

#Alpha: Extracting Market Sentiment From 140 Characters

Introducing social media indicators derived from tweets filtered to estimate sentiment

We present indicators that gauge investor outlook on firms utilizing aggregate tweet data to identify potential buy and sell candidates. The factors classify the text content in daily Twitter posts to construct sentiment and volume signals.

- Tweets identified as relevant to a particular stock are scored for sentiment to produce a cutting-edge measure of social media sentiment
- We report robust *S-Score™* and *Normalized Volume Adjusted Sentiment Score* average daily decile return spreads of 0.1680% and 0.140% for high-frequency tweeted names, persistent to 5-day periods
- Correlation analysis confirms unique signal content versus standard sentiment metrics from equity, options and short interest markets

INTRODUCTION

With the widespread use of the internet in everyday life, the availability of search engine and social media data has opened up a new discipline for research information. With its worldwide popularity, Twitter is one such application that enables timely tracking of public sentiment. Twitter is a microblogging service in which tweets, i.e., text messages limited to 140 characters, are posted conveying information describing what users are thinking, doing and feeling.

Since the launch of Twitter in 2006, the number of users contributing tweets on the social networking website has grown exponentially worldwide. Social media is growing quickly and is soon expected to reach one billion tweets per day. Additionally, over 60% of adults worldwide now use social media and over 100,000 users are added each day¹.

With this vast access to public sentiment, Twitter has been used in many industries to replace resource-intensive surveys by gathering timely feedback for market research strategy. Conversely, the information content in tweets can also be used to research public mood on subjects ranging from politics to retail and, in our case, investments.

Social Market Analytics, Inc. (SMA) operates in data services and provides analysis of social media data streams to estimate market sentiment at the stock level. Through

our partnership with the firm, we are introducing a suite of social media indicators constructed to capture timely information gleaned from Twitter posts. The measures are based on analysis of the text content in daily Twitter posts. Tweets are filtered for financial trading relevance and scored for market sentiment content. Tweet scores are then aggregated for each stock to produce a sentiment measurement from which the indicators are derived.

The remainder of this report provides an introduction to the Markit Research Signals social media indicators suite. We begin with background detail describing this unique data source and the factors built upon it. Next we present descriptive statistics describing several representative measures and round out the report with performance results for select key indicators.

LITERATURE REVIEW

The use of survey data and social mood has been a growing discipline in the prediction of financial markets, particularly since the behavioral finance field has become widely accepted, and online data sets have opened up a much more large-scale and expeditious resource for analysis. Mao *et al.* (2011) conducted a comprehensive study, building upon earlier work by Zhang *et al.* (2011),

¹<http://www.statisticbrain.com/twitter-statistics>

of a range of online data sets and sentiment tracking methods to compare their predictability of financial market indicators. By utilizing surveys, news headlines, search engine data and Twitter feeds, they compute sentiment indicators including Survey Investor Sentiment, Negative News Sentiment, Google search volumes of financial terms, Twitter Investor Sentiment and Tweet volumes of financial terms. Google Insights for Search, a service providing search volume data, revealed a significant correlation between financial term searches and Dow Jones Industrial average closing values, trading volume and VIX values, while Investor Intelligence surveys did not. They also found that an indicator of Twitter Investor Sentiment and the frequency of occurrence of financial terms on Twitter measured over recent days are also statistically significant predictors of daily market returns, while Daily Sentiment Index readings are not.

DATA AND METHODOLOGY

Our social media data is sourced from SMA, which analyzes social media data streams to estimate market sentiment. More specifically, metrics are estimated from analysis of Twitter message stream that are converted into actionable indicators in their family of measures called S-Factors™, designed to capture the signature of financial market sentiment. However, not all tweets for a particular stock are useful. Only tweets that pass SMA's filtering processes and identified as "indicative" tweets posted by confirmed accounts are used in sentiment estimates.

The methodology involves a 3-step process:



Exhibit 1: SMA extraction process

1. The **Extractor** collects tweet content and source information using a source agnostic retrieval platform that extracts all signals for designated financial terms and symbols. SMA's servers poll API's of Twitter and GNIP with access to over 500 million daily tweets.
2. The **Evaluator** filters the tweets to only include "indicative" tweets, those with relevant sentiment to the particular stock. The process utilizes established Natural Language Processing algorithms, enhanced and tuned for performance in the domain of financial markets. The Evaluator identifies words, phrases and stock symbols in the captured tweets, then removes duplicates and applies re-tweet policies to reduce the noise level of the tweet stream from sources such as "spamming" users. Lastly, it analyzes the set of relevant tweets with respect to SMA ratings for the Twitter accounts that are the originators of the captured tweets.
3. The **Calculator** analyzes the tweet language using a Sentiment Dictionary tuned for performance in the financial market domain with relevant, industry-specific terms. Sentiment level for each word parsed from a tweet is obtained from the dictionary. SMA's Sentiment Dictionary currently has about 18,000 words (uni-grams) and 400 two-word phrases (bi-grams) that have content and sentiment levels of relevance to financial market activity. Raw sentiment level is the simple aggregate of all indicative tweet sentiment levels captured during the prior 24 hours. Lastly, a normalization and scoring process calculates the final sentiment measures.

With this, we introduce 22 social media indicators on the Markit Research Signals platform using SMA sentiment data. Factors cover the following broad categories:

- **Tweet sentiment** quantifies alpha-generating sentiment from a previously untapped source of information flow of tweets filtered for financial trading relevance and scored for market sentiment content
- **Tweet volume** identifies increased interest in a stock
- **Relative value** computes scores relative to the market and/or the stock's recent history and provides a clearer view of sentiment levels
- **Changing sentiment** measures 1-day to 20-day look back signals to identify trends in the sentiment signals
- **Dispersion** assesses the number of unique tweet sources to gauge the validity of a signal

For the full factor list and brief definitions, please see the Appendix (on Pages 10 and 11).

Several weighting methodologies are utilized for various measures with details as follows:

- **Unweighted** metrics simply score and aggregate all tweets captured during the sample window to produce the sentiment estimate
- **Exponential weights** scale a tweet's sentiment score prior to aggregation by an exponential weight function that varies as a function of the arrival time of the tweet, with a maximum at the time of the stock's sentiment estimate and decreasing smoothly to a minimum at the start of the sample window
- **Normalized values** are computed as the z-score normalization of the time series

Data coverage begins December 1st 2011. While data and factor scores are available on an intraday basis, performance analytics in this report are based on pre-open data, published at 9:00 AM EST prior to August 8th 2012 and at 8:55 AM EST subsequently. However, updates may at times be sparse as messaging on Twitter expressing market sentiment for any stock may be variable in time and volume. Thus, for our coverage universe Markit US Total Cap, representing 98% of cumulative market cap or 3,000+ stocks, we find on average 1,200 names covered daily. For example, Figure 1 presents the time series trend in coverage for *S-Score*[™], a representative factor discussed in more detail below. We also present a count for names with a minimum of three tweets, which tends to be just under half of the coverage universe.




Figure 1: US Total Cap *S-Score*[™] coverage, Dec 2011 – Nov 2013

DESCRIPTIVE STATISTICS

We begin by presenting descriptive statistics for a handful of representative factors to demonstrate their underlying characteristics:

- *Raw-S*[™] is computed as an unweighted aggregation of sentiment score of all indicative tweets captured during the prior 24-hour window
- *S-Volume*[™] measures the number of tweets used to calculate the sentiment score
- *Volume Adjusted Sentiment Score* gauges the sentiment score per indicative tweet




Figure 2: US Total Cap Raw-S™ aggregate statistics, Dec 1 2011 – Nov 30 2013




Figure 3: US Total Cap scores, Nov 29 2013

First, we view a time series display of aggregate *Raw-S™* averages and standard deviations over the analysis period (Figure 2). We observe that averages are positive implying that sentiment values tend to be positive; however, the series is marked by measurable variation.

S-Volume™ also exhibits considerable aggregate variability around an average of 6.4 daily tweets. Furthermore, positive sentiment tends to be skewed to names with the most tweets. Indeed, a sample scatter plot of *Raw-S™* and *S-Volume™* scores on November 29th 2013 (Figure 3) confirms a clustering of values among positive sentiment scores with a positive linear slope of 1.6 associated with the main group of names with <20 tweets. Furthermore, we note a limited number of outlier volumes, particularly that of Apple Inc. Drilling down further into this name (Figure 4) we observe daily tweet volume exceeds 1,000 frequently with significant levels even on non-trading days.

Lastly, we review *Volume Adjusted Sentiment Score*, a computed measure which adjusts sentiment score for tweet volume. A time series plot of aggregate averages and standard deviations (Figure 5) demonstrates the greater stability achieved by this indicator versus the raw sentiment score (see Figure 2), particularly in the standard deviation statistics.




Figure 4: Apple Inc sentiment scores, Dec 1 2011 – Nov 30 2013

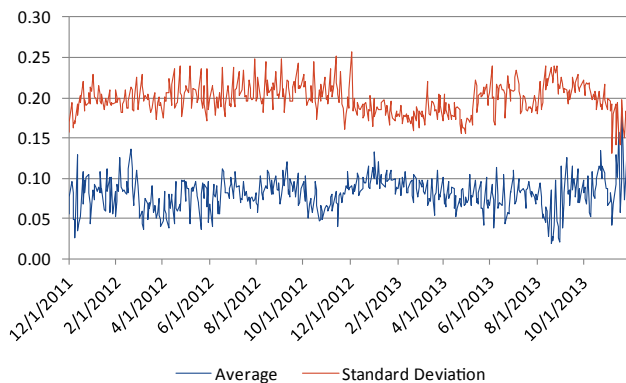


Figure 5: US Total Cap Volume Adjusted Sentiment Score aggregate statistics, Dec 1 2011 – Nov 30 2013

PERFORMANCE RESULTS AND ATTRIBUTION

S-Score™

Turning to performance analysis, we first present a detailed review of a key SMA factor, S-Score™, which measures a stock's aggregated raw sentiment score normalized to the average and standard deviation of the past 20 days. It is a gauge of the deviation of a stock's sentiment intensity level from a normal state computed as a z-score of weighted raw sentiment scores. While the underlying measure is highly dependent on tweet volume, the z-score normalization adjusts for this impact. Positive (negative) scores are aligned with positive (negative) sentiment and receive a top (bottom) rank.

We analyze several strategies for performance analytics of S-Score™. First, we detail results at the tails of the underlying factor distribution with a score >3 (<-3), in other words, a current relative sentiment score in excess of 3 (-3) standard deviations away from the normal aggregate sentiment level, indicating positive (negative) sentiment and considered a buy (sell) signal. We will detail performance on a decile ranking basis as well, but results indicate that the factor's strength is in identifying names at the extreme tails of the distribution. We present a time series of cumulative 1-day returns for a strategy to buy at the open and sell at the close (Figure 6) compared to the market return proxied by the SPDR S&P 500 ETF (SPY). Based on our empirical results, we report a cumulative (average) return of 76% (0.12%) for the buy portfolio compared to a 14% loss (-0.03%) for the sell portfolio and an open-to-close market return of 20% (0.04%).




Figure 6: US Total Cap S-Score™ open-to-close daily returns, Dec 1 2011 – Nov 30 2013

Cumulative returns show a very strong divergence of the positive and negative sentiment stocks, relative to the market. One drawback of this strategy is that we are assuming we will reinvest all capital into the strategy the next day. However, we will show that there will be dates with a low number of stocks meeting our thresholds. Moreover, we know that there tends to be more positive sentiment signals than

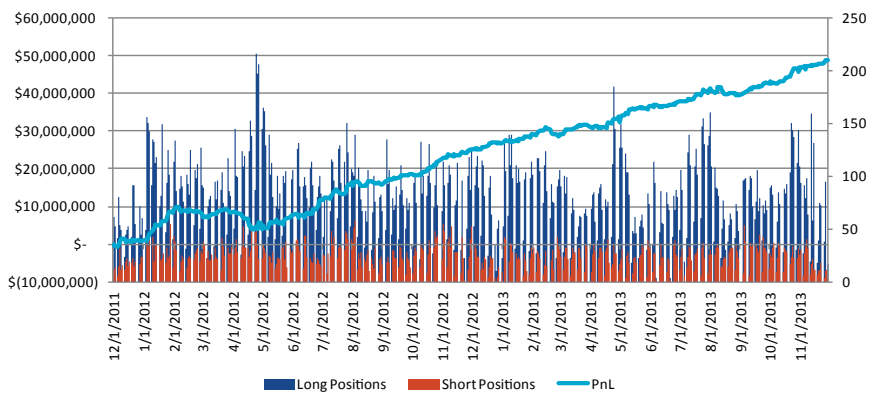


Figure 7: US Total Cap S-Score™ fixed position PnL, Dec 1 2011 – Nov 30 2013

negative. Taking into account issues of sparse signals resulting in large swings in the number of buy and sell candidates each day, we also present a fixed position size strategy which holds constant the dollar amount invested in each position. Here we go long S-Score™ >3 , putting \$10,000 in each position while hedging the same amount of capital by shorting SPY. Likewise, we go short S-Score™ <-3 by shorting \$10,000 in each position and buying the same amount of capital in SPY. The profit and loss (PnL) chart for the strategy is displayed in Figure 7. This strategy shows remarkable consistency with minimal drawdowns. We also plot the number of positions on the long and short side to demonstrate the skew towards positive sentiment signals. Our

	Open-to-Close (%)	1-day (%)	2-day (%)	3-day (%)	4-day (%)	5-day (%)	Size
D1	0.112	0.120	0.220	0.316	0.412	0.515	118
S-Score™ >2	0.098	0.124	0.220	0.317	0.415	0.511	160
S-Score™ >3	0.117	0.124	0.215	0.310	0.407	0.513	81
D10	-0.012	0.080	0.181	0.271	0.367	0.456	121
S-Score™ <-2	-0.014	0.076	0.182	0.299	0.385	0.498	54
S-Score™ <-3	-0.025	0.032	0.126	0.217	0.278	0.393	24
D1-D10 Spread (IR)	0.124 (0.04)	0.040 (0.12)	0.039 (0.08)	0.045 (0.08)	0.045 (0.07)	0.059 (0.08)	
S-Score™ 2 Spread (IR)	0.113 (0.26)	0.048 (0.11)	0.038 (0.06)	0.019 (0.02)	0.030 (0.03)	0.013 (0.01)	
S-Score™ 3 Spread (IR)	0.142 (0.19)	0.092 (0.02)	0.089 (0.10)	0.093 (0.08)	0.129 (0.10)	0.121 (0.08)	

Table 1: US Total Cap S-Score™ average returns, Dec 1 2011 – Nov 30 2013

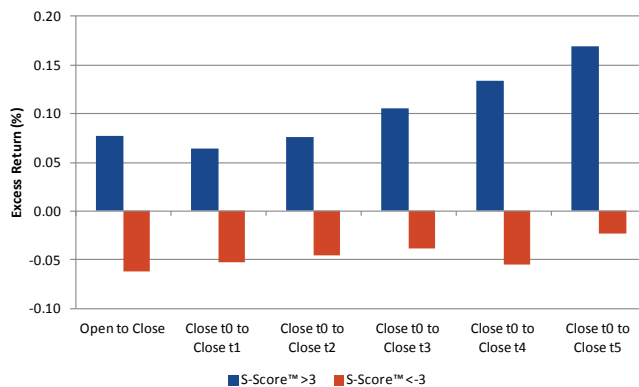


Figure 8: US Total Cap S-Score™ average excess returns, Dec 1 2011 – Nov 30 2013

results not only maintain attractive returns to this more robust performance simulation, they also demonstrate exceptional excess returns during a period of a particularly strong bull market run.

Next, we consider signal persistence beyond the open-to-close period by examining cumulative returns out to 5-days to confirm the robustness of S-Score™ signals given it is a known high-turnover strategy (92% and 90% 1-day turnovers for top and bottom decile names, respectively). In other words, the *n*-day return is computed from the day 0 close to the day *n* close. Here we also expand the tail portfolios for further robustness checks. Table 1 summarizes results of decile ranks along with underlying S-Score™ levels >2 and >3, while Figure 8 focuses further on S-Score™ >3 and < -3 strategy excess returns over the SPY ETF.

Our results confirm that the signal power persists beyond the open-to-close period. Indeed, we report buy portfolio 5-day returns (D1: 0.515%; S-Score™ >2: 0.511%; S-Score™ >3: 0.513%) in excess of open-to-close returns (D1: 0.112%; S-Score™ >2: 0.098%; S-Score™ >3: 0.117%). Persistent positive spreads are also recorded across all strategies with affirmation from the information ratio (IR), which is a volatility-adjusted statistic computed as the average divided by the standard deviation. However, we remark that, while the overall open-to-close hit rate (percent of firms with positive returns) for S-Score™ >3 (>2) is 51.4% (51.2%) versus 49.2% (49.4%) for S-Score™ <-3 (<-2), outcomes at the extended holding periods are more similar at both tails, suggesting that the signal outperformance is achieved more so from return differentials.

To ascertain market capitalization biases, we further break out performance results for US Large Cap and Small Cap universes (see Appendix Tables A1 and A2, respectively on Page 10). We focus on decile results as the z-score portfolios become moderate in size and returns are to some extent sparser. Overall, we find that large cap spread outperformance is more concentrated in the open-to-close period, while small caps respond more robustly to the signal out to the 5-day holding period, suggesting that market participants are faster to react to information content from the indicative tweets of large cap names.

S-Score™ >3 excess returns also confirm that the alpha continues to improve five days out from the original signal, though the S-Score™ <-3 excess returns tend to mean revert more quickly. We also find that the skew to positive sentiment (as observed in Figure 2 on Page 4) produces a higher number of securities on average for S-Score™ >2 or >3 than <-2 or <-3, a nuance to be keenly aware of when implementing the signal in a portfolio. Furthermore, we remark that overall hit rates for excess returns are similar to those previously cited, again suggesting that the signal outperformance at longer holding periods is more return based. Lastly, large and small cap results (see Appendix Tables A3 and A4, respectively on Page 10) confirm more robust spreads to the latter.

In studying which stocks register signals, we find some cases where top or bottom scoring names only have one or two tweets. While we can make a case that these are potentially important tweets – say Carl Icahn announces a new position in an otherwise unmentioned stock – we study the signal strength after removing names with sparse tweet volumes. For this we apply a filter for

a minimum *S-Volume*TM of 3 to adjust the *S-Score*TM signal for confirmed information content. Updated results are listed in Table 2. In general, outcomes are robust to this additional check, with the largest impact associated with deeper underperformance for the sell portfolios.

	Open-to-Close (%)	1-day (%)	2-day (%)	3-day (%)	4-day (%)	5-day (%)	Size
D1	0.123	0.127	0.216	0.316	0.407	0.505	54
S-ScoreTM >2	0.094	0.123	0.198	0.290	0.391	0.471	86
S-ScoreTM >3	0.124	0.125	0.174	0.243	0.350	0.456	42
D10	-0.045	0.067	0.172	0.268	0.377	0.435	53
S-ScoreTM <-2	-0.061	0.064	0.191	0.306	0.392	0.417	23
S-ScoreTM <-3	-0.080	-0.013	0.132	0.188	0.243	0.256	9
D1-D10 Spread	0.168	0.060	0.044	0.048	0.030	0.070	
S-ScoreTM 2 Spread	0.155	0.058	0.007	-0.016	-0.001	0.054	
S-ScoreTM 3 Spread	0.205	0.138	0.043	0.055	0.107	0.201	

Table 2: US Total Cap with S-VolumeTM filter S-ScoreTM average returns, Dec 1 2011 – Nov 30 2013

Rounding out the *S-Score*TM analysis, we drill down to attribution of the factor scores. First, we test for uniqueness of the signal content by examining its daily percentile rank correlation with several standard short-term technical indicators from the equity, options and short interest markets. We follow this with an analysis of factor rank exposures to several systematic risk indicators.

We verify a very low rank correlation with other “sentiment” based factors (Table 3) indicating the unique nature of the twitter sentiment. The highest absolute correlations are associated with 5-day *Industry Relative Return* (-0.08) and *Implied Volatility* (0.06), confirming that sentiment is not captured by equity price movement, options implied volatility or the short interest market.

Factor	Correlation
60-Month Beta	0.00
5-day Industry Relative Return	-0.08
Most Recent Earnings Surprise	-0.01
Net # of Revisions for Fiscal Year 1	0.00
ATM Put Volatility - ATM Call Volatility	-0.02
Implied Volatility	0.06
Short Interest	0.00
1-Month Change in Short Interest	0.01

Table 3: US Total Cap S-ScoreTM rank correlations, Dec 1 2011 – Nov 30 2013

Furthermore, we report only minor active exposures, in general, to a representative list of systematic risk indicators (Table 4) gauged by the average factor percentile scores for the *S-Score*TM universe. For context, we also include scores for top (D1) and bottom (D10) names of the distribution where we observe more positive sentiment toward relatively smaller cap, higher volatility names. Not surprisingly, the names at the top and bottom deciles tend to have more exposure to *Implied Volatility* than the overall universe. However, the exposures are nowhere near as extreme as we would expect if this was merely a proxy for option implied volatility.

	S-Score TM universe	D1	D10	Interpretation
60-Month Beta	52	52	52	low beta (1) – high beta (100)
Book-to-Market	54	51	52	undervalued (1) – overvalued (100)
Natural Logarithm of Market Capitalization	62	52	60	small cap (1) – large cap (100)
Implied Volatility	57	49	55	high volatility (1) – low volatility (100)

Table 4: US Total Cap with S-ScoreTM percentile factor exposures, Dec 1 2011 – Nov 30 2013

We observe skewness in tweet volume and aggregate sentiment to popular names such as Apple, but are there consistent biases towards particular sectors, such as technology and retail? Table 5 on Page 8 lists active equal-weight exposures of names with *S-Score*TM ranks and those at the tails of the factor distribution versus the Total Cap universe on average over the analysis period. In general, active exposures are modest suggesting that, while daily exposures are sporadic, the average over time is in-line with market sector exposures as no one sector is favored by tweeters. While slight, the largest positive exposures tend to be Cyclical Goods & Services and Healthcare and the largest negative exposures are associated with Energy and Financials.

	S-Score™ universe	S-Score™ >2	S-Score™ >3	S-Score™ <-2	S-Score™ <-3
Energy	0%	-1%	-3%	-2%	-4%
Basic Materials	-1%	-1%	-1%	0%	0%
Industrials	-1%	1%	1%	1%	2%
Cyclical Goods & Services	3%	1%	0%	2%	2%
Non-cyclical Goods & Services	1%	0%	0%	0%	0%
Financials	-5%	-2%	-1%	-3%	0%
Healthcare	2%	2%	2%	1%	0%
Technology	1%	0%	1%	1%	0%
Telecommunication Services	0%	0%	0%	0%	0%
Utilities	0%	0%	0%	0%	0%

Table 5: US Total Cap S-Score™ average active equal-weight sector exposures, Dec 1 2011 – Nov 30 2013

Normalized Volume Adjusted Sentiment Score

We now turn to a Markit-specific Social Media indicator, *Normalized Volume Adjusted Sentiment Score* (S-VolAdj), computed as the z-score normalization of the volume-relative sentiment score per indicative tweet over a 20-day period. By adjusting the sentiment score by volume, we add a second step, along with the z-score normalization, to take out the impact of tweet volume on the scoring system. This 2-step process applies a more robust methodology to address the bias in positive sentiment to names with the most tweets. While similar in construction to S-Score™, we remark that the rank correlation between the two factors is 0.85.

Focusing again on tail performance, we once more apply a filter for a minimum S-Volume™ of 3 to adjust the S-Score™ signal for confirmed information content in addressing random effects that may arise from sparse tweets. Updated spread results for decile ranks along with underlying S-VolAdj levels are listed in Table 6. Our results similarly confirm outperformance over the open-to-close period that persists to the 5-day period. We report buy portfolio 5-day returns (D1: 0.456%; S-VolAdj >2: 0.445%; S-VolAdj >3: 0.654%) in excess of open-to-close returns (D1: 0.107%; S-VolAdj >2: 0.127%; S-VolAdj >3: 0.144%) along with persistent positive spreads across all strategies confirming robustness for this high-turnover signal. Furthermore, outcomes are in general robust to this additional check, with the largest impact again associated with deeper underperformance to the sell portfolios.

	Open-to-Close (%)	1-day (%)	2-day (%)	3-day (%)	4-day (%)	5-day (%)	Size
D1	0.107	0.098	0.213	0.294	0.378	0.456	54
S-VolAdj >2	0.127	0.094	0.196	0.256	0.333	0.445	27
S-VolAdj >3	0.144	0.132	0.235	0.392	0.534	0.654	10
D10	-0.033	0.083	0.204	0.308	0.388	0.445	53
S-VolAdj <-2	-0.084	0.042	0.108	0.122	0.080	0.128	11
S-VolAdj <-3	-0.150	-0.068	0.074	0.108	0.142	0.156	4
D1-D10 Spread	0.140	0.015	0.009	-0.015	-0.009	0.011	
S-VolAdj 2 Spread	0.211	0.053	0.087	0.134	0.253	0.317	
S-VolAdj 3 Spread	0.294	0.200	0.160	0.284	0.393	0.497	

Table 6: US Total Cap S-VolAdj with S-Volume™ filter average returns, Dec 1 2011 – Nov 30 2013

The overall open-to-close hit rate for S-VolAdj >3 (>2) is 50.2% (50.0%) versus 48.3% (48.5%) for S-VolAdj <-3 (<-2), while outcomes at the extended holding periods are more similar at both tails, once again suggesting that the signal outperformance is achieved more so from return differentials. Additionally, large and small cap results (see Appendix Tables A5 and A6, respectively on Page 10) confirm healthier spreads to the latter, persistent out to the 5-day holding period.

Lastly, we present S-VolAdj attribution analysis. First, we examine S-VolAdj rank correlations with the aforementioned short-term technical indicators (see Table 7 on Page 9). We again verify very low rank correlations, with the highest in magnitude associated with *5-day Industry Relative Return* (-0.07) and *Implied Volatility* (0.05), confirming that sentiment is not captured by equity price movement, options implied volatility or the short interest market. Additionally, factor exposures, proxied by a representative list

of systematic risk indicators (Table 8), once more confirm only minor active exposures with more positive sentiment toward relatively smaller cap, higher volatility names.

Factor	Correlation
60-Month Beta	0.00
5-day Industry Relative Return	-0.07
Most Recent Earnings Surprise	0.00
Net # of Revisions for Fiscal Year 1	0.01
ATM Put Volatility - ATM Call Volatility	-0.01
Implied Volatility	0.05
Short Interest	0.00
1-Month Change in Short Interest	0.01

Table 7: US Total Cap S-VolAdj rank correlations, Dec 1 2011 – Nov 30 2013

	S-VolAdj universe	D1	D10	Interpretation
60-Month Beta	52	52	52	low beta (1) – high beta (100)
Book-to-Market	54	50	52	undervalued (1) – overvalued (100)
Natural Logarithm of Market Capitalization	62	50	60	small cap (1) – large cap (100)
Implied Volatility	57	50	55	high volatility (1) – low volatility (100)

Table 8: US Total Cap S-VolAdj percentile factor exposures, Dec 1 2011 – Nov 30 2013

Lastly, average active equal-weight sector exposures versus the Total Cap universe (Table 9) are moderate once again. Overall, the largest positive (negative) exposures are associated with Industrials (Energy).

	S-VolAdj universe	S-VolAdj >2	S-VolAdj >3	S-VolAdj <-2	S-VolAdj <-3
Energy	0%	-2%	-4%	-2%	-3%
Basic Materials	-1%	0%	0%	0%	0%
Industrials	-1%	1%	2%	1%	2%
Cyclical Goods & Services	3%	0%	-1%	2%	1%
Non-cyclical Goods & Services	1%	0%	0%	0%	0%
Financials	-5%	-1%	3%	-3%	0%
Healthcare	2%	1%	0%	1%	0%
Technology	1%	0%	0%	1%	1%
Telecommunication Services	0%	0%	0%	0%	0%
Utilities	0%	1%	1%	0%	0%

Table 9: US Total Cap S-VolAdj average active equal-weight sector exposures, Dec 1 2011 – Nov 30 2013

CONCLUSION

In this research note we introduce our social media indicators constructed from tweets conveying stock-level sentiment. Twitter is one such application that enables timely tracking of public sentiment via posts conveying information on what is going on in the user's life. As such, Twitter feeds are a vast repository of public mood data that are leveraged in numerous applications to quickly and easily gather feedback and, in our case, investor sentiment.

Our social media indicators are sourced in partnership with SMA which analyzes social media data streams to estimate market sentiment. More specifically, metrics are estimated from analyzing Twitter message streams that are then converted into actionable indicators. However, not all tweets are useful; therefore, only tweets that pass SMA's filtering processes and identified as "indicative" tweets posted by confirmed accounts are used in sentiment estimates. Thus, the full process involves extracting relevant tweets, validating the source, evaluating the meaning and calculating factors. The resulting indicators provide real time sentiment tracking and public mood modeling.

With 22 metrics added to our factor suite, we focus our comments on several key indicators. We begin with descriptive statistics of a few representative measures. In general, we observe measurable variability in *Raw-STM* and *S-VolumeTM* scores with a

positive skew in sentiment to names with the most tweets, while our proprietary Volume Adjusted Sentiment Score adjusts for this bias demonstrating a more stable time series.

We round out the report with performance analytics of two key indicators. For a buy (sell) portfolio based on *S-Score*[™] valued >3 (<-3), we report a cumulative (average) 1-day return of 76% (0.12%) for the buy portfolio compared to a 14% loss (-0.03%) for the sell portfolio and a market return of 20% (0.04%). Outcomes are persistent out to a 5-day horizon with healthier return spreads to small cap names and low rank correlations with other sentiment factors. *Normalized Volume Adjusted Sentiment Score* implements a more robust 2-step process to take out the impact of tweet volume on the sentiment scoring system and similarly confirms outperformance over the open-to-close period that also persists to the 5-day horizon.

Lastly, nuances to the overall strategy to be aware of when implementing the signals in a portfolio include its high-turnover nature and bias to positive sentiment. Future research will focus on these features as well as signal performance related to market interactions surrounding events such as earnings release dates and price momentum. Stay tuned for subsequent publications.

APPENDIX

Performance Results

	Open-to-Close (%)	1-day (%)	2-day (%)	3-day (%)	4-day (%)	5-day (%)	Size
D1	0.098	0.096	0.176	0.272	0.353	0.449	61
D10	0.033	0.107	0.205	0.294	0.379	0.486	60
D1-D10 Spread	0.065	-0.011	-0.029	-0.021	-0.027	-0.037	

Table A1: US Large Cap S-Score[™] average returns, Dec 1 2011 – Nov 30 2013

	Open-to-Close (%)	1-day (%)	2-day (%)	3-day (%)	4-day (%)	5-day (%)	Size
D1	0.110	0.145	0.258	0.362	0.467	0.576	57
D10	-0.041	0.061	0.149	0.254	0.349	0.434	60
D1-D10 Spread	0.151	0.085	0.109	0.107	0.119	0.142	

Table A2: US Small Cap S-Score[™] average returns, Dec 1 2011 – Nov 30 2013

	Open-to-Close (%)	1-day (%)	2-day (%)	3-day (%)	4-day (%)	5-day (%)	Size
D1	0.099	0.104	0.187	0.294	0.370	0.475	34
D10	0.007	0.099	0.202	0.290	0.396	0.494	33
D1-D10 Spread	0.092	0.005	-0.015	0.004	-0.026	-0.019	

Table A3: US Large Cap S-Score[™] with S-Volume[™] filter average returns, Dec 1 2011 – Nov 30 2013

	Open-to-Close (%)	1-day (%)	2-day (%)	3-day (%)	4-day (%)	5-day (%)	Size
D1	0.132	0.123	0.189	0.298	0.390	0.459	21
D10	-0.121	0.038	0.113	0.214	0.333	0.339	20
D1-D10 Spread	0.253	0.086	0.076	0.084	0.056	0.119	

Table A4: US Small Cap S-Score[™] with S-Volume[™] filter average returns, Dec 1 2011 – Nov 30 2013

	Open-to-Close (%)	1-day (%)	2-day (%)	3-day (%)	4-day (%)	5-day (%)	Size
D1	0.085	0.100	0.197	0.271	0.325	0.410	34
D10	0.015	0.103	0.199	0.291	0.394	0.485	33
D1-D10 Spread	0.071	-0.003	-0.002	-0.020	-0.069	-0.075	

Table A5: US Large Cap S-VolAdj with S-Volume[™] filter average returns, Dec 1 2011 – Nov 30 2013

	Open-to-Close (%)	1-day (%)	2-day (%)	3-day (%)	4-day (%)	5-day (%)	Size
D1	0.114	0.110	0.239	0.358	0.506	0.606	21
D10	-0.159	0.028	0.195	0.332	0.417	0.412	20
D1-D10 Spread	0.273	0.081	0.044	0.026	0.089	0.194	

Table A6: US Small Cap S-VolAdj with S-Volume[™] filter average returns, Dec 1 2011 – Nov 30 2013

SMA Factor Definitions

Raw-S™: Unweighted aggregation of sentiment score of all indicative tweets captured during the prior 24-hour window. Markit ranks this factor in descending order.

Raw-S-Mean™: 20-day moving average of the unweighted sentiment score. Markit ranks this factor in descending order.

Raw-Volatility™: 20-day moving standard deviation of the unweighted sentiment score. Markit ranks this factor in ascending order.

Raw-Score™: Z-score normalization of the unweighted sentiment score over a 20-day period. Markit ranks this factor in descending order.

S™: Exponentially weighted aggregation of sentiment score of all indicative tweets captured during the prior 24-hour window. Markit ranks this factor in descending order.

S-Mean™: 20-day moving average of the exponentially weighted sentiment score. Markit ranks this factor in descending order.

S-Volatility™: 20-day moving standard deviation of the exponentially weighted sentiment score. Markit ranks this factor in ascending order.

S-Score™: Z-score normalization of the exponentially weighted sentiment score over a 20-day period. Markit ranks this factor in descending order.

S-Volume™: Number of indicative tweets used to calculate the sentiment score. Markit ranks this factor in descending order.

SV-Mean™: 20-day moving average of the indicative tweet volume. Markit ranks this factor in descending order.

SV-Volatility™: 20-day moving standard deviation of the indicative tweet volume. Markit ranks this factor in ascending order.

SV-Score™: Z-score normalization of the indicative tweet volume over a 20-day period. Markit ranks this factor in descending order.

S-Dispersion™: Ratio of the number of distinct sources to the number of indicative tweets, which measures the concentration level of the tweet sources contributing to a sentiment score. The higher the dispersion value, the greater the number of distinct sources. Markit ranks this factor in descending order.

S-Buzz™: Normalized indicative tweet volume relative to the universe average. Markit ranks this factor in descending order.

S-Delta™: Absolute change in the normalized weighted sentiment score over a 15-minute period. Markit ranks this factor in descending order.

Markit Research Signals: Social media indicators

1-day Change in Normalized Weighted Sentiment Score: Percentage change in the normalized weighted sentiment score over a 1-day period. Markit ranks this factor in descending order.

5-day Change in Normalized Weighted Sentiment Score: Percentage change in the normalized weighted sentiment score over a 5-day period. Markit ranks this factor in descending order.

Volume Adjusted Sentiment Score: Sentiment score per indicative tweet. Markit ranks this factor in descending order.

20-day Average of Volume Adjusted Sentiment Score: 20-day moving average of the sentiment score per indicative tweet. Markit ranks this factor in descending order.

20-day Standard Deviation of Volume Adjusted Sentiment Score: 20-day moving standard deviation of the sentiment score per indicative tweet. Markit ranks this factor in ascending order.

Normalized Volume Adjusted Sentiment Score: Z-score normalization of the sentiment score per indicative tweet over a 20-day period. Markit ranks this factor in descending order.

Relative Standard Deviation of Indicative Tweet Volume: Coefficient of variation of the indicative tweet volume over a 20-day period. Markit ranks this factor in descending order.

REFERENCES

Mao, H., S. Counts and J. Bollen. "Predicting Financial Markets: Comparing Survey, News, Twitter and Search Engine Data", 2011.

Zhang, X., H. Fuehres, P. A. Gloor. "Predicting Stock Market Indicators Through Twitter 'I hope it is not as bad as I fear' ", 2011.

About Markit

Markit is a leading global diversified provider of financial information services.

We provide products that enhance transparency, reduce risk and improve operational efficiency in the financial markets.

Our customers include banks, hedge funds, asset managers, central banks, regulators, auditors, fund administrators and insurance companies.

Founded in 2003, we employ over 3,000 people in 11 countries. For more information, please see www.markit.com.

More Information

For more information on the products and services from Markit, please contact us at sales@markit.com or call one of our regional offices:

London	+ 44 20 7260 2000
New York	+ 1 212 931 4900
Naperville	+ 1 630 637 8088
Dallas	+ 1 972 560 4420
Sydney	+ 61 2 0 89 89 80
Tokyo	+ 81 3 6402 0130
Singapore	+ 65 6499 0088
Frankfurt	+ 49 69 299 868 140
Amsterdam	+ 31 20 502 5800

These materials are confidential. Distribution is Prohibited.

Opinions, estimates and projections in this report constitute the current judgment of the author(s) at the time of writing. They do not necessarily reflect the opinions of Markit Group Limited. Markit Group Limited has no obligation to update, modify or amend this report or to otherwise notify a reader thereof in the event that any matter stated herein, or any opinion, projection, forecast or estimate set forth herein, changes or subsequently becomes inaccurate.

The content, information and any materials ("data") provided by Markit in this presentation is on an "as is" basis. Markit Group makes no warranty, expressed or implied, as to its accuracy, completeness or timeliness, or as to the results to be obtained by recipients, and shall not in any way be liable to any recipient for any inaccuracies, errors or omissions herein. Without limiting the foregoing, Markit Group shall have no liability whatsoever to a recipient of this report, whether in contract (including under an indemnity), in tort (including negligence), under a warranty, under statute or otherwise, in respect of any loss or damage suffered by such recipient as a result of or in connection with any opinions, recommendations, forecasts, judgements, or any other conclusions, or any course of action determined, by it or any third party, whether or not based on the content, information or materials contained herein.

Copyright © 2014, Markit Group Limited. All rights reserved. Any unauthorised use, disclosure, reproduction or dissemination, in full or in part, in any media or by any means, without the prior written permission of Markit Group Limited is strictly prohibited.