



COINs2010

Predicting Stock Market Indicators Through Twitter “I hope it is not as bad as I fear”

Xue Zhang^{1,2*}, Hauke Fuehres², Peter A. Gloor²

¹National University of Defense Technology, Changsha, Hunan, China

²MIT Center for Collective Intelligence, Cambridge MA, USA

Elsevier use only: Received date here; revised date here; accepted date here

Abstract

This paper describes early work trying to predict stock market indicators such as Dow Jones, NASDAQ and S&P 500 by analyzing Twitter posts. We collected the twitter feeds for six months and got a randomized subsample of about one hundredth of the full volume of all tweets. We measured collective hope and fear on each day and analyzed the correlation between these indices and the stock market indicators. We found that emotional tweet percentage significantly negatively correlated with Dow Jones, NASDAQ and S&P 500, but displayed significant positive correlation to VIX. It therefore seems that just checking on twitter for emotional outbursts of any kind gives a predictor of how the stock market will be doing the next day.

Keywords: Twitter, economic indicator prediction, Web buzz analysis, coolhunting

* Corresponding author. Tel.: +1 617 253 7018

E-mail address: xuezhang@mit.edu

1. Introduction

Twitter is a very popular microblogging website, where users can update their status in tweets, follow the people they are interested, retweet others' posts and even communicate with them directly. Since it launched in 2006, its user base has been growing exponentially. As of June 2010, about 65 million tweets are posted each day, equaling 750 tweets sent each second (<http://en.wikipedia.org/wiki/Twitter>).

Recently, Twitter's popularity has drawn more and more attention of researchers from different disciplines. There are several streams of research investigating the role of Twitter. One stream of research focuses on understanding its usage and community structure. By examining the follower network, Java et al. (2007) found that there is a great variety in users' intentions. A single user may have multiple intentions and may even serve different roles in different communities. Huberman et al. (2009) analyzed the social interaction on Twitter, revealing that the driver of usage is a sparse hidden network among friends and followers, while most of the interaction links are meaningless.

Another stream of research concentrates on influence of Twitter users and information propagation. Cha et al. (2010) compared three different measures of influence – indegree, retweets and user mentions. They found that popular users who have high indegree are not necessarily influential in terms of spawning retweets or mentions. Also, Romero et al. (2010) showed that the correlation between popularity and influence is weaker than it might be expected, because most users are passive information consumers and do not forward the content to the network. By constructing a model capturing the speed, scale and range of information diffusion, Yang et al. (2010) claimed that some properties of the tweets themselves predict greater information propagation.

Besides the general understanding of Twitter, other researchers are interested in its prediction power and potential application to other areas. Asur and Huberman (2010) used Twitter to forecast box-office revenues of movies. They showed that a simple model built from the rate at which tweets are created about particular topics could outperform market-based predictors. In their study, Tumasjan et al. (2010) analyzed Twitter messages mentioning parties and politicians prior to the German federal election 2009 and found that the mere number of tweets reflects voter preferences and comes close to traditional election polls. Other researchers speculate that Twitter also could be used in areas such as tracking the spread of epidemic disease (Lampos, V. & Cristianini, N. 2010).

There is also prior work on analyzing correlation between web buzz and stock market. Antweiler and Frank (2004) determine correlation between activity in Internet message boards and stock volatility and trading volume. Other researches employed blog posts to predict stock market behavior. Gilbert and Karahalios (2010) used over 20 million posts from the LiveJournal website to create an index of the US national mood, which they call the Anxiety Index. They found that when this index rose sharply, the S&P 500 ended the day marginally lower than is expected. Besides the posts' contents itself, other properties of communication such as the number of comments, the length and response time of comments etc. are also helpful. Choudhury et al. (2010) modeled such contextual properties as a regression problem in a Support Vector Machine framework and trained it with stock movement. Their results are promising, yielding about 87% accuracy in predicting the direction of movement.

In recent years, we have been working on trying to predict market indicators by analyzing Web Buzz, predicting who will win an Oscar, or how well movies do at the box office (Doshi et. al 2009). Among other things we have correlated posts about a stock on Yahoo!Finance and Motley's Fool with the actual stock price, predicting the closing price of the stock of the next day based on what people say today on Yahoo!Finance, on the Web and Blogs about a stock title (Gloor et al. 2009). In this paper, we describe early work trying to predict stock market indicators such as Dow Jones, NASDAQ and S&P 500 by analyzing Twitter posts.

2. Method

The rising popularity of twitter gives us a novel way of capturing the collective mind up to the last minute. In our current project we analyze the positive and negative mood of the masses on twitter, comparing it with stock market indices such as Dow Jones, S&P 500, and NASDAQ. We collected the twitter feeds from one whitelisted IP for six months from March 30, 2009 to Sept 7, 2009, ranging from 8100 to 43040 tweets per day. According to Twitter, this corresponds to a randomized subsample of about one hundredth of the full volume of all tweets, as the total volume in 2009 was about 2.5 million tweets per day.

3. Results

3.1. Measuring Investor Fear by Tracking “Fear” Words

As is well known, emotional state can influence our decisions, and no doubt such choice includes stock market investment decision (Gilbert et. al 2010). When people are pessimistic or uncertain about the future, they will be more cautious to invest and trade. So capturing the collective mind – especially people’s mood – becomes one possible way to predict the stock market movement.

Twitter is a microblogging service in which users post very short messages: less than 140 characters, averaging 11 words per message (Connor 2010). This implies that most of the tweets have simple meaning, and even just one or two key words may capture the main topic. Inspired by this property, we decided to use mood words, for example “fear”, “worry”, “hope” etc., as emotional tags of a tweet. Then we measured collective emotion each day by simply counting all tweets containing such words. Table 1 below summarizes our results. The emotional words are divided into two groups: positive ones – hope and happy, and negative ones – fear, worry, nervous, anxious, and upset. Due to the different sample size everyday, the daily amount of each emotion is also highly variable. There were 4 to 49 “fear” tweets and 5 to 51 “worry” tweets per day; for “hope” the daily tweet numbers range from 54 to 467. More interestingly, we also find that the number of positive tweets is much higher than that of negative ones, more than double on average, which might suggest that people prefer optimistic to pessimistic words.

	Average per day	Min per day	Max per day
Tweet #	29758	8100	43040
Hope #	307	54	467
Happy #	260	37	1806
Fear #	28	4	49
Worry #	27	5	51
Nervous #	13	0	36
Anxious #	4	0	9
Upset #	14	2	25
Positive #	570	91	2204
Negative #	86	11	125

Table 1. Number of Twitter Posts from March 30, 2009 to Sept 7, 2009

3.2. Selection of Baseline

Next, we investigated against which baseline the number of tweets about a certain topic such as “hope, fear, and worry” should be measured. In our work we looked at three different baselines:

1. The number of tweets per day
2. The number of followers per day
3. The number of retweets per day

First we investigated the number of tweets about a certain topic in relation to the total number of tweets. The daily total number of tweets has been growing incrementally over the last years (Figure 1).




Figure 1. Growth in tweets per day (<http://mashable.com/2010/02/22/twitter-50-million-tweets/>)

In our own data sample we were using the Twitter “public timeline” function, implemented in such a way to deliver a more or less constant stream of messages per day. This stream allowed us to measure the percentage of emotional tweets among all the tweets. Using “hope” as an example, we defined $hope\%_t$ as the ratio between the number of “hope” tweets on day t and the amount of tweets we collected that day, comparing it with the stock market indicators on day $t+1$. Table 2 displays the correlation analysis result.

	Dow	NASDAQ	S&P 500	VIX
Hope %	- 0.381**	- 0.407**	- 0.373**	0.337**
Happy %	- 0.107	- 0.105	- 0.103	0.114
Fear %	- 0.208*	- 0.238*	- 0.200	0.235*
Worry %	- 0.300**	- 0.305**	- 0.295**	0.305**
Nervous %	- 0.023	- 0.054	- 0.021	0.015
Anxious %	- 0.261*	- 0.295**	- 0.262*	0.320**
Upset %	- 0.185	- 0.188	- 0.184	0.126
Positive %	- 0.192	- 0.197	- 0.187	0.188
Negative %	- 0.294**	- 0.323**	- 0.288**	0.301**

Table 2. Correlation Coefficient of emotional tweets percentage and stock market indicators (N=93) with total number of tweets per day as a baseline

** . Correlation is significant at the 0.01 level (2-tailed). * . Correlation is significant at the 0.05 level (2-tailed).

As external benchmark of investor fear we used the Chicago Board Options Exchange Volatility Index VIX, which strongly negatively correlated with Dow, S&P 500, and NASDAQ, which is not surprising, as the spread of stock options on a given day is used to calculate VIX. Initially we expected that the correlation between optimistic mood and stock market indicators would be positive, and the pessimistic mood would negatively correlate. Surprisingly, we found positive correlation for all of them with VIX, and negative correlation with Dow, NASDAQ and S&P500. This implies that people start using more emotional words such as hope, fear and worry in times of economic uncertainty, independent of whether they have a positive or negative context.

As our second candidate for a baseline we investigated the total number of followers per day. Follower is a key concept in Twitter, it is commonly seen as a measure of popularity. It is likely that the more followers a user has, the more people s/he can affect. In particular, the bigger the audience of one pessimist is, the more people may be infected and feel the same negative way. We analyzed the correlation between percentage of potential emotional audience and stock market indicators. For instance, we added all the follower numbers of “worry” tweets of day t and divided it by the total number of followers on that day, ($worryfollower\%_t$ in Table 3) then comparing it with Dow_{t+1} , $NASDAQ_{t+1}$ and $S \& P500_{t+1}$. The correlation coefficients are 0.143, 0.149 and 0.146 separately, which are relatively lower than we expected. As can be seen in Table 3, this index is therefore not a good predictor of stock market indices.

	Dow	NASDAQ	S&P 500	VIX
Hope-followers %	-0.086	-0.048	-0.077	0.023
Happy-followers %	-0.19	-0.181	-0.188	0.178
Fear-followers %	-0.005	-0.051	0.012	0.026
Worry-followers %	0.143	0.149	0.146	-0.008
Nervous-followers %	0.074	0.089	0.076	-0.108
Anxious-followers %	-0.156	-0.177	-0.177	0.187
Upset-followers %	-0.106	-0.116	-0.103	0.074

Table 3. Correlation Coefficients of percentage of potential emotional audience and stock market indicators (N=93)

Finally we looked at the number of retweets per day, based on the hypothesis that the more a topic is being picked up and retweeted by others, the more it is relevant. In an accumulated way, the total number of retweets is a proxy for the activity of the twitter users on a particular way.

	Average per day	Min per day	Max per day
Retweet #	1083	221	1884
Hope-retweet #	7	0	15
Happy-retweet #	9	0	40
Fear-retweet #	3	0	9
Worry-retweet #	1	0	51

Table 4. Number of retweets from March 30, 2009 to Sept 7, 2009

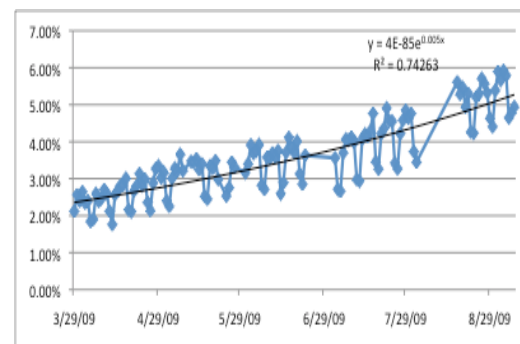


Figure 2. Percentage of retweets per day

Table 4 above illustrates the number of retweets about a certain topic per day. The retweet numbers range from 221 to 1884, nearly 3% – 5% of the tweets. As Figure 2 shows, the retweets percentage displayed an exponential growth too. We also found that there were about 40% less retweets at weekends (the nodes underneath the black line in figure 2 are weekends). We speculate that on weekends, active tweeters have the time to send more original tweets, while during the week they pick up tweets from others they find worthwhile retweeting. This means, however, that they “stake their reputation” on others’ tweets during the weekdays.

Next, we analyzed the correlation between the emotional retweets percentage and the stock market indicators. Again, taking “hope”, for example, we defined $hope - retweet\%_t$ as the ratio between the number of retweets which contain “hope” on day t and the amount of retweets on that day, then we compared it with the stock market indicators on day t+1. Table 5 below displays the correlation analysis result. Obviously, number of retweets is a better baseline than number of followers, but simply taking the total number of tweets gives the best results. This is not surprising, however, because the number of retweets containing “hope” is much lower than the number of tweets containing “hope”, which means that the fluctuation in the results is much higher, therefore leading to smaller sample size and less significant correlations. We speculate that the correlations would have been higher if we would have been able to collect a larger subsample of all the tweets.

	Dow	NASDAQ	S&P 500	VIX
Hope-retweet %	- 0.139	- 0.156	- 0.158	0.119
Happy-retweet %	- 0.011	- 0.008	0.003	- 0.015
Fear-retweet %	- 0.258*	- 0.245*	- 0.253*	0.280**
Worry-retweet %	0.037	0.036	0.047	- 0.083

Table 5. Correlation Coefficients of emotional retweets percentage and stock market indicators (N=93)
 **. Correlation is significant at the 0.01 level (2-tailed). *. Correlation is significant at the 0.05 level (2-tailed).

3.3. Time Lag of Prediction

We also investigated how much the discussion on twitter precedes fluctuations on the stock market. To put it in other words, the question is how long it takes the stock market to react on the buzz on Twitter.

All analysis in Section 3.2 focused on the Twitter buzz one day before the trading day. However, this way, a considerable portion of information has been wasted. For example, when we analyze Monday’s Dow, just the data of Sunday are used. All the data about Fridays and Saturdays becomes useless, because weekends are not trading days. In order to find out the time lag of prediction and achieve full use of resources, we created a simple twitter-volatility index averaging the buzz of day_t , day_{t-1} and day_{t-2} to predict the stock market indicators of day_{t+1} . This index displays significant negative correlations to Dow, NASDAQ and S&P500, and significant positive correlation to VIX (in this section we only use the emotional tweet percentages with the total number of tweets as baseline as predictor). Among all the emotional words, *hope*, *fear* and *worry* work best in this analysis. We also added them together to test whether the sum of them might improve results further, which it turns out it did not (see table 6).

	Dow	NASDAQ	S&P 500	VIX
Hope%	- 0.381**	- 0.407**	- 0.373**	0.337*
Hope%-2 mean	- 0.618**	- 0.631**	- 0.607**	0.518**
Hope%-3-mean	- 0.737**	- 0.738**	- 0.724**	0.621**
Fear%	- 0.208 *	- 0.238 *	- 0.2	0.235*
Fear%-2-mean	- 0.259*	- 0.285**	- 0.253*	0.312**
Fear%-3-mean	- 0.346**	- 0.368**	- 0.342**	0.403**
Worry%	- 0.3**	- 0.305**	- 0.295**	0.305*
Worry%-2-mean	- 0.421**	- 0.415**	- 0.414**	0.410**
Worry%-3-mean	- 0.472**	- 0.460**	- 0.467**	0.459**
Hope+Fear+Worry%	- 0.379**	- 0.405**	- 0.37**	0.347*
Hope+Fear+Worry%-2-mean	- 0.612**	- 0.625**	- 0.6**	0.532**
Hope+Fear+Worry%-3-mean	- 0.726**	- 0.728**	- 0.713**	0.633**

Table 6. Correlation Coefficient of average emotional tweets percentage and stock market indicators (N=93)
 **. Correlation is significant at the 0.01 level (2-tailed). * . Correlation is significant at the 0.05 level (2-tailed).

The picture below visualizes the negative correlation between Dow (blue) and “hope+fear+worry%-3-mean” (green) in the period March 30, 2009 to Sept 7, 2009.




Figure 3. Correlation between “hope, fear and worry-3 mean” and Dow Jones Industrial Average

4. Discussion

To put it in simple words, when the emotions on twitter fly high, that is when people express a lot of hope, fear, and worry, the Dow goes down the next day. When people have less hope, fear, and worry, the Dow goes up. It

therefore seems that just checking on twitter for emotional outbursts of any kind gives a predictor of how the stock market will be doing the next day.

In this paper, we have presented very preliminary results, much more work is needed to verify it further.

References

- Antweiler, W. & Frank, M.Z. (2004). Is All That Talk Just Noise? The Information Content of Internet Stock Message Boards. *Journal of Finance* Vol. 59, No. 3 (Jan., 2004), pp.1259-1294.
- Asur, S. & Huberman, B. A. (2010). Predicting the Future With Social Media. <http://arxiv.org/abs/1003.5699>.
- Boyd, danah, Scott Golder, & Gilad Lotan (2010). Tweet, Tweet, Retweet: Conversational Aspects of Retweeting on Twitter. HICSS-43. IEEE: Kauai, HI, January 6.
- Cha, M., Haddadi, H., Benevenuto, F., & Gummadi, K. P. (2010). Measuring User Influence in Twitter: The Million Follower Fallacy. 4th International AAAI Conference on Weblogs and Social Media (ICWSM), 2010.
- Choudhury, M. D., Sundaram, H., John, A. & Seligmann, D. D. (2010). Can Blog Communication Dynamics be Correlated with Stock Market Activity? Proceedings of the nineteenth ACM conference on Hypertext and hypermedia, 2010.
- Connor, B., Balasubramanian, R., Routledge, B. R. & Smith, N.A. (2010). From Tweets to Polls: Linking Text Sentiment to Public Opinion Time Series. 4th International AAAI Conference on Weblogs and Social Media (ICWSM), 2010.
- Doshi, L. Krauss, J. Nann, S. Gloor, P. Predicting Movie Prices Through Dynamic Social Network Analysis. Proceedings COINs 2009, Collaborative Innovations Networks Conference, Savannah GA, Oct 8-11, 2009
- Gilbert, E. & Karahalios, K. (2010). Widespread Worry and the Stock Market. 4th International AAAI Conference on Weblogs and Social Media (ICWSM), 2010.
- Gloor, P. & Zhao, Y. (2004) TeCFlow - A Temporal Communication Flow Visualizer for Social Networks Analysis. ACM CSCW Workshop on Social Networks. ACM CSCW Conference, Chicago, Nov. 6. 2004.
- Gloor, P., Krauss, J., Nann, S., Fischbach, K. & Schoder, D. (2009). Web Science 2.0: Identifying Trends through Semantic Social Network Analysis. IEEE Conference on Social Computing (SocialCom-09), Aug 29-31, Vancouver, 2009.
- Huberman, B. A., Romero, D. M., & Wu F. (2009). Social networks that matter: Twitter under the microscope. *First Monday*, 14(1), 2009.
- Java, A., Song, X., Finin, T. & Tseng, B. (2007). Why We Twitter: Understanding Microblogging Usage and Communities. 9th WebKDD and 1st SNA-KDD workshop on web mining and social network analysis, 2007.
- Lamos, V. & Cristianini, N. (2010). Tracking the flu pandemic by monitoring the Social Web. IAPR 2nd Workshop on Cognitive Information Processing (CIP 2010), 14-16 Jun 2010.
- Romero, D. M., Galuba, W., Asur, S. & Huberman, B. A. (2010). Influence and Passivity in Social Media. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1653135.
- Tumasjan, A., Sprenger, T. O., Sandner P. G. & Welp I. M. (2010). Predicting Elections with Twitter: Whan 140 Characters Reveal about Political Sentiment. 4th International AAAI Conference on Weblogs and Social Media (ICWSM), 2010.
- Weng, J., Lim, E., Jiang, J. & He, Q. (2010). TwitterRank: Finding Topic-sensitive Influential Twitterers. WSDM, 2010.
- Yang, J. & Counts, S. (2010). Predicting the Speed, Scale, and Range of Information Diffusion in Twitter. 4th International AAAI Conference on Weblogs and Social Media (ICWSM), 2010.