



LEUPHANA
UNIVERSITÄT LÜNEBURG

FINAL

**Mining User-Generated
Financial Content to Predict
Stock Price Movements**

Andreas Mastel
Jürgen Jacobs

Technical Reports and Working Papers
Leuphana Universität Lüneburg

Hrsg. der Schriftreihe FINAL: Ulrich Hoffmann
Volgershall 1, D-21339 Lüneburg

Mining User-Generated Financial Content to Predict Stock Price Movements

Andreas Mastel
Jürgen Jacobs

This paper presents the results of a Master's Thesis completed by Andreas Mastel under the guidance of Jürgen Jacobs.

Abstract:

Sentiment extraction from user-generated online content to predict stock price movements has become an active research field. This paper gives an overview of common approaches to this topic and analyzes the content generated by the financial social network Seekingalpha.com. The first finding is that a large proportion of users' attention is focused on only a few stocks. Regarding these stocks it can be shown that sentiment is significantly driven by past abnormal performance. Only the sentiment of premium users contains some degree of predictive power. Generally, the users' sentiment is consistent with a naïve momentum mentality.

Keywords: Stock price prediction, event study, sentiment analysis

1 Introduction

Numerous approaches to predict future stock price movements can be found in literature. So far, none of these approaches clearly refuted the implications of the efficient market hypothesis (EMH) which was first introduced by FAMA (1970) and later revised by FAMA (1991). The semi-strong form of the EMH states that it is impossible to predict price changes and outperform a buy-and-hold strategy, since all publicly available information is already incorporated in the prices and new information gets incorporated very quickly. One of the newer approaches to disprove the EMH is to verify that investors' sentiment may be used as a predictor for price movements. BAKER AND WURGLER (2006) argue that it is essential for researchers to incorporate sentiment into their descriptive models of expected returns. The authors challenge the research community to find appropriate measures for it. They identify six proxies to capture various states of investor sentiment through time and unveil that these have cross-sectional effects on price changes.

Several studies have been formed around the assumption that a good proxy for investors' sentiment can be found in user-generated content on the internet. VAN BOMMEL (2003) examines motivations and effects of spreading stock tips. He points out that even investors with a small trading capacity can influence stock prices by

information diffusion. He shows that spreading rumors makes economic sense. The rumormonger as well as followers profit from trading on rumors. In our empirical study, we do not try to differentiate between the various roles and motivations of contributors within the online community. We just differentiate between quality articles published by experts and contributions by the online community as a whole. To answer the question whether experts' or crowds' opinion will drive the stock price one has first to gather online content expressing opinions about stocks. Thereafter one has to infer a sentiment measure. This is usually done by aggregating online information about the regarded stocks. Finally, one has to test whether these measures contain valuable information for the prediction of stock prices.

Several methods have been introduced to extract and to aggregate sentiment information and also to test whether it has predictive power. Chapter 2 gives an overview of these methods. Chapter 3 first describes the data sample of user-generated content (subchapter 3.1) and stock market data (subchapter 3.2). Subchapter 3.3 provides a description of the applied sentiment extraction method and subchapter 3.4 introduces the event study methodology used. Chapter 4 shows the results of the analysis. Chapter 5 concludes the study's findings.

2 Related Work

Around the turn of the millennium, the internet developed in a way that O'REILLY (2007) summarized under the term "Web 2.0". Internet users intensively started to share content through different kinds of internet-services, e.g. photo sharing services, file sharing services, weblogs, online message boards and online social networks. With this development, several new research fields in analyzing user-generated content emerged. For example, researchers started to monitor the behavior and opinion of big crowds that were expressed in online content (PANG AND LEE (2008)). Several authors tried to find out whether crowds' opinion contains valuable information for the prediction of stock prices.

The first systematic work in this area has been published by WYSOCKI (1998). His sample is based on posts from the Yahoo! message boards. He uncovers several cross-sectional firm characteristics which have a positive effect on the volume of the message board posts. Additionally, he tests an approach to predict subsequent stock market activity, i.e. abnormal stock returns and trading volume. For 50 companies with the highest posting volume between January and August 1998, his findings show that overnight posting activity has a positive effect on the next day's trading volume, abnormal return variance and abnormal returns. Contrarily, daytime posting activity does not show this effect. He attributes his findings to the quick representation of information in stock prices. One important characteristic of his work is that he concentrates solely on message counts and not on the content of the posts. This is one constraint later authors made good for.

At about the same time BAGNOLI, BENEISH, WATTS (1999) collect and analyze earnings forecasts from Thomson's First Call service and compare them to unofficial forecasts from different sources. These unofficial forecasts are commonly referred to as "whispers". The authors collect them from several webpages, news wires and other

sources. The resulting sample data ranged from January 1995 to May 1997. The authors find that, compared with their official counterparts, whispers tend to be published more recently before the release date of the respective earnings report. Also, whispers seem to forecast earnings more accurately and tend to overestimate them, unlike the First Call forecasts which tend to underestimate them. They can show that a trading strategy based on the difference between whispers and First Call consensus yields significant returns. It is to mention that their data sample is highly biased towards firms from the technology sector. Besides, the processing of the earnings forecasts is made manually, which means that the method can't be applied to very big data samples without extraordinary effort.

In contrast to their work, a later study by DEWALLY (2008) states that whispers do not contain more information than official forecasts since they are not more accurate than the analysts' consensus. Strangely, the author still finds a profitable trading strategy of shorting stocks for which whispers forecast an outperformance.

In one of his former studies, DEWALLY (2003) analyzes stock recommendations from the two newsgroup sites "alt.invest.penny-stocks" and "misc.invest.stocks". Newsgroups are similar to online discussion boards as users can also write posts and publish them to other users. The author's data sample contains two time periods which represent two different market moods. The first period in April 1999 is labeled as "up market" and the second period in February 2001 is called "down market". His findings show that the majority of recommendations are positive. Surprisingly, this also applies to the down market. Positive recommendations are often preceded by strong prior stock performance. This behavior conforms to a naïve momentum mentality. Notably, in both market mood situations positive recommendations are followed by weak stock performance. In summary, he can't show that newsgroup recommendations contain predictive information.

TUMARKIN AND WHITELAW (2001) find no evidence that discussion board messages contain valuable information for the prediction of stock market activity. They collect a sample of posts from the popular message board RagingBull.com. The board posts in their sample contain special fields in which the users express their sentiment towards certain stocks directly. This information is processed into an opinion measure which is then used to group the board posts into positive and negative groups. Then, they conduct an event study to determine the impact of high message volume on security prices and trading volume. An event is triggered when the message volume exceeds a certain threshold. Additionally, the authors use a one-day-lagged VAR-model to explore the dependencies between the stock price, the trading volume, the number of board messages and the sentiment measure. None of these methods show that message board activity can help to predict future market activity. Yet, it is observable that trading volume remains high for one day after the event but this effect is seen as economically insignificant. In line with DEWALLY (2003), positive sentiment on stocks is preceded by a strong abnormal performance.

With the exception of WYSOCKI (1998), who makes use of message volumes, all other authors cited so far focus on users' sentiment towards stocks. However, the sentiment is not automatically derived from the texts. It is either extracted manually or predefined meta-data that already contains this information is used. There come some practical restrictions with both of these approaches. On the one hand, it becomes more or less impossible to process user-generated content manually, as it has

increased dramatically in volume. On the other hand most data sources don't provide the researcher with meta-information about the users' sentiment.

Automated opinion extraction from texts is a very complex task which needs to be performed with a lot of care. One of the reasons for this complexity is that ambiguity often impedes the analysis. In many cases, texts cannot be assigned to one distinct sentiment class because they either include many subjects with differing sentiments or the meaning of the text is completely ambiguous. Also, texts may lack any sentiment at all. Accordingly, high classification error-rates are common for sentiment detection. Many researches have accepted this fact. As a consequence, they use aggregated information, in order to reduce the impact of individual misclassifications. Besides information aggregation, a limitation to only a few sentiment classes is common practice. The most important sentiment expressions in a financial context are either positive or negative, because their polarity can be related to corresponding buying or selling decisions. Therefore, research projects usually focus on the polarity between a "buy" sentiment and a "sell" sentiment. Many researchers also incorporate a "hold" sentiment in their studies. Other mood states are excluded. There exist more obstacles in this field, but reviewing them is not the goal of this paper. PANG AND LEE (2008) give an overview of opinion mining methods and the challenges raised by opinion-oriented information-seeking systems.

Sentiment classification for stock price predictions is performed in at least two ways. One can either concentrate solely on word occurrences or also consider grammatical and contextual information. Methods in which only word counts are regarded are referred to as bag-of-words approaches. Alternatively, one can use part-of-speech tagging or try to detect contextual polarity, e.g., whether the combination of a positive expression and a negation results in a negative expression. Such structural components remain unseen with the bag-of-words approach. WILSON ET AL. (2005), for example, use a method for the extraction of contextual polarity on a sentence-basis. DAS AND CHEN (2007), for example, combine different classification algorithms to form a complex knowledge-discovery architecture which allocates message board posts to the classes buy, hold and sell. Among other methods, they include part-of-speech tagging into their method set. A particularly interesting approach in this context is to look for adjectives, adverbs and their surrounding terms only. This approach is based on the assumption that adjectives and adverbs are more expressive regarding sentiments than other parts of speech and require a greater weight in the classification process.

However, most studies use the bag-of-words approach because it is the most straightforward method and avoids strong assumptions regarding linguistic patterns. Most commonly for these methods, a term-document matrix is created which represents the occurrences of words from a word vector across all texts. On top of this information, one can either analyze word counts directly or use them to train a supervised algorithm.

In the context of sentiment analysis, the words in the word vector are often restricted to pre-classified terms for which the connotation is already known. Dictionaries for different purposes have been compiled. The Harvard Psychosociological Dictionary IV-4, for example, includes word lists for several mood dimensions which can generally be applied to extract sentiment information

from texts. LOUGHRAN AND MCDONALD (2011) argue that the negative words from this dictionary include terms which typically do not have a negative meaning in a financial context. They introduce a dictionary which is tailored to financial sentiment expressions. A good example for the use of the Harvard Dictionary IV-4 is provided by TETLOCK (2007). His work attempts to find a relationship between news media and stock market activity. He analyzes the daily Wall Street Journal (WSJ) column “Abreast of the Market” for the time between 1984 and 1999. He conducts a factor analysis and selects the categories “negative” and “weak” as the most important semantic components from the 77 categories of the Harvard Dictionary IV-4. The selected two categories form a new one that he calls “pessimism”. The words belonging to the aforementioned categories from the WSJ columns are counted and used to predict the Dow Jones Industrial Average’s (DJIA) market prices and market trading volume. The analysis is based on vector-autoregressive modeling. Similar to the concept of Granger causality (GRANGER (1969)), the explanatory power of past media pessimism on DJIA returns is measured by the significance of the estimates for the pessimism’s coefficients. The author finds that high media pessimism predicts a pattern of falling prices that is followed by a reversion. This effect is stronger for small stocks, what is consistent with the assumption that media content is linked to the behavior of individual investors, who own a disproportionate fraction of small stocks. Furthermore, unusually high or low media pessimism predicts high market trading volume. Finally, low market returns lead to high media pessimism. Based on this work, TETLOCK ET AL. (2008) extend the analysis to news messages. Their study is based on the proportion of negative words in news stories of the Wall Street Journal (WSJ) and the Dow Jones News Service (DJNS) from 1980 to 2004. The study is limited to stories about individual S&P 500 firms. The authors conduct a time-lagged OLS regression analysis and an event study to examine the predictability of stock returns and quarterly accounting earnings. In the event study, the publication of exceptionally negative or positive stories is seen as an event. Both approaches of their work provide statistically significant results and show that negative news stories predict negative firm fundamentals and the corresponding market’s reaction. Notably, negative messages that include the word stem “earn” predict earnings and returns more precisely. This effect is explained by the assumptions that earnings-related stories are better predictors for accounting measures and that the market particularly reacts to news stories that relate to fundamentals. Probably this finding also supports the assumption that the terms in a bag-of-words approach are unequally informative since the occurrence of one particular word may have major influence on the predictive power of the text. For more recent studies in the field of news sentiment aggregation to predict stock prices, one may refer to TETLOCK (2011) and SCHUMAKER AND CHEN (2009). The latter work also discusses a variety of textual analysis techniques that produce valuable results for stock price predictions.

Besides the direct analysis of word counts, there is the alternative of using supervised classifiers. A major drawback of this approach is that it requires a pre-classified training sample of texts. Pre-classification has to be achieved manually and thus is work-intensive and entails subjective human judgment.

ANTWEILER AND FRANK (2004) were among the first authors who extract sentiment by applying Naïve Bayes and Support-Vector-Machine methods to classify online discussion board posts into a buy, sell or hold tone. In financial terms, a

common optimistic opinion regarding a certain stock or the general situation in the market is often being referred to as “bullishness”. The opposite is called “bearishness”. The authors try to answer the question whether bullishness in discussion board messages helps to predict returns. They also test whether general disagreement among users is associated with more trades and whether bullishness or the level of message posting will help to predict volatility. For a sample of more than 1.5 million posts referring to 45 firms that belong to the Dow Jones Industrial Average (DJIA) and the Dow Jones Internet Commerce Index (XLK), the authors produce bullishness and agreement measures based on the counts of classified messages across time. Agreement in this case is defined as a function of the bullishness measure. They find that a positive shock to board message volume predicts negative returns on the next day but bullishness has no significant effect in this case. Message volume and disagreement helps to predict subsequent trading volume, but against expectations greater disagreement predicts fewer trades and not more. Lastly, message volume predicts volatility. This effect was found within a range of methods of modeling market volatility, including realized volatility methods and GARCH methods.

As already mentioned, DAS AND CHEN (2007) introduce an approach with several knowledge discovery algorithms in combination. Their work puts a strong focus on the development of a new natural language processing algorithm to classify stock messages. They base their analysis on posts from the Morgan Stanley High-Tech Index (MSH) message boards. Every text is classified as having a buy, hold or sell sentiment. The authors introduce a voting mechanism of 5 classifiers with the purpose of producing a better signal to noise ratio in the extraction of sentiment. This already leads to acceptable classification results. The main improvement of sentiment classification is achieved after applying another preprocessing step. Following the first classification, messages are additionally labeled with an optimism score that is based on the in-text ratio of positive to negative terms from the Harvard Dictionary IV-4. The descriptive statistics of this score shows that members of the buy class show the highest optimism score average and members of the sell class the lowest. After the authors filter out all texts that deviate strongly from the mean of their sentiment class, they classify the texts again on basis of the reduced sample. With this additional preprocessing step, the authors intend to remove ambiguous texts. Indeed, the classification results improve significantly and are comparable to widely used Bayes classifiers like the Rainbow algorithm of MCCALLUM (1996), which has been applied by ANTWEILER AND FRANK (2004). Moreover, the authors achieve the desired noise-reduction as they reduce the number of false-positives remarkably. This dive into the procedure of their work shows that also the combination of supervised classifiers and dictionary based sentiment classification has been proven to yield positive results. DAS AND CHEN (2007) build a sentiment index based on the daily counts of buy and sell messages. The levels of this index correlate to contemporaneous MSH index levels respectively individual stock price levels. Yet, the correlation between the daily changes of the sentiment index and the MSH index returns respectively individual stock returns is weak. No predictive power of sentiment can be proven in this study. Consistent with ANTWEILER AND FRANK (2004) the authors find a relationship between message board activity and contemporaneous market activity.

Most recently, the research in predicting stock prices from sentiment in user-generated content shifted to new data sources. As social networks generate a massively increasing amount of content, many researchers started to scan them for information about the opinion of big crowds. Especially the microblogging service Twitter.com got a lot of attention because it apparently became the speaking tube for a big proportion of internet users. SPRENGER AND WELPE (2010) analyze a sample of 250000 stock related Twitter messages (tweets) in the year 2010 and investigate the predictive power of their sentiment. Tweets are short messages that are limited to 140 characters. In line with ANTWEILER AND FRANK (2004) they also use a Naïve Bayes classifier to build a bullishness measure from the classified tweets. They also adopt the measure for agreement from their predecessors' work and replace message board volume by tweet volume as their third measure. They formulate several hypotheses on the effect of these three measures on market indicators like abnormal stock returns, trading volume and stock price volatility. Additionally, they analyze several aspects of information aggregation from twitter specific structures like retweets, followers and mentions. To test their market related hypotheses, they investigate on the contemporaneous and time-lagged relationships with regression models and also conduct an event study for extreme fluctuations of the bullishness index. The regression analysis shows that an increase of the bullishness index is followed by higher subsequent returns. Also, increased message volume precedes higher trading volume. Slightly significant results show that an increase in disagreement entails higher trading volume. The event study shows that stock microbloggers follow a contrarian strategy because high bullishness is preceded by low abnormal performance and low bullishness by high abnormal performance. The direction of the abnormal returns directly after the event is consistent with the Twitter.com users' expectations. A bullish event is followed by a short period of positive abnormal returns and vice versa.

Another work by BOLLEN ET AL. (2011) analyzes a sample of over 9 million tweets from the year 2008. The authors make use of two mood assessment tools which not only account for positive and negative sentiment but also for other mood dimensions. One of these tools is based on the work of WILSON ET AL. (2005). It is called OpinionFinder (OF) and provides functions for the recognition of contextual polarity, but the authors don't make use of them. They only create a sentiment index that is based on occurrences of positive or negative words from the OF sentiment lexicon. Additionally, they create a second mood analysis tool that is derived from a well-known psychometric instrument called Profile of Mood States (POMS). The authors call their derivative "GPOMS". Different from the OF approach, GPOMS can classify a text into the mood states calm, alert, sure, vital, kind and happy. For each of these mood dimensions, the authors produce a daily index. They determine if one of the mood indices Granger-causes the subsequent daily returns of the Dow Jones Industrial Average (DJIA) for a time frame from February to November 2008. The sentiment index from the OF lexicon produces significant results only for a one day lag. Other than that, the calm mood from the GPOMS tool shows highly significant results from 2 up to 6 day lags. Besides, the happy mood has significant statistics for 6 day lags. Surprisingly, the mood state calm seems to have the most significant influence on the DJIA index returns. In addition to linear time-series modeling, BOLLEN ET AL. (2011) conduct a non-linear analysis of the DJIA and the already discussed mood time series. They train a Self-organizing

Fuzzy Neural Network (SOFNN) to predict upward and downward directions for the next day's DJIA movement. The training data consists of the past three DJIA values in addition to several permutations of the past values of their mood time series. The algorithm predicts the direction of the next-days DJIA return with an impressive accuracy of 86.7 percent. One has to keep in mind that the test sample was binomially distributed with a 50.0 percent chance of success. Remarkably, the inclusion of the mood state calm not only improves the classification accuracy significantly, it also outperforms every other mood state. However, the test sample consists of only 15 trials what roughly corresponds to about 3 weeks of trading.

Similar to the aforementioned prediction of stock movement directions with SOFNN methods, the work of CHOUDHURY ET AL. (2008) shows an example of an Support Vector Machine regression framework which is trained on blog posts from a technology-oriented online community called Engadget.com. Their framework is aimed to predict the direction and the magnitude of abnormal returns for the stocks Apple, Microsoft, Google and Nokia in a time period between January 2007 and November 2007. In addition to past index values, they incorporate structural information of the weblog like, for example, the number of posts, the length and response time of comments and other variables into their training sample. The authors achieve a prediction accuracy of 86.59% for the prediction of the direction and 78% accuracy for the magnitude of the movement.

ZHANG ET AL. (2010) also analyze tweets and find a negative correlation between the volume of tweets containing words like hope, fear and worry and the subsequent values of indices like the DJIA, NASDAQ-100 and S&P 500. Moreover, they find a positive correlation to the future values of the volatility index VIX.

Twitter.com is not the only interesting online social network to be analyzed. CHEN ET AL. (2011) focus on the social network Seekingalpha.com (SA). Their work is based on user-generated articles that were published between 2006 and 2010. A more detailed description of the structures of SA will follow in the next section of this work. The authors build a negativity measure that is based on negative words which were compiled by LOUGHRAN AND McDONALD (2011). Their measure is defined as the fraction of negative words to all words across all articles that are concerning a certain company on a certain day. Additionally, they calculate the same measure for WSJ news articles to test if a similar effect can be shown for traditional media-outlets. Initially, they inspect cumulative abnormal returns for two portfolios based on stocks that were discussed in articles which belong to the daily-calculated upper (bearish) or lower (bullish) tercile of the negativity measure distribution. The direct effect of negativity on abnormal returns is investigated with a regression analysis. Among other independent variables, a high negativity score is related to lower contemporaneous returns and also lower subsequent returns. This effect is stable even after controlling for the negativity measure of WSJ articles. They show that their findings are stronger for companies with articles that are more closely followed by market participants and for companies which are held mostly by retail investors.

The literature overview presented so far demonstrates that various data sources and various methods to interpret the content concerning stocks and accounting information have been used. There are also various methods to uncover the relationship between the interpreted content and market activity or price changes. A common approach is to search for a direct relation to time-lagged returns or abnormal

returns with the method of panel regressions. In doing so, it is to consider that there are widely applied control variables which were proven to have an effect on the predictability of stock returns. For example, it is common to control for the size-effect which assumes that market capitalization explains some cross-sectional variation in stock returns (FAMA AND FRENCH (1992)). Two more well accepted control variables are the book equity to market equity ratio and the earnings-price ratio. All these variables are said to incorporate cross-sectional predictable patterns and should be considered when searching for a direct relation between aggregated online content and subsequent stock returns. Examples of the aforementioned regression method can be found in WYSOCKI (1998), ANTWEILER AND FRANK (2004), TETLOCK ET AL. (2008), SPRENGER AND WELPE (2010) or CHEN ET AL. (2011).

In some cases, more sophisticated time-series models have been applied to capture linear interdependencies of several factors including online content. For example, the vector autoregressive model (VAR) can simultaneously represent relations between different variables in a contemporaneous and a time-lagged manner. Therefore, more variables of interest like message volume and trading volume have been tested simultaneously by some authors. Most of these approaches emphasize the importance of certain model coefficients to be statistically different from zero. For the prediction of market activity, one would expect the coefficients expressing dependencies between measures of online content and variables like stock returns, abnormal stock returns or volatility to be significant. Examples of the application of VAR models are given by TUMARKIN AND WHITELAW (2001), TETLOCK (2007) or BOLLEN ET AL. (2011).

Besides the approaches to capture a linear relationship between online content and market activity, there have been some efforts to investigate on a possible non-linear relationship using, for example, Neural Networks or Support Vector Machines. Often, the authors restrain their analyses to the prediction of stock movement directions rather than to the prediction of the exact value of subsequent returns. Successful examples for improvements of the classification accuracy by the inclusion of online content parameters can be found in BOLLEN ET AL. (2011) and CHOUDHURY ET AL. (2008).

Some authors interpret different kinds of fluctuations in online content as either positive or negative events and conduct an event study. In most cases deviations in message volume or changes in the value of a sentiment index form the triggering events for this type of analysis. The usual procedure is to scan a time frame before and after the event for abnormal returns. This work will go into detail on this methodology in a later section. In some studies, the online content's fluctuations are not named as being "events", but still the common event study methodology is applied. One can refer to BAGNOLI, BENEISH, WATTS (1999), TUMARKIN AND WHITELAW (2001), DEWALLY (2003), TETLOCK ET AL. (2008), SPRENGER AND WELPE (2010) and CHEN ET AL. (2011) for examples of event studies.

The body of literature in this research area is still growing and the number of studies is steadily increasing. However, statements about the predictive power of online content are contradictory. This paper analyzes data from the online social network Seekingalpha.com. It extends the work of CHEN ET AL. (2011) by analyzing a larger sample of texts and by differentiating between texts that are written by users with different engagements in the web community. The bag-of-words approach and

the event study methodology will be applied to analyze whether texts with certain sentiment expressions can predict positive or negative abnormal returns.

3 Data and Methodology

Similar to CHEN ET AL. (2011), user-generated content has been downloaded from the online social network SeekingAlpha.com (SA). In total, texts that were written by 47072 registered users have been collected. This section will first describe the characteristics of SA network's texts and then go into detail on the market data sample for the analysis. Also, it will describe the motivation for the selection of analyzing S&P 500 stocks and it will explain the event study methodology that is applied in this study.

3.1 Seekingalpha.com

SA is a social network with a heterogeneous structure of content. The texts published on SA vary in purpose, length and quality. The online community claims that it consists of opinions rather than news¹. This section gives an overview of the text categories.

3.1.1 Articles

User-written articles form a major part of the network's content. They can be contributed to SA exclusively by registered authors but can be read by every visitor of the website. Articles represent the authors' opinion on manifold topics which range from investment ideas to macroeconomic estimates. To become an author, a registered user must submit an article which then has to be reviewed by a panel. Authors have several incentives to write quality articles. Firstly, a bad article would not pass the review and would not be published. Secondly, a certain amount of page views for so-called premium articles is being rewarded with a monetary compensation². Lastly, the network provides structures of opinion leadership measures and authors would try to achieve a good status on SA. For example, every registered user can follow or be followed by others. In this context, the best-known registered users are listed in rankings. Another difference between authors and registered users is that authors must disclose more personal information like, for example, their positions in the mentioned stocks. For all these reasons articles should contain quality information and represent the views of experts. All articles that were published between January 2004 and August 2011 have been downloaded. The sample consists of 272048 articles with an average length of 1423 words. Articles are tagged with information about the stocks mentioned in them and they link to their author's profile. The articles in the sample were written by 3295 distinct authors. This shows that articles are provided by a rather small group of the SA members. It is

¹ See http://seekingalpha.com/page/about_us (checked on 10/24/2011).

² See seekingalpha.com/premium-program/intro (checked on 10/24/2011).

important to note that the review process and the extent of content in an article suggest a delay between the completion of the final draft and the eventual publication of the article. This could be a problem for the predictive power of SA articles because the market could have already adjusted the prices during the review process.

3.1.2 Stocktalks

This type of text is very closely related to the microblogging posts on Twitter.com. Every registered user can type a text of not more than 140 characters and submit it to a message-stream. Since stocktalks do not underlie a revision process, the messages are likely to contain more noise than articles. Every stocktalk from between April 2009, the start of the stocktalk service, and August 2011 has been downloaded. This resulted in a sample of 189318 texts with an average length of 15 words. Even if this type of content does not contain as much quality information as articles might do, stocktalks are published timelier and could represent the crowds' opinions faster. Unfortunately, stocktalks are not tagged with information about the stocks discussed, so this information has to be inferred from the texts. This was possible because almost every stock mentioned was referred to in form of its ticker symbol. Since ticker symbols in SA are mostly also hyperlinks to a company profile site, the search for certain structures in the HTML form of the text provided the needed information.

3.1.3 Market currents

Whereas all other text categories represent opinions, market currents are short news messages that are published in a news stream on SA. These are composed by the SA staff and always consist of a short text that sometimes contains a link to a more detailed news article. Every market current is tagged with information about the stocks mentioned in the text. All market currents from between May 2008, the start of the market current service, and August 2011 have been downloaded. The sample size consists of 86286 messages with an average length of 32 words. The selection procedure of news on SA is not documented. News might incorporate a certain selection bias. It is not inconceivable that market currents are also affected by opinions to some degree, because their selection seems to be based upon the SA staffs' estimates of the relevance of the respective news.

3.1.4 Comments

Seekingalpha.com provides the possibility to comment on almost every text object submitted by the users. According to CHEN ET AL. (2011), most comments concerning articles are made on its day of publication, so comments should also provide timely information about users' opinions on the respective topic. Since they are not tagged with information about which stocks are mentioned, it is assumed that the comments always concern the same stocks as the underlying texts they were made for. 750448 comments concerning articles and 130704 comments that were made on market currents have been downloaded. The former type of comments consists of 95 words and the latter of 54 words on average.

3.1.5 Descriptive statistics of SA texts

Because the downloaded sample of texts comprises almost the full range of content that exists on SA, it is worth examining which stocks, funds, etc. lay in the focus of attention. According to the web traffic tracking site Quantcast.com, SA is very popular in the USA³. The following statistics will show that due to the regional location of the majority of SA users, the stocks mostly followed by them are US stocks. More precisely, companies that are listed in the S&P 500 index got a disproportionately high fraction of the users' attention. When drawing conclusions from the analysis, this bias should be considered.

For texts in which several stocks are mentioned, it is hard to determine the exact allocation of meaning to stocks. In order to avoid allocation errors, only texts that mention one single stock are considered in this study. Whenever a stock is referred to somewhere in the text we call this a stock mention. Table 1 shows the numbers of texts for each category, the number of stock mentions, as well as the number of distinct ticker symbols per category.

	Total #	# of texts (all mentions)	# of texts (single mentions)	# of texts (single mentions of S&P 500 stocks)	# of distinct stock symbols mentioned
Articles	272048	230357	105323	36107	11992
Comments on articles	750339	458227	166654	52591	9918
Market currents	86286	47521	29046	14855	3941
Comments on market currents	130704	41335	22520	14429	1858
Stocktalks	189318	100593	68371	22159	6432

Table 1. Descriptive statistics of text counts per category and number of stock mentions⁴

First of all, it can be seen that roughly one third of all texts mention single stocks, and roughly one third of texts with single mentions refer to S&P 500 stocks. If one compares the number of S&P 500 stocks included in this analysis (roughly 500) with the total number of distinct stocks mentioned in the respective texts, one can assume an overweight in attention for S&P 500 stocks. More generally, the attention for stocks among SA users seems to be unequally distributed. The next section will go more into detail on this.

³ See quantcast.com/seekingalpha.com (checked on 10/24/2011).

⁴ A list of S&P 500 constituents used in this work is provided in the appendix 7.1.

3.1.6 Distribution of stock mentions in SA texts

All of the aforementioned texts refer to a total of 14497 distinct ticker symbols of stocks, ETFs, funds, etc. Several tickers with special symbols, meaning that the stock was delisted or is traded “over-the-counter” are included in this number. Omitting these special symbols, the total number of distinct ticker symbols still is 11062. The distributions of stock mentions show that public attention is mainly focused on only a small fraction of stocks. Fig. 1 shows the distribution of stock mentions in articles on a log-log scale. Additionally, it shows the cumulative distributions of S&P 500 stock mentions (red) and all other stock mentions (green).

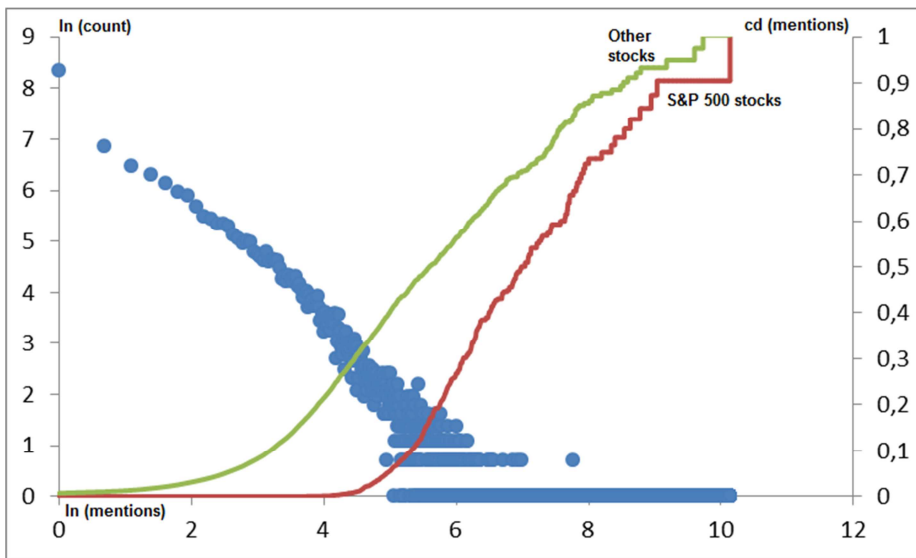


Fig. 1. Log-log diagram of the quantity distribution of stock mentions (blue) and cumulative quantity distribution of S&P 500 stocks (red) and all other stocks (green) for articles

The distribution of mentions is highly skewed. A relatively small fraction of stocks gets a disproportionately high fraction of attention. Fig. 1 shows approximately a straight line that gets noisy at the end. The distribution dwindles in this region, meaning that each bin only has a few samples in it, if any. This straight line indicates a power-law distribution of stock mentions. The respective distributions of stocktalks, market currents and comments show even clearer lines. For those, the assumption of power-law behavior is stronger. The plots of their distributions are reported in appendix 7.2.

Many characteristics of big online social networks seem to be distributed by a power-law. For example, KUMAR ET AL. (2010) report that the degree distributions, meaning the number of friends per user, of the networks Flickr and Yahoo! 360 follow a power-law. Furthermore, CHA ET AL. (2010) find that the numbers of retweets and mentions per user on Twitter.com also follow such a law. Hence, it is not surprising to see power-law behavior for SA users’ attention to stocks. For articles, the ratio of S&P stock mentions (not reported in Table 1) to all mentions is 31.4%.

For stocktalks it is 34.5%, for comments on articles it is 36.5%, for market currents it is 46.6% and for comments on market currents it is even 50.2%.

To illustrate the overweight in attention to S&P 500 stocks, the red line in Fig. 1 shows the cumulative distribution of their mentions. The green line illustrates the same for all other stock mentions. One can see that S&P 500 stocks have on average more mentions than other stocks. This is noteworthy because these stocks comprise roughly 500 of the 11992 different ticker symbols in articles. Despite that, even the less discussed S&P 500 stocks have more mentions than a big fraction of all other stocks. This is shown by the late incline of the red line compared to the green one. For instance, less than only 10% of the S&P stocks have fewer than 200 mentions. For all other stocks, it is less than 45%. As a result, one can say that the attention of the SA users has a strong bias towards members of the S&P 500 index. This statement is also valid for the other text categories.

In line with these observations, the symbol that gets the most attention in articles is also S&P 500 related. It is the SPDR S&P 500 Trust ETF (SPY) which is an exchange traded fund that replicates the movement of the S&P 500 index. It gets a total number of 25189 mentions, including those from texts with multiple mentions. Next to the SPY, the overall top ranked symbols are ETFs that replicate the movements of the Nasdaq-100 index, the Dow Jones Industrial Average and some symbols of funds that are active in commodities and currencies. The prevalence of attention for a small group of stocks may be due to regional aspects. Since most of the users on SA seem to come from the USA, their attention might not be equally spread on international markets. According to these findings, we will analyze only the impact of S&P 500 related SA content on the market activity of the underlying stocks. Even if it is overconfident to say that the stocks which are discussed more frequently are also discussed with more quality, one can suppose that SA users follow the stock price and company development of those popular stocks more closely. The next section will describe the sample of stock market data for the selected set of stocks.

3.2 Stock price data

Stock prices have been downloaded from the Yahoo! Finance service for all actual S&P 500 constituents and some stocks that were members recently (see appendix 7.1). The daily adjusted closing prices have been used to calculate returns. These closing prices are adjusted for stock splits and dividend multipliers, adhering to Center of Research in Security Prices (CRSP) standards⁵. The sample period spans from January 2006 to August 2011. This is due to the content volume from SA. Even if the first articles go back to the year 2004 a volume increase was observable after 2006. This is also reported by CHEN ET AL. (2011). The market current service and the stocktalk go back to the years 2008 and 2009 respectively. Only trading days are regarded. So, for example, the return from Friday to Monday is treated as a one day return.

⁵ See help.yahoo.com/l/us/yahoo/finance/investing/fitahistquo.html (checked on 10/24/2011).

3.3 Sentiment measure

In line with studies that utilize dictionary-based approaches for their sentiment analysis, the negative words compiled by LOUGHRAN AND McDONALD (2011) are used to determine the negativity of a text. The list of positive words is not considered because a positive meaning might be more frequently negated than a negative one (CHEN ET AL. (2011)). A positive tone could be best captured by including methods of contextual polarity detection. Capturing the negativity of a text by using the counts of negative words has already been successfully implemented by TETLOCK (2007), TETLOCK ET AL. (2008) and CHEN (2011). The negativity measure is defined as the ratio of negative words to all words written about a stock on a certain day:

$$neg_{i,t,m} = \frac{neg_w_{i,t,m}}{total_w_{i,t,m}}$$

Here, $neg_{i,t,m}$ stands for the negativity for company i on day t in text category m , $neg_w_{i,t,m}$ denotes the number of negative words that were found in all texts concerning the respective company on that day and, $total_w_{i,t,m}$ denotes the total word count of all these respective texts. This procedure is similar to the one introduced by CHEN ET AL. (2011) and can be assigned to the bag-of-words approaches. One major difference to their work is that almost all text categories that exist on SA are analyzed. Admittedly, the dictionary-based procedure might not be equally well suited for all types of texts. Short texts might cause problems. As already mentioned, the average text length across categories varies from 15 to 1423 words. Another problem might be the vocabulary. Articles and market currents use finance specific terminology, whereas ordinary language predominates in comments and stocktalks.

In the following it will be explained how the scores for each stock are classified into two opposite sentiment groups. The negativity scores for each day are distinguished into a positive (bullish) and a negative (bearish) class by splitting the set of scores for each day into an upper and a lower tercile. Texts which belong to the lower third of the negativity score are labeled as bullish and those in the upper third are labeled as bearish. Texts belonging to negativity scores of zero are always labeled as bullish, because text categories with only a few words on average quite often do not contain negative words at all. In the end, this approach results in sets of days for each stock on which the negativity scores indicate a bullish or a bearish sentiment for each text category.

Table 2 shows some descriptive statistics for the negativity measure among the bullish and the bearish class. The analysis only considers texts that refer to S&P 500 stocks. Hence, the average number of words per text is shown again for this smaller sample.

	avg. words per text (stdev.)	avg. negative words per text (stdev.)	avg. negativity score (stdev.)	avg. negativity score for bullish texts (stdev.)	avg. negativity score for bearish texts (stdev.)
Articles	2953.07 (4154.09)	24.39 (36.05)	0.0088 (0.0071)	0.0031 (0.0025)	0.0153 (0.0076)
Comments on articles	315.27 (650.15)	3.52 (8.96)	0.0100 (0.0150)	0.0001 (0.0007)	0.0168 (0.0163)
Market currents	35.67 (28.39)	0.58 (1.05)	0.0157 (0.0266)	0.0001 (0.0016)	0.0441 (0.0272)
Comments on Market currents	153.55 (301.67)	2.07 (4.81)	0.0133 (0.0283)	0.0000 (0.0006)	0.0260 (0.0351)
Stocktalks	27.45 (39.35)	0.31 (0.81)	0.0112 (0.0280)	0.0002 (0.0029)	0.0523 (0.0394)

Table 2. Descriptive statistics of the negativity measure

On average comments on articles are longer than comments on market currents. This might indicate that articles are discussed in more detail. Market currents and stocktalks are both relatively short. This might intimidate sentiment extraction. Nevertheless, their average negativity scores are similar to those of the other classes. The very low average negativity scores for bullish texts across all text categories except for articles indicate that many texts contain no negative words at all. At a first glance, the similarity between the negativity scores across text categories motivates the application of the dictionary-based approach even for short texts. In the following, the methodology to analyze the influence of SA sentiment on the stock market will be described.

3.4 Event studies

The EMH has sustained many empirical tests. Several methods have been tried to prove that stock prices don't react to new information as quickly as stated by the EMH. In this context, event studies have become a popular method for testing the speed of stock price adjustments to new information. Common events of such studies are investment decisions, financing decisions and changes in corporate control (FAMA (1991)). Event studies examine the assumption that certain events are associated with abnormal behavior in stock returns during the time around the event. From the viewpoint of EMH supporters, new price-relevant information should be incorporated in prices quickly after the event. Though many authors claim to have found anomalies which contradict the EMH, FAMA (1991) summarizes that most literature provides evidence for market efficiency. The first event study was performed by FAMA ET AL. (1969) using stock splits as events. Recently several authors have used this methodology to capture the effect of sentiment in online content on stock market activity. There is not a clear definition of what to define as events in this context. MACKINLAY (1997) states that an event study usually consists of three steps: determination of the selection criteria for the inclusion of a given firm, the definition of the events of interest, and the definition of abnormal returns.

The first task has already been described in subsection 3.1.6. Only S&P 500 stocks are included in our study, because these stocks are frequently discussed on SA. A long time span and a high amount of messages are important to avoid spurious results which may arise due to data-dredging and chance sample-specific conditions as was pointed out by FAMA (1991).

An event for a stock is defined as the day at which texts concerning this stock are published and the negativity score for these texts classifies them as either bullish or bearish. In case that a stock is discussed over several consecutive days with the same sentiment, only the starting point of the stream of texts is selected. Moreover, events occurring in-between 10 trading days after another event of the same sentiment class are omitted. This definition is applied to each text category separately. Distinguishing between bullish and bearish events does not only allow interpreting abnormal returns in any direction but also testing whether bullish events are associated with positive abnormal returns and vice versa. Admittedly, this is an interpretation of events in a very broad sense of the event study's original meaning but the methodology seems to be very suitable for this kind of analysis. Assuming one would always find abnormal returns after the publication of certain texts, this would allow conclusions on the predictive power of online sentiment. Similarly, abnormal returns before the publication of certain texts could give information on the reasons for the authors to write certain kinds of content.

The third task of defining abnormal returns can be done in several ways (MACKINLAY (1997)). All of them have in common that the abnormal return is defined as the actual return of the stock minus its expected return:

$$AR_{it} = R_{it} - E(R_{it})$$

R_{it} denotes the actual return of stock i at day t , $E(R_{it})$ denotes the expected return of the stock at day t . According to MACKINLAY (1997), there are several ways to estimate $E(R_{it})$. He mentions at least three different alternatives – the constant mean return model, the market model and the market-adjusted return model. BROWN AND WARNER (1985) point out that the market model and the market-adjusted return model have similar power regarding daily data. We will use market-adjusted returns:

$$AR_{it} = R_{it} - (R_{mt})$$

R_{mt} denotes the return of the market portfolio at day t . Broad based stock indices like the S&P 500 index, the CRSP value weighted index or the CRSP equal weighted index are often used as proxies for the market portfolio. It is also common to use logarithmic returns for R_{it} and R_{mt} because it is assumed that continuously compounded returns better conform to the assumptions of the model (FAMA ET AL. (1969)).

Aggregation can be performed over time and across companies. Aggregation across companies is achieved by building the average:

$$AAR_t = \frac{1}{N} \sum_{i=1}^N AR_{it}$$

N denotes the number of companies affected by events at time t . Aggregation through time is obtained by summing up the ARs for several days around the event resulting in cumulative abnormal returns (CAR):

$$CAR_{i\tau_1\tau_2} = \sum_{\tau=\tau_1}^{\tau_2} AR_{i\tau}$$

Here, τ_1 is the start of the event window and τ_2 denotes the end. If there is more than one event for the respective stock, one also has to aggregate over these events. Finally, the CARs are aggregated across companies:

$$ACAR_{\tau_1\tau_2} = \frac{1}{N} \sum_{i=1}^N CAR_{i\tau_1\tau_2}$$

It is common to test the null hypothesis that AAR and ACAR are normally distributed with zero mean and known variance. Since market-adjusted returns are assumed to have a zero mean, significant deviations from zero will indicate that the corresponding returns are abnormal. The calculated variances have to be interpreted with caution. Firstly, the true variance of the AARs and ACARs is unknown so it is necessary to find a good estimate. MACKINLAY (1997) gives some examples for that. Secondly, these estimators often imply distributional assumptions which have to be considered. One of them is that ARs and CARs need to be independent across securities. More precisely, their covariances across stocks need to be zero. This becomes an issue when event windows for different stocks overlap. According to BROWN AND WARNER (1985) this leads to underestimations of the variance and hence to more rejections of the null hypothesis. Knowing that clustering is problematic when all events are at the same day, the restriction that two events from the same sentiment class must not occur within 10 trading days was made to slightly accommodate for this problem. Even with this restriction, the event windows occasionally overlap. This means that variances might be underestimated in some cases. For this reason, only strongly significant results will be used to draw conclusions from.

4 Results

Our results are separately described for each of the five text categories, first for bullish thereafter for bearish contents. Conclusions are mainly drawn from the patterns of average ARs (ARRs). Therefore results for a two-tailed test with the null hypothesis that the ARs are normally distributed with zero mean and given variance will be listed. To extend the insights into the results, also average CARs (ACARs) and the ratio of positive to negative ARs across companies and events will be shown. Along with each result set, the average CARs will be displayed in diagrams. Ideally, CARs should increase after the event of positive mentions and decrease after the event of negative mentions. Finally, results regarding the S&P 500 index returns will be presented.

4.1 Results for articles




Fig. 2. Average CAR around positive and negative events in articles

	number of events			7408	avg. events per stock			15.12			
day	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	
AAR	0.04%	0.03%	0.05%**	0.03%	0.01%	0.01%	0.04%	0.04%	0.04%	0.16%**	
t-stat.	1.553	1.035	1.999	1.098	0.444	0.224	1.431	1.446	1.494	5.057	
pos./neg.	49.39%	50.58%	50.23%	49.47%	49.42%	49.62%	50.80%	50.24%	51.19%	52.36%	
ACAR	0.04%	0.07%	0.12%	0.15%	0.16%	0.16%	0.20%	0.24%	0.28%	0.45%	
0	1	2	3	4	5	6	7	8	9	10	
0.18%**	0.10%**	-0.02%	0.04%*	0.04%*	-0.02%	0.01%	0.05%**	-0.01%	0.00%	-0.02%	
4.608	2.775	-0.723	1.718	1.661	-0.630	0.243	1.968	-0.232	-0.197	-0.943	
51.46%	50.99%	48.61%	49.85%	49.39%	49.23%	49.65%	50.96%	49.27%	49.14%	49.49%	
0.63%	0.73%	0.71%	0.75%	0.79%	0.78%	0.78%	0.83%	0.82%	0.82%	0.80%	

** significance at the 5% level

* significance at the 10% level

Table 3. Average ARs and average CARs for positive events (articles)

It is first to note that the average ARs for bullish articles (left-hand side of Fig. 2; Table 3) are positive from the beginning of the event window until the day when the event occurs. There are significant values on days -8, -1 and 0. The latter two days are most significant and the positive-negative ratios of avg. ARs take their highest values. These observations indicate that articles with a positive tone are preceded by positive avg. ARs of the respective stocks. Normally one would infer that a positive sentiment in articles is due to previous positive performance. However, one has to keep in mind the aforementioned review process of articles. It is not clear how much time has passed between the creation of the text and its publication. The avg. ARs after the event show significant values on days 1 and 7 and slightly significant values on days 3 and 4. The most significant value is observed right after the event. It is considered strong enough to be interpreted as a sign for the predictive power of bullish articles for a one day horizon.

Going long on a portfolio of stocks with a positive tone in SA articles might have yielded a gross performance of 20 basis points above the index considering a 7 days holding period. Accordingly, one might have got the chance to beat the market in the years 2006 until 2011. But one still has to remember that trading costs and other mitigating factors could completely absorb this outperformance. For defenders of the buy-and-hold strategy and followers of the EMH, these findings might be reason to investigate further on this topic.

	number of events			7487	avg. events per stock						15.13
day	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	
AAR	-0.02%	-0.02%	0.01%	0.00%	0.05%*	0.00%	-0.08%**	-0.06%*	-0.11%**	-0.19%**	
t-stat.	-0.627	-0.793	0.489	0.109	1.678	-0.029	-2.583	-1.835	-3.225	-4.322	
pos./neg.	48.64%	48.38%	48.35%	48.62%	49.19%	48.27%	48.14%	49.00%	48.76%	48.78%	
ACAR	-0.02%	-0.04%	-0.03%	-0.02%	0.03%	0.03%	-0.05%	-0.11%	-0.22%	-0.41%	
0	1	2	3	4	5	6	7	8	9	10	
-0.13%**	-0.06%	-0.03%	0.02%	-0.01%	0.00%	0.00%	0.02%	-0.04%	-0.03%	0.01%	
-2.617	-1.405	-0.821	0.638	-0.329	0.054	0.006	0.599	-1.350	-0.895	0.496	
47.87%	48.51%	47.59%	48.26%	48.56%	49.02%	49.27%	49.71%	48.63%	48.98%	49.57%	
-0.54%	-0.60%	-0.62%	-0.60%	-0.61%	-0.61%	-0.61%	-0.59%	-0.63%	-0.66%	-0.64%	

** significance at the 5% level

* significance at the 10% level

Table 4. Avg. ARs (AAR) and avg. CARs (CAR) for negative events (articles)

Bearish articles (right-hand side of Fig. 2; Table 4) are generally showing the opposite effect of bullish articles. Except for one slightly significant positive value on day -6, all days from day -4 until day 0 are negative and either slightly or strongly significant. Also the positive-negative ratio is smaller than for positive events in most cases. For the post-event window, the biggest negative avg. AR can be seen again one day after the event but this time it is not significant. Even though the prevalent direction of the post-event CARs is negative after all, no significance can be spotted. Again, it is to question whether the predictive power would have been greater without the delay from the review process.

The observed positive avg. ARs before bullish and negative avg. ARs before bearish articles suggest a naïve momentum mentality of SA authors. Their sentiment on stocks is based on past price trends.

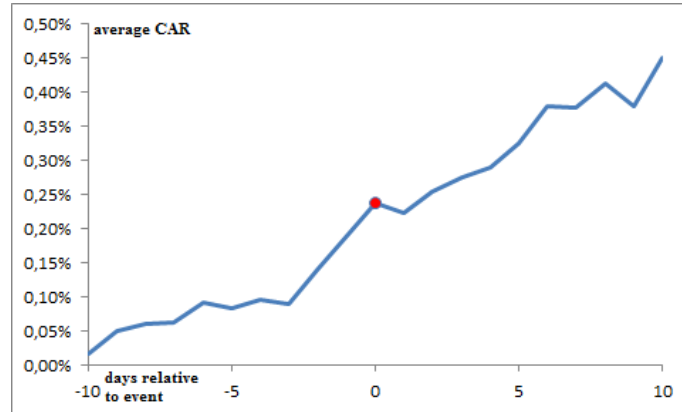


Fig. 3. Average abnormal return and average cumulative abnormal return around the publication of neutral articles

	number of events					avg. events per stock				
	7685					15.34				
day	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1
AAR	0.02%	0.03%	0.01%	0.00%	0.03%	-0.01%	0.01%	-0.01%	0.05%*	0.05%
t-stat.	0.663	1.262	0.415	0.087	1.177	-0.368	0.486	-0.274	1.746	1.448
pos./neg.	49.80%	50.25%	49.67%	49.66%	50.10%	49.85%	50.54%	49.81%	51.33%	51.48%
ACAR	0.02%	0.05%	0.06%	0.06%	0.09%	0.08%	0.10%	0.09%	0.14%	0.19%
0	1	2	3	4	5	6	7	8	9	10
0.05%	-0.01%	0.03%	0.02%	0.02%	0.04%	0.05%**	0.00%	0.04%	-0.03%	0.07%**
0.943	-0.424	1.129	0.843	0.642	1.413	2.223	-0.118	1.378	-1.378	2.718
50.51%	48.34%	48.64%	49.99%	50.08%	49.89%	50.02%	48.95%	49.76%	49.72%	50.24%
0.24%	0.22%	0.25%	0.27%	0.29%	0.33%	0.38%	0.38%	0.41%	0.38%	0.45%

** significance at the 5% level

* significance at the 10% level

Table 5. Average ARs and average CARs around the publication of neutral articles

Articles that are neither labeled as bearish nor as bullish also reveal interesting insights. The performance around the publication of these articles can be seen in Fig. 3 and Table 5. According to the definition of bullish and bearish texts in chapter 3.3, these neutral articles contain a certain amount of negative words but don't belong to either the upper or lower tercile of the negativity measure on their publication day. The mean of the average abnormal return is positive from day 2 until day 8, and significantly deviating from zero on days 6 and 10. Even though the post-event cumulative abnormal returns indicate an outperformance of 21 basis points for a ten-days holding period, the pre-event performance is not as high as for bullish events. Here, the most significant positive avg. ARs are located at the end of the event window. This positive outperformance is not really surprising when looking at the average negativity scores of the underlying articles. The direction seems to be consistent with expectations. Although "neutral" articles were not assigned to the bullish or the bearish class, the average negativity score for them is 0.0074 with a

standard deviation of 0.0024. According to Table 2, this is not only below the average of the overall sample; it is also much closer to the average of the bullish group than to the average of the bearish one. This encourages a revision of the class boundaries of the negativity score. Positive sentiment thereby would not only be revealed in the lower third of the negativity score. Further investigations on this effect have not been conducted in this study.

In summary, there is only little evidence that articles may be used for stock price predictions. Although one can observe that the attention of authors with non-negative sentiments is mainly focused on stocks which perform better than the market in the 10 days after publication, the significance of this observation might fall under acceptable confidence levels if the underestimation of the variance due to clustering will be taken into account. The opposite conclusion that attention to stocks is driven by extraordinary prior performance will resist a much stronger mitigation of significance.

4.2 Results for comments on articles




Fig. 4. Average CAR around positive and negative events for comments on articles

Before we go into detail on the results for comments on articles, it is to mention that the average negativity score of articles with bullish comments is 0.0089 with a standard deviation of 0.0078, whereas for articles with bearish comments, the average score is 0.0101 with a standard deviation of 0.0085. The first average is higher and the second one is lower than the respective average for the whole sample of bullish and bearish articles (see Table 2). In other words, the articles that were commented on with a positive tone were not the ones with the most bullish content and vice versa.

		number of events			4928	avg. events per stock			11.51		
day	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	
AAR	0,01%	0,03%	0,00%	-0,01%	0,01%	-0,02%	0,07%*	-0,07%	-0,01%	-0,02%	
t-stat.	0,153	0,769	-0,065	-0,228	0,215	-0,519	1,694	-1,574	-0,234	-0,353	
pos./neg.	48,74%	48,90%	48,54%	48,23%	48,54%	48,76%	50,67%	49,59%	50,39%	49,70%	
ACAR	0,01%	0,03%	0,03%	0,02%	0,03%	0,01%	0,07%	0,01%	0,00%	-0,02%	
	0	1	2	3	4	5	6	7	8	9	10
	0,05%	0,10%**	0,06%*	0,05%	-0,02%	0,02%	0,08%**	0,07%*	0,05%	-0,02%	0,09%**
	1,033	2,495	1,650	1,550	-0,448	0,482	2,287	1,910	1,427	-0,471	2,596
	48,72%	50,53%	48,80%	49,86%	49,23%	49,84%	50,32%	49,94%	50,85%	49,51%	51,24%
	0,02%	0,12%	0,18%	0,24%	0,22%	0,24%	0,32%	0,38%	0,43%	0,42%	0,50%

** significance at the 5% level

* significance at the 10% level

Table 6. Average ARs and average CARs for positive events (comments on articles)

Table 6 shows that most AARs for bullish comments on articles are positive and significant after the event has occurred. Positive texts of this text category conform better to the expectation of having predictive power than any other category. The ACARs for the post-event phase are higher than those of bullish and neutral articles. Most interestingly, the strong positive performance of 48 basis points in 10 days is not preceded by high avg. ARs before the event.

		number of events			4485	avg. events per stock			11.24		
day	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	
AAR	-0.01%	0.01%	-0.02%	0.03%	0.04%	-0.01%	-0.04%	0.06%	-0.17%**	-0.01%	
t-stat.	-0.306	0.157	-0.519	0.751	0.855	-0.243	-0.791	1.364	-3.404	-0.148	
pos./neg.	47.91%	48.96%	49.08%	49.35%	47.84%	48.57%	48.55%	51.01%	47.91%	48.73%	
ACAR	-0.01%	-0.01%	-0.03%	0.00%	0.04%	0.03%	-0.01%	0.05%	-0.12%	-0.13%	
	0	1	2	3	4	5	6	7	8	9	10
	-0.03%	0.08%*	0.04%	0.00%	0.01%	0.00%	0.06%	0.08%**	0.04%	0.03%	-0.03%
	-0.599	1.772	0.862	0.031	0.252	0.100	1.527	2.134	0.972	0.787	-0.866
	48.85%	49.65%	49.67%	49.42%	49.05%	49.83%	51.49%	50.86%	49.24%	48.57%	49.10%
	-0.16%	-0.08%	-0.05%	-0.04%	-0.03%	-0.03%	0.03%	0.11%	0.15%	0.18%	0.15%

** significance at the 5% level

* significance at the 10% level

Table 7. Average ARs and average CARs for negative events (comments on articles)

The results for negative comments (Table 7) are disappointing as the post-event phase until day 9 is also characterized by positive AARs – two of them being significant. It seems that many negative comments refer to articles discussing stocks with positive abnormal returns after the day of the comment's creation. This is all the more surprising as negative comments refer to articles with a lower negativity score average than those in the class of all negative articles. Additionally, bearish comments on articles have the second lowest average negativity score compared to the classes of

all other bearish texts (Table 2). This can be interpreted either as a problem with the sentiment measure or the commenters are simply too pessimistic and wrong in their opinion.

4.3 Results for market currents




Fig. 5. Average CAR around positive and negative events for market currents

	number of events			6485	avg. events per stock				12.92		
day	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	
AAR	0.03%	0.04%	0.05%	0.06%**	-0.01%	0.04%	0.01%	-0.02%	0.05%	0.46%**	
t-stat.	1.031	1.124	1.586	2.078	-0.169	1.026	0.376	-0.571	1.401	8.219	
pos./neg.	49.82%	50.62%	50.13%	50.50%	49.82%	50.13%	48.93%	50.84%	52.03%	54.48%	
ACAR	0.03%	0.07%	0.12%	0.18%	0.17%	0.21%	0.22%	0.20%	0.25%	0.71%	
0	1	2	3	4	5	6	7	8	9	10	
0.24%**	-0.01%	-0.06%*	0.02%	0.02%	-0.03%	0.02%	0.02%	-0.04%	-0.02%	0.00%	
4.515	-0.420	-1.731	0.785	0.747	-0.989	0.591	0.535	-1.349	-0.834	0.111	
50.13%	48.56%	48.51%	49.13%	49.99%	49.08%	50.49%	49.58%	48.79%	48.64%	48.91%	
0.95%	0.94%	0.88%	0.90%	0.92%	0.90%	0.91%	0.93%	0.89%	0.86%	0.87%	

** significance at the 5% level

* significance at the 10% level

Table 8. Average ARs and average CARs for positive events (market currents)

Positive news in market currents are preceded by strongly significant positive avg. ARs one day before the event and at the event day (see Table 9). The post-event performance is neither positive nor negative. Only one slightly significant value on day 2 can be observed. These findings imply that market currents are mainly driven by past performance and don't contain predictive information. Since the avg. ARs on the event day are highly significant, it would be interesting to examine whether stock prices adjust to the new information in market currents on an intra-day basis.

	number of events			3313	avg. events per stock			7.02			
day	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	
AAR	-0.04%	-0.04%	-0.02%	0.03%	-0.01%	0.02%	-0.06%	0.00%	-0.15%**	-0.77%**	
t-stat.	-0.724	-0.748	-0.355	0.500	-0.256	0.393	-1.126	0.020	-2.613	-7.850	
pos./neg.	48.42%	48.57%	49.11%	49.98%	48.17%	49.05%	50.14%	48.29%	48.72%	42.86%	
ACAR	-0.04%	-0.08%	-0.10%	-0.07%	-0.08%	-0.06%	-0.12%	-0.12%	-0.27%	-1.04%	
0	1	2	3	4	5	6	7	8	9	10	
-0.57%**	-0.06%	-0.04%	-0.01%	0.00%	-0.05%	0.03%	-0.09%*	-0.04%	0.02%	0.08%	
-6.687	-1.123	-0.665	-0.219	-0.060	-0.942	0.617	-1.828	-0.781	0.468	1.426	
45.00%	47.90%	48.81%	48.78%	47.33%	47.93%	49.11%	47.66%	49.32%	49.41%	49.98%	
-1.61%	-1.67%	-1.71%	-1.73%	-1.73%	-1.77%	-1.74%	-1.83%	-1.87%	-1.85%	-1.77%	

** significance at the 5% level

* significance at the 10% level

Table 9. Average ARs and average CARs for negative events (market currents)

Negative market currents are preceded by extremely significant negative avg. ARs (see Table 9). The magnitude and significance of the avg. ARs on day -1 is the highest in the whole study. Also the positive-negative ratio has its lowest value on this day. The post-event performance again does not show any mentionable significance. Only the value on day 7 is slightly significant and negative.

In summary, market currents seem to react strongly to past performance rather than to contain valuable predictive information. The assumption that this might be due to a possible delay of news on SA compared to other news streams cannot be confirmed. A small amount of market currents has been compared to other news wires on the internet and only short time lags of not more than a couple minutes have been observed.

4.4 Results for comments on market currents




Fig. 6. Average CAR around positive and negative events for comments on market currents

		number of events			1551	avg. events per stock			5.26		
day	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	
AAR	-0.01%	0.14%**	0.01%	0.02%	0.09%	0.07%	-0.04%	-0.07%	0.17%**	-0.19%**	
t-stat.	-0.095	1.989	0.236	0.325	1.341	0.954	-0.639	-0.921	2.434	-2.347	
pos./neg.	49.07%	49.71%	48.55%	49.39%	49.45%	49.00%	50.87%	47.71%	50.03%	47.32%	
ACAR	-0.01%	0.13%	0.14%	0.17%	0.26%	0.33%	0.29%	0.22%	0.39%	0.20%	
0	1	2	3	4	5	6	7	8	9	10	
0.10%	0.09%	0.00%	-0.08%	0.01%	-0.04%	-0.01%	0.16%**	-0.07%	0.00%	-0.08%	
0.914	1.000	-0.033	-1.263	0.094	-0.604	-0.221	2.350	-1.172	-0.053	-1.390	
52.29%	47.20%	49.77%	47.39%	48.10%	47.52%	49.13%	49.84%	45.91%	47.97%	46.68%	
0.30%	0.39%	0.38%	0.30%	0.31%	0.26%	0.25%	0.41%	0.35%	0.34%	0.26%	

** significance at the 5% level

* significance at the 10% level

Table 10. Average ARs and average CARs for positive events (comments on market currents)

The avg. ARs around the publication of positive comments on market currents show some single significant values but overall there is no clear picture (see Table 10). Only a slight positive tendency before the event can be observed. In the post-event period there is a significant positive avg. AR on day 7, but there are a lot of negative post-event AARs. Even if they all lack significance, they sometimes have high negative values.

		number of events			1262	avg. events per stock			5.32		
day	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	
AAR	0.10%	0.00%	0.01%	-0.01%	0.02%	0.03%	-0.05%	-0.09%	-0.13%	0.00%	
t-stat.	1.288	-0.033	0.141	-0.085	0.201	0.508	-0.785	-1.146	-1.613	-0.002	
pos./neg.	50.63%	49.37%	48.81%	49.21%	49.13%	49.21%	48.34%	46.67%	48.26%	48.18%	
ACAR	0.10%	0.10%	0.11%	0.10%	0.12%	0.15%	0.10%	0.01%	-0.12%	-0.12%	
0	1	2	3	4	5	6	7	8	9	10	
0.11%	0.00%	-0.12%	0.01%	-0.05%	0.12%	-0.11%	0.08%	-0.10%	-0.06%	0.15%*	
0.814	-0.014	-1.523	0.112	-0.552	1.343	-1.466	1.096	-1.386	-0.794	1.666	
51.43%	49.29%	45.32%	48.42%	47.70%	51.51%	44.45%	48.10%	47.70%	48.81%	48.73%	
-0.01%	-0.01%	-0.13%	-0.12%	-0.17%	-0.05%	-0.16%	-0.08%	-0.18%	-0.23%	-0.09%	

** significance at the 5% level

* significance at the 10% level

Table 11. Average ARs and average CARs for negative events (comments on market currents)

The results for negative comments on market currents also show an unclear picture. There is only one slightly significant positive avg. AR on day 10 (see Table 11).

4.5 Results for stocktalks




Fig. 7. Average CAR around positive and negative events for stocktalks

	number of events				5726				avg. events per stock				11.88								
day	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6	7	8	9	10
AAR	0.04%*	0.03%	0.05%**	0.08%**	0.05%**	0.02%	0.00%	0.07%**	0.05%**	0.12%**	0.33%**	0.03%	-0.02%	0.00%	-0.03%	0.00%	-0.03%	0.02%	0.01%	0.00%	0.03%
t-stat.	1.740	1.364	2.148	3.531	2.210	0.861	-0.062	3.142	2.108	4.291	8.335	1.058	-0.980	0.086	-1.472	-0.086	-1.374	0.720	0.560	0.041	1.203
pos./neg.	49.39%	49.44%	49.39%	50.38%	50.05%	49.88%	48.60%	49.28%	50.00%	51.29%	55.26%	49.30%	47.61%	49.13%	47.52%	48.97%	47.38%	50.00%	49.60%	48.86%	48.83%
ACAR	0.04%	0.07%	0.12%	0.20%	0.24%	0.26%	0.26%	0.33%	0.38%	0.51%	0.84%	0.87%	0.85%	0.85%	0.81%	0.81%	0.78%	0.80%	0.81%	0.81%	0.84%

** significance at the 5% level

* significance at the 10% level

Table 12. Average ARs and average CARs for positive events (stocktalks)

Bullish stocktalks have strong positive avg. ARs before the event (see Table 12) and thus conform to a naïve momentum mentality. On the event day the deviation is strongly significant. This again motivates an intra-day analysis to check whether stocktalks precede strong abnormal performances or vice versa. Bullish stocktalks cannot be used for predictions as there is no significant AAR after the event.

	number of events			1653	avg. events per stock				5.07		
day	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	
AAR	0.01%	-0.02%	-0.05%	-0.02%	0.08%*	-0.02%	0.02%	0.09%*	-0.01%	-0.07%	
t-stat.	0.262	-0.329	-1.268	-0.455	1.821	-0.533	0.467	1.919	-0.150	-1.275	
pos./neg.	48.94%	48.09%	47.01%	48.28%	50.39%	47.49%	49.06%	48.52%	49.30%	49.73%	
ACAR	0.01%	0.00%	-0.06%	-0.08%	0.00%	-0.02%	0.00%	0.09%	0.08%	0.01%	
0	1	2	3	4	5	6	7	8	9	10	
-0.39%**	-0.02%	0.04%	-0.02%	0.00%	-0.05%	0.00%	0.01%	0.01%	0.09%**	0.00%	
-4.146	-0.266	0.865	-0.468	-0.014	-0.941	-0.064	0.170	0.180	2.044	-0.092	
47.19%	47.43%	49.00%	47.79%	49.00%	45.92%	48.40%	48.82%	48.34%	50.15%	49.97%	
-0.38%	-0.40%	-0.36%	-0.38%	-0.38%	-0.43%	-0.43%	-0.42%	-0.42%	-0.33%	-0.34%	

** significance at the 5% level

* significance at the 10% level

Table 13. Average ARs and average CARs for negative events (stocktalks)

Bearish stocktalks show a significant negative avg. AR only on the event day (see Table 13). This motivates an intra-day inspection once again. The values before the event day vary from positive to negative and thus are not conclusive. There is only one significant AAR in the post-event period. Therefore also bearish stocktalks are not predictive.

Stocktalks can be expected to be the most recent proxy for the sentiment of SA users because they can be immediately posted to the stocktalks stream. Unfortunately they have a limitation of 140 characters. This implies some linguistic complications for sentiment classification. Furthermore, stocktalks contain common language and ironical formulations what impedes textual analysis even more. It is not clear whether the lack of predictive power has to be attributed to the limitations of the sentiment extraction method.

4.6 S&P 500 related texts and index returns

In addition to texts dealing with individual S&P 500 stocks, there are also texts that deal with the index itself. These texts include mentions of either the symbol SPX, the ticker of the S&P 500 index, or the symbol SPY, the ticker of an ETF which tracks the performance of this index. Since abnormal returns are not applicable in this case, the index returns for several periods of time around the publication of the respective texts are pooled and compared to the negativity score of the texts. This additional analysis was motivated by DAS AND CHEN (2007) who found that the relation from sentiment to price changes was stronger for index returns than for individual stock returns.

Scatterplots of the negativity measures and the index returns for different time periods (1-day, 2-day, 3-day, 5-day, and 10-day return before and after the day of publication) are analyzed. Since market currents and comments on market currents contain less than 20 texts, these categories are not analyzed.

The results of this additional analysis can be summarized as follows. Stocktalks and comments on articles show no significant relation to any of the pre- or post-

publication returns. For articles, the negativity scores are related to pre-publication returns only. The strongest relation is observed for the 10-day returns before the event. The scatterplot containing 969 observations is shown in Fig. 8. Additionally, the equation and statistics of an OLS regression with past returns as the independent and negativity as the dependent variable is shown below.

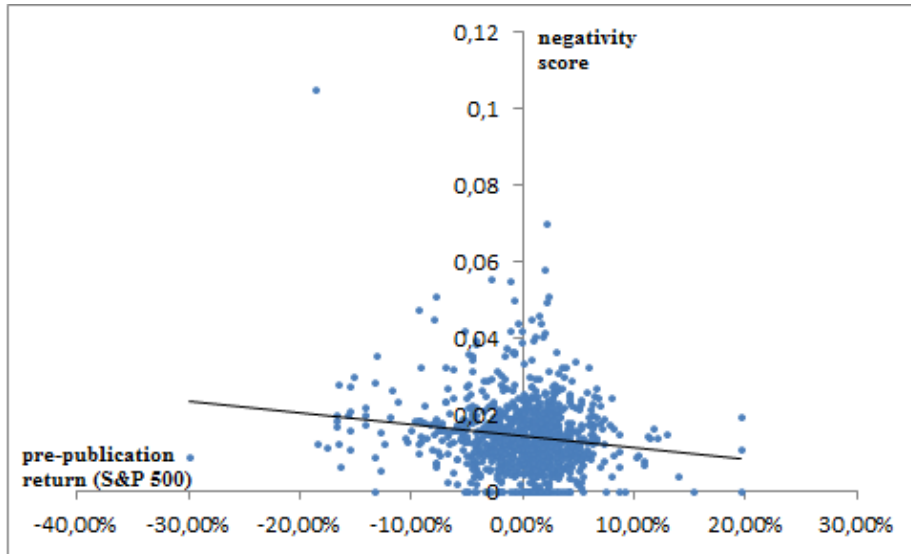


Fig. 8. Scatter plot of S&P 500 10-days return pre-publication of S&P 500 related articles and the negativity score on the respective publication days⁶

$$\begin{aligned} neg^{SP500} &= 0.0144 (0.000) - 0.0302(0.000) \cdot R_{-10,0}^{SP500} \\ R^2 &= 0.022 \\ \text{p-value of F-Test} &= 0.000 \end{aligned}$$

The values in brackets denote the p-values of the t-statistics for the coefficients of the regression. As one can see, there is a significant relation between the past index performance and the negativity in articles. The R-squared statistic is rather weak but not unusual when compared to the other studies about stock price changes. The significance of pre-event abnormal returns is consistent with our previous findings. The same is true regarding the predictive power of texts as none of the OLS regressions shows significant results.

⁶ For everyone who wonders what the titles of the articles were that are described by the topmost point in the plot; the point stands for only one article with the title “It’s Not You, It’s the Market - Now Officially the Worst S&P Decline in History”. Interestingly, two days after the publication of this very negative article on 8th October 2008, a much worse decline occurred and four articles were written which are described by the leftmost point in the plot. Their titles were “Investor Sentiment: Bullish or Bearish?”, “Too Late to Short SPY? An Historical Perspective”, “Fear the Market or Fear Yourself?” and “Reading the S&P 500’s Crashing Waves”.

5 Conclusion

The first finding is that crowds' attention focuses on a small fraction of stocks only. As a consequence, the prediction of stock price movements by aggregating multiple opinions is only feasible for a few stocks.

It appears that online users follow a naïve momentum mentality and that their sentiment is influenced by the past performance of stocks. With only few exceptions, positive sentiment is preceded by positive abnormal performance and negative sentiment by negative abnormal performance. This effect can be observed clearly in user-generated articles, news messages and microblogging posts. No clear statement can be derived for comments on these texts. An additional regression of the past performance of the S&P 500 index and the negativity in articles concerning that index further confirms the naïve momentum assumption. Here, a high negativity in the texts can be associated with poor prior index performance. These findings are consistent with the results from TUMARKIN AND WHITELAW (2001), DEWALLY (2003), TETLOCK ET AL. (2008) but are not in line with SPRENGER AND WELPE (2010) who ascertain that online users follow a contrarian strategy.

Overall, no convincing results have been found for the predictive power of SA content. User-generated articles which are written by a small group of very active users are to some extent predictive for positive abnormal returns, but not for negative abnormal returns. Authors of articles can be regarded as being experts because of high quality requirements for this text category. Comments on articles are followed by significant positive abnormal returns. Surprisingly, this is the case for bullish as well as for bearish comments. Bullish comments are showing the most significant after-event abnormal performance of all text categories. Besides articles and comments on articles, no other text category contains significant predictive information. Neither the market currents news stream nor microblogging posts can be proven to precede abnormal stock returns. As a result, the opinion of premium users is more valuable than the opinion of regular users and news.

Taking into account that date clustering might reduce the significance of our observations, the results on the predictive power of online content are rather weak. The influence of past abnormal performance on the sentiment of online users is highly significant, though, and will withstand certain significance reductions.

As the literature suggests, the application of other methods for the sentiment extraction or the analysis of the influence on stock market activity may be worthwhile and will possibly allow deeper inspections.

6 References

- Antweiler**, Werner; **Frank**, Murray Z. (2004): Is All That Talk Just Noise? The Information Content of Internet Stock Message Boards. In: *Journal of Finance* 59 (3), 1259–1294.
- Bagnoli**, Mark; **Beneish**, Messod D.; **Watts**, Susan G. (1999): Whisper forecasts of quarterly earnings per share. In: *Journal of Accounting and Economics* 28 (1), 27–50.
- Baker**, Malcolm; **Wurgler**, Jeffrey (2006): Investor Sentiment and the Cross-Section of Stock Returns. In: *Journal of Finance* 61 (4), 1645–1680.
- Bollen**, Johan; **Mao**, Huina; **Zeng**, Xiaojun (2011): Twitter mood predicts the stock market. In: *Journal of Computational Science* 2 (1), 1–8.
- Brown**, Stephen J.; **Warner**, Jerold B. (1985): Using daily stock returns. In: *Journal of Financial Economics* 14 (1), S. 3–31.
- Cha**, Meeyoung; **Haddadi**, Hamed; **Benevenuto**, Fabricio; **Gummadi**, Krishna P. (2010): Measuring User Influence in Twitter. The Million Follower Fallacy. In: In Fourth International AAAI Conference on Weblogs and Social Media.
- Choudhury**, Munmun de; **Brusilovsky**, Peter; **Davis**, Hugh; **Sundaram**, Hari; **John**, Ajita; **Seligmann**, Dorée Duncan (2010): Can blog communication dynamics be correlated with stock market activity? In: *Proceedings of the nineteenth ACM conference on Hypertext and hypermedia*.
- Chen**, Hailiang; **De**, Prabuddha; **Hu**, Yu Jeffrey; **Hwang**, Byoung-Hyoun (2011): The Customer as Advisor: The Role of Social Media in Financial Markets. Online link: <http://ssrn.com/abstract=1807265>, checked on 10/24/2011.
- Das**, S. R.; **Chen**, M. Y. (2007): Yahoo! for Amazon: Sentiment Extraction from Small Talk on the Web. In: *Management Science* 53 (9), 1375–1388.
- Dewally**, Michaël (2003): Internet Investment Advice: Investing with a Rock of Salt. In: *Financial Analysts Journal* (59(4)), 65–77.
- Dewally**, Michaël (2008): The Informational Value of Earnings Whispers. In: *American Journal of Business* 23 (1), 37–52.
- Fama**, Eugene F. (1970): Efficient Capital Markets. In: *J Finance* 25 (2), 383–417.
- Fama**, Eugene F. (1991): Efficient Capital Markets II. In: *Journal of Finance* 46 (5), 1575–1617.
- Fama**, Eugene F.; **French**, Kenneth R. (1992): The Cross-Section of Expected Stock Returns. In: *Journal of Finance* 47 (2), 427–465.

- Fama**, Eugene F.; **Fisher**, Lawrence; **Jensen**, Michael C.; **Roll**, Richard W. (1969): The Adjustment of Stock Prices to New Information. In: *International Economic Review* 10, 1–21.
- Granger**, Clive William John (1969): Investigating Causal Relations by Econometric Models and Cross-spectral Methods. In: *Econometrica* 37 (3).
- Loughran**, Tim; **McDonald**, Bill (2011): When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks. In: *Journal of Finance* 66 (1), 35–65.
- MacKinlay**, Craig A. (1997): Event Studies in Economics and Finance. In: *Journal of Economic Literature* 35 (1), 13–39.
- McCallum**, Andrew (1996): A Toolkit for Statistical Language Modeling, Text Retrieval, Classification and Clustering. Carnegie-Mellon University, Pittsburgh, PA. School of Computer Science. Online link: <http://www.cs.umass.edu/~mccallum>, checked on 10/24/2011.
- O'Reilly**, Tim (2007): What is Web 2.0. Design Patterns and Business Models for the Next Generation of Software. In: *Communications & Strategies* (1), 17.
- Pang**, Bo; **Lee**, Lillian (2008): Opinion Mining and Sentiment Analysis. In: *FNT in Information Retrieval* 2 (1–2), 1–135.
- Schumaker**, Robert P.; **Chen**, Hsinchun (2009): Textual analysis of stock market prediction using breaking financial news: The AZFin text system. In: *ACM Transactions on Information Systems (TOIS)* 27 (2).
- Sprenger**, Timm O.; **Welpe**, Isabell M. (2010): Tweets and Trades: The Information Content of Stock Microblogs. Working Paper. Technical University of Munich, Munich.
- Tetlock**, Paul C. (2007): Giving Content to Investor Sentiment: The Role of Media in the Stock Market. In: *Journal of Finance* 62 (3), 1139–1168.
- Tetlock**, Paul C.; **Saar-Tsechansky**, Maytal; **Macskassy**, Sofus (2008): More Than Words: Quantifying Language to Measure Firms' Fundamentals. In: *Journal of Finance* 63 (3), 1437–1467.
- Tetlock**, P. C. (2011): All the News That's Fit to Reprint: Do Investors React to Stale Information? In: *Review of Financial Studies* 24 (5), 1481–1512.
- Tumarkin**, Robert; **Whitelaw**, Robert F. (2001): News or Noise? Internet Postings and Stock Prices. In: *Financial Analysts Journal* (57(3)), 41–51.
- Kumar**, Ravi; **Ungar**, Lyle; **Craven**, Mark; **Gunopulos**, Dimitrios; **Eliassi-Rad**, Tina; **Novak**, Jasmine; **Tomkins**, Andrew (2010): Structure and evolution of online social networks. In: *Link Mining: Models, Algorithms, and Applications* 4, 337–357.
- van Bommel**, Jos (2003): Rumors. In: *Journal of Finance* 58 (4), 1499–1520.

Wilson, Theresa; **Mooney**, Raymond J.; **Wiebe**, Janyce; **Hoffmann**, Paul (2005): Recognizing contextual polarity in phrase-level sentiment analysis. In: Human Language Technologies Conference/Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP 2005), 347–354.

Wysocki, Peter D. (1998): Cheap Talk on the Web: The Determinants of Postings on Stock Message Boards. Working Paper. University of Michigan, Ann Arbor.

Zhang, Xue; **Fuehres**, Hauke; **Gloor**, Peter (2010): Predicting stock market indicators through Twitter. “I hope it is not as bad as I fear”. In: Collaborative Innovations Networks Conference (Savannah), 1–8.

7 Appendix

7.1 List of ticker symbols (Standardandpoors.com checked on 10/14/2011)

A	ARG	CCE	CVC	EQR	GNW	IRM
AA	ATI	CCL	CVH	EQT	GOOG	ISRG
AAPL	AVB	CEG	CVS	ERTS	GPC	ITT
ABC	AVP	CELG	CVX	ESRX	GPS	ITW
ABT	AVY	CEPH	D	ETFC	GR	IVZ
ACE	AXP	CERN	DD	ETN	GS	JBL
ACN	AZO	CF	DE	ETR	GT	JCI
ADBE	BA	CFN	DELL	EW	GWV	JCP
ADI	BAC	CHK	DF	EXC	HAL	JDSU
ADM	BAX	CHRW	DFS	EXPD	HAR	JEC
ADP	BBBY	CI	DGX	EXPE	HAS	JNJ
ADSK	BBT	CINF	DHI	F	HBAN	JNPR
AEE	BBY	CL	DHR	FAST	HCBK	JNS
AEP	BCR	CLF	DIS	FCX	HCN	JOYG
AES	BDX	CLX	DISCA	FDO	HCP	JPM
AET	BEN	CMA	DNB	FDX	HD	JWN
AFL	BHI	CMCSA	DNR	FE	HES	K
AGN	BIG	CME	DO	FFIV	HIG	KEY
AIG	BIIB	CMG	DOV	FHN	HNZ	KFT
AIV	BJS ⁷	CMI	DOW	FII	HOG	KIM
AIZ	BK	CMS	DPS	FIS	HON	KLAC
AKAM	BLK	CNP	DRI	FISV	HOT	KMB
AKS	BLL	CNX	DTE	FITB	HP	KMX
ALL	BMC	COF	DTV	FLIR	HPQ	KO
ALTR	BMS	COG	DUK	FLR	HRB	KR
AMAT	BMV	COH	DV	FLS	HRL	KSS
AMD	BRCM	COL	DVA	FMC	HRS	L
AMGN	BSX	COP	DVN	FO ⁸	HSP	LEG
AMP	BTU	COST	EBAY	FRX	HST	LEN
AMT	BXP	COV	ECL	FSLR	HSY	LH
AMZN	C	CPB	ED	FTI	HUM	LIFE
AN	CA	CPWR	EFX	FTR	IBM	LLL
ANF	CAG	CRM	EIX	GAS	ICE	LLTC
ANR	CAH	CSC	EK	GCI	IFF	LLY
AON	CAM	CSCO	EL	GD	IGT	LM
APA	CAT	CSX	EMC	GE	INTC	LMT
APC	CB	CTAS	EMN	GILD	INTU	LNC
APD	CBE ⁹	CTL	EMR	GIS	IP	LO
APH	CBG	CTSH	EOG	GLW	IPG	LOW
APOL	CBS	CTXS	EP	GME	IR	LSI

⁷ Excluded on 04/29/2010 due to acquisition

(<http://www.standardandpoors.com/ratings/articles/en/us/?assetID=1245210962207>).

⁸ Excluded on 09/29/2011 due to spin off

(http://www.msnbc.msn.com/id/44721788/ns/business-stocks_and_economy/t/sp-backs-fortune-brands-rating-spin-off-nears/#.TpvwnEfeVnM).

⁹ Excluded on 08/28/2009

(<http://news.pnnewswire.com/DisplayReleaseContent.aspx?ACCT=104&STORY=/www/story/08-28-2009/0005084772&EDATE=>).

LTD	MWV	PBCT	RDC	STJ	UNH	XEL
LUK	MWW	PBI	RF	STR ¹⁰	UNM	XL
LUV	MYL	PCAR	RHI	STT	UNP	XLNX
LXK	NBL	PCG	RHT	STZ	UPS	XOM
M	NBR	PCL	RL	SUN	URBN	XRAY
MA	NDAQ	PCLN	ROK	SVU	USB	XRX
MAR	NE	PCP	ROP	SWK	UTX	YHOO
MAS	NEE	PCS	ROST	SWN	V	YUM
MAT	NEM	PDCO	RRC	SWY	VAR	ZION
MCD	NFLX	PEG	RRD	SYK	VFC	ZMH
MCHP	NFX	PEP	RSG	SYMC	VLO	
MCK	NI	PFE	RTN	SYU	VMC	
MCO	NKE	PFG	S	T	VNO	
MDT	NOC	PG	SAI	TAP	VRSN	
MEE ¹¹	NOV	PGN	SBUX	TDC	VTR	
MET	NRG	PGR	SCG	TE	VZ	
MHP	NSC	PH	SCHW	TEG	WAG	
MHS	NSM ¹²	PHM	SE	TER	WAT	
MI ¹³	NTAP	PKI	SEE	TGT	WDC	
MJN	NTRS	PLD	SHLD	THC	WEC	
MKC	NU	PLL	SHW	TIE	WFC	
MMC	NUE	PM	SIAL	TIF	WFM	
MMI	NVDA	PNC	SI ¹⁴	TJX	WFR	
MMM	NVLS	PNW	SJM	TLAB	WHR	
MO	NWL	POM	SLB	TMK	WIN	
MOLX	NWSA	PPG	SLE	TMO	WLP	
MON	NYT ¹⁵	PPL	SLM	TROW	WM	
MPC	NYX	PRU	SNA	TRV	WMB	
MRK	ODP ¹⁶	PSA	SNDK	TSN	WMT	
MRO	OI	PWR	SNI	TSO	WPI	
MS	OKE	PX	SO	TSS	WPO	
MSFT	OMC	PXD	SPG	TWC	WU	
MSI	ORCL	QCOM	SPLS	TWX	WY	
MTB	ORLY	QEP	SRCL	TXN	WYN	
MU	OXY	R	SRE	TXT	WYNN	
MUR	PAYX	RAI	STI	TYC	X	

¹⁰ Excluded on 06/23/2010 due to replacement
(<http://www.s-p-500.com/qep-resources-replaces-questar-corp-on-sp-500/>).

¹¹ Excluded on 05/01/2011 due to replacement
(<http://www.s-p-500.com/anr-to-replace-mee-on-sp-500/>).

¹² Excluded on 09/23/2011 due to acquisition
(http://newscenter.ti.com/Blogs/newsroom/archive/2011/09/23/texas-instruments-completes-acquisition-of-national-semiconductor-813808.aspx?DCMP=National_branding_NS&HQS=more-bhprn).

¹³ Excluded on 06/28/2011 due to replacement
(<http://www.usatoday.com/money/perfi/columnist/krantz/story/2011-09-22/is-standard-poors-500-index-misleading/50514580/1>).

¹⁴ Excluded on 08/10/2010 due to replacement
(http://markets.financialcontent.com/stocks/news/read/14385256/tyco_to_replace_smith_international_on_s&p_500).

¹⁵ Excluded on 12/10/2010 due to replacement
(http://money.cnn.com/2010/12/10/markets/SP_500_new_companies/index.htm).

¹⁶ Ib.

7.2 Log-log plots of mentions distributions

7.2.1 Market currents




Fig. 9. Log-log diagram of quantity distribution of stock mentions for market currents

7.2.2 Comments on market currents




Fig. 10. Log-log diagram of quantity distribution of stock mentions for comments on market currents

7.2.3 Stocktalks




Fig. 11. Log-log diagram of quantity distribution of stock mentions for stocktalks

In der Reihe FINAL sind bisher erschienen:

1. Jahrgang 1991:

1. Hinrich E. G. Bonin; Softwaretechnik, Heft 1, 1991 (ersetzt durch Heft 2, 1992).
2. Hinrich E. G. Bonin (Herausgeber); Konturen der Verwaltungsinformatik, Heft 2, 1991 (überarbeitet und erschienen im Wissenschaftsverlag, Bibliographisches Institut & F. A. Brockhaus AG, Mannheim 1992, ISBN 3-411-15671-6).

2. Jahrgang 1992:

1. Hinrich E. G. Bonin; Produktionshilfen zur Softwaretechnik --- Computer-Aided Software Engineering --- CASE, Materialien zum Seminar 1992, Heft 1, 1992.
2. Hinrich E. G. Bonin; Arbeitstechniken für die Softwareentwicklung, Heft 2, 1992 (3. überarbeitete Auflage Februar 1994), PDF-Format.
3. Hinrich E. G. Bonin; Object-Orientedness --- a New Boxologie, Heft 3, 1992.
4. Hinrich E. G. Bonin; Objekt-orientierte Analyse, Entwurf und Programmierung, Materialien zum Seminar 1992, Heft 4, 1992.
5. Hinrich E. G. Bonin; Kooperative Produktion von Dokumenten, Materialien zum Seminar 1992, Heft 5, 1992.

3. Jahrgang 1993:

1. Hinrich E. G. Bonin; Systems Engineering in Public Administration, Proceedings IFIP TC8/ WG8.5: Governmental and Municipal Information Systems, March 3--5, 1993, Lüneburg, Heft 1, 1993 (überarbeitet und erschienen bei North-Holland, IFIP Transactions A-36, ISSN 0926-5473).
2. Antje Binder, Ralf Linhart, Jürgen Schultz, Frank Sperschneider, Thomas True, Bernd Willenbockel; COTEXT --- ein Prototyp für die kooperative Produktion von Dokumenten, 19. März 1993, Heft 2, 1993.
3. Gareth Harries; An Introduction to Artificial Intelligence, April 1993, Heft 3, 1993.
4. Jens Benecke, Jürgen Grothmann, Mark Hilmer, Manfred Hölzen, Heiko Köster, Peter Mattfeld, Andre Peters, Harald Weiss; ConFusion --- Das Produkt des AWÖ-Projektes 1992/93, 1. August 1993, Heft 4, 1993.
5. Hinrich E. G. Bonin; The Joy of Computer Science --- Skript zur Vorlesung EDV ---, September 1993, Heft 5, 1993 (4. ergänzte Auflage März 1995).
6. Hans-Joachim Blanke; UNIX to UNIX Copy --- Interactive application for installation and configuration of UUCP ---, Oktober 1993, Heft 6, 1993.

4. Jahrgang 1994:

1. Andre Peters, Harald Weiss; COMO 1.0 --- Programmierumgebung für die Sprache COBOL --- Benutzerhandbuch, Februar 1994, Heft 1, 1994.
2. Manfred Hölzen; UNIX-Mail --- Schnelleinstieg und Handbuch ---, März 1994, Heft 2, 1994.
3. Norbert Kröger, Roland Seen; EBrain --- Documentation of the 1994 AWÖ-Project Prototype ---, June 11, 1994, Heft 3, 1994.
4. Dirk Mayer, Rainer Saalfeld; ADLATUS --- Documentation of the 1994 AWÖ-Project Prototype -- -, July 26, 1994, Heft 4, 1994.
5. Ulrich Hoffmann; Datenverarbeitungssystem 1, September 1994, Heft 5, 1994. (2. überarbeitete Auflage Dezember 1994).
6. Karl Goede; EDV-gestützte Kommunikation und Hochschulorganisation, Oktober 1994, Heft 6 (Teil 1), 1994.
7. Ulrich Hoffmann; Zur Situation der Informatik, Oktober 1994, Heft 6 (Teil 2), 1994.

5. Jahrgang 1995:

1. Horst Meyer-Wachsmuth; Systemprogrammierung 1, Januar 1995, Heft 1, 1995.
2. Ulrich Hoffmann; Datenverarbeitungssystem 2, Februar 1995, Heft 2, 1995.
3. Michael Guder / Kersten Kalischefski / Jörg Meier / Ralf Stöver / Cheikh Zeine; OFFICE-LINK --- Das Produkt des AWÖ-Projektes 1994/95, März 1995, Heft 3, 1995.
4. Dieter Riebesehl; Lineare Optimierung und Operations Research, März 1995, Heft 4, 1995.
5. Jürgen Mattern / Mark Hilmer; Sicherheitsrahmen einer UTM-Anwendung, April 1995, Heft 5, 1995.
6. Hinrich E. G. Bonin; Publizieren im World-Wide Web --- HyperText Markup Language und die Kunst der Programmierung ---, Mai 1995, Heft 6, 1995.
7. Dieter Riebesehl; Einführung in Grundlagen der theoretischen Informatik, Juli 1995, Heft 7, 1995.
8. Jürgen Jacobs; Anwendungsprogrammierung mit Embedded-SQL, August 1995, Heft 8, 1995.
9. Ulrich Hoffmann; Systemnahe Programmierung, September 1995, Heft 9, 1995 (ersetzt durch Heft 1, 1999).
10. Klaus Lindner; Neuere statistische Ergebnisse, Dezember 1995, Heft 10, 1995.

6. Jahrgang 1996:

1. Jürgen Jacobs / Dieter Riebesehl; Computergestütztes Repetitorium der Elementarmathematik, Februar 1996, Heft 1, 1996.
2. Hinrich E. G. Bonin; "Schlanker Staat" & Informatik, März 1996, Heft 2, 1996.
3. Jürgen Jacobs; Datenmodellierung mit dem Entity-Relationship-Ansatz, Mai 1996, Heft 3, 1996.
4. Ulrich Hoffmann; Systemnahe Programmierung, (2. überarbeitete Auflage von Heft 9, 1995), September 1996, Heft 4, 1996 (ersetzt durch Heft 1, 1999).
5. Dieter Riebesehl; Prolog und relationale Datenbanken als Grundlagen zur Implementierung einer NF2-Datenbank (Sommer 1995), November 1996, Heft 5, 1996.

7. Jahrgang 1997:

1. Jan Binge, Hinrich E. G. Bonin, Volker Neumann, Ingo Stadtsholte, Jürgen Utz; Intranet-/Internet- Technologie für die Öffentliche Verwaltung --- Das AWÖ-Projekt im WS96/97 --- (Anwendungen in der Öffentlichen Verwaltung), Februar 1997, Heft 1, 1997.
2. Hinrich E. G. Bonin; Auswirkungen des Java-Konzeptes für Verwaltungen, FTVI'97, Oktober 1997, Heft 2, 1997.

8. Jahrgang 1998:

1. Hinrich E. G. Bonin; Der Java-Coach, Heft 1, Oktober 1998, (CD-ROM, PDF-Format; aktuelle Fassung).
2. Hinrich E. G. Bonin (Hrsg.); Anwendungsentwicklung WS 1997/98 --- Programmierbeispiele in COBOL & Java mit Oracle, Dokumentation in HTML und tcl/tk, September 1998, Heft 2, 1998 (CD-ROM).
3. Hinrich E. G. Bonin (Hrsg.); Anwendungsentwicklung SS 1998 --- Innovator, SNIFF+, Java, Tools, Oktober 1998, Heft 3, 1998 (CD-ROM).
4. Hinrich E. G. Bonin (Hrsg.); Anwendungsentwicklung WS 1998 --- Innovator, SNIFF+, Java, Mail und andere Tools, November 1998, Heft 4, 1998 (CD-ROM).
5. Hinrich E. G. Bonin; Persistente Objekte --- Der Elchtest für ein Java-Programm, Dezember 1998, Heft 5, 1998 (CD-ROM).

9. Jahrgang 1999:

1. Ulrich Hoffmann; Systemnahe Programmierung (3. überarbeitete Auflage von Heft 9, 1995), Juli 1999, Heft 1, 1999 (CD-ROM und Papierform), Postscript-Format, zip-Postscript-Format, PDF-Format und zip-PDF-Format.

10. Jahrgang 2000:

1. Hinrich E. G. Bonin; Citizen Relationship Management, September 2000, Heft 1, 2000 (CD-ROM und Papierform), PDF-Format.
2. Hinrich E. G. Bonin; WI>DATA --- Eine Einführung in die Wirtschaftsinformatik auf der Basis der Web_Technologie, September 2000, Heft 2, 2000 (CD-ROM und Papierform), PDF-Format.
3. Ulrich Hoffmann; Angewandte Komplexitätstheorie, November 2000, Heft 3, 2000 (CD-ROM und Papierform), PDF-Format.
4. Hinrich E. G. Bonin; Der kleine XMLer, Dezember 2000, Heft 4, 2000 (CD-ROM und Papierform), PDF-Format, aktuelle Fassung.

11. Jahrgang 2001:

1. Hinrich E. G. Bonin (Hrsg.): 4. SAP-Anwenderforum der FHNON, März 2001, (CD-ROM und Papierform), Downloads & Videos.
2. J. Jacobs / G. Weinrich; Bonitätsklassifikation kleiner Unternehmen mit multivariater linear Diskriminanzanalyse und Neuronalen Netzen; Mai 2001, Heft 2, 2001, (CD-ROM und Papierform), PDF-Format und MS Word DOC-Format
3. K. Lindner; Simultantestprozedur für globale Nullhypothesen bei beliebiger Abhängigkeitsstruktur der Einzeltests, September 2001, Heft 3, 2001 (CD-ROM und Papierform).

12. Jahrgang 2002:

1. Hinrich E. G. Bonin: Aspect-Oriented Software Development. März 2002, Heft 1, 2002 (CD-ROM und Papierform), PDF-Format.
2. Hinrich E. G. Bonin: WAP & WML --- Das Projekt Jagdzeit ---. April 2002, Heft 2, 2002 (CD-ROM und Papierform), PDF-Format.
3. Ulrich Hoffmann: Ausgewählte Kapitel der Theoretischen Informatik (CD-ROM und Papierform), PDF-Format.
4. Jürgen Jacobs / Dieter Riebesehl; Computergestütztes Repetitorium der Elementarmathematik, September 2002, Heft 4, 2002 (CD-ROM und Papierform), PDF-Format.
5. Verschiedene Referenten; 3. Praxisforum "Systemintegration", 18.10.2002, Oktober 2002, Heft 5, 2002 (CD-ROM und Papierform), Praxisforum.html (Web-Site).

13. Jahrgang 2003:

1. Ulrich Hoffmann; Ausgewählte Kapitel der Theoretischen Informatik; Heft 1, 2003, (CD-ROM und Papierform) PDF-Format.
2. Dieter Riebesehl; Mathematik 1, Heft 2, 2003, (CD-ROM und Papierform) PDF-Format.
3. Ulrich Hoffmann; Mathematik 1, Heft 3, 2003, (CD-ROM und Papierform) PDF-Format und Übungen.
4. Verschiedene Autoren; Zukunft von Verwaltung und Informatik, Festschrift für Heinrich Reinermann, Heft 4, 2003, (CD-ROM und Papierform) PDF-Format.

14. Jahrgang 2004:

1. Jürgen Jacobs; Multilayer Neural Networks; Heft 1, 2004, (CD-ROM und Papierform) PDF-Format.

15. Jahrgang 2005:

1. Ulrich Hoffmann; Mathematik für Wirtschaftsinformatiker; Heft 1, 2005, (CD-ROM und Papierform) PDF-Format.
2. Ulrich Hoffmann; Übungen & Lösungen zur Mathematik für Wirtschaftsinformatiker; Heft 1, 2005, (CD-ROM und Papierform) PDF-Format.
3. Ulrich Hoffmann; Datenstrukturen & Algorithmen; Heft 2, 2005, (CD-ROM und Papierform) PDF-Format.

16. Jahrgang 2006:

1. Hinrich E. G. Bonin; Systemanalyse für Softwaresysteme; Heft 1, August 2006, (CD-ROM und Papierform) PDF-Format.
2. Hinrich E. G. Bonin; Faszination Programmierung; Heft 2, August 2006, (CD-ROM und Papierform) PDF-Format.
3. Dieter Riebesehl; Strukturanalogien in Datenmodellen, Heft 3, Dezember 2006, (CD-ROM und Papierform) PDF-Format.

17. Jahrgang 2007:

1. Ulrich Hoffmann; Ausgewählte Kapitel der Theoretischen Informatik; Heft 1, August 2007, (CD-ROM und Papierform) PDF-Format.
2. Ulrich Hoffmann; Mathematik für Wirtschaftsinformatiker und Informatiker; Heft 2, August 2007, (CD-ROM und Papierform) PDF-Format.
3. Hinrich E. G. Bonin; Der Java-Coach, Heft 3, September 2007, (CD-ROM und Papierform) PDF-Format.
4. Jürgen Jacobs; Dichteprognoze autoregressiver Zeitreihen, Heft 4, September 2007, (CD-ROM und Papierform) PDF-Format.

18. Jahrgang 2008:

1. Verschiedene Autoren; Festschrift für Prof. Dr. Meyer-Wachsmuth; Heft 1, Juli 2008, (CD-ROM und Papierform) PDF-Format.
2. Ulrich Hoffmann; Ausgewählte Kapitel der Theoretischen Informatik; Heft 2, Dezember 2008, (CD-ROM und Papierform) PDF-Format.

19. Jahrgang 2009:

1. Verschiedene Autoren; Festschrift für Prof. Dr. Goede; Heft 1, August 2009, (CD-ROM und Papierform) PDF-Format.

20. Jahrgang 2010:

1. Hinrich E. G. Bonin; Konstrukte, Konstruktionen, Konstruktionsempfehlungen – Programmieren in LISP; Heft 1, März 2010, (CD-ROM und Papierform) PDF-Format.
2. Verschiedene Autoren; Festschrift für Prof. Dr. Bonin; Heft 2, April 2010, (CD-ROM und Papierform) PDF-Format.
3. Verschiedene Autoren; Frühwarnindikatoren und Risikomanagement,
1. Forschungssymposium an der Leuphana Universität Lüneburg, Oktober 2009; Heft 3, April 2010, (CD-ROM und Papierform) PDF-Format.

21. Jahrgang 2011:

1. Verschiedene Autoren; Frühwarnindikatoren und Risikomanagement,
2. Forschungssymposium an der Leuphana Universität Lüneburg, November 2010; Heft 1, Februar 2011, (CD-ROM und Papierform) PDF-Format.

22. Jahrgang 2012:

1. Andreas Mastel, Jürgen Jacobs; Mining User-Generated Financial Content to Predict Stock Price Movements, Heft 1, Dezember 2012, (CD-ROM und Papierform) PDF-Format.

Herausgeber der Schriftenreihe FINAL:

Prof. Dr. Ulrich Hoffmann

Leuphana Universität Lüneburg, Volgershall 1, D-21339 Lüneburg, Germany

email: ulrich.hoffmann@uni.leuphana.de

Verlag:

Eigenverlag (Fotographische Vervielfältigung), Leuphana Universität Lüneburg
(vormals Fachhochschule Nordostniedersachsen)

Erscheinungsweise:

ca. 4 Hefte pro Jahr.

Für unverlangt eingesendete Manuskripte wird nicht gehaftet. Sie sind aber willkommen.

Digitales FInAL-Archiv:

<http://www.leuphana.de/institute/iwi/final.html>

Copyright:

All rights, including translation into other languages reserved by the authors. No part of this report may be reproduced or used in any form or by any means --- graphic, electronic, or mechanical, including photocopying, recording, taping, or information and retrieval systems --- without written permission from the authors, except for noncommercial, educational use, including classroom teaching purposes.

Copyright: Hoffmann Apr-1995,..., Dezember 2012, all rights reserved