risks

Article

Initial Investigations of Intra-Day News Flow of S&P500 Constituents

Jim Kyung-Soo Liew [1],* and Zhechao Zhou [2]

[1] Finance Department, The Johns Hopkins Carey Business School Baltimore, MD 21202, USA

[2] Investment Technology Group, Inc., New York, NY 10006, USA; E-Mail: zhechao.zhou@itg.com

* Author to whom correspondence should be addressed; E-Mail: kliew1@jhu.edu;
Tel.: +410-234-9401.

Abstract: In this work, we examine Thomas Reuters News Analytics (TRNA) data. We found several fascinating discoveries. First, we document the phenomenon that we label "Jam-the-Close": The last half hour of trading (15:30 to 16:00 EST) contains a substantial and statistically significant amount of news sentiment releases. This finding is robust across years and months of the year. Next, upon further investigations we found that the "novelty" score is on average 0.67 in this period *vs.* 2.09 prior to midday. This indicates that "new" news is flowing at a rapid pace prior to the close. Finally, we discuss the implication of such phenomena in the context of existing financial literature.

Keywords: TRNA; news sentiments; intra-day prices; S&P500

1. Introduction

Any stock trader will attest to the influence of "news" on securities prices. In just a few months of trading, most neophyte traders will experience the impact of a market-moving news event that spooks the entire market into a downward panic, or an idiosyncratic news shock that crashes an individual stock. Frustrating as it may be, the monitoring of news releases related to stock positions is a must for any seasoned trader.

One of the most exciting times to witness reactions of stock prices to news events, occurs in the earnings seasons. Whether a company misses or beats its expected "numbers" will have a profound effect on the security prices behavior, especially when the "earnings surprise" is highly unexpected.

For example, when a company misses earnings for the first time, after consistently beating consensus estimates quarter-after-quarter, this will have serious consequences on the price behavior, especially if they were expected to beat consensus estimates and trade higher. Panic selling occurs as fear strikes deep into the trader's heart and the "unknown-unknowns" have manifested in such a negative fashion. Thoughts of "How could we have missed this?" and "Are our views completely incorrect here?" surface from the logical part of the mind and many times emotional panic sets in: "Sell. Sell. Sell!"

Such occurrences happen all the time during earnings seasons. As such, news stories become very important as traders are constantly looking to update their expectations of company's fundamentals which include expected future cash flow generating abilities.

Additionally, traders make use of news to stay informed about the future prospects of the companies that they trade. In today's world, where news is abundant and easily accessible, the trader's problem is not so much finding the news, but making sense of the blizzard of new stories generated each day.

With the advent of fast computers and powerful sentiment processing engines which enables the instantaneous transformation of text news stories into sentiment values, we have recently embarked on the "new-new" idea in finance: the birth and continuous growth of the language recognition algorithms that quickly interpret news flows. We have reached a point where machines can analyze the tone and relevance of news stories in practically real time (milliseconds to microseconds delay) and, ultimately, automatically trade news flows at the intraday frequency.

While understanding of news sentiment has only just begun, we have heard that many high-frequency trading shops have already mined this data and have made many of the easy trading strategies obsolete. In this work, we begin to examine, for mass consumption, the behavior of such sentiment data and stock prices.

We have already seen some industrial research papers that intended to propagate such products and services. However caution should be exercised as many of these papers have biased scientists frantically searching for results that would provide added attractiveness for the vendor's product. As such, we believe such biases creep into the analyses as finding such trading strategies would help the sales efforts of these vendors but unfortunately in practice such results may not come into fruition. The markets move towards efficiency very quickly and such "edges" should and do disappear very quickly in the modern marketplace. Not to mention the many data-snooping biases that, all too quickly, enter the realm of research crunching massive amounts of data with little investigational intuition to guide these works.

In this paper, we slowly and thoughtfully attempt to first understand the news sentiment data. Specific attention to detail is performed to discover empirical differences of the data's behavior from other studies. Additionally, we examine the relationship between the sentiment data and changes in security prices. In particular we are interested in applying our intuition to formulate hypotheses and test these hypotheses on the data. Some ideas worked the way we thought it would *a priori*, while others go counter to our intuition. In both cases we document our results accordingly.

Most seasoned quants in industry understand the limitations of stock market anomalies found in the academic literature. Once real-world frictions enter the extraction of these anomalies, profits becomes unattainable. It is well-known that many of the previous financial anomalies are driven by those stocks that are small (as defined by market capitalization) and thinly traded (as measured by volume).

Additionally, many prior studies have not properly accounted for the distinction between normal trading hours of 09:30 to 16:00 *vs.* after-market trading hours 16:00 to the following day 09:29:59.999. We find both the quantity and quality of news releases are very distinct during those time periods. We find also that there is a significant increase in the number of news sentiment releases in the last half hour of trading. We name this phenomenon "Jam-the-Close." We also find that these sentiment data released in this period have a highly "novel" meaning as they contain new pieces of information.

We review the academic literature and place our work in the context of prior works, as well as the growing body of news sentiment research. Next, we discuss our data and document relatively well-known behaviors seen within our data set. Afterwards, we show our main results after thorough documentation and investigations into the "Jam-the-Close" phenomenon. Finally, we conclude our paper with thoughts about future research.

2. Literature Review

The theoretical justification of this empirical work is best cast with the assertions of Grossman and Stiglitz [1], and more recently updated by Stein [2]. That is, even with a large number of institutional traders or information intermediaries, this does not guarantee that markets will be efficient. We have seen many prior works identifying stock anomalies and questioning the efficient market hypothesis.

For example, Jegadeesh and Titman [3] argue that momentum is a result of under-reaction to information. Momentum has been documented to also exist in international stocks [4,5]. Additionally, Asness, Moskowitz, and Pedersen [6] have summarized much of the existing literature on momentum across stocks, bonds, currencies, and commodities. Recently, Sinha [7] documents that the tone of news articles can predict future returns over the next 13 weeks. He concludes that the market underreacts to news releases.

Other anomalies have been documented that further question the market's ability of being efficient such as work by Banz [8] which documents the "size effect", wherein using a security's market capitalization one can cross-sectionally distinguish abnormal returns. In our work, we purposely restrain our analysis to a survivorship bias-free news sentiment dataset of the S&P500 constituents, so as to discover results that are robust to such size critics. Practitioners and academics readily understand the limitation of attempting to extract trading strategies based on fundamentals in the small-cap area [8–10].

Other fundamental variables, such as book-to-market and price-to-earnings, have been well documented as anomalies [5,10–14]. Fama and French [12,15,16] have even adjusted the CAPM of Sharpe [17] and Lintner [18] to include some of these factors. Moreover, Carhart [19] has extended Fama and French's 3-factor model to include momentum as defined along the lines of Asness [20].

Building upon such strong foundational research, we have witnessed a continued growth in the literature with a focus on better understanding news sentiment data. Much of this work attempts to uncover relationships between news sentiment data and security prices, see [5-7,21–30].

3. The Data

In this project we employ Thomson Reuters News Analytics (TRNA) data to explore the relationship between company-related news and price dynamics. Traditionally, news items are fundamental and qualitative in nature, consequently it is nearly impossible to investigate news flow from a systematic perspective. Thomson Reuters has however developed a language recognition algorithm that quantifies news items and informs investors of how relevant and how positive or negative the content truly is about the companies of interest. The Bayesian algorithm developed analyses news provided by more than 60 reporting sources (news agencies, magazines, *etc.*) and browses through a dictionary of about 20,000 words which have been segregated and defined as relevant amongst corporate news ("take-over," "profit warning," "beat," "disappoint," *etc.*). The outcome produced by the algorithm is a set of quantitative scores based on the presence and the position of the words in the text. The scores broadly fall into the following categories:

1. Relevance: A number of measures of how relevant the news item is to the asset. It ranges from 0 to 1, with 1 being the most relevant.
2. Sentiment: Whether the news item talks about the asset in a positive, neutral or negative manner. Every new story has three scores, the positive, neutral and negative scores. Each one is regarded as the "probability" that the nature of the story falls into the positive/negative/neutral category. The three scores of a story add up to 1.
3. Novelty: A measure of similarity of this item to previously seen news items.
4. Volume: Counts of the number of recent items mentioning the asset.
5. Headline Classification: Specific analysis of the headline.

A snapshot of the sample data is provided in Figure 1. For example, if news is released about company X "beating the analysts' estimation on earnings this season", it is then likely that the algorithm will rate the sentiment of the news as extremely positive (>95%) and the relevance of the news as very high (>95%). If the news announcement is about the company "sponsors a new research building in a university overseas", the algorithm is likely to rate the sentiment of the news as slightly positive (let's say +12%) and the relevance of the news as fairly low (let's say 5%).

Figure 1. A snapshot of the news sentiment data.

IDN_TIME	BCAST_REF	STOCK_RIC	RELEVANCE	SENTIMENT	SENT_POS	SENT_NEUT	SENT_NEG	LNKD_CNT1	LNKD_CNT2	LNKD_CNT3	LNKD_CNT4	LNKD_CNT5	ITEM_TYPE	BCAST_TEXT
10/1/12 0:00:03.03	600362.SS	600362.SS	0.123091	−1	0.218358	0.120649	0.661093	0	1	1	1	1	ARTICLE	RPT-UPDATE 1-INTERVIEW-Exxon yet to inspect Afghan
10/1/12 0:00:03.03	XOM.N	XOM.N	1	1	0.413991	0.284191	0.301818	0	3	3	3	3	ARTICLE	RPT-UPDATE 1-INTERVIEW-Exxon yet to inspect Afghan
10/1/12 0:00:15.15	ADNC.O	ADNC.O	1	1	0.79934	0.140003	0.0606572	0	0	0	0	0	ARTICLE	Audience earSmart Technology Delivers Industry Lea
10/1/12 0:15:58.58	ASL.AX	ASL.AX	1	−1	0.0556283	0.125228	0.819144	0	0	0	0	0	ALERT	AUSDRILL LTD - TRADING HALT <ASL.AX>
10/1/12 0:33:56.56	ARI.AX	ARI.AX	1	1	0.455159	0.167824	0.377016	1	1	1	1	1	ARTICLE	Australia's Arrium rejects $1 bln bid from Noble,
10/1/12 0:33:56.56	005490.KS	005490.KS	1	−1	0.203471	0.160732	0.635796	0	0	0	0	0	ARTICLE	Australia's Arrium rejects $1 bln bid from Noble,
10/1/12 0:33:56.56	NOBG.SI	NOBG.SI	1	−1	0.203471	0.160732	0.635796	0	0	0	0	0	ARTICLE	Australia's Arrium rejects $1 bln bid from Noble,
10/1/12 0:36:53.53	BP.L	BP.L	0.707107	−1	0.0556283	0.125228	0.819144	0	0	0	0	0	ARTICLE	BUZZ-AUD/USD-Steadies around 1.0350/TDMI inflation
10/1/12 0:39:13.13	ARI.AX	ARI.AX	1	1	0.77273	0.14842	0.07885	5	5	5	5	5	ARTICLE	Singapore's Noble Group confirms interest in Aust
10/1/12 0:39:13.13	005490.KS	005490.KS	0.288675	1	0.628861	0.171895	0.199244	1	1	1	1	1	ARTICLE	Singapore's Noble Group confirms interest in Aust
10/1/12 0:39:13.13	NOBG.SI	NOBG.SI	1	1	0.785504	0.144584	0.0699114	1	1	1	1	1	ARTICLE	Singapore's Noble Group confirms interest in Aust
10/1/12 0:40:46.46	SIFI.MI	SIFI.MI	0.25	1	0.688999	0.16586	0.145141	0	0	0	0	0	ARTICLE	PRESS DIGEST - British business - Oct 1
10/1/12 0:40:46.46	WMH.L	WMH.L	0.5	1	0.418416	0.163552	0.418033	0	0	0	0	0	ARTICLE	PRESS DIGEST - British business - Oct 1
10/1/12 0:40:46.46	SBT.L	SBT.L	0.5	1	0.43358	0.165463	0.400956	0	0	0	0	0	ARTICLE	PRESS DIGEST - British business - Oct 1

Whilst many fields are available through the data, we decided to focus our attention on two fields: the news relevance and the news sentiment of S&P 500 constituent stocks. In the future, we intend to include additional fields such as novelty of the news, as well as on different asset classes (commodities, FX, *etc.*). Here are some basic statistics about the data from 1 January 2003 to 31 October 2012. There are 4,003,206 stories related to S&P 500 stocks during that period.

We notice that the number of news stories on Saturday and Sunday is much less than the number on weekdays (as shown in Figure 1) and relatively insignificant, therefore we filtered out the stories on weekends. After the filtering, there are 3,935,936 stories left for the period. The number of stories reported outside of trading hours (defined here 09:30 EST to 16:00 EST referring to the hours of the New York Stock Exchange (NYSE) is 2,232,088, while 1,703,848 stories are reported during trading hours.

Many of the stories in the data are highly related to the entities mentioned. The mean relevance of the stories is 0.6770. The amount of 2,146,792 stories have a relevance score of 1; whereas 2,183,603 events have a relevance of >0.75 (highly relevant). Out of all the stories, 892,655 stories are categorized ALERTs, which are single-line texts, 2,464,740 stories are ARTICLEs. ALERTS usually come out ahead of ARTICLEs which have both headlines and texts. The lags between an ALERT and the following ARTICLE range from a few minutes to a few hours. The positiveness/negativeness of an ARTICLE is possibly more accurate than an ALERT. However, the ALERTs have a very high relevancy score: of all the ALERTs, the average relevance score is 0.9765, while mean relevance score for all ARTICLEs is 0.620.

4. News Flows Seasonality

A fundamental principle of market efficiency is that investors react to new information as it arrives, resulting in price changes that reflect investors' expectations of risk and return. Previous literature has documented that public information arrival is non-constant [30]. The first question to be asked about our data is whether the rate of public information flow displays any seasonal and intraday patterns. We define the rate of public information flow as the number of news stories per unit of time in our data, and document the number of news over various time intervals, such as when the market is open *vs.* when it is closed, or during a regular month *vs.* an earning month.

Table 1 displays our data organized by month of the year. It is apparent from the table that there is significantly more news released during the earnings announcement months (January, April, July, and October). Those four months have the most news releases among all twelve and in total account for 38% of the stories in the data, despite some irregular news bursts like the one during the September 2008 crisis which boosts the average count for September. The occurrence of news items aggregated by month has a cyclical pattern as shown in Figure 2. After an earning-announcement, more news is released about the company, and there is a peak around that time. As time progresses less news is released. The month prior to the earning announcement usually has the least amount of news, possibly due to quiet periods.

These first-stage results on news flow seasonality suggest that a statistical test of the news counts in earning months and regular months is warranted Figure 3. Accordingly, we form the null hypothesis as the news counts of earning announcement months and the news counts of regular months having the same mean value. In order to test this hypothesis, we divide the data into two groups, Group 1 with the news counts aggregated by weeks in regular months, Group 2 with the news story counts aggregated by weeks from earning announcement months. We run the *t*-test on two groups to examine if the counts have the same mean value.

Figure 2. Total counts of new stories by day of the week from 1 January 2003 to 31 October 2012. The numbers of news on Saturday and Sunday are much less than on weekdays.
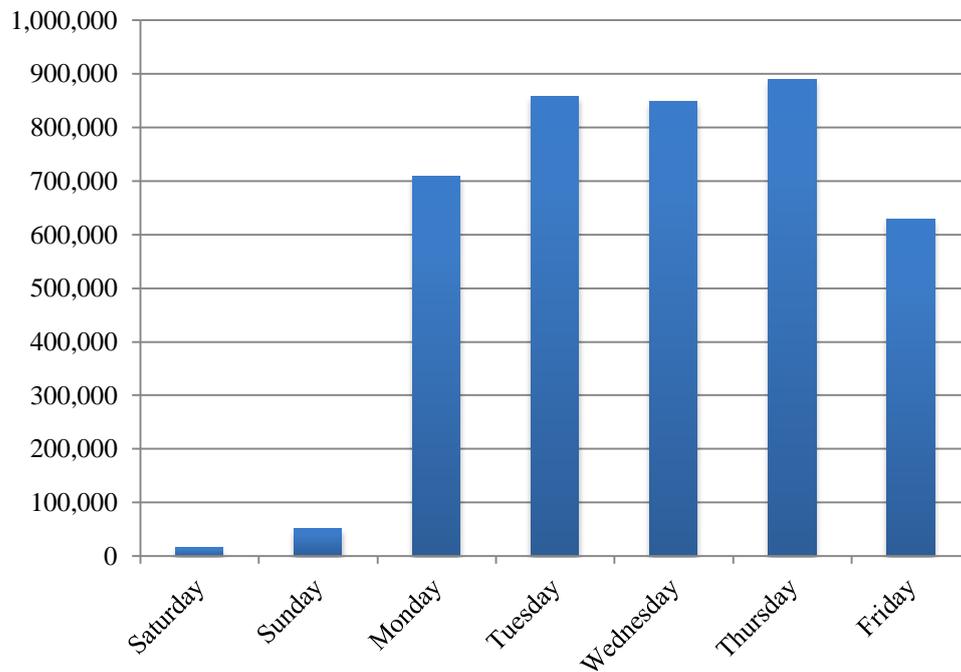

Table 1. News counts by month from 1 January 2003–31 October 2012. For every month during January 2001–October 2012, we count the number of news stories by binning the timestamps into the monthly bins, and we aggregate the counts by months of the year (Std. Dev. = standard deviation).

	Sum of Count	Average Monthly Count	Std. Dev. of Monthly Count	Average Daily Count	Std. Dev. of Daily Count
January	351,560	35,156	10,412	1,598	799
February	332,991	33,299	10,833	1,648	698
March	305,808	30,581	8,986	1,371	517
April	351,902	35,190	10,080	1,644	794
May	322,270	32,227	11,206	1,465	710
June	306,258	30,626	9,321	1,418	472
July	384,936	38,494	11,860	1,750	874
August	296,520	29,652	9,402	1,336	566
September	323,386	32,339	12,502	1,504	725
October	423,092	42,309	15,157	1,950	834
November	290,295	32,255	9,790	1,504	690
December	246,918	27,435	7,942	1,222	651

Figure 3. Percentage of news stories out of total by month, from 1 January 2003 to 30 September 2012. The occurrence of news items aggregated by month has a cyclical pattern.
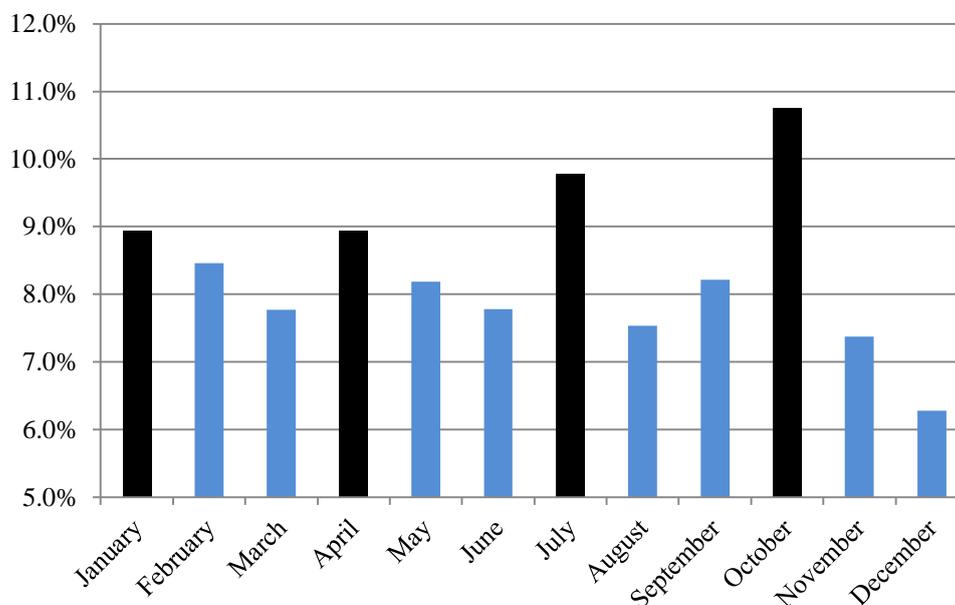
The result of the *t*-test is displayed in Table 2. The null hypothesis that the news counts from earning announcement months and regular months have the same mean is rejected at 99% confidence level. The news counts in earning announcement months are significantly greater than the counts in regular months.

Table 2. *T*-test on the same mean value for news counts in regular months and earning announcement months, assuming unequal variance.

|	Group 1 (Regular Month)	Group 2 (Earning Month)
Mean	7,020	8,890
Variance	7,383,273	11,469,717
Observations	339	175
Hypothesized Mean Difference	0	
df	292	
t-Stat	6.327	
P(T ≤ t) one-tail	4.672×10^{-10}	
t-Critical one-tail	2.345	
P(T ≤ t) two-tail	9.344×10^{-10}	
t-Critical two-tail	2.600	

5. News Intraday Patterns: "Jam-the-Close"

Various documents have found that news arrival displays distinct intraday patterns around the clock. In our data, the news flow shows a substantial increase beginning at 06:00, and continues increasing in volume until the market opens; the fact that Europe has been open prior to the opening of US markets could explain this increase. The flow is relatively stable throughout the trading hours, and

then it jumps again and peaks between 15:30 and 16:00, right before the market close. We suspect news agencies hold out on information and release them prior to the close (16:00). We identify this jump in frequency of news releases as "Jam-the-Close" and this effect was seen to be most significant during 2007–2010 as in Figure 4.

After the trading hours, the flow remains high between 16:00 and 18:00 due to company news that cannot be announced during regular trading hours for regulatory reasons. Generally speaking, the total amount of news during the trading hours and the after-hours are roughly the same, however the patterns of each are distinct. After-hour news releases are more symmetrical whereas in the normal-trading day, news release exhibits very "skewed" behavior with a massive peak in news prior to the close—"Jam-the-Close."

Figure 4. Total number of news stories by time of day. Time interval 0:00 represents 0:00 to 00:29 and interval 23:00 represents 23:00 to 23:29. Trading hours are highlighted in red. Bin 15:30 to 16:00 has the most news count.
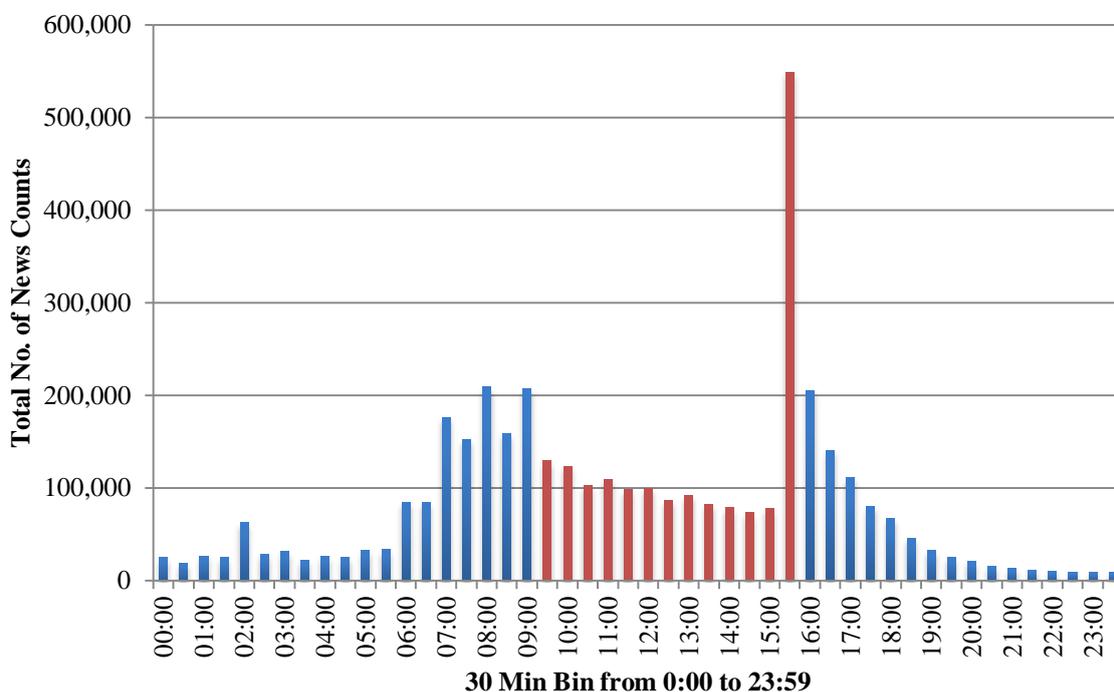
Interestingly enough, during the "Jam-the-Close" period a very distinct pattern occurs. At the time period of 15:40, 15:45, and 15:50, the proportion of market commentaries peaks significantly, which indicates that, systematically, market commentary news are released strategically, namely at 15:40, 15:45 and 15:50 please see Figure 5. Prior to those periods, the frequency of market commentary releases appears to taper, waiting to "time" the close at exactly 20, 15, and 10 minutes prior to the close at 16:00. Are such delays in affect attempting to concentrate the news impact on the closing price? Possibly, and definitely it warrants further investigations. "Jam-the-Close" maybe due to the delayed effect of market news items strategically released at time intervals prior to the closing of normal trading hours.

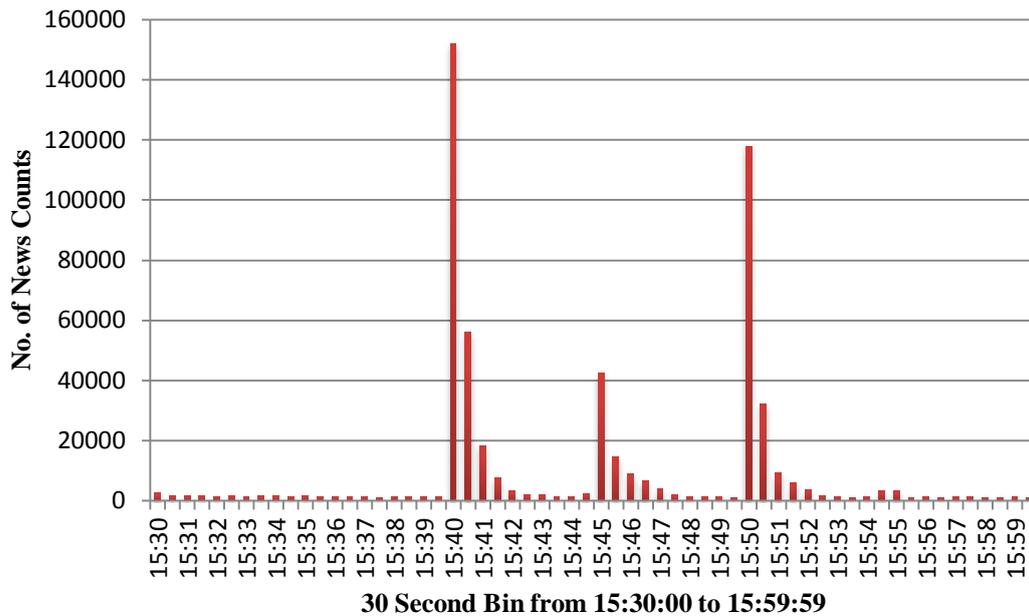**Figure 5.** Total number of news stories each 30 s from 15:30:00 to 15:59:59.

Table 3. Percentages of selected topic codes associated with news stories.

Code	Description	Definition	% in All News	% in All News between 15:30–16:00	Difference
LEN	English	Stories in English	6.87	14.49	7.62%
US	United States	United States	6.02	14.34	8.32%
RTRS	Reuters News	Reuter News	5.68	14.25	8.57%
BACT	Corporate Events	All business events relating to companies and other issuers of securities.	1.88	7.19	5.31%
STX	Equity Markets	Share market trading news, including the performance of share market indexes and individual shares.	1.21	7.17	5.96%
CORA	Corporate Analysis	Analyses about a company or group of companies.	0.46	6.80	6.34%

To better understand the nature and the content of the news stories placed between 15:30 and 16:00, we examine the distributions of associated topic codes for the stories. Table 3 compares the distributions of selected topic codes associated with all the news stories and the stories between 15:30 and 16:00. The percentages of news stories with topic codes LEN or US during 15:30–16:00 are significantly higher than the overall pool, as the US market remains very active during that time. Higher occurrence of topic code BACT and STX during 15:30–16:00 suggests that there are

significant amounts of company and listing equities news scheduled and released just prior to the close.

To better understand how the normal trading-hour news flow differs from the after-hour news, and how the market reacts to the news flow, we present a methodology to calculate a news sentiment indicator at the market level in this section. For each new story, we use (positive score − negative score) × (1 − neutral score) as the measurement for that story. The range of that measurement is [−1, 1], 1 meaning the most positive story, −1 meaning the most negative story and 0 being neutral. One company may have a few news stories on a single day. To aggregate the news flow of one company on one day, we take the arithmetic mean of the measurements from the news stories and get a sentiment score for that company on that day. The aggregation is performed at two time intervals, the regular trading hours (9:30–16:00, current day) and the after-hours (16:00–9:30, after current day to the beginning of next trading day), as well as the 24-h period.

For example, suppose stock ABC on a regular trading day (Day 1) has five news stories.

Story 1: 11:00 on Day 0, measure 0.53
Story 2: 02:00 on Day 0, measure −0.12
Story 3: 09:34 on Day 1, measure −0.25
Story 4: 15:24 on Day 1, measure −0.10
Story 5: 16:46 on Day 1, measure −0.13

Our method will produce 0.205 (= (0.53 − 0.12)/2) as the after-hour (16:00 Day 0–09:30 Day 1) sentiment score, −0.175 (=(−0.25 − 0.10)/2) as the trading-hour score (09:30 Day 1–16:00 Day 1) and the overall score 0.015 (average the first four signals) for this stock on Day 1 (16:00 Day 0–16:00 Day 1). Story 5 has no impact on the scores for Day 1. Instead it will be included in the score calculation for Day 2 (16:00 Day 1–16:00 Day 2).

After the step described above, we now have an aggregated score ranging from −1 to 1 for every ticker that gets mentioned in the news flow in a day. We count the number of the tickers that have positive or negative aggregated scores during the whole day/trading hour/after hour, across all the S&P500 stocks, and propose them as the market sentiment indicators. For each set of indicators on one time interval, we calculate the difference between the number of stocks with positive aggregated sentiment scores and the number of stocks with negative flows, and compare the results against the S&P500 index in Figure 4. The three indicators appear to experience big drops when the index drops, and the magnitude of the indicator drop is much bigger than that of the index. Moreover, the three indicators tend to move together. The all-day and market indicators are most prominent. The all-day indicator and the market-hour indicator exhibit higher volatility and correlation, while the after-hour indicator is less volatile and barely goes below zero.

We test the null hypothesis that three indicators on different time intervals are the same. To test this hypothesis, we run the linear regression on the scatter plots of the three indicators. If two indices are roughly the same on the same day, then we expect most points lie along the diagonal. The regression test shows that the all-day indicator can be represented roughly by either the market-hour indicator or the after-hour indicator. However, the market-hour and the after-hour indicators are rather different, so the researcher may want to break the sentiment analysis into trading hours and after-hours separately.

Finally, we use the Granger causality test to examine whether the knowledge of the news sentiment indicators can improve short-term forecast of the future S&P500 index movement. We perform the test on the lagged terms on the changes in the S&P500 index (close to open, open to close, close to previous close and open to previous open) and the three news indicators up to 20 days window period.

For each test, the null hypothesis is that the lagged terms of the indicator do not provide any statistically significant information about the S&P500 index change. The outputs are displayed as in Table 4. We find that the all-day and the market-hour indicators at Day t are both significant for predicting the S&P500 index movements between Close(t) to Open(t + 1), Close(t) to Close(t + 1), Open(t) to Open(t + 1). Based on the statistics, the null hypothesis that the indicators have no prediction power on the S&P500 index is rejected.

Table 4. The correlation between all-day and market-hour indicators is 0.9381. The correlation between all-day and after-hour indicators is 0.7857. The correlation between the after-hour and market-hour indicators is 0.5391.

	Slope	Interval
All-day *vs.* market-hour	1.2355	[1.1965, 1.2745]
All-day *vs.* after-hour	1.1582	[1.1393, 1.1772]
After-hour *vs.* market-hour	0.6172	[0.5738, 0.6606]

Table 5. Granger causality test on the all-day/market-hour/after-hour indicators. Test statistics that are significant at the 10% level are in bold.

	After-Hour Indicator (t + 1)	Market-Hour Indicator (t)	All-Day Indicator (t)
Close(t + 1)–Open(t + 1)	f = 0.9065	f = 0.9435	f = 0.7192
	p = 0.5789	p = 0.5305	p = 0.8099
Open(t + 1)–Close(t)	f = 0.9012	**f = 2.7722**	**f = 1.9163**
	p = 0.5858	**p = 4.0817e-05**	**p = 0.0085**
Close(t + 1)–Close(t)	f = 0.9787	**f = 2.6454**	**f = 1.8505**
	p = 0.4853	**p = 9.5325e-05**	**p = 0.0121**
Open(t + 2)–Open(t + 1)	f = 0.8815	**f = 2.6445**	**f = 1.6868**
	p = 0.6116	**p = 9.5893e-05**	**p = 0.0288**

Among the three indicators, the market-hour and the all-day indicators have more prediction power than the after-hour indicator. The after-hour indicator between 16:00 on Day t and 09:30 on day t + 1 is not significant on any daily index change. On the other hand, the market-hour indicator from the news stories between 09:30 and 16:00 on Day t and the all-day indicator between 16:00 on Day t − 1 and 16:00 on the next day predicts the index change between the close price on Day t and the open price on Day t + 1. This suggests that the effect of the news stories spills over to after-hour trading until the next opening.

To test the spillover effect of news flow, we ran the Granger causality regression between the three indicators at time t and the future S&P500 index change at time T (T > t). The test is performed with increasing T on all the three indicators until the significance level moved beyond 10%. As shown in Tables 5 and 6, both the market-hour and the all-day indicators are significant on predicting the index

daily changes up to one calendar week. This suggested a possible weighted scheme including both the current news events and the past news stories when traders wish to use this news flow as part of their trading signals.

Table 6. Granger causality test on the all-day/market-hour/after-hour indicators at time t and the S&P500 index daily change on day T (T > t). Test statistics that are significant at the 10% level are in bold.

	After-Hour Indicator (t + 1)	MarketHhour Indicator (t)	All-Day Indicator (t)
Close(t + 1)–Close(t)	f = 0.9787	**f = 2.6454**	**f = 1.8505**
	p = 0.4853	**p = 9.5325e-05**	**p = 0.0121**
Close(t + 2)–Close(t + 1)	f = 0.9352	**f = 2.6369**	**f = 1.8572**
	p = 0.5413	**p = 1.0086e-04**	**p = 0.0117**
Close(t + 3)–Close(t + 2)	f = 1.0073	**f = 2.6761**	**f = 1.8393**
	p = 0.4493	**p = 7.7735e-05**	**p = 0.0129**
Close(t + 4)–Close(t + 3)	f = 0.9514	**f = 2.6832**	**f = 1.7314**
	p = 0.5203	**p = 7.4194e-05**	**p = 0.0229**
Close(t + 5)–Close(t + 4)	f = 1.0391	**f = 1.5611**	**f = 1.5672**
	p = 0.4108	**p = 0.0534**	**p = 0.0519**
Close(t + 6)–Close(t + 5)	f = 0.9573	**f = 1.6144**	**f = 1.7635**
	p = 0.5126	**p = 0.0413**	**p = 0.0621**
Close(t + 7)–Close(t + 6)	f = 1.0715	**f = 1.4810**	**f = 1.4980**
	p = 0.3732	**p = 0.0775**	**p = 0.0717**
Close(t + 8)–Close(t + 7)	f = 1.0702	f = 1.2042	f = 1.1816
	p = 0.3746	p = 0.2401	p = 0.2601
Close(t + 10)–Close(t + 9)	f = 1.0347	f = 0.9950	f = 0.9835
	p = 0.4161	p = 0.4647	p = 0.4792

6. Conclusions

News has been known to drive securities prices since the days of the formulation of the Efficient Market Hypothesis. As technology advances and the ability of converting news stories into sentiment data continues, the call for researching into this exciting realm of finance beckons us all to join the fray. In this work, we begin our investigations into news sentiment data by examining sentiment data, limiting our investigation to a sample of constituents of the S&P500.

We made some fascinating discoveries. First, we document that there is a distinct behavior of sentiment data around earnings cycles. Not surprisingly as companies release their "numbers", news releases capture such information and quickly disseminate it across the market place. Next, we uncover the distinct behavior of sentiment data across after-hours *vs*. normal trading-hours trading. "Jam-the-Close" becomes increasingly important as we learn more about the relationship between sentiment and future market prices. In this work, we attempt to scratch the surface and peak the interest of those wanting to exploit potential predictability of news sentiment. As we continue our research into news sentiment data, we hope to inspire other young researchers to join us in our quest of uncovering

new discoveries and building creative stories that, we hope, may hold up to the scrutiny of other scientists and the test of time.

Author Contributions

Both authors contributed to all aspects of this work.

Conflicts of Interest

The authors declare no conflict of interest.

References

1. Grossman, S.J.; Stiglitz, J.E. On the impossibility of informationally efficient markets. *Am. Econ. Rev.* **1980**, *70*, 393–408.
2. Stein, J. Presidential address: Sophisticated investors and market efficiency. *J. Financ.* **2009**, *64*, 1517–1548.
3. Jegadeesh, N.; Titman, S. Returns to buying winners and selling losers: Implications for stock market efficiency. *J. Financ.* **1993**, *48*, 65–91.
4. Rouwenhorst, K.G. International momentum strategies. *J. Financ.* **1998**, *53*, 267–284.
5. Liew, J.; Vassalou, M. Can book-to-market, size, and momentum be risk factors that predict economic growth? *J. Financ. Econ.* **2000**, *57*, 221–245.
6. Asness, C.S.; Moskowitz, T.J.; Pedersen, L.H. *Value and Momentum Everywhere*; SSRN Working Paper Series; Social Science Electronic Publishing, Inc.: Rochester, NY, USA, 2013; doi:10.2139/ssrn.1363476.
7. Sinha, N.R. *Underreaction to News in the US Stock Market*; SSRN Working Paper Series; SSRN Working Paper Series; Social Science Electronic Publishing, Inc.: Rochester, NY, USA, 2010; doi:10.2139/ssrn.1572614.
8. Banz, R. The relation between return and market value of stocks. *J. Financ. Econ.* **1981**, *38*, 269–296.
9. Kyle, A.S. Continuous auctions and insider trading. *Econometrica* **1985**, *53*, 1315–1335.
10. Fama, E.F.; French, K.R. Dissecting anomalies. *J. Financ.* **2008**, *63*, 1653–1678.
11. Ball, R.; Brown, P. An empirical evaluation of accounting income numbers. *J. Account. Res.* **1968**, *6*, 159–178.
12. Fama, E.F. Market efficiency, long-term returns, and behavioral finance. *J. Financ. Econ.* **1998**, *49*, 283–306.
13. Rosenberg, B.; Reid, K.; Lanstein, R. Persuasive evidence of market inefficiency. *J. Portf. Manag.* **1985**, *11*, 9–17.
14. Lakonishok, J.; Shleifer, A.; Vishny, R.W. Contrarian investment, extrapolation, and risk. *J. Financ.* **1994**, *49*, 1541–1578.
15. Fama, E.F.; French, K.R. Multifactor explanations of asset pricing anomalies. *J. Financ.* **1996**, *51*, 55–84.

16. Fama, E.F.; French, K.R. Common risk factors in the returns on stocks and bonds. *J. Financ. Econ.* **1993**, *33*, 3–56.

17. Sharpe, W.G. Capital Asset Pricing: A theory of market equilibrium under conditions of risk multifactor explanations of asset pricing anomalies. *J. Financ.* **1964**, *19*, 425–462.

18. Lintner, J. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Rev. Econ. Stat.* **1995**, *47*, 13–37.

19. Carhart, M. On persistence in mutual fund performance. *J. Financ.* **1997**, *52*, 57–82.

20. Asness, C.S. Variables that Explain Stock Returns. Ph.D. Thesis, University of Chicago, Chicago, IL, USA, 1994.

21. Akbas, F.; Kocatulum, E.; Sorescu, S.M. *Mispricing Following Public News: Overreaction for Losers, Underreaction for Winners*; SSRN Working Paper Series; Social Science Electronic Publishing, Inc.: Rochester, NY, USA, 2008; doi:10.2139/ssrn.1107690.

22. Chan, W. Stock price reaction to news and no-news: Drift and reversal after headlines. *J. Financ. Econ.* **2003**, *70*, 223–260.

23. Hong, H.; Lim, T.; Stein, J.C. Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies. *J. Financ.* **2000**, *55*, 265–295.

24. Hong, H.; Stein, J.C. A unified theory of underreaction, momentum trading and overreaction in asset markets. *J. Financ.* **1999**, *54*, 2143–2184.

25. Mitchell, M.; Mulherin, J. The impact of public information on the stock market. *J. Financ.* **1994**, *49*, 923–950.

26. Mitchell, M.; Stafford, E. Managerial decisions and long-term stock price performance. *J. Bus.* **2000**, *73*, 287–329.

27. Mullainathan, S.; Shleifer, A. The market for news. *Am. Econ. Rev.* **2005**, *95*, 393–408.

28. Solomon, D. Selective publicity and stock prices. *J. Financ.* **2012**, *67*, 599–638.

29. Tetlock, P.C. Giving content to investor sentiment: The role of media in the stock market. *J. Financ.* **2007**, *62*, 1139–1168.

30. Tetlock, P.C.; Saar-Tsechansky, M.; Macskassy, S. More than words: Quantifying language to measure firms' fundamentals. *J. Financ.* **2008**, *63*, 1437–1467.