Social media indicators in the UK

February 2019

Research Signals

In March 2014, we introduced a set of social media indicators for US markets, in partnership with Social Market Analytics, Inc., that classify the text content in daily Twitter posts to construct a family of social media signals. We now expand our coverage to the UK market using a similar factor structure.

- For names at the extreme tails (2 standard deviations) of the factor distribution, we report notable S-Score™ average daily return spreads of 0.097% since August 2015, with robustness out to longer 10- and 20-day holding periods
- When focusing on frequently tweeted names, average 20-day return spreads improve to 0.383% from 0.298% for the stand-alone strategy, while, for long-only strategies, our empirical results again demonstrate positive performance for names at the 2-standard-deviation tail, with average excess returns of 0.049% on an open-to-close basis, extending to 0.389% out to 20 days
- For Relative Standard Deviation of Indicative Tweet Volume, a Research Signals defined measure of the tweet volume volatility of a stock, we find average daily return spreads of 0.224%, which reached 0.342% at the 20-day horizon

Contacts

Research Signals · MK-ResearchSignals@ihsmarkit.com

Confidential. © 2019 IHS Markit™. All rights reserved.

Introduction

Social media influences the lives of people all over the world and it was just a matter of time before its influence would extend to other disciplines such as stock market pricing. Coupled with the emerging fields of search engines, artifical intelligence, machine learning and natural language processing, social media data has opened up a new discipline for research information and equity signals. With its worldwide popularity, Twitter is one such application that enables timely tracking of public sentiment.

Through our partnership with Social Market Analytics, Inc. (SMA), we provide a suite of social media indicators constructed to capture timely information gleaned from Twitter posts. SMA operates in data services and provides analysis of social media data streams to estimate market sentiment at the stock level.With a track record of successfully capturing social media sentiment in the US, SMA has expanded their coverage to cover stocks trading in the UK using the same robust process. Sentiment 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.

In this report, we extend our methodology underpinning the Research Signals Social Media factor suite to stocks trading on the London Stock Exchange. We begin with background detail describing this unique data source and the factors built upon it. We then turn to performance results for select key indicators, including both long-short and long-only strategies.

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 (Figure 1) involves a 3-step extraction process:

Figure 1

Source: IHS Markit

© 2019 IHS Markit

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 nearly 19,000 words (uni-grams) and over 18,000 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 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

The full factor list and brief definitions are provided in the Appendix. Several weighting methodologies are used 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 for the analytics begins 24 August 2015 and runs through 31 December 2018. While data and factor scores are available on an intraday basis, performance analytics in this report are based on pre-open data, published at 7:40 AM UK local time, and pre-close data, published at 4:10 PM UK local time. Daily average open-to-close returns are reported, along with varying holding periods ranging from 1 to 20 business days (overlapping periods) to check for persistence beyond the open-to-close period given that the signal is a known high-turnover strategy.

Updates may at times be sparse as messaging on Twitter expressing market sentiment for any stock may be variable in time and volume for a particular stock. Thus, for our coverage universe of stocks that trade on the London Stock Exchange, we find on average 330 names covered daily. For example, Figure 2 presents the time series trend in average monthly coverage for S-Score™, a representative factor discussed in more detail below.

Confidential. © 2019 IHS Markit. All rights reserved.

Figure 2

S-Score™ coverage

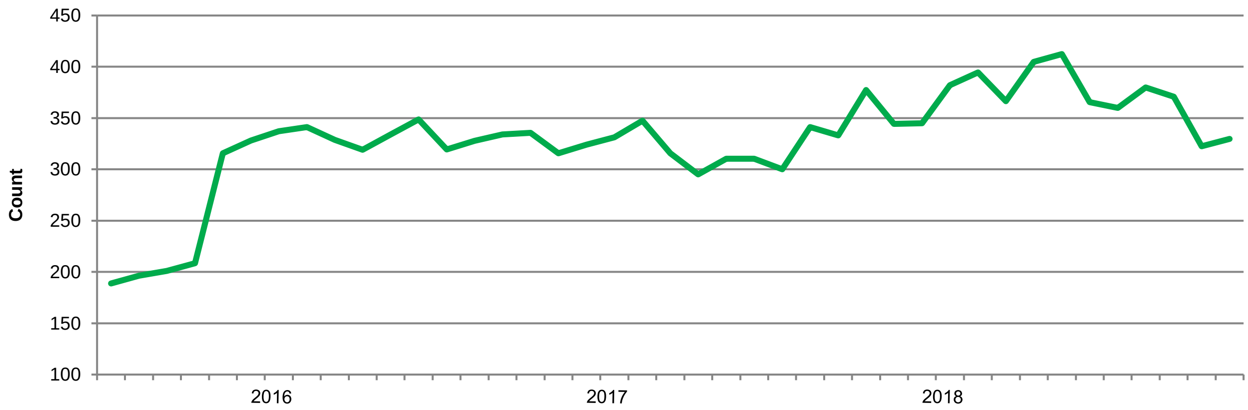

Source: IHS Markit

© 2019 IHS Markit

Performance results

We report performance results for select key indicators, beginning with detailed statistics for S-Score™, a key factor in our Social Media factor suite. We also include summary results for a Research Signals defined factor, Relative Standard Deviation of Indicative Tweet Volume.

S-Score™

S-Score™ 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 considered a buy (sell) signal.

We analyze several strategies for performance analytics of S-Score™. First, we detail results at the tails of the underlying factor distribution, for example, with a score >=3 (<=-3), in other words, a current relative sentiment score of at least 3 (-3) standard deviations away from the normal aggregate sentiment level, indicating positive (negative) sentiment. We also detail performance at scores of +/- 2 and on a quintile ranking basis providing larger targeted sets. While, the factor's strength is in identifying names at the extreme tails of the distribution, we focus results on the +/-2 tail, as the low counts at the more extreme +/-3 tail render the performance summary results susceptible to spurious outcomes (see Table A1 in the Appendix for portfolio sizes).

Updates may at times be sparse as messaging on Twitter expressing market sentiment for any stock may be variable in time and volume, with some top or bottom scoring names having just one tweet. While we can make a case that these are potentially important tweets, we study the signal strength after removing names with sparse tweet volumes. For this we apply a filter for a minimum S-Volume™ (number of indicative tweets) of 2 to adjust the S-Score™ signal for confirmed information content (see Table A1 in the Appendix for the impact on portfolio sizes).

First, we analyze spread performance which is the difference between equal-weight average returns for the buy portfolio less the sell portfolio (Table 1) and feature the results at the extreme +/-2 tail (Figure 3), where the benchmark is the average return of the underlying UK social media universe constituents with at least 1 tweet on that day. S-Score™ open-to-close return spreads at the +/-2 tail averaged 0.097% since August 2015. Signal power persisted out to three days

Confidential. © 2019 IHS Markit. All rights reserved.

(0.040%), though not exceeding the open-to-close average. However, longer holding period spreads picked back up, nearly tripling that of the open-to-close average, with 15- and 20-day average spreads of 0.272% and 0.298%, respectively.

At the more restrictive +/- 3 tail, performance tended to be weaker at the short time horizons, but we did see stronger spreads at longer holding periods from 10 days (0.111%) to 20 days (0.195%). At the broader quintile tails, performance was weaker in general across all time horizons, consistent with expectations of increased factor strength at the extreme tails.

We also detail daily cumulative spreads at the +/-2 tail over the full analysis period (Figure 4), comparing close-to-close returns for signals generated pre-open, pre-close and our benchmark universe. Our results demonstrate a cumulative return spread of 57% for longs versus 41% for shorts for the pre-open signals and 9% for the market. As expected, considering the importance of acting upon the signals with low latency, we note that the pre-close signal works better than the pre-open signal when trading at the close, with a cumulative return of 75% for longs versus just 6% for shorts.

Figure 3

S-Score™ +/-2 tail average return spreads, 24 Aug 2015 - 31 Dec 2018

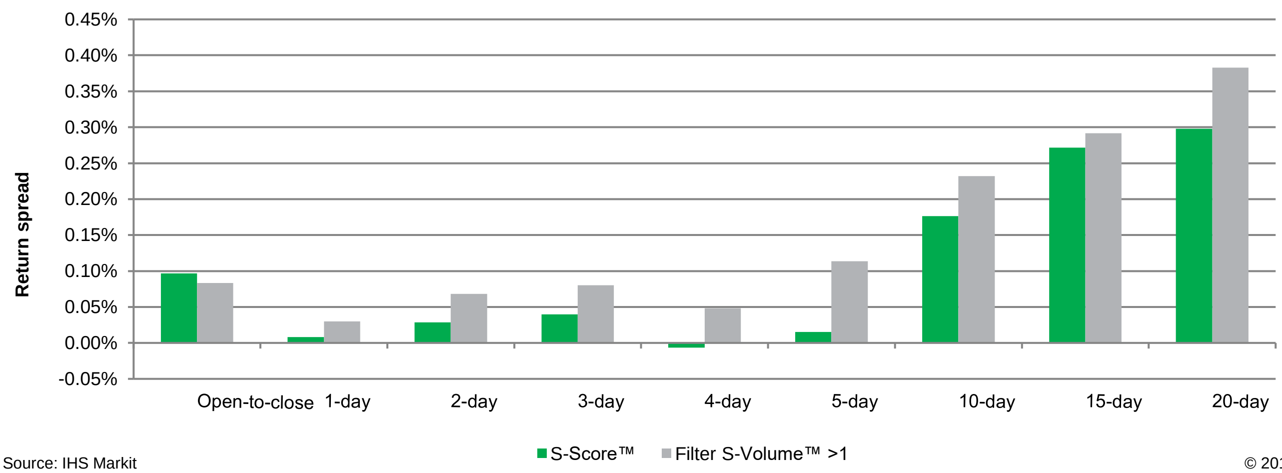

Source: IHS Markit

© 2019 IHS Markit

Table 1

S-Score™ average return spreads, 24 Aug 2015 – 31 Dec 2018

Strategy	Open-to-close (%)	1-day (%)	2-day (%)	3-day (%)	4-day (%)	5-day (%)	10-day (%)	15-day (%)	20-day (%)
S-Score™									
+/-3 tail spread	0.004	0.012	-0.027	0.015	-0.029	-0.068	0.111	0.131	0.195
+/-2 tail spread	0.097	0.008	0.028	0.040	-0.006	0.015	0.177	0.272	0.298
Quintile spread	0.017	0.011	0.013	0.015	0.006	0.014	0.065	0.038	0.013
Filter S-Volume™ >1									
+/-3 tail spread	-0.045	0.012	-0.020	0.031	-0.022	0.027	0.097	0.104	0.248
+/-2 tail spread	0.083	0.030	0.068	0.080	0.048	0.113	0.232	0.292	0.383
Quintile spread	0.033	0.009	0.010	0.021	0.008	0.027	0.097	0.091	0.041

Source: IHS Markit

© 2019 IHS Markit

Figure 4

S-Score[TM] +/-2 tail daily cumulative return

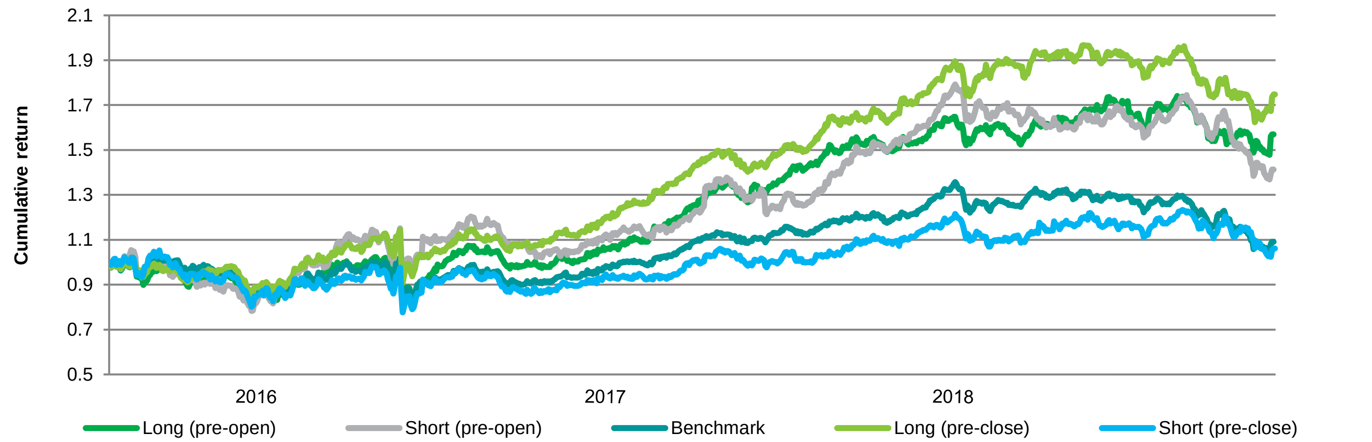

Source: IHS Markit © 2019 IHS Markit

Turning to the filtered strategies, outcomes were robust to these additional checks, under the caveat of more limited portfolio sizes. With the filter for S-Volume™>1, average open-to-close spreads at the +/- 2 tail averaged 0.083%, and exceeded the stand-alone factor results at each of the longer holding periods. For example, the 10-day average spread improved to 0.232% from 0.177% while the 20-day average spread jumped to 0.383% from 0.298%.

Overall, our empirical results suggest an immediate response to the signal at the +/-2 tail with some signs of short-term decay, though with the potential for success at longer holding periods, while the less stringent quintile tail results peaked around the 10- to 15-day horizons. Furthermore, we again remark on the skew to positive sentiment (as observed in Table A1 in the Appendix) producing a higher number of securities on average for S-Score™ >2 than <-2, a matter to be acutely aware of when implementing the signal in a portfolio.

Thus, given these findings that the distribution of scores demonstrates a skew toward positive sentiment signals, we now take into account long-only strategies. We remark that the portfolios remain reasonably sized for our analysis. For this study, we focus on performance for just the buy portfolio (Table 2) and again feature the results at the extreme +2 tail (Figure 5). Performance is reported as returns in excess of the average return of the underlying UK social media universe (LSE).

S-Score™ open-to-close excess returns at the +2 tail averaged 0.044%, exceeding both the +3 tail (0.022%) and quintile (-0.001%) strategies. At longer holding periods of 10 and 20 days, average excess returns reached 0.241% and 0.296%, respectively. At the broader quintile tail, excess returns were somewhat muted across each holding period, while, for the most stringent +3 tail, average excess returns tended to only slightly trail the +2 tail strategy.

As with our spread results cited above, excess returns for the +2 tail applying the S-Volume™ >1 filter were stronger than the stand-alone strategy out to the 30-day horizon. However, in this case, results increased in general with each incremental extension in holding period. Overall, the results confirm the benefits of the signal for long-only portfolio managers as a short term signal.

Figure 5

S-Score™ +2 tail excess returns, 24 Aug 2015 - 31 Dec 2018

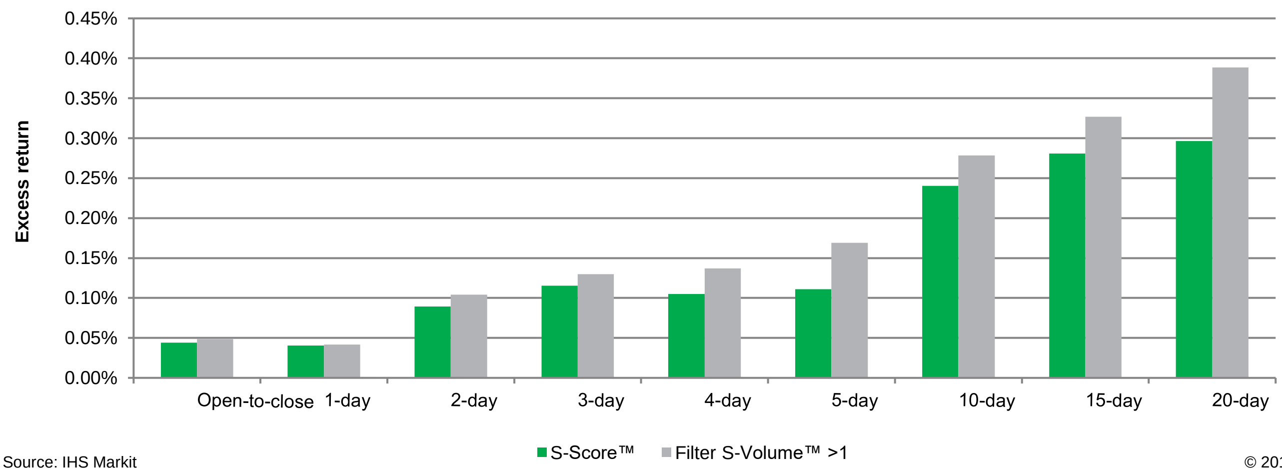

Source: IHS Markit

© 2019 IHS Markit

Table 2

S-Score™ average excess returns, 24 Aug 2015 – 31 Dec 2018

Strategy	Open-to-close (%)	1-day (%)	2-day (%)	3-day (%)	4-day (%)	5-day (%)	10-day (%)	15-day (%)	20-day (%)
S-Score™									
+3 - LSE	0.022	0.064	0.059	0.096	0.078	0.064	0.188	0.158	0.310
+2 - LSE	0.044	0.040	0.089	0.115	0.105	0.111	0.241	0.281	0.296
Quintile 1 - LSE	-0.001	0.012	0.017	0.023	0.022	0.027	0.064	0.064	0.064
Filter S-Volume™ >1									
+3 - LSE	0.046	0.022	-0.002	0.059	0.045	0.071	0.156	0.130	0.308
+2 - LSE	0.049	0.042	0.104	0.130	0.137	0.169	0.278	0.327	0.389
Quintile 1 - LSE	0.023	0.010	0.017	0.028	0.033	0.045	0.094	0.118	0.096

Source: IHS Markit

© 2019 IHS Markit

Relative Standard Deviation of Indicative Tweet Volume

Next we report summary performance results for our proprietary factor Relative Standard Deviation of Indicative Tweet Volume, a gauge of the tweet volume volatility of a stock. It is defined as the coefficient of variation of the indicative tweet volume over a 20-day period; in other words, the 20-day tweet volume standard deviation as a ratio of the 20-day tweet volume average, where a higher volatility is preferred whether associated with positive or negative sentiment.

Average spreads and excess returns for the quintile tail strategy are plotted in Figure 6 and summarized in Table 3. We observe a notable open-to-close spread of 0.224%, though spreads tend to tail off in the following days. However, over longer holding periods, the signal picks back up for the 10-day (0.159%), 15-day (0.229%) and 20-day (0.342%) time horizons, demonstrating persistence in the signal content after the initial pullback.

We also evaluate buy rated stocks for the case of long-only portfolios. On an open-to-close basis, we report an average excess return of 0.070% and again find an immediate reversion to more neutral results. Yet, the signal persistence out to longer holding periods is evident, with a 20-day average of 0.196%, again well in excess of the open-to-close results.

In summary, irrespective of sentiment, we find stocks with volatile tweet volumes over the past 20 days tend to outperform open-to-close, perhaps reflecting bias in the social media sentiment towards positive events that drive stocks to positive returns. In addition, these stocks outperform over the 20-day horizon, reflecting that the anomaly is not fully captured on day one and there is persistence in the returns.

Figure 6

Relative Standard Deviation of Indicative Tweet Volume quintile average returns, 24 Aug 2015 - 31 Dec 2018

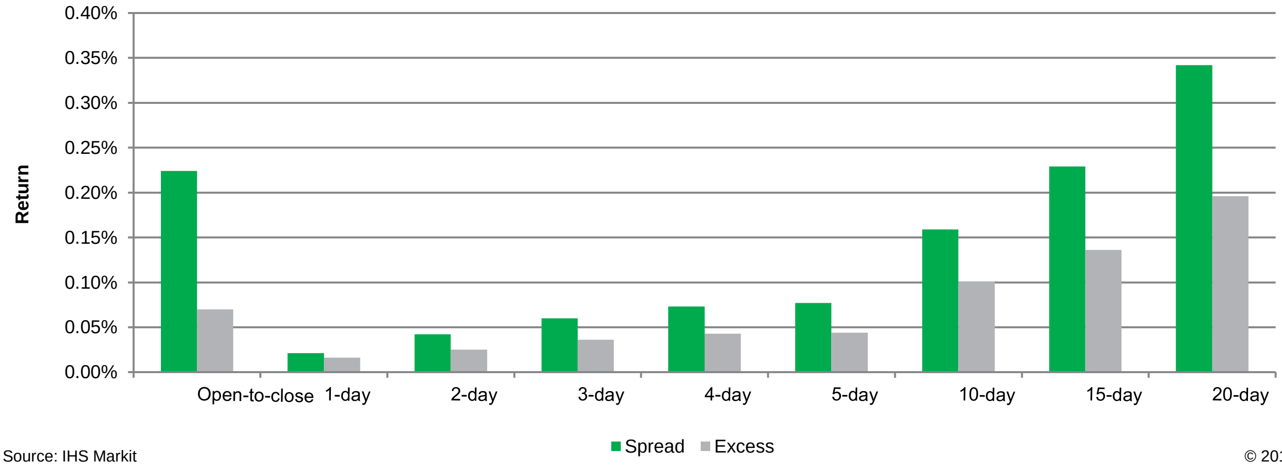

Source: IHS Markit © 2019 IHS Markit

Table 3

Relative Standard Deviation of Indicative Tweet Volume quintile average returns, 24 Aug 2015 – 31 Dec 2018

Strategy	Open-to-close (%)	1-day (%)	2-day (%)	3-day (%)	4-day (%)	5-day (%)	10-day (%)	15-day (%)	20-day (%)
Quintile spread	0.224	0.021	0.042	0.060	0.073	0.077	0.159	0.229	0.342
Quintile 1 - LSE	0.070	0.016	0.025	0.036	0.043	0.044	0.101	0.136	0.196

Source: IHS Markit © 2019 IHS Markit

Conclusion

We extend the Research Signals Social Media factor suite, originally introduced for US stocks, to the UK market. Our social media indicators are sourced in partnership with SMA, which analyzes social media data streams to estimate market sentiment. The factors are based on analysis of the text content in daily Twitter posts which are filtered for financial trading relevance and scored for market sentiment content.

We focus particularly on the main predictive factor, S-Score™, along with applying a filter for a minimum S-Volume™ of 2 to adjust the S-Score™ signal for confirmed information content in addressing random effects that may arise from sparse tweets. S-Score™ open-to-close return spreads at the +/-2 tail averaged 0.097%, persistent to 10-day (0.177%) and 20-day (0.298%) holding periods. On a cumulative basis, we report a pre-close spread return of 75% for buy-rated stocks versus 6% for sell-rated stocks and 9% for the market.

Results were robust to filters on minimum tweets and to long-only strategies. Applying the S-Volume™ >1 filter, open-to-close spreads for the +/-2 tail strategy averaged 0.083%, and exceeded the stand-alone factor results at each of the longer holding periods. For buy portfolios, S-Score™ (with S-Volume™>1) open-to-close excess returns at the +2 tail averaged 0.044% (0.049%) and increased in general with each incremental extension in holding period reaching 0.296% (0.389%) at 20 days, confirming the benefits of the signal for long-only portfolio managers.

Lastly, we report on one of our proprietary measures, Relative Standard Deviation of Indicative Tweet Volume, where we also find strong results. Stocks with volatile tweet volume pre-market tend to outperform open-to-close (spread: 0.224%; excess return: 0.070%), while these stocks also outperform at longer time horizons, reaching a spread of 0.342% out to 20 days (excess return: 0.196%).

Appendix

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.

Confidential. © 2019 IHS Markit. All rights reserved.

Research Signals factor definitions

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.

Tables

Table A1

S-Score™ average counts, 24 Aug 2015 – 31 Dec 2018

Strategy	Buy portfolio	Sell portfolio
S-Score™		
+/-3 tail	9	5
+/-2 tail	19	12
Quintile	65	65
Filter S-Volume™ >1		
+/-3 tail	6	4
+/-2 tail	14	9
Quintile	47	46

Source: IHS Markit

© 2019 IHS Markit

References

Deutsche Bank Securities Inc. “Translating Words to Numbers”, 16 June 2017.

IHS Markit Customer Support:

Support@ihsmarkit.com
Americas: +1 877 762 7548
Europe, Middle East, and Africa: 00800 6275 4800
Asia and the Pacific Rim: +65 6922 4210

Disclaimer

The information contained in this presentation is confidential. Any unauthorized use, disclosure, reproduction, or dissemination, in full or in part, in any media or by any means, without the prior written permission of IHS Markit Ltd. or any of its affiliates ("IHS Markit") is strictly prohibited. IHS Markit owns all IHS Markit logos and trade names contained in this presentation that are subject to license. Opinions, statements, estimates, and projections in this presentation (including other media) are solely those of the individual author(s) at the time of writing and do not necessarily reflect the opinions of IHS Markit. Neither IHS Markit nor the author(s) has any obligation to update this presentation in the event that any content, opinion, statement, estimate, or projection (collectively, "information") changes or subsequently becomes inaccurate. IHS Markit makes no warranty, expressed or implied, as to the accuracy, completeness, or timeliness of any information in this presentation, and shall not in any way be liable to any recipient for any inaccuracies or omissions. Without limiting the foregoing, IHS Markit shall have no liability whatsoever to any recipient, whether in contract, in tort (including negligence), under warranty, under statute or otherwise, in respect of any loss or damage suffered by any recipient as a result of or in connection with any information provided, or any course of action determined, by it or any third party, whether or not based on any information provided. The inclusion of a link to an external website by IHS Markit should not be understood to be an endorsement of that website or the site's owners (or their products/services). IHS Markit is not responsible for either the content or output of external websites. Copyright © 2019, IHS Markit™. All rights reserved and all intellectual property rights are retained by IHS Markit.

IHS Markit