Singapore Management University
Institutional Knowledge at Singapore Management University

Dissertations and Theses Collection (Open Access) Dissertations and Theses

2010

News which Moves the Market: Assessing the Impact of Published Financial News on the Stock Market

Yu Chiang Soon
Singapore Management University, ycsoon.2008@smu.edu.sg

Follow this and additional works at: http://ink.library.smu.edu.sg/etd_coll

Part of the Finance and Financial Management Commons, and the Portfolio and Security Analysis Commons

Recommended Citation
Soon, Yu Chiang. News which Moves the Market: Assessing the Impact of Published Financial News on the Stock Market. (2010). Dissertations and Theses Collection (Open Access). .
Available at: http://ink.library.smu.edu.sg/etd_coll/58

This Master Thesis is brought to you for free and open access by the Dissertations and Theses at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Dissertations and Theses Collection (Open Access) by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email libIR@smu.edu.sg.

NEWS WHICH MOVES THE MARKET:

ASSESSING THE IMPACT OF PUBLISHED FINANCIAL NEWS ON THE STOCK MARKET

SOON YU CHIANG

SINGAPORE MANAGEMENT UNIVERSITY

2010

NEWS WHICH MOVES THE MARKET:

ASSESSING THE IMPACT OF PUBLISHED FINANCIAL NEWS ON THE STOCK MARKET

SOON YU CHIANG

SUBMITTED IN PARTIAL FURFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF MASTER OF SCIENCE IN FINANCE

Supervisor: Prof. Chiraphol New Chiyachantana

SINGAPORE MANAGEMENT UNIVERSITY

2010

©2010
Soon Yu Chiang
All Rights Reserved

News Which Moves the Market:

Assessing the Impact of Published Financial News on the Stock Market

Soon Yu Chiang

ABSTRACT

Recent years have seen a large increase in the volume of financial news available to investors daily. What has traditionally been restricted to print media has now evolved to include the internet and satellite television as important media sources for financial news. With this overwhelming flow of information available to investors, the impact of financial news on market prices is at best uncertain. In this paper, a computational text-scoring methodology will be employed to uncover behavioral responses by investors to negative news.

The empirical methodology employed in this paper will consist of three parts. Firstly, through the General Inquirer (GI) content analysis software, a sentiment score is derived from daily news articles published in the Wall Street Journal. The second part will be an analysis of the sentiment time series which was obtained, where comparison will be made to existing barometers of market sentiment and market volatility. The final part of the modeling methodology which will be presented is a predictive model of market implied volatility using daily news scores as the main input.

In conclusion, it is found that high negative news scores do not necessarily predict negative abnormal returns in the S&P 500 across a 1-day to 5-day window. However, high negative news scores are highly correlated with higher market volatility. Given that the negative news is published prior to the market's trading start in the morning; we are able to utilize this information to construct a predictive model of the CBOE VIX index.

Table of Contents

1. Introduction: Investor Sentiment and News Page 1
2. Text Analytics: Potential Source for Alpha? Page 9
3. Data and Research Methodology Page 14
4. Text Analysis of Wall Street Journal Articles Page 16
5. Correlation Study on Negative News Page 19
6. Negative News and Market Returns: A Trading Strategy Study Page 21
7. Does Negative News Cause Market Volatility? Page 25
8. Introduction to the VIX Index Page 28
9. Modeling the VIX using Negative News Scores Page 30
10. Conclusions Page 34
11. Bibliography and References Page 37
12. Tables and Charts Page 40
13. Appendix A Page 58
14. Appendix B Page 60

Acknowledgments

Best wishes to all my fellow course mates, professors and thesis advisor at Singapore Management University. All your help has always been invaluable and appreciated.

My gratitude to Alexander Weaver of the General Inquirer Webserver, for server access.

Introduction: Investor Sentiment and Market Responses to News

This part of the paper introduces previous research which covers theories of investor sentiment and examines the motivation behind the research put forward in the rest of the paper. The first part will cover a theoretical aspect of investor behavior responses to news and propose the motivation of research. The second will examine past research which contribute towards the subject. The third section will cover the latest innovations in text analysis and computational scoring of text articles. The final part will present a formal overview of the methodology employed and a "road-map" that defines the paper's research.

Proposition: How Does News Affect the Market ?

Classical economic thought and traditional financial theory often require a stable equilibrium point defined by a rational utility maximising agent. Without which there can be no stable equilibrium, and therefore the risk-based pricing of securities will eventually break down. Recent studies in behavioral science have shown that investor rationality can be challenged. Investors have a tendency to be irrational more often than not. Yet in classical finance, there is no scope for investor sentiment as any mispricing in securities by irrational investors would immediately be arbitraged away, and hence there is minimal market impact.

Decisions in financial markets often involve respectable sums of money and therefore put a significant amount of an investor's financial capital at risk. The process of making transactional decisions in the stock market is often complicated given the amount of information available to each individual investor. To complicate matters further, the pace and rate of information arrival has intensified due to technological innovations such as internet based trading and news reports. Information has become much more rapid in nature to investors as compared to what was decades ago.

Behavioral studies have shown that investors are subject to an affect heuristic; especially when overwhelmed by information and required to make decisions which are inherently risky. Affect in a psychological state refers towards a state of emotion or mood, which can either be positive or negative. Epstein (1994) documents that humans often have a dual process of thought; whereby emotions play an important role. Humans tend to use affective reasoning rather than analytical reasoning when emotionally shocked.

Cognitive judgment and decision making then falters. In short, human behavior often tends to irrationality when faced with a highly negative mood or emotion which can be influenced by the news and market surroundings.

If emotions do play a part in the financial market, then which emotions can be stated as the dominant ones which affect investors? Recall that an investor's main purpose is to achieve profit. Investment profit is made through making decisions which maximise any gain, while minimising any losses made. Through this, it can be proposed that investors are fearful of large losses and at the same time greedy for gains to their portfolio. Fear and greed are very natural human characteristics and it can be postulated that these are the two main emotional states which are embedded within the stock market at any given time.

This tendency towards irrationality by investors has its implications. Irrational trading behavior implies that there will be either over reaction or under reaction in security prices. This deteriorates the pricing quality or signal as sentiment fluctuates. Investor confidence has been studied extensively over the past decade. Odean (1998) Daniel, Hirshleifer, Subrahmayman (1998) Barbara, Shleifer, Visny (1998) build theoretical models of where investors are biased to their own private signals, and underweight the public signal. These models of investor over confidence point towards decreasing price quality, increased volatility, lower trading profits and hence lower utility, as investor over confidence increases. In summary, we have over reactions in prices, which later reverse. Daniel, Hirshleifer, Subrahmayman (1998) provide another model where investor confidence is dynamic; private signals which are later confirmed by public signals are

strengthened, and vice versa, this leads to variations in investor confidence. In example, if an investor's buy signal is later confirmed by a public signal, his confidence is likely to increase.

The papers stated above try to explain the origin of over confidence; touching on representativeness, natural selection bias and many heuristic biases. However, very little has been added on the affect heuristic with regard towards investor sentiment, or perceived market sentiment through news. This paper proposes that published financial news is the main driver of time varying investor sentiment. In essence we only consider the existence of one investor class. This investor class perceives the market sentiment through the arrival of financial news. Market sentiment in this setting has an effect on market prices and hence proposes perceived risk-based pricing.

How investor sentiment affects price volatility is also of great interest to investors and portfolio managers. Spikes in volatility are often catastrophic for risk management systems. Portfolio managers employ models which rely on the conditional variance of market returns; such as Value-at-Risk and other portfolio risk management systems. Volatility spikes can disrupt the risk measures and risk prudence values that these models provide to the portfolio managers. These volatility spikes could be influenced by investor sentiment which may arise from the tone of financial news. Historically spikes in market volatility has led to some major fund blowups (Long Term Capital Management). As such, spikes in volatility would undoubtedly have some impact on the stability of the international financial system as a whole.

The main goal of this paper is to estimate any impact that published financial news might have on market returns, volatility and other market characteristics. Following that, if any relationship is found – a model of market volatility through financial news will be proposed. This could help improve risk management measures and in turn aid financial market stability. The following section attempts to summarise leading research papers that have covered on investor sentiment, reactions to news and models of market volatility.

Background: A Review of Past Research of Importance

I postulate that the stock market is largely driven by these sentiment cycles that are caused primarily (but not solely) by news. This news can be either financial news or even from the geo-political front. Inextricably the relationship between what we (as human beings) visualise through the media plays a huge part in our sentiment. Good news would tend to lift the market, while on the other hand bad news would tend to dampen the market's growth. However, this news effect is not symmetric. Good news may not lift the market as much as bad news might be reacted upon.

Previous studies on investor sentiment and news have been numerous but most studies had mainly focused on stock specific effects. Baker & Wurgler (2006) examined sentiment effects by using a time series of investor sentiment constructed from various sentiment proxies such as new initial public offerings and mutual fund flows. They find

a large effect on sentiment for stocks with speculative valuation which is significant in Fama-French valuation models. The authors conclude that a better understanding of investor sentiment is required to understand patterns in stock valuation. While the authors do not attempt to dissect exactly how investor sentiment is motivated by published financial news or other sources, their research provides a basis for recognizing the role of investor sentiment.

Engle (1990) and Engle & Ng (1993) assess that there exists a asymmetric impact on the stock market from news. Campbell and Hentschel (1992), Pagan and Schwert (1990) also document a similar skewness in market volatility. News might also have different effects on the market return depending on the overall state of the market. Good news in a bull market may be reacted on very differently from perhaps if the same piece of news had arrived during a bear market. The mechanics of market reaction towards such dynamics is complicated to understand fully and is hence difficult to model. Hamilton and Susmel (1994) introduced regime switching GARCH models which attempt to capture these effects, but the threshold point for switching is often determined in a case-by-case basis.

The models used to generate the News Impact Curve defined by Engle & Ng (1993) are generally GARCH type variants, where good news is defined to be observed if the last period's stock return is positive and vice-versa. This assumes a zero-mean market return and proposes a time series approach to volatility modeling. However this assumption that a non-negative market return implies the arrival of good news in the model can be reliant on time series properties (estimating a conditional variance). In this paper, a different

approach to quantifying good or bad news is undertaken using a computational word scoring method.

The analysis of textual data is not new in financial research. Textual analysis has become significantly prominent within the domain of finance and investment (This will be discussed more within the following section). Antweiler and Frank (2001) use a text mining approach to analyse messages posted on internet forums pertaining to stock investment (Yahoo Finance and Ragingbull forums). The authors in their paper use a classification methodology using a Naive-Bayesian classifier which learns from a training set to classify between good and bad news, depending on the textual content of specific posts. This classification system is similar to that used by spam email filters utilised by various email service providers. The main results found by Antweiler and Frank suggest that there is a significant relation between message volume (how active are the forums on a particular day) and market volatility. The authors also find that the degree of bullishness implied by the message boards have a direct effect on trading volume. This was perhaps one of the first research papers which contribute towards textual analysis of financial news. The methodology employed however does have its flaws, as Naive-Bayesian classification often relies on human judgment during the training process, whereby the labeling of news in the "example set" required for machine-learning is subject to human bias. More recently another dictionary predetermined "bag of words" approach was employed to decompose financial news articles.

Tetlock (2007) uses the Harvard-IV dictionary and General Inquirer[1] content analysis system, to examine and score news articles from the Wall Street Journal's "Abreast of the Market" column. The column is published daily and generally is a summary of the previous day's market performance. This methodology is more appealing than that used by Antweiler and Frank as it removes the human bias towards constructing the "bag of words" that is required for analysis. The Harvard-IV dictionary presents predetermined words which define negativity, weakness and other emotional themes (The following section will explain further on the dictionary methodology); thus removing any need to construct a new dictionary which may be subject to bias. The author finds that high media pessimism[2] predicts downward pressure on market prices and a subsequent reversal which implies that the market overreacts to negative news. According to the paper, a one standard deviation increase in the bad news measure predicts a 3.7 basis point decline in the Dow Jones Industrial Average the following day, and this is significant to the 1% level. In the trading window 2 to 5 days after the news event, a subsequent reversal of 5.6 basis points is observed over the same event. The author subsequently finds that the same media pessimism measure used weakly predicts higher market volatility at times of high pessimism, and also slightly higher market volume. The articles used for text analysis by the author are not recent (1984-1999); and often reflect the market's performance on the previous day. Arguably, the results may reflect trend behavior in market returns, where the next day abnormal returns may be a result of trend continuation or momentum. Borrowing on the author's main methodology, I aim to

[1]An overview of the General Inquirer text scoring system is presented in Appendix A.

[2]Tetlock uses the term frequency of Negative words defined by the Harvard-IV dictionary to proxy for media pessimism. A sample of the words is presented in Appendix B.

correct for this by using a combination of articles in the Wall Street Journal and use a data sample that reflects on more recent times.

In this section, a summary of past important research on investor sentiment, models of volatility and textual financial news analysis has been presented. Table 1(a) and 1(b) below summarizes the research papers discussed and their relevant contributions to our knowledge. The following section will expand further on the growing importance of textual analysis in finance and a brief overview of other research papers which explore the subject.

(Insert Table 1 about Here)

Text Analytics : A Potential Source for Alpha ?

The analysis of financial news through machine learning or dictionary bases methodologies is a upcoming concept. Outperforming the market requires constant innovation of available data by investors. News is a potential new data source that is not fully exploited as yet. Proprietary trading desks and hedge funds now use news based trading algorithms to add an edge to their investment strategies. Many financial data service providers such as Thomson Reuters, Bloomberg and Dow Jones are now providing text analytic algorithms to read, score or rate financial news, with a view to

provide sentiment signals to investors. A column from the Financial Times (January 28th 2010) states the following on the rise of text analytics;

"The arms race in trading technology is set to intensify this week as Thomson Reuters, the news and market data company, on Monday unveils a service for "high-frequency" traders allowing them to make split-second trading decisions based on news articles "before the information moves the market" . . .

So-called "machine readable news" services, such as the new Thomson Reuters product, have grown up in parallel with the emergence of high-frequency and algorithmic trading, which depend on lightning-fast delivery of data and news to traders specialising in such computer-driven trading strategies.

NewsScope Analytics by Reuters measures the sentiment, or tone, of the article. A very positively or very negatively-toned article might suggest a bias in the price to the upside- or downside, or to help predict a spike in volatility, which can then be incorporated into your strategies."

The scoring or determination of market sentiment from news articles is difficult to quantify compared to technical charts and other traditional market signals. The complexity involved poses several challenges to the news analyst. The first challenge relates to the massive amount of qualitative news articles and events available. The

existence of many news sources and articles can be overwhelming to analyse. The news analyst must first decide which news is relevant, and which news source should be used for analysis. I propose that consistency, reliability and choosing only reputable news sources should be paramount. This would ensure unbiased analysis.

The second challenge posed relates to the methodology to be used by the analyst to decipher the news content into quantifiable variables which are easy to interpret into a model of sentiment, volatility or returns. Natural language processing or "text mining" has advanced to include many techniques capable of analysing and classification of textual data. The techniques relevant to financial news can be broadly be divided into two main categories. The first relies on human supervised machine learning to train classification models from training sets (sorted out usually by human interaction), and henceforth capable of extracting sentiment values categorically from new articles fed into the model. An example of this methodology was used by Antweiler and Frank (2001), where a naïve Bayesian classifier was used to sort news into Positive, Neutral and Negative categories through conditional probability of word occurrences within the articles. Classification techniques create a dictionary of words based on human interaction and are often difficult to train.

The second technique uses a predetermined dictionary of words which gives scores based on the term frequency[3] of the words categories. Each individual observation or news

[3]Given a group of words W = $, w_2, w_3, \ldots, w_n)$, which belong to category A. The term frequency of category A is given by $= \frac{\sum c_A}{\sum c_{A,B,C,\ldots,Z}} = \frac{\sum w_{1,2,3,\ldots,n}}{\sum c_{A,B,C,\ldots,Z}}$. The term frequency is the sum of all words that define a category, divided by the total number of words that define all categories which appear in the document.

event is scored based on several categories which can be used to define a sentiment score. The work by Tetlock (2007, 2008) uses this methodology and presents a simple count of Negative words to define market sentiment. The categories defined in his study are through the Harvard-IV dictionary (which was designed by psychologists), relate to positive, weak, negative or strong words. This methodology hence does not require a supervised initial training set; as the dictionary has already been provided. This eases the analytical process and requires far less news data to provide a very consistent sentiment score. A potential downside to this form of analysis is that the analyst relying on a static "bag of words" which may not reflect time trends in linguistic content of the news.

In determining which methodology which should be employed, the analyst should again aim for consistency and reliability depending on his available resources. If large amounts of presorted news data[4] and ample computational power is available then a classification methodology may be suitable. However if the resources are unavailable to the analyst, a simpler approach must be employed to ensure consistency in the results. Table 2 below documents the differences between the two approaches available to the news analyst:

(Insert Table 2 about Here)

[4]This sort of data may be available to analysts at large financial services institutions. Brokerages and news vendors have access to datasets of news articles which contain end user opinions on each specific article. The end users are often professional fund managers and investors.

Data and Research Methodology

In this paper, news articles collected from the daily Wall Street Journal column "Ahead of the Tape" are used for analysis. I have used the Wall Street Journal as the main source for the news articles used to construct my sentiment proxy of the market. The Wall StreetJournal has an estimated distribution 2.1 million copies[5] which represents the highest distributed newspaper in the United States.

The "Ahead of the Tape" column was started on 9th April 2002 and provides daily opinions and expert commentary on the trading day ahead. The Wall Street Journal's description of the column is as follows;

"Located on the front of the Money & Investing section, Ahead of the Tape is like having breakfast with the Journal's financial editors: It explains in plain language what to watch for today that may move the markets. We expect it will become essential reading for professional investors -- and be even more helpful for the rest of us. It has become an instant touchstone for market professionals and interested amateurs, highlighting crucial opportunities and pitfalls before the opening of each market."

[5]Estimate provided by the Audit Bureau of Circulations as of October 2009.

Recall that the first challenge to providing a well-structured analysis of textual data is to choose a reliable, consistent and reputable source for news. The Wall Street Journal's "Ahead of the Tape" column fulfills that criteria and is relevant, given that it is read by professional investors on a daily basis. The column is published in the morning and is received by investors before each trading day commences. The column thus provides suitable qualitative content that is required for textual analysis, and subsequently sentiment extraction. The news articles were compiled from Dow Jones Factiva, with a time range starting from 9th April 2002 to 31st December 2009.

In selecting a suitable methodology to analyse the news articles collected; I use the General Inquirer[6] content analysis software, aided by the Harvard-IV dictionary to score the daily articles and provide a sentiment score daily. This sentiment score is represented by the Negative[7] category which consists of 2291 words that convey negativity (weakness, lack of confidence, suffering, ect). The analysis provided by the General Inquirer is easy to understand, and can be interpreted without having to construct a dictionary of words through classification. This should provide a consistent sentiment reading of the market without any computational or human biases in classification. This methodology has also been previously used by Tetlock (2007, 2008) in his research.

A graphical evaluation and correlation study will then be done between the negative scores, risk-adjusted market returns and various previous proxies for sentiment and market volatility. The proxies for sentiment and market volatility will include absolute

[6]A web-version of the General Inquirer is available at http://www.webuse.umd.edu. In this paper, the server based General Inquirer was used (with courtesy from the Stone Center EU).

[7]An example of some of the words contained in the Negative Category can be found in Appendix C.

market returns from the S&P 500, daily trading range from the S&P 500 index, S&P 500 trading range, CBOE VIX daily closing value and the S&P 500 daily trading volume. The motive here is to determine if the Negative scores obtained are suitable to explain either risk-adjusted market returns, market trading volume or market volatility. This first initial step will provide a rough guide to determining if any relationship exists between the studied variables and is the basis for any further predictive analysis in this paper.

Predictive modeling in this paper will follow a standard cross-sectional regression approach; where market volatility, returns and trading volume will be the main forecast variables. The regression methodology will be explained further in section 2 of this paper, where a regression specification is used to determine market volatility reaction to negative news. Graph 1 below presents the road-map or overview of the research methodology in this paper:

(Insert Graph 1 about Here)

Textual Analysis of Daily Wall Street Journal Articles

In this second part of the paper, I detail the results from the textual analysis on the Wall Street Journal articles collected. The section will be divided mainly into three sections. The first section will provide the summary statistics of the text data analysis. The second section will provide a correlation study between, our negative news score, market returns, market volatility and trading volume. In the final section, an interpretation of the statistics acquired will be presented. This final section will also provide the basis for further statistical modeling, and attempt to propose a linkage between market volatility and the negative tone of financial news.

Summary Statistics of "Ahead of the Tape" Articles from the Wall Street Journal

In our initial sample of the news articles merged to daily trading days, a total of 1645 observations were obtained between 9th April 2002 and 31st December 2008. This provides a sufficient amount of observations to analyse the effect of daily news on market characteristics. The news articles from 1st January 2009 - 31st December 2009 are left out for an out-of-sample evaluation in the third part of this paper. The summary statistics are summarised in Table 3 below:

(Insert Table 3 about Here)

In the analysis provided in this paper, two versions of the negative news series will be considered. The summary statistics reveal high deviation or high daily volatility within our negative news data series. The logged negative scores (*lognegative*) exhibit a high standard deviation of 0.42 across its mean of 1.22. This could be due to the nature of the daily fluctuations in linguistic content of news articles captured by the General Inquirer program.

The negative scores thus will be smoothed to a 30-day simple moving average[8], this is to ensure that the series captures trend movements rather than daily spikes in negativity, which are noisy and could be misleading to analysis. This smoothed series is denoted by *lognegative30* in Table 3 above. The variables can be understood further through yearly subsets and box-whisker plots. Table 4 below provides yearly analysis on the distributions of our factor variables, supplemented by Graph 2 and Graph 3 which provide graphical analysis.

(Insert Table 4 , Graph 2, Graph 3 about here)

The two graphs and statistics in Table 4 above show that the distribution is not highly varied across our whole sample for our negative news scores. The *lognegative*and *lognegative30* measure's distribution remains stable across the seven year sample period. The non-smoothed*lognegative* score however does exhibit more extreme observations; which can either be corrected through removal or using the smoothed series for further analysis.

[8] It is defined as $gative_{30DAYSMA} = \frac{(Negative_t + Negative_{t-1} + \ldots + Negative_{t-29})}{30}$

Correlation Study between Negative News Score and Market Variables

This section will study the correlation between the studied variables and provides a visual representation into any relationships between our negative news score and market variables. As with most financial data series; calculated correlation coefficients are often time varying. The sample will be split to report yearly subset correlations and that of the full sample as well. This will present a detailed and concise background to any predictive modeling that will be performed. Table 5 below reports the Pearson correlation coefficients between the *lognegative* and *lognegative30* measures and the market variables.

(Insert Table 5 about Here)

The correlation analysis from the table above depicts the yearly Pearson correlation coefficient between the negative news score and the market variables which are of interest in this study. The first result that can be seen is that there is no direct or conclusive evidence of any correlation between the negative news score and S&P 500 market returns (Full sample correlation of -0.01 between the *lognegative30* score and S&P 500 market returns). This suggests that our negative news score cannot be used to predict the market accurately. However there seems to be a significant correlation

between the *lognegative30* score and the other market variables such as S&P 500 volume *(0.5064)*, absolute market returns *(0.2956)*, trading range *(0.4563)* and the daily closing value of the CBOE VIX index *(0.5941)*. The negative news score appears to be highly positively correlated with the indicators of market volatility, and this suggests that the negative news appear to increase the uncertainty of market returns; and also to increase the range of market returns. The *lognegative30* score represents the negative news compiled from the last 30-days; and thus could be a strong measure of market sentiment given its effect on market volatility. The correlation from the non-smoothed negative news series lognegative is also significantly positively correlated with the market volatility indicators and S&P daily trading volume. However the correlation is not as strong as that of the smoothed series; and this may be due to lagged reaction to news or other factors. Graph 4 below presents a scatter correlation graph between all the market variables and our negative news scores.

(Insert Graph 4 about here)

Through graphical analysis the correlation between the *lognegative30* score and the daily closing value of the CBOE VIX index seems to be validated. The market returns appear to increase in volatility and range as the *lognegative30* score increases. The increase in market volatility is not highly skewed towards negative or positive returns. There are several uncertainties that still remain; is it market volatility, trends and stock market

returns that drive the negativity of published news; or is it the negative news which drives up market volatility? Another question worth asking is if the negative news scores are reacted to preemptively prior to publication, immediately after publication or is there a lagged reaction from the market to the negative news scores that were derived?

Does Negative News Score Have Economic Value? A Study of Market Returns and a Simple Trading Strategy Analysis

To obtain further understanding of this; a percentile analysis of the negative news score and subsequent market reaction is presented in Table 6 below.

(Insert Table 6 about here)

The analysis presented in Table 6 suggests that the negative news score does have a significant effect on the risk-adjusted market returns over a 5-day window (the cumulative market returns across 5-days). The market returns do not seem to have a lagged reaction to negative news clustered on a single day in the 5-day window. Instead the negative news effect appears to be spread out across all 5-days. The total market return for negative news above the 75% percentile averages at *-0.5%*; while the total market return for negative news below the 25% percentile averages at *+0.2%*. Through

using a market direction measure[9], it is also observable that the S&P 500 index is more likely to regress across a 5-day period, when the negative news score is above the 75% percentile. Conversely the S&P 500 index is more likely to appreciate if the negative news score is below the 25% percentile.

The results for the market direction and returns however have a high variation and therefore are unsuitable for any trading strategy or prediction modeling. The standard deviations of returns and market direction seem to increase across negative news percentiles. The total market return above the 75% percentile exhibiting a *0.4%* standard deviation; while in comparison the total market return below the 25% percentile only exhibits a *0.1%* standard deviation. This increase in market volatility is consistent with our correlation study stated in Table 5. The logged daily closing value of the CBOE VIX index also seems to increase in both volatility and value across percentiles. Below the 25% percentile; the logged daily VIX closing value averages at *2.644* with a standard deviation of *0.209*. Comparatively when the negative news score is above its 75% percentile; the logged daily VIX closing value averages at *3.300*, with a standard deviation of *0.373*.

Using Tetlock's (2007) methodology, the negative news score can be examined to determine if there is any predictability or forecast value towards the S&P 500 index returns. Tetlock uses a Vector Autoregression methodology, with the following specification:

[9]A score of 1 is assigned to a day where the S&P 500 index experiences a gain; while a score of -1 is assigned to a day where the S&P 500 index experiences a loss. The market direction measure is the 5-day average of these scores; and captures the directional trend of the stock index.

$$\text{Mktrf} = \alpha_1 + \beta_1 \,.\; \text{L5 (Mktrf)} + \gamma_1 . \text{L5 (Log.negative30)} + \delta_1 \,.\, \text{L5 (Log.SPvolume)} + \varepsilon_{1t}$$

Where L5 are lag operators of the vector coefficients (β, δ, γ). Table 7(a) below details the results:

(Insert Table 7(a) about here)

It can be seen that contrary to Tetlock's initial result in his paper, there seems to be no direct forecast value of the negative news score on S&P 500 market returns. All the coefficients γ reported in Table 7(a) are not significant and this rules out the possibility that our negative news score can provide information pertaining to the trading day ahead. This result may be driven by the updated sample and differences in article choice from Tetlock's initial study.

To provide some further analysis on the economic significance of the negative news score, a simple buy and hold trading strategy will be tested between the S&P 500 E-Mini futures returns and the negative news score (as an event trigger). The trading strategy is defined as such:

If the lognegative30 score is above its 75% percentile, this will be defined as a negative news event, and correspondingly the e-Mini futures will be shorted at opening price, and sold at the closing price after t days. The returns are then calculated.

If the lognegative30 score is below its 25% percentile, this will be defined as a positive news event, and correspondingly the e-Mini futures will be longed at opening price, and sold at the closing price after t days. The returns are then calculated.

(Insert Table 7(b)about here)

The results in Table 7(b) show the two trading portfolio returns on the e-Mini futures. It can be seen that both the Long and Short portfolio exhibit positive returns both the 5-day and 15-day holding period. The average return per trade is positive for both portfolios; but only the 5-day holding period short portfolio average return is significant; at 0.75% average returns per trade. This dampens any possibility that the negative news series can be used to trade successfully in the futures market. However, again the standard deviations of returns within the Short portfolio (where negative news is high) is significantly higher than if compared to the Long portfolio (where negative news is low); by about 250 basis points.

While the negative news score that was obtained may not be suitable for predictive modeling and strategic trading; what can be observed is that negative news increases the spectrum of returns in both the positive and negative territory. This strengthens the case for the negative news score being a suitable predictor of market volatility (as it is published prior to market trading), a barometer of market sentiment, and a risk management tool for portfolio managers.

Does Negative News Cause Market Volatility?

One question remains on the causal relationship between market volatility and the negative news score. Table 8 below presents a Granger Causality analysis between market volatility indicators and the negative news score.

(Insert Table 8about Here)

The Granger Causality test between the market variables and the negative news score was performed with lags of 1 and 5 respectively (denoting daily and weekly lags), and in both directions of causality. The results in Table 8 are mixed, but there is some evidence of causality (within 1 lag) between market volatility and the negative news score. The test statistics indicate a high degree of causality arising from the negative news score, towards the CBOE VIX daily closing value, S&P 500 trading range and S&P 500 trading volume. However, there is no result or evidence of causality between S&P 500 market returns and the negative news score. There is however a caveat to this methodology which should be documented. Granger causality tests within this context are usually difficult to evaluate and interpretation of the test statistics can vary across lags. Therefore, causality tests on financial time series are often sample specific; and must be supplemented with further analysis and economic interpretation.

In this section, a basic statistical study has been concluded to examine the properties of the negative news series and its distribution effects on market returns, volatility and trading volume. The main results that can be concluded is that market returns are often not predictable from our negative news series alone. The negative news series however does seem to have an effect on market volatility and based on this initial finding; a model of market volatility can be proposed.

Modeling Market Implied Volatility through Textual News Data

This next section will first provide a background to the CBOE VIX index and how it relates to market sentiment. A literature review of past research focusing on predictive models of the CBOE VIX index will then be outlined. Following which, a predictive model of the CBOE VIX index utilizing our negative news score will be proposed. Interpretation of the results of the Vector Autoregressive Model (VAR) will then conclude this section.

The CBOE VIX Index: Background and Past Research

The Chicago Board Options Exchange (CBOE) first introduced its volatility index (VIX) in 1993, where its value is calculated through a weighted blend of S&P index options (eight at-the-money call and put options for 30 days ahead). The VIX index value hence represents the market expectation of 30 day volatility of the S&P 500 index, and is often quoted as the implied volatility of the S&P 500. When first implemented in 1993, the VIX utilized at-the-money calls and puts for the S&P 100 index, and this was updated in 2003 to use options on the S&P 500 index for a broader market implied volatility interpretation. Graph 5 below shows the CBOE VIX since its implementation in 1993, and the VIX graphed with the Negative News Score since 2002:

(Insert Graph 5 about here)

An interesting characteristic of the CBOE VIX index is that its return series has a very strong negative correlation with the S&P 500 market returns series. In the days where there are upward spikes in the CBOE VIX index, the S&P 500 market index tends to decline with it. This has led to popular reference to the VIX as the so-called 'Fear Gauge' of the market, or a leading barometer of market sentiment and portfolio insurance by fund managers. However, high VIX values do not necessarily mean a purely 'bearish' outlook

for stocks, but does imply a likely large movement in either direction (sharp reversals to a bearish trend). On the flipside when the VIX is at a low value, investors perceive relatively low downside or upside risk to the S&P 500 market index. This property can be seen to be fairly similar to how the negative news score relates to S&P 500 market returns, and forms the basis of this section.

Past research on modeling the CBOE VIX index has been fairly limited to time-series approaches utilizing lagged values of the VIX index and other exogenous variables (market returns, volatility forecasts). Most previous research on the subject had focused on using the logged closing VIX value as the main forecast target.

Engle & Gallo (2003) in their paper first introduce a multiplicative error model, which is used to provide intraday volatility forecasts (trading range, absolute returns and realized volatility). They then proceed to introduce a Vector Autoregressive (VAR1) model with the intraday volatility forecasts as exogenous variables.In a more recent paper, Ahoneimi (2006) used a moving average time-series approach to model the VIX index. Using an ARIMA model with S&P 500 market returns as an exogenous variable, the author obtains an accurate forecast model of the VIX index up to 3-days ahead. Ahoneimi (2006) did not obtain any significance for GARCH terms which were added to the model.

Modeling the CBOE VIX Index using Negative News Score

In this paper, the VIX index will be modeled through a Vector Autoregression (VAR) methodology, and using 6 different model specifications for robustness. Similarly to Engle & Gallo (2003), the logged VIX closing value will be the main forecast target in all model specifications. Table 9 below depicts the modeling process and equations:

(Insert Table 9 about here)

The modeling process will begin using a basic 1-factor model of the daily negative news score (*lognegative30*) (Model 1). Lagged values of market volatility, trading volume and market returns will then be added to the model to control for past market conditions (*sprange*, *mktrf*, *spvolume*) (Models 2, 3). An AR(1) correction term on *logvixclose* is added to control for autocorrelation and persistence of market volatility (Model 4). Model 5is similar to Model 4 but includes an AR (2) correction term. Model 6 is a first difference model containing all the variables in Model 4, included for completeness and comparison reference.

It should be noted that no lags of the negative news score series have been used in the modeling methodology. This is as the news is published prior to the trading day commencing, and the *lognegative* score hence being representative of the news implied sentiment of the current day. Adding lags of *lognegative* do not improve the forecast in

the extended models and do not have an effect on the significance of the other variables. They have been omitted as a result of that. The results from all 6 model specifications are detailed below in Table 10:

(Insert Table 10 about here)

From Table 10, we can observe that the *lognegative30* score variable is positively significant in all model specifications. This would imply that an instance where negative news or news pessimism is higher would result in a higher closing VIX index on that day. This fits perfectly into the theoretical construct poised in the introduction of the paper. Models 4 & 5 are the best fitting models (R-squared of 0.97) and even with an AR (t) correction; the negative news score still seems to have some forecast value towards the VIX daily closing value. The t-statistics of the AR terms in Models 4 & 5 are large. This is mainly due to the high persistence of the VIX index and is a common characteristic (Engle 2003) when modeling volatility, and the VIX index.

Past values of the other market characteristic variables (*spvolume and sprange*) do not appear to have any effect on the daily closing VIX value once autocorrelation is controlled for. Past market returns *mktrf* are negatively significant in all models except Model 5, which includes an AR (2) term. This implies that there is a degree of persistency in market volatility given a large negative market movement (possible reversals to market returns).

While past market returns appear to affect the daily VIX closing value; this effect is subsumed when we account for past volatility. This is similar to what was found in Ahoneimi (2006) where the MSCI EAFE and S&P 500 returns did not improve forecasts once an ARIMA model of order 2 was used. The negative news score in this paper however remains robust to this effect, and is independent of past volatility in its forecast value.

Model 6 is a first difference model included for comparative reference. The *lognegative30* score remains positively significant at a 10% level in this model. This provides supportive evidence that the negative news score is robust across various model specifications, in providing forecast value to market implied volatility

.

The negative news score or media pessimism can be concluded as such to have an effect on market volatility. An explanation as to why this occurs could be traced back to the behavioral reasoning of investors. Overreaction to news is often the case; whereby investors are overwhelmed by negativity embedded within financial press, and act beyond analytical reasoning in their transactions. Oversold or overbought positions in securities occur, and subsequently price reversals occur to correct for the initial mistake. Debondt and Thaler (1990) in their statistic examination of portfolios also documents an overreaction by security analysts to visible stock returns. This overreaction behavioral characteristic leads to heightened market volatility surrounding such negative news events, and affirms that linguistic content in news is an influential factor in the stock market.

This section has provided a detailed time-series analysis of the CBOE VIX index and its relationship with the negative news score which proxies media pessimism. In the final section, a conclusion to the paper will be presented, which will focus on reviewing the results found, highlighting possible usage of textual analysis in finance and finally identifying possible future research in this topic.

Conclusion: News Which Moves the Market – Assessing the Impact of Published Financial News on the Stock Market

Through analysis of over 2000 news articles published in the Wall Street Journal, this paper has attempted to uncover the market reactions to linguistic content of financial news. The results obtained through the statistical studies seem to support the notion that investors are subject to an affect heuristic and are prone to overreact to security prices based on the content of financial news. The main findings of the paper can be summarized as below:

1) Stock market prices (S&P 500) exhibit an average positive (negative) trend when investors are subject to low (high) negative language tone in published financial news. However, this may not be profitable to trade on; as the directional accuracy of trading through this system is low.

2) Stock market prices exhibit low (high) volatility when investors are subject to low (high) negative language tone in published financial news. This characteristic is consistent with other research that has been done previously Tetlock (2007) Antweiler Frank (2001).

3) Through the negative news series; a predictive model of the CBOE VIX index was proposed. The negative news score variable has a positive significant

relationship (forecast value); towards the daily closing value of the CBOE VIX index. This is result is robust to model selection, inclusion of other market variables and autocorrelation in volatility.

What are the implications and significance of the results?

Given the volatility prediction properties of the negative news series, it is possible to create a risk-management measure or "early warning system" of an imminent extreme market movement through fine-tuning our negative news measure. This of course requires more research into the subject, methodology improvements and significantly more test data requirements. If such a system is able to be modeled with a good degree of accuracy, portfolio risk management and general financial sector stability can be improved.

The negative news score is arrived at through computational analysis of Wall Street Journal articles, and hence can be replicated for any newspaper or text articles from any region/country/market. This leads towards the possibility of creating a sentiment index that relates to a certain exchange, which mimics the CBOE VIX index. This is intuitive, as the negative news series can forecast the VIX to a certain degree of accuracy. This could have economic value in countries where there are no options exchange security markets, in which an option implied volatility index can be constructed.

In conclusion, text analysis of financial market news is a growing subject and is a possible source of alpha and a powerful risk measurement tool if calibrated accurately to market variables. Future research would require a more flexible approach to text-scoring news articles, using a combination of machine learning and a predefined "financial news dictionary" which will cater solely to the field of finance.

As financial information technology will always advance to distribute more sources of media and information to investors, research too has to innovate and understand how new information sources can affect investor reaction. Often, an investor searches for the fundamentals within media sources which he has access, to formulate investment ideas and get confidence for his trades. Without perceived sentiment from media sources, financial markets would perhaps only exhibit a fraction of the volatility that we see today on a daily basis. Understanding how media factors such as language play a part in financial markets is a new dimension in research, and will only become more important as time progresses.

-THANK YOU-

Bibliography and References

Ahoniemi, Katja 2008, Modeling and Forecasting the VIX index, Working Paper, Aalto University School of Economics; Helsinki Center of Economic Research (HECER).

Antweiler Werner, and Murray Z. Frank, 2004, Is all that talk just noise? The information content of Internet stock message boards, Journal of Finance 59, 1259-93.

Antweiler Werner, and Murray Z. Frank, 2006, Do U.S. stock markets typically overreact to corporate news stories? Working paper, University of British Columbia.

Baker Malcolm, and Jeffrey Wurgler, 2005, Investor sentiment and the cross-section of stock returns, Journal of Finance 61, 1645-80.

Bjerring James, Lakonishok Josef and Vermaelen, Theo, 1983, Stock Prices and Financial Analyst Recommendations, Journal of Finance 18, 129-73.

Boyd John, Hu Jian and Jagannathan Rav,i 2005, The Stock Market's Reaction to Unemployment News: Why Bad News Is Usually Good for Stocks, Journal of Finance 60, 649-72.

Campbell John and Hentschel L 1992, No news is good news, open: An asymmetric model of changing volatility in stock returns, Journal of Financial Economics 31, 281-318.

Chan, L. and N. Jegadeesh, et al. 1996.Momentum Strategies. The Journal of Finance 51, 1681-1713.

Chan and Wesley S, 2003, Stock price reaction to news and no-news: drift and reversal after headlines, Journal of Financial Economics 70, 223-260.

Cutler David, Poterba James and Summers Lawrence, 1989, What moves stock prices? Journal of Portfolio Management 15, 4-12.

Daniel Kent, Hirshleifer David and Subrahmanyam Avanidhar, 1998, Investor Psychology and Security Market Under- and Overreactions, Journal of Finance 53, 1839-85.

Das, Sanjiv, and Mike Chen, 2006, Yahoo! for Amazon: sentiment extraction from small talk on the web, Working paper, Santa Clara University.

De Bondt, Werner F.M., and Thaler Richard, 1990, Do security analysts overreact? American Economic Review 80, 52-57.

Engle Robert and Ng Victor, 1993, Testing and Measuring the Impact of News on Volatility, Journal of Finance 36, 1749-87.

Engle Robert and Gallo Giampiero, 2006, A multiple indicators model for volatility using intra-daily data, Journal of Econometrics 131, 3-27.

Epstein Seymore, 1994 , Integration of the cognitive and the psychodynamic unconscious, American Psychologist 49, 709-724.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, Journal of Financial Economics 33, 3-56.

Granger, Clive W.J., 1969, Investigating causal relations by econometric models and cross spectral methods, Econometrica 37, 424-38.

Grossman, Sanford J., and Joseph E. Stiglitz, 1980, On the impossibility of informationally efficient markets, American Economic Review 70, 393-408.

Hamilton James and Susmel Raul, 1994, Autoregressive conditional heteroskedasticity and changes in regime, Journal of Econometrics 64, 307-33.

Kim Oliver and Verrechia Robert 1991, “Trading Volume and Price Reactions to Public Announcements” Journal of Accounting Research 29. Issue 2, 301-329.

Mittermayer, M. 2004. Forecasting Intraday Stock Price Trends with Text Mining Techniques.Hawaii International Conference on System Sciences, Kailua-Kona, HI.

Mittermayer, Knolmayer 2006, "Text Mining Systems for Market Response to News: A Survey", Working Paper, The University of Bern.

Odean Terrence 1998, Volume, Volatility, Price, and Profit When All Traders Are Above Average, Journal of Finance 53, 1887-1934.

Pagan Adrian and Schwert William 1990, Alternative Models for Conditional Stock Volatility, NBER Working Paper Series.

Riffe, Daniel, Stephen Lacy, and Frederick G. Fico, 1998, Analyzing media messages: using quantitative content analysis in research (Lawrence Erlbaum Associates: London).

Robert P. Schumaker and Hsinchun Chen, 2007.A Quantitative Stock Prediction System based on Financial News. Working Paper.The University of Arizona.

Tetlock, Paul., 2007, Giving content to investor sentiment: the role of media in the stock market, Journal of Finance 62, 1139-1168.

Tetlock Paul, Saar-Tsechansky.M, and Macskassy.M 2008, "More than Words: Quantifying Language to Measure Firm Fundamentals", Journal of Finance, Forthcoming

Wall Street Journal, 9 April 2002 – 31 December 2008, "Ahead of the Tape" Articles, (Dow Jones Company, NY, New York).

Whaley, Robert E., 2000, The investor fear gauge, Journal of Portfolio Management 26, 12-17.

Wüthrich, et al 1998, Daily Prediction of Major Stock Indices from Textual WWW Data, Proceedings 4th ACM SIGKDD Int. Conference on Knowledge Discovery and Data Mining. New York 1998, S. 364-368.

Table 1(a): Summary of Influential Research Papers on Structured News

Authors	Studies	Focus Area	Results
Bjerring, Lakonishok, Vermaelen (1980)	Stock price reaction to financial analyst recommendations	Analyst recommendations	Positive abnormal return to recommendations
Goh, Edenington (1993)	Stock price reaction to bond rating downgrades	Bond rating downgrades	Mixed results. Most downgrades are followed by negative stock reaction. Has scope for alternative reaction due to reason of downgrade
Chan (2003)	Stock price reaction to degree of media coverage	Media coverage of firms	Stocks covered by media tend to underreact, while stock movements fueled by private information tend to reverse
Boyd, Hu, Jagannathan (2005)	Stock market reaction to economic news	Economic news	Mixed results. Suggests bad employment news during economic expansion has good impications to stocks

Table 1(b): Summary of Influential Research Papers on Sentiment and News

Research Paper	Studies	Methodology	Conclusion
Baker &Wurgler (2006)	Investor sentiment and market returns. Sentiment proxy used is a mixture of various measures.	Sentiment Index Time Series, Fama French Valuation Models	Sentiment effect on speculative stocks, unable to be explained by Fama French Model. States a role for investor sentiment in stock market patterns.
Engle (1990), Engle & Ng (1993), Campbell &Hentschel (1992)	Market volatility, volatility clustering, time-series approaches.	GARCH & GARCH variants. Past market return represents news inflow to market. Introduces the News Impact Curve.	Asymmetric distribution of volatility response to good or bad news.
Antweiler& Frank (2001)	Stock market response to textual data in internet investor forums.	Text classification and event studies on abnormal returns, volatility and volume.	Not conclusive on stock returns. Finds a relationship between market volatility and internet message volume. Weak relationship between degree of bullishness and market trading volume.
Tetlock (2007)	Stock market response to Wall Street Journal column	Text scoring via General Inquirer. Time series and VAR approach to response variables, Fama French Valuations	Media pessimism or negative news predicts downward market movement. Weak relationship to trading volume and volatility.

Table 2: Sentiment Analysis via Textual Analysis – Two Main Approaches

	Text Classification Methodology	**Predetermined Dictionary Methodology**
Text Data Required for Modelling	Large amount. Training data required to construct dictionary of linguistic features	Small amount, No training data required
Preprocessing	Textual data must be simplified using existing dictionary stemming algorithms	No simplification required, only requires removal of non-identified text
Model Output	Classifies documents to sentiment categories; many documents needed to define a daily sentiment score.	Provides a simple negative or positive score to a single document.
Suitability	Analysts with large data availability and computational power	Small sentiment or survey analysts with limited data and computational power

Graph I : Research Road-Map

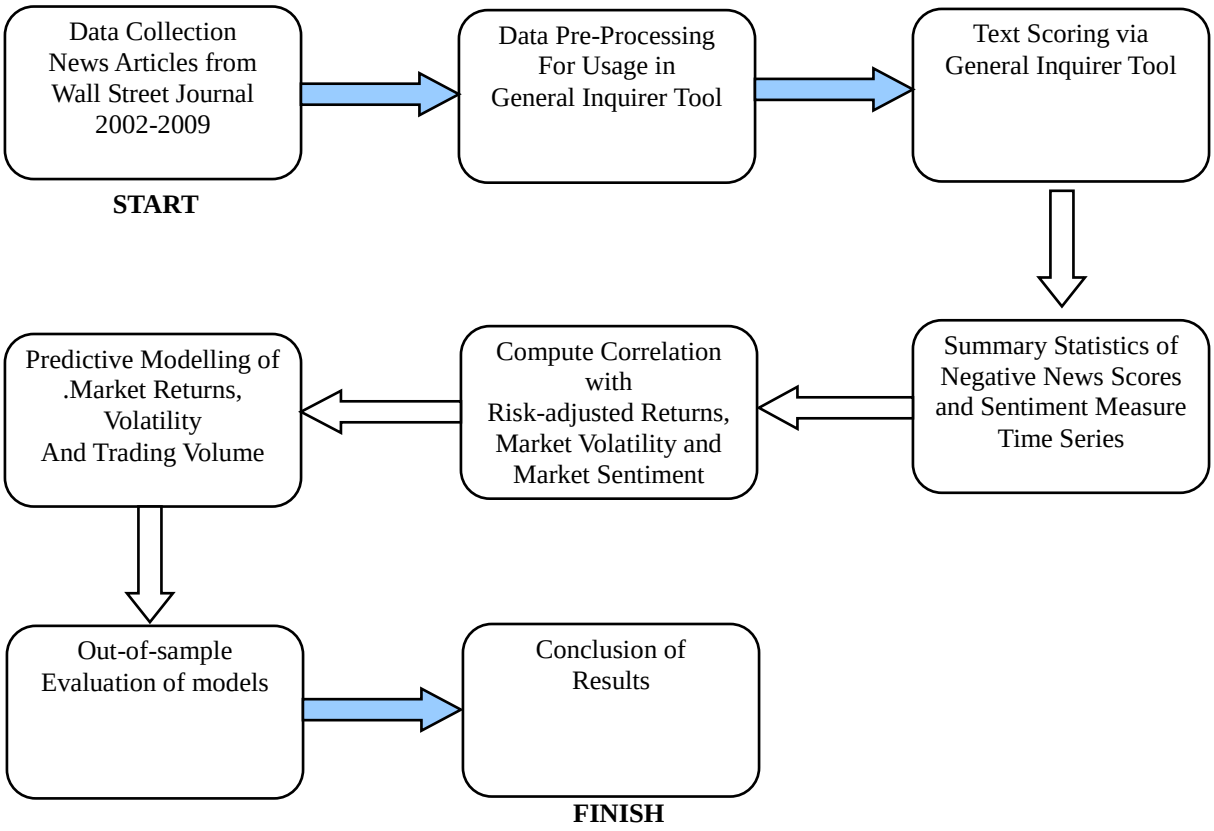

Table 3: Summary Statistics of Negative News Data Series, S&P 500 and CBOE VIX

(9th April 2002 - 31st December 2008)

The table data comes from CRSP, NYSE and textual numerical data compiled from the General Inquirer program using news articles from the Wall Street Journal.

Mktrf is the risk adjusted market return for the S&P 500 index, logspvolume is the logged daily trading volume of the S&P 500, sprange is the daily trading range of the S&P 500, logvixclose is the logged daily closing value of the CBOE VIX index, lognegative is the logged negative news score derived from textual analysis of Wall Street Journal articles, lognegative30 is the logged 30-day simple moving average of the negative news score.

Variable	N	Mean	StdDev	Minimum	Maximum	Median
lognegative	1645	1.217	0.428	-1.386	2.297	1.261
lognegative30	1616	1.289	0.141	0.867	1.595	1.280
logvixclose	1645	2.901	0.411	2.291	4.392	2.820
mktrf	1645	-0.0001	0.013	-0.090	0.115	0.001
sprange	1645	16.313	11.881	2.900	102.650	12.98
logspvolume	1645	21.493	0.495	19.690	23.161	21.416

Table 4: Means and Standard Deviations of All Variables Sorted by Year

		VARIABLES				
		lognegative	**lognegative30**	**logsprange**	**logspvolume**	**logvixclose**
YEAR						
2002	**Mean**	1.22	1.30	2.95	21.07	3.35
	Std	0.40	0.08	0.42	0.24	0.23
2003	**Mean**	1.23	1.31	2.47	21.05	3.02
	Std	0.41	0.10	0.41	0.19	0.23
2004	**Mean**	1.09	1.17	2.28	21.09	2.67
	Std	0.45	0.08	0.41	0.19	0.14
2005	**Mean**	1.07	1.15	2.28	21.45	2.52
	Std	0.48	0.12	0.40	0.16	0.11
2006	**Mean**	1.18	1.26	2.36	21.66	2.55
	Std	0.42	0.08	0.46	0.19	0.18
2007	**Mean**	1.40	1.45	2.96	22.01	3.05
	Std	0.34	0.09	0.50	0.24	0.23
2008	**Mean**	1.42	1.47	3.38	22.40	3.54
	Std	0.32	0.04	0.59	0.31	0.48
All	**Mean**	1.22	1.29	2.61	21.49	2.90
	Std	0.43	0.14	0.58	0.50	0.41

Graph 2: Box-Whisker Plot of all Variables (9th April 2002 - 31st December 2008)

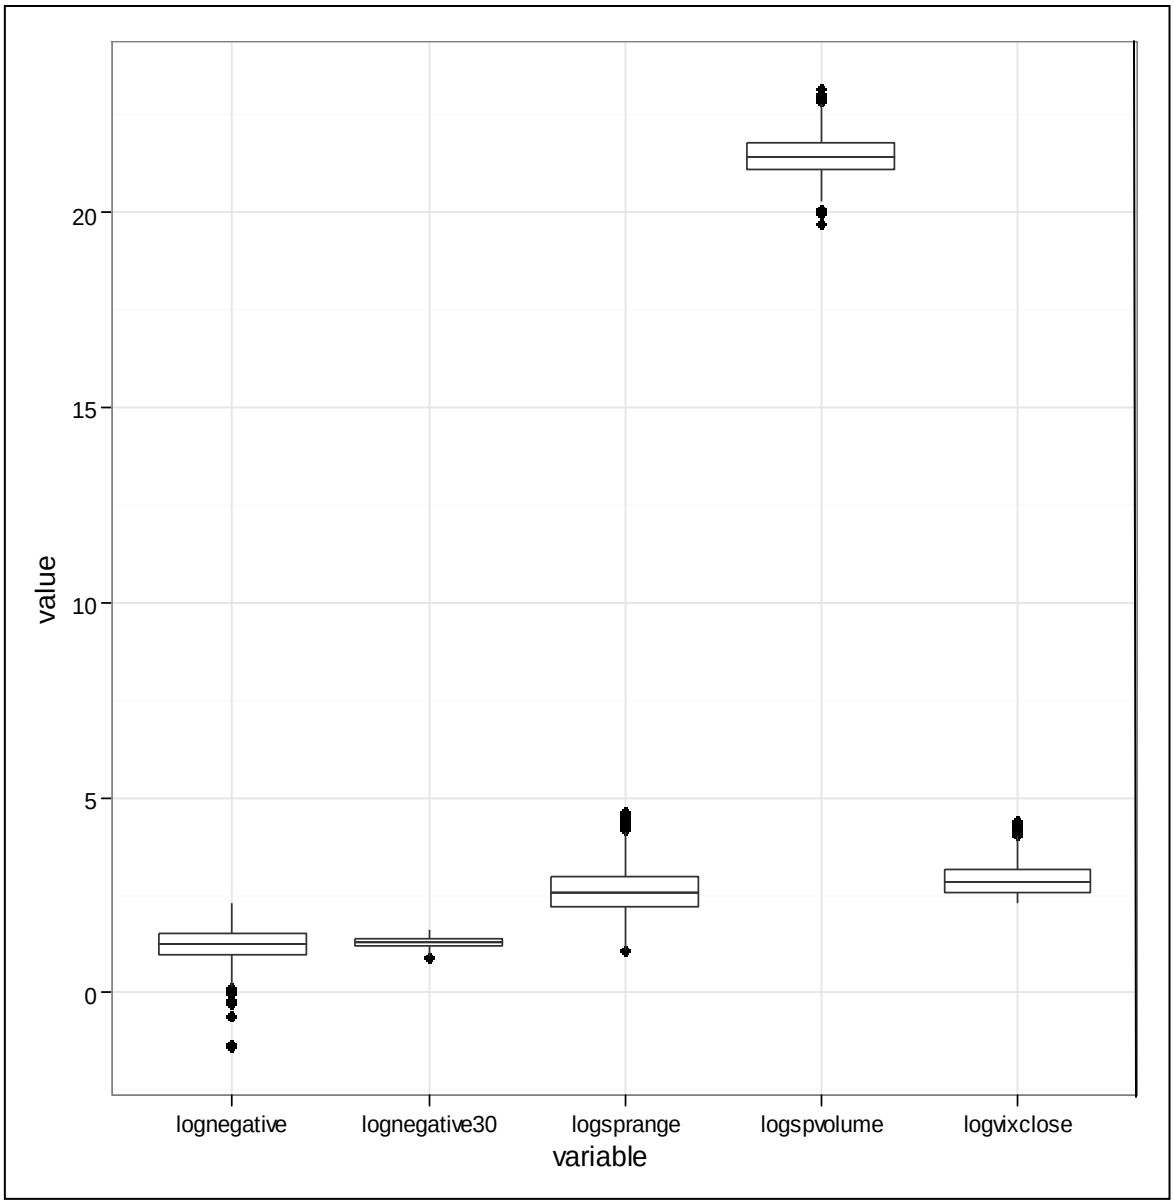

Graph 3: Box-Whisker Plot of all Variables Sorted by Year (2002-2008)

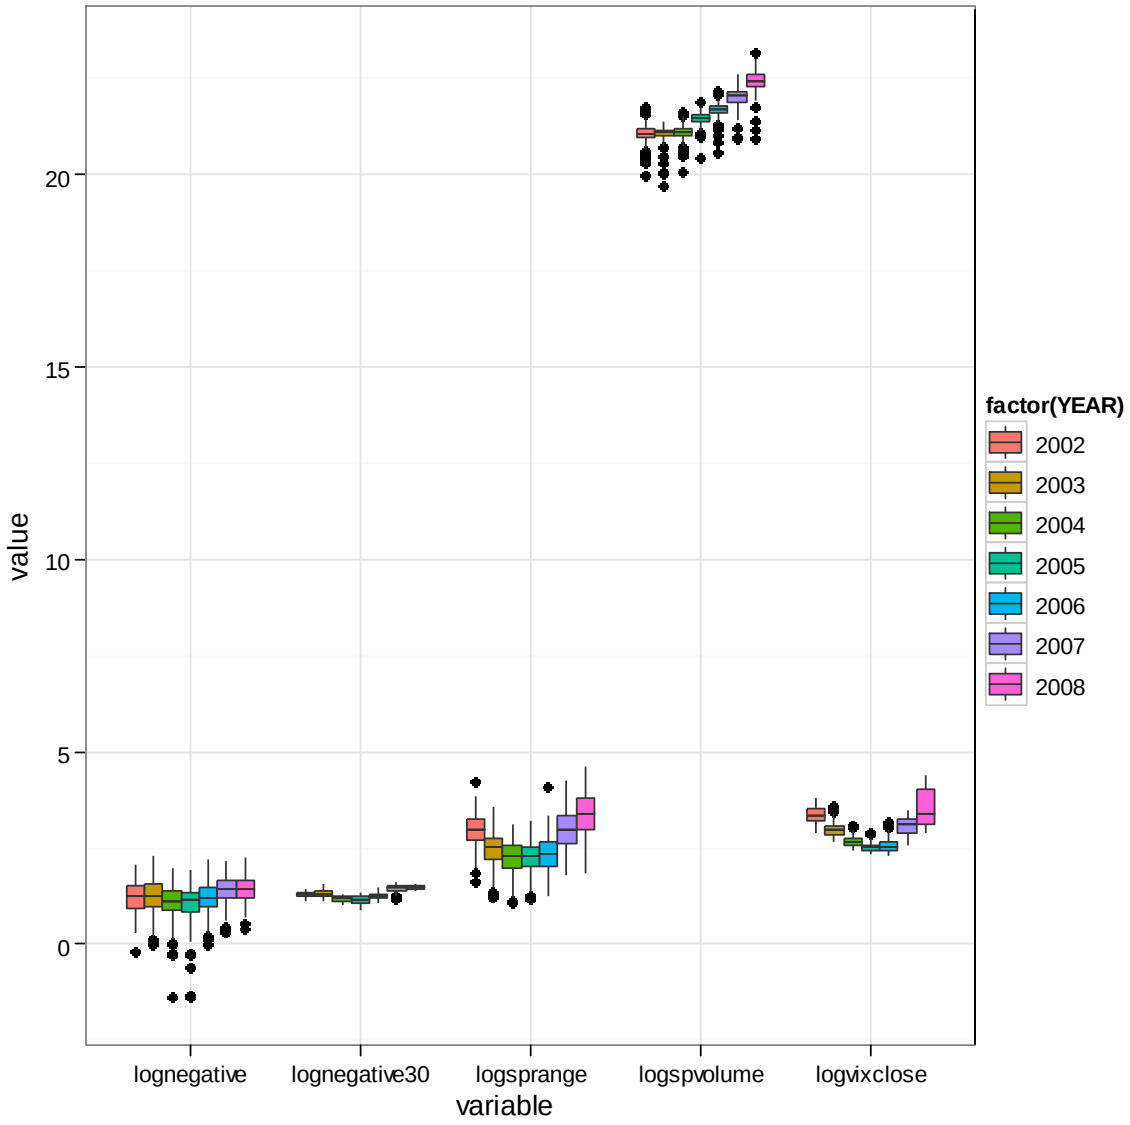

Table 5: Yearly Correlation between Negative News Scores and Market Variables

*Correlation coefficients reported are from the Pearson methodology. Numbers in parentheses denote Prob> |r| under H_0: r = 0 (where *, **, *** denote significance at 10%, 5% and 1% respectively).*

Mktrf is the risk adjusted market return for the S&P 500 index, logspvolume is the logged daily trading volume of the S&P 500, sprange is the daily trading range of the S&P 500, absmktrf is the absolute risk adjusted market returns for the S&P 500 index, logvixclose is the logged daily closing value of the CBOE VIX index, lognegative is the logged negative news score derived from textual analysis of Wall Street Journal articles, lognegative30 is the logged 30-day simple moving average of the negative news score.

		Mktrf	**Logvixclose**	**Sprange**	**Logspvolume**	**Absmktrf**
2002	*Lognegative*	0.06721 *(0.3556)*	0.14738 *(0.0419) ***	0.06063 *(0.4048)*	-0.01982 *(0.7855)*	0.05970 *(0.4120)*
	Lognegative30	0.09622 *(0.1854)*	0.40299 (<.0001) ***	0.02153 *(0.7675)*	0.16765 *(0.0204) ***	0.03224 *(0.6579)*
2003	*Lognegative*	-0.01598 *(0.7972)*	0.14714 *(0.0174) ***	0.04762 *(0.4436)*	-0.10728 *(0.0836) **	0.06816 *(0.2726)*
	Lognegative30	0.04073 *(0.5123)*	0.58484 *(<.0001) ****	0.28555 *(<.0001) ****	-0.07850 *(0.2062)*	0.20497 *(0.0009) ****
2004	*Lognegative*	0.00591 *(0.9242)*	0.02985 *(0.6306)*	-0.05436 *(0.3808)*	-0.05436 *(0.3808)*	0.01820 *(0.7693)*
	Lognegative30	-0.00240 *(0.9691)*	0.05737 *(0.3550)*	-0.07569 *(0.2221)*	-0.10215 *(0.0990) **	-0.00802 *(0.8972)*
2005	*Lognegative*	0.00395 *(0.9495)*	0.17561 *(0.0045) ****	0.16060 *(0.0095) ****	0.01438 *(0.8175)*	0.10924 *(0.0787) **
	Lognegative30	0.00999 *(0.8726)*	0.54410 *(<.0001) ****	0.15349 *(0.0132) ****	-0.02976 *(0.6329)*	0.09504 *(0.1264)*
2006	*Lognegative*	-0.00306 *(0.9608)*	0.00829 *(0.8942)*	0.07466 *(0.2302)*	0.05710 *(0.3591)*	0.07417 *(0.2333)*
	Lognegative30	0.00348 *(0.9555)*	-0.08360 *(0.1790)*	0.11235 *(0.0705) **	0.25177 *(<.0001) ****	0.00079 *(0.9899)*
2007	*Lognegative*	-0.04501 *(0.4691)*	0.25568 *(<.0001) ****	0.08347 *(0.1788)*	0.13492 *(0.0293) ***	0.05380 *(0.3867)*
	Lognegative30	-0.00916 *(0.8829)*	0.76325 *(<.0001) ****	0.35915 *(<.0001) ****	0.40886 *(<.0001) ****	0.24208 *(<.0001) ****
2008	*Lognegative*	0.14560 *(0.0754) **	-0.03407 *(0.6790)*	0.08715 *(0.2889)*	0.10921 *(0.1834)*	0.14646 *(0.0737) **
	Lognegative30	-0.00396 *(0.9616)*	-0.15390 *(0.0601) **	-0.12955 *(0.1141)*	0.05242 *(0.5241)*	-0.11043 *(0.1785)*
Full Sample	*Lognegative*	0.01599 *(0.517)*	0.22808 *(<.0001) ****	0.19481 *(<.0001) ****	0.18580 *(<.0001) ****	0.15295 *(<.0001) ****
N=1645	*Lognegative30*	-0.01144 *(0.6428)*	0.59415 *(<.0001) ****	0.45638 *(<.0001) ****	0.50645 *(<.0001) ****	0.29561 *(<.0001) ****

Graph 4: Scatter Plot and Correlation Matrix for Risk-adjusted Market Returns, Negative News Score and Market Volatility Indicators

*Correlation coefficients reported are from the Pearson methodology. Numbers in parentheses denote Prob> |r| under H_0: r = 0 (where *, **, *** denote significance at 10%, 5% and 1% respectively).*

Mktrf is the risk adjusted market return for the S&P 500 index, logspvolume is the logged daily trading volume of the S&P 500, sprange is the daily trading range of the S&P 500, absmktrf is the absolute risk adjusted market returns for the S&P 500 index, logvixclose is the logged daily closing value of the CBOE VIX index, lognegative30 is the logged 30-day simple moving average of the negative news score.

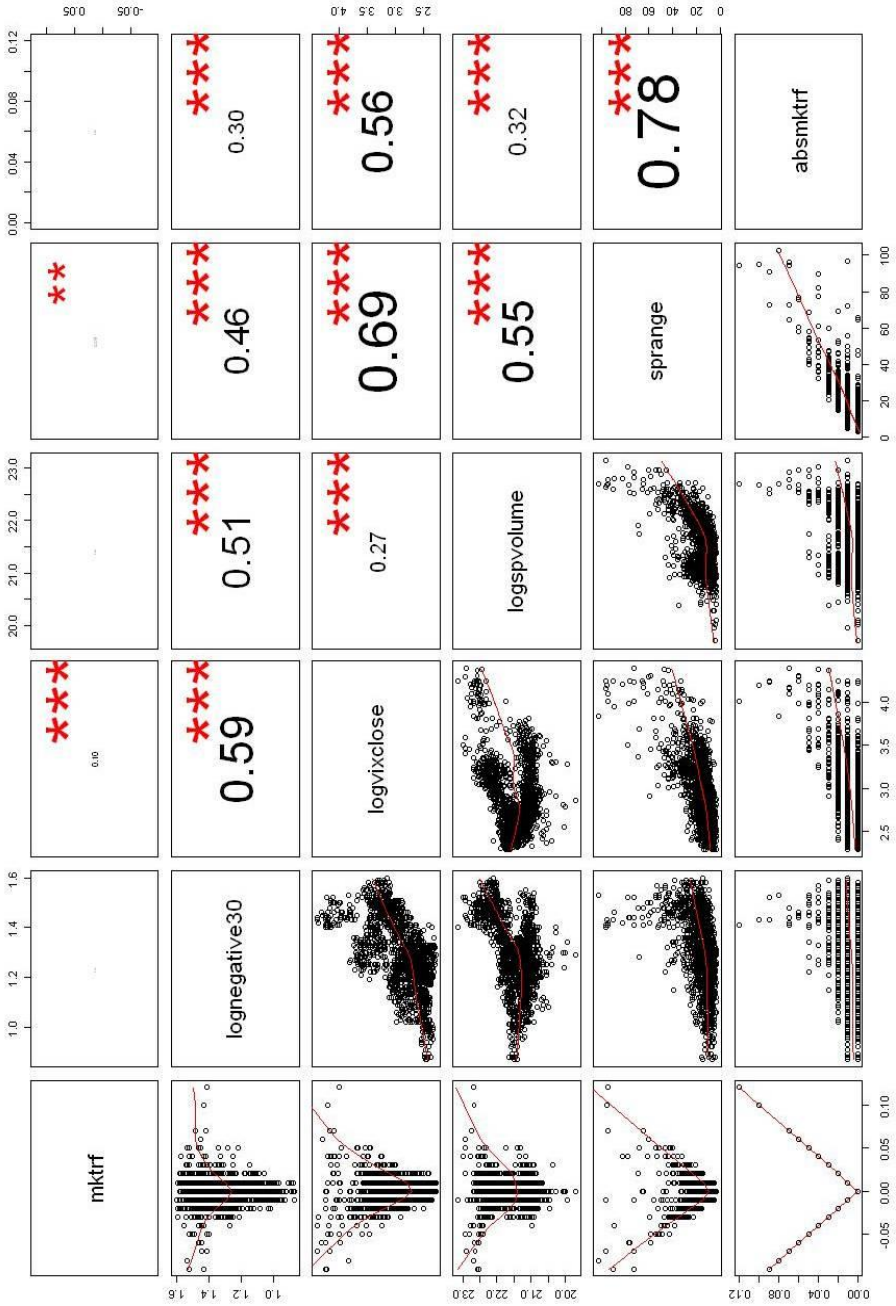

Table 6: Percentile Analysis of Negative News Score on Market Returns, Direction and Volatility*

Lognegative30 Percentile	Variables	N	T-value	Mean	Std
Above 90%	Mktrf	165	0.34	0.0005	0.0195
	$Mktrf_1$	165	0.06	0.0001	0.0197
	$Mktrf_2$	165	-0.25	-0.0004	0.0200
	$Mktrf_3$	165	-0.20	-0.0003	0.0201
	$Mktrf_4$	165	-0.59	-0.0009	0.0193
	$Mktrf_5$	165	-0.37	-0.0006	0.0190
	Average Mktrf	165	-0.37	-0.0002	0.0068
	Total Mktrf	165	-0.37	-0.0010	0.0340
	Market Direction	165	-0.77	-0.0255	0.4244
	Logvixclose	165	152.15	3.3151	0.2799
Above 75%	Mktrf	414	-0.65	-0.0007	0.0214
	$Mktrf_1$	413	-1.01	-0.0011	0.0213
	$Mktrf_2$	412	-0.80	-0.0008	0.0215
	$Mktrf_3$	411	-0.66	-0.0007	0.0215
	$Mktrf_4$	410	-0.86	-0.0009	0.0216
	$Mktrf_5$	409	-0.61	-0.0006	0.0214
	Average Mktrf	410	-2.30	-0.0009	0.0080
	Total Mktrf	410	-2.30	-0.0046	0.0400
	Market Direction	410	-1.49	-0.0302	0.4112
	Logvixclose	414	179.97	3.3003	0.3731
Below 25%	Mktrf	411	1.09	0.0004	0.0073
	$Mktrf_1$	411	0.88	0.0003	0.0075
	$Mktrf_2$	411	0.51	0.0002	0.0075
	$Mktrf_3$	411	0.80	0.0003	0.0075
	$Mktrf_4$	411	1.83	0.0007	0.0072
	$Mktrf_5$	411	1.21	0.0004	0.0074
	Average Mktrf	411	2.41	0.0004	0.0031
	Total Mktrf	411	2.41	0.0019	0.0156
	Market Direction	411	1.51	0.0307	0.4117
	Logvixclose	411	255.78	2.6443	0.2096

*$Mktrf_t$ refers to S&P market returns on day t. Average Mktrf is the average market returns across all 6 days ahead. Total Mktrf is the total market directions of all 6 days. Market Direction is a measure of average market directional movement which ranges from -1 (market declined across all days) to 1 (market moved upward across all days)

Table 6: Percentile Analysis of Negative News Score on Market Returns, Direction and Volatility (Continued)

Lognegative30 Percentile	Variables	N	T-value	Mean	Std
Below 10%	Mktrf	164	0.04	0.0000	0.0065
	$Mktrf_1$	164	-0.12	-0.0001	0.0069
	$Mktrf_2$	164	0.21	0.0001	0.0066
	$Mktrf_3$	164	0.08	0.0000	0.0066
	$Mktrf_4$	164	0.66	0.0003	0.0068
	$Mktrf_5$	164	0.68	0.0004	0.0068
	Average Mktrf	164	0.41	0.0001	0.0028
	Total Mktrf	164	0.41	0.0005	0.0142
	Market Direction	164	-0.49	-0.0146	0.3796
	Logvixclose	164	193.79	2.5709	0.1699

Table 7(a): Predicting S&P 500 Returns Using Negative News Score

Using Tetlock (2007) methodology, we examine if S&P 500 returns are predictable using the negative news score updated. The VAR model is defined as below:

$$Mktrf = \alpha_1 + \beta_1 . \ L5\,(Mktrf) + \gamma_1 . \ L5\,(Log.negative30) + \delta_1 . \ L5\,(Log.SPvolume) + \varepsilon_{1t}$$

Variable	**Coefficient** γ	**Std.Dev**	**P-Value**
Constant	0.002	0.00317	0.49
Negative News(t)	0.014	0.0190	0.44
Negative News(t-1)	-0.019	0.0268	0.46
Negative News(t-2)	0.003	0.02683	0.90
Negative News(t-3)	-0.023	0.01684	0.17
Negative News(t-4)	-0.024	0.01909	0.19

Table 7(b): Buy and Hold Returns for S&P 500 Long and Short Trading Strategy using Negative News Score from 9th April 2002 - 31st December 2008

If the negative news score is below its 25% percentile (proxy for low negativity in financial news), the S&P 500 e-mini futures is bought at opening price and position held for a time period t.Position is then theoretically sold at closing price.

If the negative news score is above its 75% percentile (proxy for high negativity in financial news), the S&P 500 e-mini futures is sold and position held for a time period t. Position is then theoretically sold at closing price.

Transaction costs and other market costs are not taken into account in this table. T-statistics are reported in the parentheses, all returns are risk-adjusted returns (portfolio returns minus risk free rate).

Long Portfolio Where negative news score < 25% percentile					
	Total Returns of all Trades	**Average Returns on Each Trade**	**Standard Deviation of Returns**	**Max Loss on a Single Trade**	**Max Gain on a Single Trade**
Holding Period t = 5	27.00%	0.101% *(1.21)*	1.50%	-5.00%	5.00%
Holding Period t = 15	16.00%	0.058% *(0.72)*	3.00%	-8.00%	8.00%
Short Portfolio Where negative news score > 75% percentile					
	Total Returns of all Trades	**Average Returns on Each Trade**	**Standard Deviation of Returns**	**Max Loss on a Single Trade**	**Max Gain on a Single Trade**
Holding Period t= 5	260.00%	0.75% *(3.25)*	4.00%	-19.00%	19.00%
Holding Period t = 15	11.00%	0.04% (0.58)	5.50%	-19.00%	33.00%

Table 8: Granger-Causality Wald Test on Negative News Score and other Market Variables

A Granger-Causality Wald-Test was run on the Negative News Score (lognegative30) and all other variables through a vector autoregressive VAR(2,L) model. Lag lengths (L) of 1 and 5 were selected for comparison purposes. Variables indicated before arrow imply the granger causality direction to subsequent variables. The test statistics reported below are Chi-squared values from the Wald test and p-values are in the parentheses (, **, *** denote significance at a 10%, 5% and 1% level respectively).*

Variables and Direction of Causality	**Test Statistics – χ^2 (Pr> χ^2)**	
	Lag = 1	**Lag = 5**
Negative News → Risk-Adjusted Returns *(Mktrf)*	0.41 *(0.5195)*	2.45 *(0.7843)*
Negative News → VIX *(Logvixclose)*	5.01 *(0.0252)* **	6.76 *(0.2382)*
Negative News → S&P 500 Trading Range *(Sprange)*	18.92 *(0.0001)* ***	13.40 *(0.0199)* **
Negative News → S&P 500 Trading Volume *(Logspvolume)*	24.67 *(0.0001)* ***	6.41 *(0.2687)*
Risk-Adjusted Returns *(Mktrf)* → Negative News	0.13 *(0.7219)*	4.13 *(0.5303)*
VIX *(Logvixclose)* → Negative News	4.91 *(0.0403)* **	9.57 *(0.0883)* *
S&P 500 Trading Range *(Sprange)* → Negative News	2.29 *(0.1309)*	9.23 *(0.1002)*
S&P 500 Trading Volume *(Logspvolume)* → Negative News	3.85 *(0.0550)* **	8.83 *(0.l160)*

Graph 5: The Chicago Board Options Exchange Volatility Index (1993 – 2008) and VIX – Negative News Score (April 2002 – December 2008)

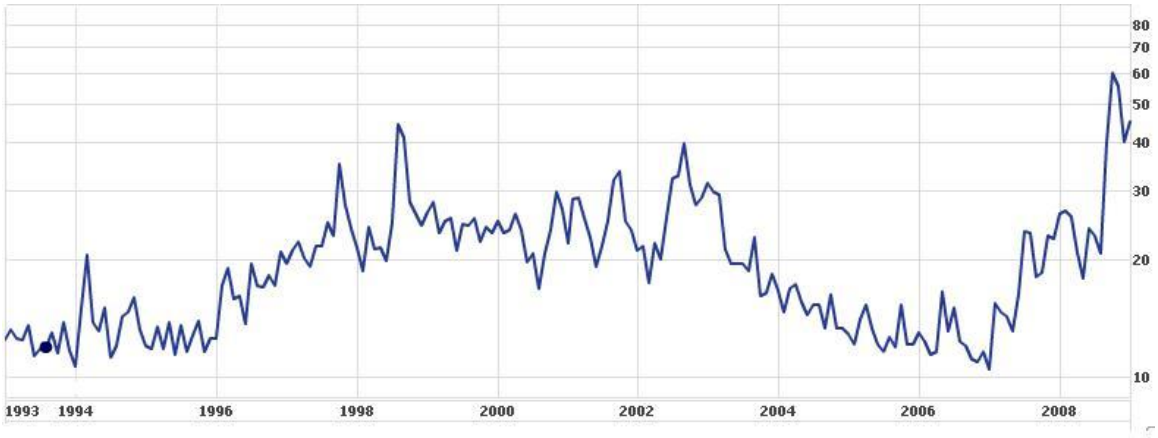

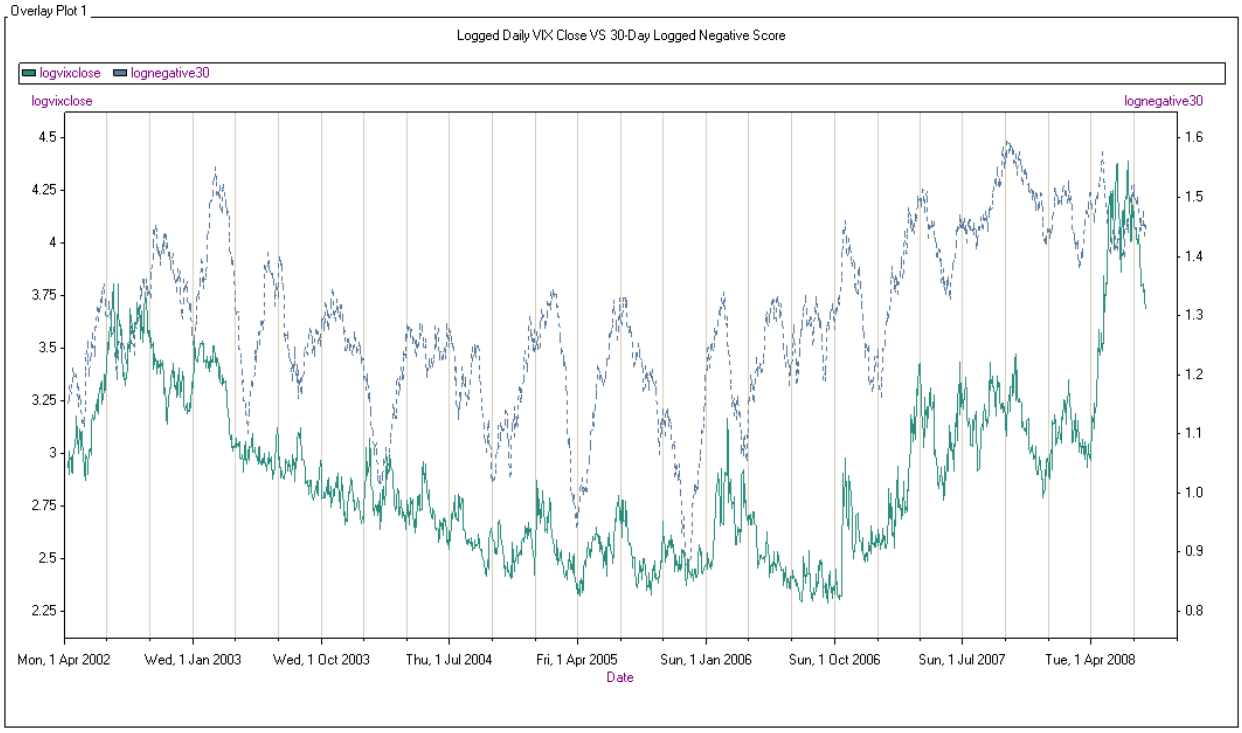

Table 9: VAR Model Specifications

Model	Specification
Model1: ARX(0, 1)	$\log(vixclose)_t =$ $\alpha + \beta_1.\log(negative30)_t + \epsilon_t$
Model 2: ARX(0, 1)	$\log(vixclose)_t =$ $\alpha + \beta_1.\log(negative30)_t + \beta_2.\mathrm{Mktrf}_{(t-1)} + \epsilon_t$
Model 3: ARX(0, 1)	$\log(vixclose)_t =$ $\alpha + \beta_1.\log(negative30)_t + \beta_2.\mathrm{Mktrf}_{(t-1)} + \beta_3.\mathrm{SPVolume}_{(t-1)} + \beta_4.\mathrm{SPRang}$ $+ \epsilon_t$
Model 4: ARX(1, 1)	$\log(vixclose)_t =$ $\alpha + \beta_1.\log(negative30)_t + \beta_2.\mathrm{Mktrf}_{(t-1)} + \beta_3.\mathrm{SPVolume}_{(t-1)} + \beta_4.\mathrm{SPRang}$ $+ \beta_5.\log(\mathrm{vixclose})_{(t-1)} + \epsilon_t$
Model 5: ARX(2, 1)	$\log(vixclose)_t =$ $\alpha + \beta_1.\log(negative30)_t + \beta_2.\mathrm{Mktrf}_{(t-1)} + \beta_3.\mathrm{SPVolume}_{(t-1)} + \beta_4.\mathrm{SPRang}$ $+ \beta_5.\log(vixclose)_{(t-1)} + \beta_6.\log(vixclose)_{(t-2)} + \epsilon_t$
Model 6: ARX(1, 1) Diff(1)	$\delta(1).\log(vixclose)_t =$ $\alpha + \beta_1.\delta(1).\log(negative30)_t + \beta_2.\delta(1).\mathrm{Mktrf}_{(t-1)} + \beta_3.\delta(1).\mathrm{SPVolume}_{(t}$ $+ \beta_4.\delta(1).\mathrm{SPRange}_{(t-1)} + \beta_5.\delta(1).\log(vixclose)_{(t-1)}$ $+ \beta_6.\delta(1).\log(vixclose)_{(t-2)} + \epsilon_t$

Table 10: Vector Auto Regressions of Logged VIX Closing Values on Negative News Score and Control Variable Vector

*Numbers reported are coefficients with t-statistics in parentheses. *, ** and *** denote significance at a 10%, 5% and 1% level respectively.*

Model Parameters	**Model 1 ARX(0,1)**	**Model 2 ARX(0,1)**	**Model 3 ARX(0,1)**	**Model 4 ARX(1,1)**	**Model 5 ARX(2,1)**	**Model 6 ARX(1,1) Diff (1)**
Constant	0.67321 *(8.99) ****	0.67500 *(9.06) ****	0.25005 *(2.47) ***	0.05985 *(0.66)*	0.08051 *(0.90)*	0.00048 *(0.31)*
Lognegative30	1.72721 *(29.94) ****	1.72531 *(30.01) ****	1.34611 *(25.80) ****	0.03356 *(2.17) ***	0.03203 *(2.08) ***	-
Mktrf(t-1)	-	-2.40492 *(-4.16) ****	-1.21833 *(-2.75) ****	-0.41405 *(-3.75) ****	0.04262 *(0.27)*	-
LogSpvolume(t-1)	-	-	-0.26893 *(-16.89) ****	-0.00256 *(-0.59)*	-0.00375 *(-0.78)*	-
LogSprange(t-1)	-	-	0.022338 *(34.69) ****	0.00008 *(0.39)*	0.00028 *(1.27)*	-
LogVixclose(t-1)	-	-	-	0.98323 *(156.80)****	0.85896 *(23.14) ****	-
LogVixclose(t-2)	-	-	-	-	0.00535 *(0.13)*	-
ΔLognegative30	-	-	-	-	-	0.15435 *(1.63) **
ΔMktrf(t-1)	-	-	-	-	-	-0.21861 *(-2.41) ***
ΔLogSpvolume(t-1)	-	-	-	-	-	-0.01290 *(-1.59)*
ΔLogSprange(t-1)	-	-	-	-	-	-0.00003 *(-0.19)*
Δ LogVixclose(t-1)	-	-	-	-	-	-0.15649 *(-5.12) ****
ΔLogVixclose(t-2)	-	-	-	-	-	-
AIC	-2.20846	-2.21768	-2.76815	-5.5414	-5.55788	-5.55605
SBC	-2.20189	-2.20783	-2.75173	-5.52169	-5.53159	-5.53304
R^2	0.35	0.36	0.63	0.98	0.98	-
Final Prediction Error	0.1098	0.1088	0.06277	0.00392	0.00385	0.00386

APPENDIX A: The General Inquirer Content Analysis System

Abstract from "How the General Inquirer is used and a comparison with other text analysis procedures" found at http://www.wjh.harvard.edu/~inquirer/3JMoreInfo.html

The General Inquirer Content Analysis System is a computational text mapping tool. It maps each text file with counts on dictionary-supplied categories. The currently distributed version combines the "Harvard IV-4″ dictionary content-analysis categories, the "Lasswell" dictionary content-analysis categories, and five categories based on the social cognition work of Semin and Fiedler, making for 182 categories in all. Each category is a list of words and word senses. A category such as "self-references" may contain only a dozen entries, mostly pronouns. Currently, the category "negative" is our largest with 2291 entries. Users can also add additional categories of any size.

Because overlaps exist among content-analysis tools as well as with "qualitative analysis", text-management, "natural language processing" and some artificial-intelligence software, it is important to have realistic expectations about what each content-analysis tool, including the General Inquirer, can readily provide. For some research projects, especially those involving intensive analyses of modest quantities of text, other text-analysis software may be much more appropriate.

In order to map category assignments with reasonable accuracy, the General Inquirer software spends most of its processing time identifying commonly used word senses. For example, it distinguishes between "race" as a contest, "race" as moving rapidly, "race" as

a group of people of common descent, and "race" in the idiom "rat race". The General Inquirer also cautiously removes common regular suffixes so that one entry in a category can match several inflected word forms. A category entry can be an inflected word (for example, "swimming"), a root word ("swim" would match "swimming", if "swimming" is not a separate entry) or a word sense (for example, "swim#1") identified by our disambiguation routines of an inflected or root word form. These English stemming procedures, integrated with English dictionaries and routines for disambiguating English word senses, limit the current Inquirer system to English text applications.

As a tool that maps pre-specified categories, the General Inquirer differs from purely inductive mapping tools, such as the so-called neural-net building procedures that are now included in several software packages. However, as the correspondence analysis example above indicates, some inductive tools may be applied to Inquirer-produced spreadsheets of category counts, mapping either the relationships between categories or the relationships between documents into a multidimensional space.

Unlike some artificial intelligence programs that can be applied to texts within limited topic domains, the General Inquirer simply maps text according to categories and does not search after meaning. General Inquirer mappings have proven to supply useful information about a wide variety of texts. But it remains up to the researchers, not the computer, to create knowledge and insight from this mapped information, usually situating it in the context of additional information about the texts' origins.

APPENDIX B: Sample of Words in Harvard-IV Dictionary's Negative Category

ABATE	DERISIVE	FAIL	INFLATION
ABATED	DEROGATORY	FAILED	INFLICT
ABATEMENT	DESERT	FAILING	INFLICTED
ABATEMENTS	DESERTED	FAILS	INFLICTING
ABATES	DESERTING	FAILURE	INFLICTS
ABATING	DESERTION	FAILURES	INFRACTION
ABDICATE	DESERTS	FAINT	INFRACTIONS
ABDICATED	DESIRE	FAINTED	INFRINGE
ABDICATES	DESIRED	FAINTING	INFRINGED
ABDICATING	DESIRES	FAINTS	INFRINGEMENT
ABDICATION	DESIRING	FAKE	INFRINGEMENT
ABDICATIONS	DESOLATE	FAKED	INFRINGES
ABHORS	DESOLATION	FAKES	INFRINGING
ABJECT	DESPAIR	FAKING	INFURIATE
ABJECTION	DESPERATE	FALL	INFURIATED
ABJECTIONS	DESPERATION	FALLACIES	INFURIATES
ABJECTLY	DESPICABLE	FALLACY	INFURIATING
ABJECTNESS	DESPISE	FALLING	INGRATITUDE
ABNORMAL	DESPISED	FALLOUT	INHIBIT
ABNORMALITIES	DESTITUTE	FALLS	INHIBITED
ABNORMALITY	DESTROY	FALSE	INHIBITING
ABNORMALLY	DESTROYED	FALSEHOOD	INHIBITION
ABOLISH	DESTROYING	FALSEHOODS	INHIBITS
AGGRAVATED	DESTROYS	FALSELY	INHUMANE
AGGRAVATES	DESTRUCTION	FALSIFICATION	INJUNCTION
AGGRAVATING	DESTRUCTIVE	FALTER	INJUNCTIONS
AGGRAVATION	DETACHMENT	FALTERED	INJURE
AGGRAVATIONS	DETAIN	FALTERING	INJURED
AGGRESSION	DETAINED	FALTERS	INJURES
AGGRESSIVE	DETAINING	FAMINE	INJURIES
AGGRESSIVELY	DETAINS	FAMISHED	INJURING
AGGRESSIVENESS	DETER	FANATIC	INJURIOUS
AGGRIEVE	GLOAT	FANATICAL	INJURY
AGGRIEVED	GLOATED	FANATICS	INSANE
AGHAST	GLOATING	FARCE	INSECURE
AGITATE	GLOATS	FARCES	INSECURITIES
AGITATED	GLOOM	FASCIST	INSECURITY
AGITATES	GLOOMY	FAT	INSENSIBLE
AGITATING	GLUM	FATAL	INSIDIOUS
AGITATION	GODDAMN	FATALISTIC	INSIGNIFICANT
AGITATIONS	GRAB	FATIGUE	INSINUATE
AGITATOR	GRABBED	FAULT	INSINUATED
AGITATORS	GRABBING	FAULTS	INSINUATES
AGONIZE	GRABS	FAULTY	INSINUATING
AGONIZED	GRAPPLE	FEAR	INSOLENCE
AGONIZES	GRAPPLED	FEARED	INSOLENT
AGONIZING	GRAPPLES	FEARFUL	INSTABILITIES
AGONIZINGLY	GRAPPLING	FEARS	INSTABILITY
AGONY	GRATUITOUS	FEARSOME	INSTABLE
AIL	GRAVE	FRANTIC	INSUFFICIENCY
AILED	GRIEF	FRANTICALLY	INSUFFICIENT