How News Events Impact Market Sentiment

Contact Info:
support@ravenpack.com
+1 (646) 216-2140

Terms of Use

This White Paper is not intended for trading purposes. The White Paper is not appropriate for the purposes of making a decision to carry out a transaction or trade. Nor does it provide any form of advice (investment, tax, legal) amounting to investment advice, or make any recommendations regarding particular financial instruments, investments or products. RavenPack may discontinue or change the White Paper content at any time, without notice. RavenPack does not guarantee or warrant the accuracy, completeness or timeliness of the White Paper

You may not post any content from this White Paper to forums, websites, newsgroups, mail lists, electronic bulletin boards, or other services, without the prior written consent of RavenPack. To request consent for this and other matters, you may contact RavenPack at support@ravenpack.com.

THE WHITE PAPER IS PROVIDED “AS IS”, WITHOUT ANY WARRANTIES. RAVENPACK AND ITS AFFILIATES, AGENTS AND LICENSORS CANNOT AND DO NOT WARRANT THE ACCURACY, COMPLETENESS, CURRENTNESS, TIMELINESS, NONINFRINGEMENT, TITLE, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OF THE WHITE PAPER, AND RAVENPACK HEREBY DISCLAIMS ANY SUCH EXPRESS OR IMPLIED WARRANTIES. NEITHER RAVENPACK NOR ANY OF ITS AFFILIATES, AGENTS OR LICENSORS SHALL BE LIABLE TO YOU OR ANYONE ELSE FOR ANY LOSS OR INJURY, OTHER THAN DEATH OR PERSONAL INJURY RESULTING DIRECTLY FROM USE OF THE WHITE PAPER, CAUSED IN WHOLE OR PART BY ITS NEGLIGENCE OR CONTINGENCIES BEYOND ITS CONTROL IN PROCURING, COMPILING, INTERPRETING, REPORTING OR DELIVERING THE WHITE PAPER. IN NO EVENT WILL RAVENPACK, ITS AFFILIATES, AGENTS OR LICENSORS BE LIABLE TO YOU OR ANYONE ELSE FOR ANY DECISION MADE OR ACTION TAKEN BY YOU IN RELIANCE ON SUCH WHITE PAPER. RAVENPACK AND ITS AFFILIATES, AGENTS AND LICENSORS SHALL NOT BE LIABLE TO YOU OR ANYONE ELSE FOR ANY DAMAGES (INCLUDING, WITHOUT LIMITATION, CONSEQUENTIAL, SPECIAL, INCIDENTAL, INDIRECT, OR SIMILAR DAMAGES), OTHER THAN DIRECT DAMAGES, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL THE LIABILITY OF RAVENPACK, ITS AFFILIATES, AGENTS AND LICENSORS ARISING OUT OF ANY CLAIM RELATED TO THIS AGREEMENT EXCEED THE AGGREGATE AMOUNT PAID BY YOU FOR THE WHITE PAPER IN THE 12 MONTHS IMMEDIATELY PRECEDING THE EVENT GIVING RISE TO SUCH CLAIM. BECAUSE SOME STATES OR JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF LIABILITY FOR DAMAGES OR THE EXCLUSION OF CERTAIN TYPES OF WARRANTIES, PARTS OR ALL OF THE ABOVE LIMITATION MAY NOT APPLY TO YOU.

These Terms of Use, your rights and obligations, and all actions contemplated by these Terms of Use will be governed by the laws of New York, NY, USA and You and RavenPack agree to submit to the exclusive jurisdiction of the New York Courts. If any provision in these Terms of Use is invalid or unenforceable under applicable law, the remaining provisions will continue in full force and effect, and the invalid or unenforceable provision will be deemed superseded by a valid, enforceable provision that most closely matches the intent of the original provision.

How News Events Impact Market Sentiment

Peter Ager Hafez*

May 3, 2010

Abstract

News sentiment is shown to outperform one-month price momentum when predicting future returns of the S&P500. Market and industry-level sentiment indexes are constructed based on a bottom-up approach considering the impact of company-specific news events and their corresponding sentiment. As part of constructing the indexes, I show that company relevance and event novelty are important elements of a news-based strategy, since including only the most relevant and novel news stories result in improved Information Ratios. From May 2005 through December 2009, the strategies tested deliver double digit positive returns in out-of-sample testing. In addition, I show how industry sentiment can add value when constructing market-neutral strategies taking long and short positions in top-ranked and bottom-ranked industries, respectively. Finally, I show that targeted directional exposures to top-ranked and bottom-ranked industries can improve a trading strategy beyond simple S&P500 index exposures.

1 Introduction

It is broadly accepted that financial news moves stock prices either through a direct impact on a company's expected future cash flow, the discount factor that one uses, or through more behavioral or sentiment-based mechanisms. Even though news-based trading has a long history of being part of investment decision-making, only in recent years has it been possible to test "quantitatively" the impact of news events on individual stock prices or markets. Extensive academic and industry research has shown that news, particularly stories conveying sentiment, can add value in both high and low frequency trading and investment strategies - improving the prediction of price direction, volatility and trading volume.

In a previous study, I made the case for applying a news sentiment index in predicting future returns of the Dow Jones Industrial Average, and showed that taking this approach significantly outperformed a price momentum strategy [Hafez, 2009a]. Tetlock looks at a regression model and finds that ten-day reversals are reduced following company specific news, which indicates that public news is a proxy for information that has not yet been incorporated into the stock price [Tetlock, 2009]. Engelberg et.al. find that short sellers trades are more than twice as profitable in the presence of recent news, which provides strong evidence in favor of the idea that news presents profitable trading opportunities for skilled information processors [Engelberg et al., 2010]. Mitra et.al. included news sentiment as part of the construction of forward-looking covariance matrices [Mitra et al., 2008]. Interestingly, they found that sentiment could add value to the volatility prediction process beyond what could be captured by option-implied volatility. Also, Zhang and Skiena show that news is significantly correlated with both trading volume and stock returns [Zhang & Skiena, 2010].

While studies show the impact of scheduled news events can be measured in milliseconds, signals from unscheduled news events, can be measured in minutes, days, weeks, and months. For example, the intra-day abnormal return impact of positive and negative sentiment events

*RavenPack International S.L.

can be measured in minutes and hours when looking at intra-day abnormal returns (see [Hafez, 2009c]). Focusing on longer horizons, Cahan et.al. found that the effect could be measured in days and weeks [Cahan et al., 2009a, Cahan et al., 2009b]. In addition, using a one year investment horizon, Kittrell found value in using net sentiment as a measure for long-term stock selection [Kittrell, 2010].

Applying structured news data or *News Analytics* in a trading model allows for the possibility to not only react in real-time to scheduled and unscheduled news events in a fully or semi-automated fashion, but also to consider the prevailing sentiment trend on a given market. Such trends can be captured by looking at aggregated news sentiment on single companies, sectors, industries or even on broader equity portfolios. As part of previous research, a methodology was presented on how to construct market and sector sentiment indexes that were used as part of a directional sector-rotation type strategy [Hafez, 2009d]. To address news flow seasonality, the indexes were based on counts of positive vs. negative sentiment news stories that were considered to be highly relevant to one of the index constituents. As part of the study, I find that considering a Company Relevance metric is an important element in constructing sentiment-based strategies as the out-of-sample return correlation improves by a factor of 3 after filtering for relevance. In this study, I take relevance filtering a step further and include only news that is contextually relevant to the companies in the S&P500. That is, where a company has been detected to be playing a prominent role in the news story and has been involved in some type of categorized event (e.g. earnings announcement, analysts rating, product recall, etc.), and therefore has received a relevance score of 100. For more information on relevance, see Appendix B. Furthermore, I consider how it may be desirable to treat stories differently in terms of sentiment impact depending on the detected event category. That is, a bankruptcy story should count more towards a sentiment score than a story about a product or marketing campaign. Finally, I consider how event novelty may influence the construction of sentiment indexes, where novelty in this case represents how "new" or novel a news story is over a 24 hour time window.

Generally, I find that considering the impact of different company events adds value when constructing market-level sentiment indexes. For industry-level indexes, I noticed that the total number of company-specific events varied depending on the industry. In order to improve the confidence around the sentiment estimates, I apply a slightly less restrictive relevance score moving from 100% to 90% relevant. This permitted the use of other sentiment analytics available from RavenPack which provide more information by examining various aspects of each story (i.e. events, language tone, story type). Here I consult 5 different sentiment scores that classify each news story as being either positive, negative or neutral. The same approach was considered in a previous study [Hafez, 2009d]. Rather than normalizing only for news flow, I consider a normalization for changes in the event category characteristics, which seems to bring further value in the sentiment ranking of industries.

The study proceeds as follows. In Section 2, I provide an overview of the methodology on how to construct market-level sentiment indexes considering an Event Sentiment Score. Furthermore, I consider a simple trading strategy based on a US Market-level Sentiment Index. In Section 3, I describe how to construct industry sentiment indexes based on aggregated news sentiment. Using an industry rank, I first consider a simple market-neutral strategy, followed by a targeted directional strategy based on industry rather than broad market exposures. Finally, in Section 4, I present the conclusion of the study.

2 Market-Level Sentiment

Generally, news sentiment indexes try to capture the prevailing sentiment trend for a particular market or sector based on news information. In order to capture such trends, it seems reasonable to consider an aggregation of news sentiment over well-defined moving time intervals to capture the general "mood" of the market. News sentiment indexes have been useful when constructing simple investment strategies that consistently outperform similar strategies based on price-momentum [Hafez, 2009a]. Previous results have shown to be resistant to different sentiment aggregation windows, investment horizons, and for different investment timing.

2.1 Data & News Analytics

In order to measure sentiment for a particular equity index, I use news analytics data from RavenPack going back to 2005. The data set includes tens of thousands of records per day each representing a company reference in a financial news story. Currently, RavenPack tracks around 27 thousand companies globally, which represent more than 98% of the investable global market. Each record comes with a millisecond timestamp and data for sentiment, novelty, relevance, event categories, among other news analytics. One of the advantages with RavenPack's news analytics is that the data is free of survivorship bias. That is, each company is identified systematically using its respective point-in-time ticker symbols and/or other company identifiers or aliases, and both "dead" and "survivor" companies are included in the data set.

Whenever, RavenPack is able to detect one of more than 160 company related event categories as well as the role a company plays in a news story, these elements are tagged as part of the company-specific news record. For example, in a news story with the headline "IBM Completes Acquisition of Telelogic AB", the event would be identified as "acquisition-acquirer" since IBM is involved in an acquisition and is the acquiring company. Telelogic would receive the "acquisition-acquiree" tag in its corresponding record since the company is also involved in the acquisition, but as the acquired company. This applies to other events like lawsuits where it makes sense to differentiate between the "plaintiff" and the "defendant". Another example, "Toyota Files Voluntary Safety Recall on Select Toyota Division Vehicles for Sticking Accelerator Pedal" is tagged as a "product-recall" since Toyota is involved in a product recall.

Previous studies indicate that *spill-over effects* are present between company-specific news events and sector index price moves. Cahan et al. find that the sector excess returns following company-specific news events are smaller than the market excess return indicating that the sector also moves on the news event[1] [Cahan et al., 2009a]. Patton and Verardo find that news releases have an important impact on the risk and covariance of stocks, which suggests that there is contagion in the information content of news releases. In other words, new information for one stock impacts the trading of other stocks, causing the stocks to move together a little more than might be expected under normal conditions [Patton & Verardo, 2009]. Taking this into consideration, it seems reasonable to expect that such spill-over effects are also present at the market-level.

For the construction of market-level news sentiment indexes I use RavenPack's Event Sentiment Score (ESS), which indicates how event categories are typically rated by financial experts as having positive or negative share price impact. To capture news events specifically related to S&P500 companies, I use the RavenPack Company Relevance Score. This metric provides a way to capture stories that are actually relevant to the S&P500 constituents and not mere mentions in the text. The numerical score indicates "how" relevant the story is to the company and assigns higher values based on the context of the news using semantic analysis. In a previous study, I found that only 20% of all news records are relevant; hence, 80% could simply be adding noise[2] [Hafez, 2009a]. In many cases, companies are mentioned in passing and are not the central theme of the story. Filtering based on Company Rel-

[1] This seems to be more pronounced for negative than for positive sentiment events.
[2] At least 73% of stories contain a one highly relevant company.

evance, ensures that only records that have been categorized as being strongly related to one of the companies belonging to the universe of stocks are being considered. It should be mentioned that companies not detected as explicitly mentioned in a story are not given a relevance score. While a story about Yahoo! might be considered in some other context to be relevant to Google, the company Google will not be given a relevance score unless that story explicitly mentions Google. Finally, I use an Event Novelty Score (ENS), which represents how "new" or novel a news story is over a 24 hour time window. The first story disclosing an event about a company is considered to be the most novel and receives a score of 100. In Appendix A, I have included further information on Company Relevance, ESS, and ENS. CSI Market Data is the source for corporate action-adjusted pricing data.

2.2 Market-Level Index Calculation

Having described what elements are to be considered when constructing market-level news sentiment indexes, it is now possible to describe the methodology in more detail.

Let $\mathcal{N}$ denote the universe of all news stories from the RavenPack data set. Fix a company C that is mentioned within some news story from $\mathcal{N}$ with $E_C(N)$ and $D_C(N)$ representing the Event Sentiment Score and Event Novelty Score of company C for story $N \in \mathcal{N}$; respectively.

Definition 2.1 Let $\mathcal{U}$ be a universe of companies, and C be a company such that $C \in \mathcal{U}$. Let p be a time period denoting a certain number of days. Let P_N be the stories published within p days before publication of news story N up to and including N such that $\forall N_i \in P_N$ with $E_C(N) \neq \emptyset$ and $D_C(N) = 100$. In other words, every story in P_N has an Event Sentiment Score for some $C \in \mathcal{U}$ and an Event Novelty Score of 100. Finally, let $m = |P_N|$. The trailing *sentiment index*, $I_{\mathcal{U}}(N, E)$, for $\mathcal{U}$ is the quantity

$$I_{\mathcal{U}}(N, E) \quad = \quad \frac{1}{m} \sum_{i=1}^{m} E_C(N_i). \tag{1}$$

Remark 2.1 The *universe of companies* is meant here to represent the constituents of a broader equity index i.e. the S&P500 or Russell1000.

Considering the constituents of the S&P500, I construct a US Market-Level Sentiment Index applying a 90 day trailing window ($P = 90$). One of the advantages of aggregating over such period is that I capture an entire "quarterly season" in each trailing window, thereby ensuring that similar news flow characteristics are represented. Equity news flow is very much characterized by seasonality, where a quarterly pattern is evident and likely caused by the repeated earnings reporting season [Hafez, 2009b]. Figure 1 depicts the US Market-Level Sentiment Index vs. the S&P500 cumulative index log-returns covering the period March 2005 through December 2009.

Figure 1: *Market-Level Sentiment Index (**solid**-line) - primary-axis; vs. the S&P500 cumulative log-returns (**dashed**-line) - secondary-axis. The sentiment index has been constructed based on the average Event Sentiment Score of all S&P500 companies over a 90 day trailing window covering the period January 2005 through December 2009, and has been scaled, for visualization purposes, to take values between 0 and 1. In addition, events have been filtered only to include the most novel stories ($ENS = 100$).*

2.3 Strategy & Empirical Results

Under the assumption that market returns are likely to move in the same direction as market sentiment, I base my trading decision on the sentiment index delta, $\triangle_t$. Focusing on the index delta will capture the sentiment of the most recent period (i.e. one month), but also include the sentiment change from the previous period which may have similar characteristics (i.e. as captured by the "same" month in the earnings season cycle).

Definition 2.2 Let I_t be the trailing *sentiment index* value at time t, then the *sentiment index delta*, $\triangle_t$, at time t is the quantity

$$\triangle_t = I_t - I_{t-1}. \tag{2}$$

In order to construct a simple news-based strategy, whenever $\triangle_t > 0$ I take a long position in the S&P500 in the following period. Likewise, when $\triangle_t < 0$ I take a short position. More specifically, I decide at the end of each month the direction to take in the S&P500 in the following month. The S&P500 index returns are calculated as the monthly *close-to-close* log-returns.

In Figure 2, I include the cumulative return of a trading strategy based on the US Market Sentiment Index with and without applying an Event Novelty Score filter. Furthermore, I include the one-month price-momentum strategy return for benchmark purposes.

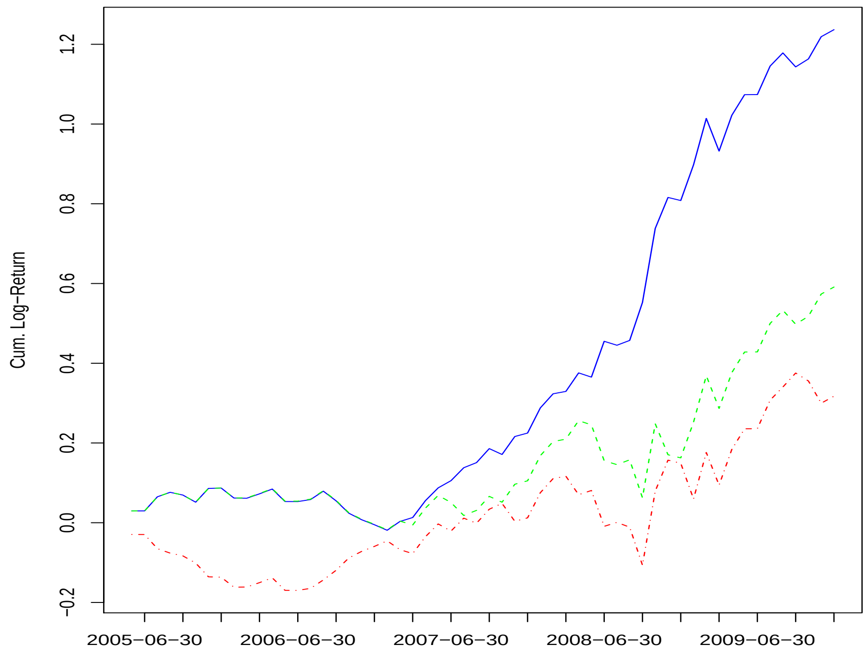

Figure 2: *Cumulative strategy returns covering the out-of-sample period May 2005 through December 2009 for the Market-Level Sentiment Index (**solid**-line) with an Event Novelty Score filter, the Market-Level Sentiment Index (**dashed**-line) without an Event Novelty Score filter, and the one-month price-momentum strategy based on the S&P500 (**dotdashed**-line). The sentiment indexes have been constructed based on a 90 day trailing window, and a trading decision was made based on the monthly index delta. A positive delta resulted in a long position, and a negative delta in short position.*

As can be observed in Figure 2, not only do both sentiment-based strategies outperform one-month price-momentum, but filtering based on event novelty seems to add significant value in predicting the future price direction of the S&P500, as given by an improvement of the Information Ratio of more than 100%. In Table 1, I have included a performance summary for the different strategies. The Event Novelty filtered strategy would have obtained an overall Information Ratio[3] of 1.75 over the period with the values 1.02 and 2.47 pre- and post the *market high* of the test period in October 2007, respectively. Overall, the event novelty filtered strategy would have realized an annualized return of 26.5% with a Hit Ratio[4] of almost 70%. Interestingly, both sentiment indexes not only deliver significantly better returns than price momentum, but do so with lower volatilities. Also, a significant improvement can be observed in the Hit Ratios.

[3]The Information Ratio is calculated as the annualized return divided by the annualized volatility.
[4]The Hit Ratio represents the percentage of months that were profitable over the period.

		Sentiment Index (ENS=100)	Sentiment Index	Price Momentum (one-month)
Information Ratio	Total	1.75	0.76	0.40
Information Ratio	Pre-Oct.	1.02	0.35	0.18
Information Ratio	Post-Oct.	2.47	1.05	0.55
Ann. Return	Total	26.50%	12.67%	6.80%
Ann. Return	Pre-Oct.	7.69%	2.74%	1.40%
Ann. Return	Post-Oct.	46.70%	23.33%	12.60%
Ann. Volatility	Total	15.16%	16.61%	16.90%
Ann. Volatility	Pre-Oct.	7.56%	7.85%	7.88%
Ann. Volatility	Post-Oct.	18.88%	22.32%	23.05%
Hit Ratio	Total	69.64%	62.50%	55.36%
Hit Ratio	Pre-Oct.	62.07%	58.62%	48.28%
Hit Ratio	Post-Oct.	77.78%	66.67%	62.96%

Table 1: *Performance statistics covering the out-of-sample period May 2005 through December 2009 for the strategies based on the Market-Level Sentiment Index (column 1) with an Event Novelty Score filter (ENS=100), the Market-Level Sentiment Index (column 2) without an Event Novelty Score filter, and the one-month price-momentum strategy based on the S&P500 (column 3). Pre- and Post Oct. refer to the market high in October 2007.*

Considering the per-year annualized return of the different strategies from Table 2, the novelty-filtered sentiment-based strategy delivers double digit positive returns in all years except 2006, which was the only year that resulted in a loss. Interestingly, price momentum outperforms the sentiment-based strategies only in 2006.

	Sentiment Index (ENS=100)	Sentiment Index	Price Momentum (one-month)
2005	13.05%	13.05%	-20.52%
2006	-9.20%	-9.09%	7.68%
2007	22.99%	10.91%	7.20%
2008	58.32%	5.75%	13.68%
2009	42.85%	42.85%	16.80%

Table 2: *Yearly annualized return covering the out-of-sample period May 2005 through December 2009 for the strategies based on the Market-Level Sentiment Index (column 1) with an Event Novelty Score filter (ENS=100), the Market-Level Sentiment Index (column 2) without an Event Novelty Score filter, and the one-month price-momentum strategy based on the S&P500 (column 3).*

Overall, it seems that the construction of bottom-up sentiment indexes outperform price momentum in predicting future returns of the S&P500. In addition, I find that Event Novelty is an important element of a news-based strategy, since including only the most novel news stories result in significantly improved Information Ratios both pre- and post- the market high in October 2007. Finally, considering the impact of different company events add value to the construction of market-level news sentiment indexes.

3 Industry-Level Sentiment

In the previous section, I demonstrated how to construct a market-level news sentiment index that can be used to predict the future direction of the S&P500. Here, I will focus on constructing market-neutral strategies taking long and short positions in the top- and bottom-ranked industries, respectively. Being able to successfully rank industries based on sentiment should enhance the performance of a strategy, since more targeted investments can be made towards expected out- and under-performers during periods with long and short exposures, respectively.

3.1 Data & News Analytics

In order to construct industry-level sentiment indexes based on news, I use five sentiment analytics available from RavenPack. Each of these analytics has been calculated using a different linguistic technique. For example, some analytics are based on keyword and phrase detection, optimized to capture key financial language. Other analytics are derived using classifiers or algorithms trained to emulate how financial experts would react to different types of news e.g. earnings and announcements, editorial and commentary, corporate actions, or stories about mergers and acquisitions. For more information on the different RavenPack classifiers, see Appendix B.

To capture news events specifically related to S&P500 companies, I also use the RavenPack Company Relevance Score. Finally, as with the market-level indexes, CSI Market Data is the source for corporate action adjusted pricing data.

3.2 Industry-Level Index Calculation

When constructing industry-level indexes, I noticed that the total number of company-specific events varied depending on the industry. In order to improve the confidence around the sentiment estimates, I apply a slightly less restrictive Company Relevance Score moving from 100% to 90% relevant. This also permits the use of other sentiment analytics available from RavenPack which provide more information by examining various aspects of each story (i.e. events, language tone, story type). Here I consult five different sentiment scores that classify each news story as being either positive, negative or neutral; and thereby evaluate the changing relationship between the count of positive and negative sentiment stories. Previous studies shows that consulting multiple classifiers to determine the sentiment of a given story or event can add significant value when trying to predict stock price direction [Cahan et al., 2009a].

Let $\mathcal{N}$ denote the universe of all news stories from the RavenPack data set. Fix a company C that is mentioned within some news story from $\mathcal{N}$ with sentiment analytics $q_i \in q$, where $i \in \{1, \ldots, 5\}$ represents each of the five sentiment analytics.

Definition 3.1 Call the function $S_C : \mathcal{N} \to \{-1, 0, 1\}$, the story sentiment indicator relative to company C where

$$S_C(N) = \begin{cases} -1, & \text{if } C \text{ receives a score } Q(N) < 50; \\ 0, & \text{if } C \text{ receives a score } Q(N) = 50; \\ 1, & \text{if } C \text{ receives a score } Q(N) > 50, \end{cases} \tag{3}$$

with

$$Q(N) = \text{avg}(q) \tag{4}$$

When I say that a news story $N \in \mathcal{N}$ is about the company C, it is assumed that N has Relevance≥ 90 for C.

Remark 3.1 A relevance value of 90 indicates that the company is referenced in the main title or headline of the news story. A company will be assigned a high mark of 100 if it plays a main role in these types of stories (context-aware).

Having classified all stories for the targeted universe of stocks as being either positive, negative, or neutral; I'm now able to define the Sentiment Ratio.

Definition 3.2 Let $\mathcal{U}$ be a universe of companies, and C be a company such that $C \in \mathcal{U}$. Let p be a time period denoting a certain number of days. Let P_N be the stories published within p days before publication of news story N up to and including N such that $\forall N_i \in P_N$ with Relevance$\geq$ 90 for some $C \in \mathcal{U}$. In other words, every story in P_N is about some company $C \in \mathcal{U}$. Finally, let $n = |P_N|$. The trailing *sentiment ratio*, $R_{\mathcal{U}}(N, S)$, for $\mathcal{U}$ is the quantity

$$R_{\mathcal{U}}(N, S) \quad = \quad \frac{\sum_{i=1}^{n} S_C(N_i)_{[S_C(N_i)=1]}}{\sum_{i=1}^{n} |S_C(N_i)_{[S_C(N_i)=-1]}|}. \tag{5}$$

In order to compare and rank industries according to sentiment, the Sentiment Ratio needs to be normalized expecting that the levels or changes in the Sentiment Ratio may depend on the industry. Therefore, I suggest performing some adjustment that considers the standard deviation of the industry-specific Sentiment Ratio. Such normalization could depend on a trailing volatility measure, or on the empirical distribution of the Sentiment Ratio of the entire back-testing period. In order to center the normalized values, I map the mean of the empirical distribution into a Sentiment Index value of 50, and apply a stepwise linear mapping of the remaining values based on standard deviations.

Let B_i be some mapping function that depends on σ, where the family of B are the number of standard deviations away from the mean. Furthermore, let Y_i be a function of B_i such that more extreme Sentiment Ratios, in terms of distance to the mean, will receive sentiment index values closer to either 0 or 100 depending on whether the Sentiment Ratio is below or above its mean value, respectively.

Definition 3.3 Let σ be the standard deviation of $R_{\mathcal{U}}$ over the entire back-testing period. Fix $B_i(\sigma)$ for $i \in \{0, \ldots, k\}$ the cut-off points for a mapping function, and let $Y(B_i)$ be the mapping score at the cut-off B_i with $Y(B_0) = 0$ and $Y(B_k) = 100$, then the *sentiment index value* for $\mathcal{U}$ at N will be the quantity

$$V_{\mathcal{U}}(R) = \begin{cases} Y(B_0), & \text{if } R_{\mathcal{U}}(N,S) < B_0(\sigma); \\ (1 - W(R)) * Y(B_i) + W(R) * Y(B_{i+1}), & \text{if } B_0(\sigma) \leq R_{\mathcal{U}}(N,S) < B_k(\sigma); \\ Y(B_k), & \text{if } R_{\mathcal{U}}(N,S) \geq B_k(\sigma); \end{cases} \tag{6}$$

with

$$W(R) = \frac{R_{\mathcal{U}}(N,S) - B_i(\sigma)}{B_{i+1}(\sigma) - B_i(\sigma)}, \tag{7}$$

and where $B_i(\sigma) \leq R_{\mathcal{U}}(N,S) < B_{i+1}(\sigma)$ and $i \in \{0, \ldots, k-1\}$.

Based on the normalized indexes, it's possible to rank industries for instance on a monthly basis in order to construct long and short portfolios of the top- and bottom-ranked industries. In the following section, I will present an empirical study focusing on the constituents of the S&P500 equity index.

3.3 Strategy & Empirical Results

As part of constructing a set of industry sentiment indexes, I map each company belonging to the S&P500 to their corresponding industry[5]: (1) Oil & Gas, (2) Basic Materials, (3) Industrials, (4) Consumer Goods, (5) Health Care, (6) Consumer Services, (7) Telecommunications, (8) Utilities, (9) Financials, and (10) Technology. Each of these indexes are mapped to an iShares industry exchange traded fund (ETF) for testing purposes.

The indexes are constructed based on a 90 day trailing window ($P = 90$), which ensures similar news flow characteristics in each window. In order to map the calculated sentiment ratio into a sentiment index value, I consider the empirical distribution over the period March 2005 through December 2009. Finally, at each month-end I decide which industries to hold long and short in the following one-month period depending on an *industry sentiment delta rank*, see Definition 2.2.

In Figure 3, I show the cumulative log-return spread between the top 5 and bottom 5 industries according to a sentiment ranking. As can be observed, the *out-of-sample* results indicate that a positive spread can be realized taking long and short positions in the top and bottom ranked industries.

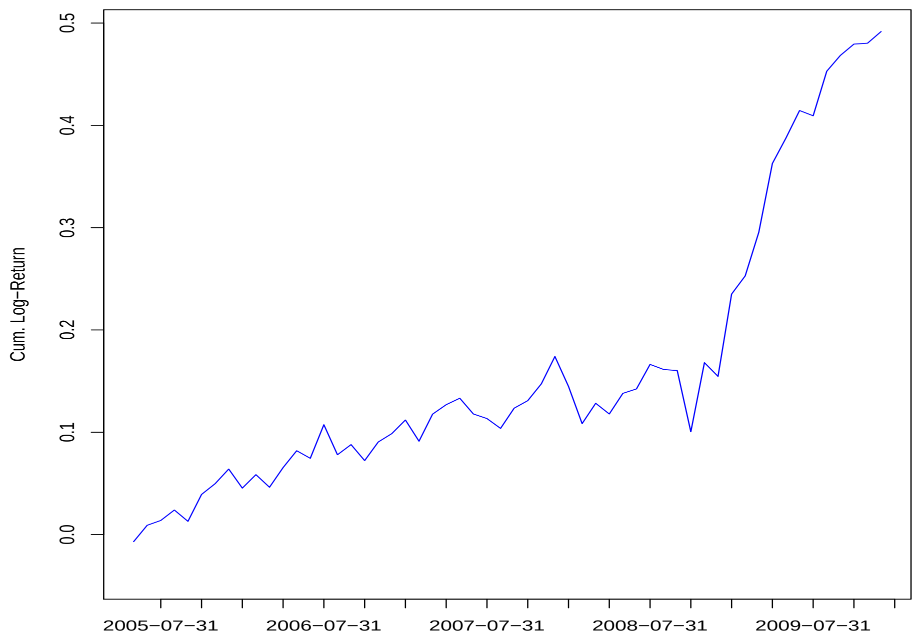

Figure 3: *Cumulative industry return spreads covering the out-of-sample period May 2005 through December 2009 for the top 5 and bottom 5 ranked industries. The sentiment indexes have been constructed based on a 90 day trailing window, and a monthly industry rank was made based on the monthly index delta.*

In Table 3, I have included a performance statistics summary for the long-short strategy on the top 5 and bottom 5 ranked industries. Over the back-testing period, the strategy yields an Information Ratio of 1.23 with values of 0.92 and 1.53 pre- and post- the market high in October 2007. In addition, the Hit Ratio reaches about 66% with positive returns in 4 out of 5 years, see Table 4.

[5]The industry mapping has been based on the ICB industry classification.

		Top/Bottom 5
Information Ratio	Total	1.23
Information Ratio	Pre-Oct.	0.92
Information Ratio	Post-Oct.	1.53
Ann. Return	Total	10.54%
Ann. Return	Pre-Oct.	5.11%
Ann. Return	Post-Oct.	16.37%
Ann. Volatility	Total	8.54%
Ann. Volatility	Pre-Oct.	5.58%
Ann. Volatility	Post-Oct.	10.73%
Hit Ratio	Total	66.07%
Hit Ratio	Pre-Oct.	62.07%
Hit Ratio	Post-Oct.	70.37%

Table 3: *Performance statistics covering the out-of-sample period May 2005 through December 2009 for the strategies based on the top 5 and bottom 5 ranked industries. Pre- and Post Oct. refer to the market high that took place in October 2007.*

	Top/Bottom 5
2005	9.60%
2006	3.46%
2007	7.54%
2008	-1.94%
2009	33.71%

Table 4: *Yearly annualized return covering the out-of-sample period May 2005 through December 2009 for the top 5 and bottom 5 ranked industries.*

Overall, it seems that going long the top-ranked and short the bottom-ranked industries based on news sentiment can add value to a market-neutral strategy.

3.4 A Directional Industry Strategy

Based on the results of the previous section, it seems reasonable to assume that taking targeted industry exposures rather than investing in a broader market index would yield improved results. Hence, instead of investing in the entire S&P500 let us consider taking long positions in the top-ranked industries when the market-level sentiment delta is positive, and short positions in the bottom-ranked industries when the market-level sentiment delta is negative. In Figure 4, I have depicted the cumulative returns of the "basic" strategy investing in the S&P500 and of the strategy taking exposures in the top 5 and bottom 5 ranked industries. As can be observed, the latter outperforms the former over the back-testing period.

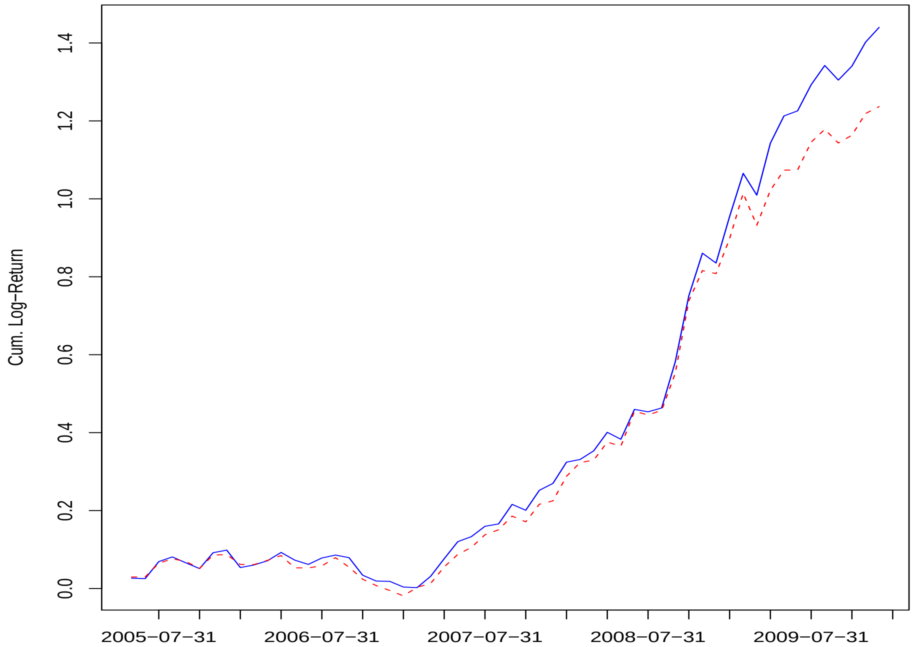

Figure 4: *Cumulative strategy returns covering the out-of-sample period May 2005 through December 2009 taking long positions in the top-ranked industries when the Market-Level Sentiment Index delta is positive and short positions in the bottom-ranked industries when the Market-Level Sentiment Index delta is negative. Depicted in the graph is a strategy based on the top 5 and bottom 5 ranked industries (**solid**-line), and on S&P500 market exposures (**dashed**-line). The sentiment indexes have been constructed based on a 90 day trailing window, and a monthly industry rank was made based on the monthly index delta.*

Considering in more detail the performance statistics depicted in Table 5, it can be observed that taking a targeted industry approach outperforms the strategy based on broad market exposures both pre- and post the market high in October 2007. Overall, the Information Ratio increases from 1.75 to 1.91 with most of the improvement obtained post the market high as indicated in the greater post-October 2007 Information Ratio. For this period, the Information Ratio has risen from 2.47 to 2.79.

		Top/Bottom 5	Market-level
Information Ratio	Total	1.91	1.75
Information Ratio	Pre-Oct.	1.03	1.02
Information Ratio	Post-Oct.	2.79	2.47
Ann. Return	Total	30.86%	26.50%
Ann. Return	Pre-Oct.	8.94%	7.69%
Ann. Return	Post-Oct.	54.40%	46.70%
Ann. Volatility	Total	16.19%	15.16%
Ann. Volatility	Pre-Oct.	8.66%	7.56%
Ann. Volatility	Post-Oct.	19.52%	18.88%
Hit Ratio	Total	67.86%	69.64%
Hit Ratio	Pre-Oct.	58.62%	62.07%
Hit Ratio	Post-Oct.	77.78%	77.78%

Table 5: *Performance statistics covering the out-of-sample period May 2005 through December 2009 taking long positions in the top-ranked industries when the Market-Level Sentiment Index delta is positive and short positions in the bottom-ranked industries when the Market-Level Sentiment Index delta is negative. Depicted in the graph is a strategy based on the top and bottom 3 ranked industries (column 1), and the S&P500 (column 2). Pre- and Post Oct. refer to the market high that took place in October 2007.*

Interestingly, it can be observed from Table 6 that the targeted industry strategy outperforms the market-level strategy in terms of annualized returns in 4 out of 5 years with 2008 being the only exception. Especially, 2009 seems to stand out.

	Top/Bottom 5	Market-level
2005	14.78%	13.05%
2006	-8.02%	-9.20%
2007	25.14%	22.99%
2008	56.56%	58.32%
2009	60.47%	42.85%

Table 6: *Yearly annualized return covering the out-of-sample period May 2005 through December 2009 taking long positions in the top-ranked industries when the Market-Level Sentiment Index delta is positive and short positions in the bottom-ranked industries when the Market-Level Sentiment Index delta is negative. Depicted in the graph is a strategy based on the top 5 and bottom 5 ranked industries (column 1), and the S&P500 (column 2).*

All in all, it seems possible to enhance the market-level strategy by taking targeted long and short positions in the top-ranked and bottom-ranked industries rather than taking broader market exposures.

4 Conclusion

Considering the relevance and impact of different company events is an important element when constructing market and industry sentiment indexes. Furthermore, applying event novelty as a filter is shown to bring significant value in predicting future returns of the S&500. Especially, I find that the market-level sentiment strategies significantly outperform one-month price momentum with Information Ratios of 1.75 vs. 0.40 and annualized returns of 26.5% vs. 6.8%. In addition, the market-level sentiment strategy delivers double digit positive returns in 4 out of 5 years. In order to measure the impact of different company events, I consider RavenPack's Event Sentiment Score, which indicates how event categories are typically rated by financial experts as having positive or negative share price impact. Also, I use RavenPack's Event Novelty Score to measure how "new" or novel a news story is over a 24 hour time window. Rather than trading based on sentiment index levels, I find value in trading on deltas or monthly index changes. By aggregating news sentiment data over a trailing period of 3 months, similar news flow characteristics are represented in each window, addressing seasonality caused by quarterly earnings reporting. Beyond capturing market-level sentiment, I calculate industry-level sentiment indexes using multiple news classifiers which provide diversity and more sentiment data. I apply a slightly less restrictive company relevance criteria while still ensuring that only high relevance news stories are considered. Based on a set of industry sentiment indexes, a positive return spread can be obtained based on a market-neutral strategy taking long and short positions in the top- and bottom-ranked industries, respectively. I show that a positive spread can be realized taking long and short positions in the top 5 and bottom 5 ranked industries, thereby obtaining an Information Ratio of 1.23 over the back-testing period. Finally I show that beyond using market-level news sentiment to invest in the S&P500, it's possible to enhance a market strategy by taking long and short positions in the top and bottom-ranked industries when the general market sentiment index is positive and negative, respectively. This approach improves the Information Ratio from 1.75 to 1.91 with outperformance in 4 out of 5 years.

References

[Cahan et al., 2009a] Cahan, R., Jussa, J., & Luo, Y. (2009a). Breaking news: How to use news sentiment to pick stocks. *Macquarie US Equity Research.*

[Cahan et al., 2009b] Cahan, R., Jussa, J., & Luo, Y. (2009b). Eventful Investing: Harnessing the Power of Event-Driven Strategies. *Macquarie US Equity Research.*

[Engelberg et al., 2010] Engelberg, J. E., Reed, A. V., & Ringgenberg, M. C. (2010). How Are Shorts Informed? Short Sellers, News, And Information Processing. *Kenan-Flagler Business School, University of North Carolina.*

[Hafez, 2009a] Hafez, P. A. (2009a). Construction of market sentiment indices using news sentiment. *RavenPack International S.L.*

[Hafez, 2009b] Hafez, P. A. (2009b). Detection Of Seasonality Patterns in Equity News Flow. *RavenPack International S.L.*

[Hafez, 2009c] Hafez, P. A. (2009c). Investigation of the Impact of News Sentiment on Abnormal Stock Return. *RavenPack International S.L.*

[Hafez, 2009d] Hafez, P. A. (2009d). Sector Rotation Strategies Driven By News Sentiment Indices. *RavenPack International S.L.*

[Kittrell, 2010] Kittrell, J. (2010). Sentiment Reversals as Buy Signals. *Knightsbridge Asset Management.*

[Mitra et al., 2008] Mitra, L., Mitra, G., & diBartolomeo, D. (2008). Equity portfolio risk (volatility) estimation using market information and sentiment. *CARISMA, Brunel University.*

[Patton & Verardo, 2009] Patton, A. & Verardo, M. (2009). Does beta move with firm news? Systematic risk and firm-specific information flows. *Duke University.*

[Tetlock, 2009] Tetlock, P. C. (2009). Does Public Financial News Resolve Asymmetric Information? *Yale University.*

[Zhang & Skiena, 2010] Zhang, W. & Skiena, S. (2010). Trading Strategies To Exploit News Sentiment. *Stony Brook University.*

A Market-Level Sentiment Data

I use the following RavenPack data to construct the market-level index.

A.1 Company Relevance

The score between 0-100 that indicates how strongly related the company is to the underlying news story, with higher values indicating greater relevance. For any story that mentions a company, RavenPack provides a relevance score. A score of 0 means the company was passively mentioned once in a story. A score of 100 means the company was predominant in the story and played a well defined role in the article. The greater the score between 0-100, the higher the relevance of the story to the company.

A.2 ESS: Event Sentiment Score

A granular score between 0 and 100 that represents the news sentiment for a given company by measuring various proxies sampled from the news. The score is determined by systematically matching stories typically rated by financial experts as having short-term positive or negative share price impact. The strength of the score is derived from training sets where financial experts classified company-specific events and agreed these events convey positive or negative sentiment and to what degree. Their ratings are encapsulated in an algorithm that generates a score range between 0-100 where higher values indicate more positive sentiment while lower values below 50 show negative sentiment.

A.3 ENS: Event Novelty Score

A granular score between 0 and 100 that represents how new or novel a news story is over a 24 hour time window. The first story disclosing an event about a company is considered to be the most novel and receives a score of 100. Subsequent stories about the companys event receive lower scores following a decay function (100, 75, 56, xxxxx). Stories outside the 24 hour window but similar to a story in a chain of events receive a score of 0.

B Industry-Level Sentiment Data

I use the following RavenPack data to construct the industry-level indexes.

B.1 COMPANY RELEVANCE

The score between 0-100 that indicates how strongly related the company is to the underlying news story, with higher values indicating greater relevance. For any story that mentions a company, RavenPack provides a relevance score. A score of 0 means the company was passively mentioned once in a story. A score of 100 means the company was predominant in the story and played a well defined role in the article. The greater the score between 0-100, the higher the relevance of the story to the company.

B.2 WLE: Word & Phrase Detection

A score that represents the news sentiment of the given news item according to the WLE classifier, which specializes in identifying positive and negative words and phrases in articles about global equities. Scores can take values of 0, 50, or 100 indicating negative, neutral or positive sentiment, respectively. This sentiment score is based on RavenPack's Traditional Methodology.

B.3 PCM: Projections, Corporate News

A score that represents the news sentiment of the given story according to the PCM classifier, which specializes in identifying the sentiment of stories that are only about earnings, developments, and projections news. Scores can take values of 0, 50, or 100 indicating negative, neutral or positive sentiment, respectively. This sentiment score is based on RavenPack's Expert Consensus Methodology.

B.4 ECM: Editorials, Commentary News

A score that represents the news sentiment of the given story according to the ECM classifier, which specializes in short commentary and editorials on global equity markets. Scores can take values of 0, 50, or 100 indicating negative, neutral or positive sentiment, respectively. This sentiment score is based on RavenPack's Expert Consensus Methodology.

B.5 RCM: Reports, Corporate Actions News

A score that represents the news sentiment of the given news story according to the RCM classifier, which specializes in reports on corporate action announcements. Scores can take values of 0, 50, or 100 indicating negative, neutral or positive sentiment, respectively. This sentiment score is based on RavenPack's Expert Consensus Methodology and has been trained on stories that lead up to a pre-identified corporate action announcement.

B.6 VCM: Merger, Acquisitions, & Takeover News

A score that represents the news sentiment of the given story according to the VCM classifier, which specializes in news stories about mergers, acquisitions and takeovers. Scores can take values of 0, 50, or 100 indicating negative, neutral or positive sentiment, respectively. This sentiment score is based on RavenPack's Expert Consensus Methodology, and has been trained on stories that lead up to a pre-identified mergers, acquisitions and takeover event.