Contents lists available at ScienceDirect



Journal of Financial Markets



journal homepage: www.elsevier.com/locate/finmar

Investor sentiment and stock returns: Wisdom of crowds or power of words? Evidence from Seeking Alpha and Wall Street Journal^{star}



Ioanna Lachana*, David Schröder

Birkbeck College, University of London, United Kingdom

ARTICLE INFO

JEL classification: G12 G41 D53 Keywords: Investor sentiment Text mining Financial markets Traditional media Social media Stock market returns

ABSTRACT

In light of changes in the media landscape from traditional print towards social media, in this study we compare the ability of investor sentiment measures obtained from various media sources to predict short-term market returns. We show that investor sentiment extracted from the social media platform Seeking Alpha is better in predicting market returns than investor sentiment obtained from the *Wall Street Journal*, a traditional print medium. Seeking Alpha is more suitable for the extraction of investor sentiment due to the richer language and timeliness of online media.

1. Introduction

Social media has become an important data source for financial market participants. For example, in their seminal paper, Chen et al. (2014) show that information shared by users of the investment-related social media platform Seeking Alpha can be used to predict the returns of individual stocks. However, articles and stories published on internet platforms might not only be valuable to extract information, but could also be useful to measure the more elusive concept of irrational investor mood or sentiment. According to Baker and Wurgler (2007), investment decisions are not always based on facts, but are to some extent irrationally influenced by investor mood. Shiller (2016) suggests that the media plays a vital part in capturing and propagating such market sentiment.

In light of the evolution in the media landscape from traditional print towards social media platforms, the objective of this study is to analyze and compare investor sentiment extracted from various media sources in their ability to predict short-term stock market returns, and to explore the underlying reasons for possible differences.

Using a large data set of daily articles and reader comments from 2006 to 2020, we analyze the effect of investor sentiment obtained from social media and traditional media on market returns for up to five trading days. Following previous studies, we use

E-mail addresses: ilacha01@mail.bbk.ac.uk (I. Lachana), d.schroeder@bbk.ac.uk (D. Schröder).

https://doi.org/10.1016/j.finmar.2025.100970

Available online 3 March 2025

 $[\]stackrel{\circ}{\sim}$ Earlier versions of this paper circulated under the titles "Investor sentiment, media and stock returns: The advancement of social media" and "Investor sentiment, media and stock returns: Wisdom of crowds or power of words?". We would like to thank Diego García for providing access to his print media data, and Seeking Alpha for providing access to the social media data. We would also like to thank Roberto Gomez Cram, David Hirshleifer, Byoung-Hyoun Hwang, Ron Smith, Clara Vega (the editor), conference and seminar participants at Birkbeck College (University of London), the World Finance Conference 2021, the Annual Southern Ph.D. Economics Conference 2021, and the Royal Economic Society 2022 Annual Conference. For the purposes of open access, the authors have applied a CC BY public copyright licence to any author accepted manuscript version arising from this submission. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Correspondence to: Birkbeck College, University of London, Malet Street, Bloomsbury, London WC1E 7HX, United Kingdom.

Received 12 January 2024; Received in revised form 4 February 2025; Accepted 18 February 2025

^{1386-4181/© 2025} The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

I. Lachana and D. Schröder

social media data from the Seeking Alpha (SA) platform (Chen et al., 2014) and articles published in the *Wall Street Journal* (WSJ) as a prominent example of traditional print media (Tetlock, 2007). We extract investor sentiment from the data using textual analysis based on the lexicon approach brought forward by Loughran and McDonald (2011).

Investor sentiment extracted from articles appearing in the Today's Market section of Seeking Alpha is significantly related to short-term equity market returns. The relation between investor sentiment and stock market returns is however transitory, not causing any permanent changes to asset prices. This result is in line with the hypothesis of Shiller (2016), who considers that the investment decisions driven by investor mood are largely irrational.

By contrast, investor sentiment obtained from the daily *Markets* column of the *Wall Street Journal* is only weakly related to stock returns. At first sight, this finding seems at odds with Tetlock (2007) and Garcia (2013), who document a strong relation between investor sentiment obtained from print media and stock returns during the twentieth century. Going back to the original data used in Garcia (2013), we find that the relation between sentiment measures obtained from print media and stock markets has decayed over time, thereby reconciling our results with Tetlock (2007) and Garcia (2013).

These results are obtained under the implicit assumption that articles appearing on SA and the WSJ are very similar (apart from appearing on different media websites), and can thus be directly compared. Yet, although articles appearing on SA and the WSJ both discuss developments of the financial markets, there are some considerable differences. The daily *Markets* column of the WSJ is published every trading day after the market close, and written by a rather small number of professional journalists. These articles adhere to a specific format and are of similar length, presumably due to space constraints in the corresponding print edition. By contrast, as a pure online medium, SA has no space constraints. This means there is no limit for the number of daily articles appearing in the Today's Market section, and these articles can be of any length. Articles appearing on SA are written by thousands of independent and non-professional contributors. In addition, these authors are not expected to publish articles on a regular basis, but have a choice of whether or not to contribute.

These differences might explain why sentiment extracted from SA is better in explaining short-term market returns than sentiment obtained from the WSJ. For example, it is a well-known result from the forecasting literature (Wang et al., 2023) that the average of many forecasts is usually a better predictor than a single forecast. Thus, the larger volume of SA articles relative to WSJ articles might naturally give sentiment extracted from SA an edge over the WSJ column. Another reason for the superiority of SA sentiment measures could be the large pool of different contributors on social media platforms, also known as the "wisdom of crowds" (Surowiecki, 2005): The large diversity of authors on SA might allow the aggregation of a much larger set of opinions compared to the few journalists employed at the WSJ. Furthermore, since SA authors have a choice (but not the obligation) to publish articles, this might create an endogeneity problem. For example, if SA authors chose to publish only on days with more activity in the stock market, then the higher predictive power of investor mood extracted from SA articles might be explained by such a self-selection bias. Finally, if the tone conveyed in articles varies with length and complexity, some of the stronger association between SA sentiment and stock returns might be explained by the higher variability in article length.

To address these concerns, we conduct a series of additional tests. We create various sub-samples of all SA publications that match the characteristics of the WSJ sample. Thus, if the reasons discussed above can indeed explain the difference between the sentiment measures extracted from SA and WSJ, this difference should disappear when using a sample of SA articles that is identical to the WSJ sample. Yet, when matching (a) the number of authors, (b) the article length and variance, and (c) the volume of articles, the sentiment extracted from SA articles still predicts market returns. Thus, we can rule out these different explanations to drive the results.

Then we investigate two more reasons why sentiment extracted from social media could be better than traditional media to capture investor sentiment. First, we consider differences in writing style and language. Following the lexicon approach, investor sentiment is measured as the fraction of "positive" and "negative" words in the (Loughran and McDonald, 2011) dictionary. Since social media posts use a richer language than newspapers articles (Hagvar, 2019), sentiment measures from social media could be more valuable. Indeed, when matching the fraction of dictionary words of SA articles to the ratio observed in the WSJ, SA sentiment indices lose their explanatory power for stock returns. Second, timing differences could also play an important role for the effectiveness of the sentiment indices obtained from social media. While WSJ articles appear in the hours after market close, SA publishes articles around the clock, including the time during the trading hours. Since the effect of investor sentiment is short-lived (Renault, 2017), such posts might be more relevant for asset prices. Dividing the sample of SA articles in two sets depending on the publication time, we can support this hypothesis. Sentiment extracted from articles published during trading hours have a significantly stronger relation with market returns than sentiment measured outside trading hours.

In the context of this study, a natural question arises: Is the effect of investor sentiment on short-term stock returns only driven by irrational investor mood, as suggested by Shiller (2016), or caused by new information contained in the articles as suggested by Chen et al. (2014)? While it is difficult to disentangle these two explanations perfectly, we adopt the identification strategy by Garcia (2013) to minimize possible confounding effects. Using a variety of analyses, we ensure that all WSJ and SA articles and user comments refer to events that happened prior to the extraction of investor sentiment. If any of these events were sufficiently important to affect the market, the information should be immediately impounded in prices. Thus, our identification strategy is that news affects the market immediately, while sentiment affects the stock market in the subsequent trading days. Furthermore, there is some substantial reversal of stock returns after the initial price impact. This observation is more consistent with the view that sentiment captures short-term irrational investor mood rather than long-term fundamental information. Hence, although we cannot exclude that there is some effect of information contained in the textual data we use, we believe that most of the sentiment effects are driven by naïve investor mood.

I. Lachana and D. Schröder

More generally, it is unlikely that SA readers have sufficient market power for investor mood conveyed in these articles to influence broad market returns. Thus, we do not interpret our results as evidence of a direct effect of the tone reflected in SA articles and comments on asset prices. We rather consider the SA platform as an ideal and timely medium where overall market investor mood is reflected, and can be easily extracted.

This work builds on earlier studies that explore the impact of irrational investor sentiment conveyed in the media on short-term stock market movements.¹ In his seminal paper, Tetlock (2007) is the first to show that a pessimism index constructed from a daily WSJ column can explain short-term, transitory stock market returns. Building on this insight, Garcia (2013) shows that investor sentiment obtained from daily *New York Times* articles can predict short-term returns of the Dow Jones Industrial Average for the extended time period from 1905 to 2005, especially during economic recessions.

More recently, financial economists have turned to social media to extract investor sentiment. Gan et al. (2020) analyze the effect of changes of the media landscape from 2012 to 2017 on the ability of media sentiment to predict stock returns. Using off-the-shelf sentiment indices of traditional and social media provided by Thomson Reuters, they do not find that investor sentiment has any significant predictive power for market returns.

Yet, it seems that the relation between investor sentiment and stock returns depends on the time horizon. Sun et al. (2016) and Renault (2017) show that market sentiment can predict intraday stock returns. In turn, Kraussl and Mirgorodskay (2017) construct a monthly media pessimism indicator to show that media pessimism is associated with negative market returns even in the long run.

The role of media sentiment for stock markets has been studied from increasingly different perspectives. For example, Hillert et al. (2018) use a daily sentiment indicator to capture the disagreement in the tone of journalists. They show that journalist disagreement negatively explains market returns. Next to stock returns, investor sentiment is also significantly related to trading volume (Hanna et al., 2020) and market volatility (Jiao et al., 2020).

In summary, previous papers establish that measures of irrational investor sentiment extracted from different media sources, can explain – among other aspects – short-term and transitory stock market returns. Yet, these studies are somewhat inconclusive as to whether or not sentiment measures obtained from social media are consistently better than measures extracted from traditional print media. More important, an analysis of why one or the other media source is more suitable to capture sentiment has been missing.

Comparing two manually created investor sentiment indices obtained from traditional and social media, the contribution of this study is to show that sentiment indices obtained from social media have a higher predictive ability for short-term stock returns relative to sentiment extracted from traditional media. Most important, this approach allows pinning down the reasons why sentiment based on textual analysis from social media has an edge over traditional media. Our sub-sampling methodology shows that sentiment captured from social media is more valuable due to the timeliness and richer language of its articles. In addition, we are, to our knowledge, the first to show that reader comments on social media can also be used to construct meaningful sentiment indices to predict the markets.

The literature surveyed above is in sharp contrast to a parallel stream of research that aims at extracting fundamental information (rather than irrational investor mood) of social media posts. For example, Chen et al. (2014) extract information from the SA platform to build firm-level indices to predict short- and long-term returns of individual firms. Similarly, Leung and Ton (2015) study how information provided on the Australian social media platform HotCopper can be used to predict the returns of small-cap Australian stocks. Kim and Kim (2014) use data from Yahoo message boards for the 2005–2010 period to predict stock returns, volatility, and trading volume across different horizons. More recently, Frankel et al. (2022) show that disclosure sentiment measures based on machine-learning algorithms are better at predicting firm-level stock returns on 10-K filing and conference call day relative to dictionary-based methods.²

The rest of the study proceeds as follows. In Section 2 we provide a description of the data. In Section 3 we describe the methodology to transform the raw media data into various measures of investor sentiment. Section 4 shows that sentiment proxies extracted from social media are better in explaining market returns than sentiment measures from traditional media sources. Section 5 shows that the primary reason for this result is the more expressive language used on social media platforms. Section 6 shows that (a) information contained in SA articles and (b) reverse causality from stock returns to investor sentiment is unlikely to drive the results. We conclude in Section 7.

2. Data

This study analyzes the effect of investor sentiment as reflected in both traditional and social media on the stock market from 2006 to 2020. This period covers the years of rapid growth of social media platforms, such as Twitter and Facebook. At the same time, this period saw a tremendous transition of traditional print media, moving to different business models by introducing paywalls on their websites while cutting down on their print runs. Finally, this period captures the Great Recession 2007/09, as well as the turbulence caused by the COVID-19 pandemic.

¹ Instead of relying on media sources, investor sentiment can also be extracted from the weather (Hirshleifer and Shumway, 2003), sports events (Edmans et al., 2007), or internet search queries (Da et al., 2015).

² Other important works extract information from Twitter data using textual analysis, including (Sprenger et al., 2014b), Sprenger et al. (2014a), Sul et al. (2017), Bartov et al. (2018), and Li et al. (2018). Siganos et al. (2017) use Facebook data.

2.1. Media data

Social media are relatively new internet-based mass communication portals. In the context of this study, there are two important differences between social media and traditional print media. First, most of the content in social media is generated by non-professional and independent contributors, i.e., the social media activity is not their main source of income. By contrast, print media content is usually provided by professional journalists who are often employed by a specific publishing house. Second, social media platforms allow users to interact with each other, whereas traditional print media are usually one-way communication channels.

2.1.1. Social media

In this study, we use social media data from the Seeking Alpha (SA) platform, in line with previous studies (Chen et al., 2014). The SA platform, established in 2004, is one of the largest investment communities where millions of users connect and discuss important factors influencing the market, as well as anything related to the investment world.

Articles published on SA are written by a large variety of independent contributors, such as individual and institutional investors, fund managers, but also college students and retirees. Up to March 2021, over one million individual articles had been published on SA, written by over 17,000 independent contributors.³

Articles appearing on SA discuss a wide range of topics. Since the focus of this study is on analyzing the relation between investor sentiment and stock market returns, only articles appearing in the Today's Market section are used. In this section, authors provide their views about some important changes in the corporate world, developments that affect the markets as a whole, including the economic or political environment, or they give some investment advice. It is important to note that SA articles rarely discuss current news. Instead they tend to give broader reflections and analyses of events that happened in the past. As Chen et al. (2014, p. 1368) note, SA's aim is to "provide opinion and analysis rather than news". See Fig. 3 for a typical example of a SA article.

Before an article appears on SA, it is reviewed by the SA editorial team to ensure that some minimum standards are met. Poorly written articles are not published. SA contributors can earn money based on how many SA readers read their articles.⁴ The reward system and the practice of editorial review differentiate SA from other social media platforms, where the authors may not have the motivation to write thorough analyses but might express unfounded opinions instead. Another important feature of SA is that articles can be much longer compared to other social media platforms such as Twitter and Stocktwits. As a result, SA contributors can provide more in-depth analyses compared to such platforms.

The SA platform also allows readers to leave comments after every article, including their agreement or disagreement with the opinions expressed in each article. Up to March 2021, readers left more than 24 million commentaries in response to these articles. Following (Chen et al., 2014), we also retrieve all user comments written in response to each of the Today's Market articles. The idea is that these comments are likely to reflect investor sentiment as well. See Fig. 4 for some examples of comments written in response to the SA article presented in Fig. 3. All SA articles and comments are directly provided by SA as raw text data.

2.1.2. Traditional print media

To investigate the ability of traditional print media to capture investor sentiment, we use articles published in the *Wall Street Journal* (WSJ) for the same time period, following the seminal paper by Tetlock (2007). Founded in 1889, the WSJ is one of the oldest financial newspapers, covering different industries and markets. Until today, it has one of the largest daily circulations in the United States.

The WSJ publishes a daily *Markets* column that contains a summary and discussion of events that took place mainly in the previous trading session.⁵ The *Markets* column covers mostly equities, but other asset classes like fixed income, currencies, and commodities are also discussed. There are brief references to economic news and international markets, along with commentary from market professionals, such as portfolio managers or equity strategists. The *Markets* column is written by many different journalists. Usually, one journalist is responsible for writing the column, before some colleague takes over. In our sample period, there are 152 different WSJ authors. The WSJ articles are manually extracted from the WSJ website and stored into individual text files before being processed. Fig. 5 is an example of a WSJ article.

2.1.3. Descriptive statistics

Panel A of Table 1 presents the descriptive statistics of our raw media data sample. The sample contains 3776 WSJ articles. Since the *Markets* column appears daily, this number corresponds to the number of trading days during our 15-year sample period. The total number of SA articles is 85,116, which is considerably larger, as there is no restriction of the number of articles appearing in the Today's Market section of SA. On average, there are around 23 SA articles a day, which are rather equally distributed over the trading days in our sample. Finally, the raw data sample contains around 1.57 million individual user comments written in response to the SA articles. This means that there are on average around 18 comments below each of the SA articles.

The average *Markets* column of the WSJ contains 656 words. The fairly low standard deviation shows that the length of the WSJ articles does not vary a lot, presumably because of the column's structure, as well as to ensure some homogeneity over time. Because no length constraints apply to SA contributors, the average length of SA articles (988 words) is considerably longer compared to the

³ This information is provided on the SA website.

⁴ Only articles that are exclusive to Seeking Alpha earn some compensation. The article compensation is based on criteria such as the traffic it generates, the topic it covers, and its quality.

⁵ The column has had different names, such as Main Markets, Today's Markets or U.S. Markets.

Table 1

Descriptive statistics of the media data.

I I I I I I I I I I I I I I I I I I I			
	WSJ articles	SA articles	SA comments
Panel A: Raw data sample			
Total number articles/comments	3,776	85,116	1,566,690
Mean number of words per article/comment	656	988	69
St. dev. of number of words per article/comment	168	1,004	94
Panel B: Retaining only SA comments written within 24	h of the corresponding SA	articles	
Total number of comments	-	-	1,055,615
Mean number of words per comment	-	-	66
St. dev. of number of words per comment	-	-	87
Panel C: Retaining only SA articles/comments containing	g the keyword "market(s)"	– final sample	
Total number of articles/comments	3,776	76,135	148,198
Mean number of words per article/comment	656	574	60
St. dev. of number of words per article/comment	168	502	62

The table reports the summary statistics of the *Wall Street Journal* (WSJ) and Seeking Alpha (SA) articles, as well as the comments written below the SA articles. Panel A reports the raw data sample, before any data cleaning. Panel B reports the summary statistics of the SA comments after retaining only those that were written within 24 h after the publication of the corresponding SA article. Panel C reports the summary statistics of the SA articles and comments after retaining only those SA articles and comments that contain the keyword "market(s)." This is the final sample. The sample period is from 2006 to 2020 and consists of 3776 trading days.

WSJ articles. The high standard deviation of the length of SA articles likely reflects the different writing styles of the individual SA contributors. With an average of 69 words, SA comments are much shorter compared to SA and WSJ articles. This is to be expected given the nature of comments, which tend to simply express agreement or disagreement with the main ideas expressed in the article.

2.2. Stock market data

We use the S&P 500 market composite index to capture the general trend of the U.S. stock market.⁶ The S&P 500 is based on a larger number of U.S. stocks than other market indices like the Dow Jones Industrial Average (DJIA) used by Garcia (2013). The conjecture is this index is more suitable to capture the general market movements.

We retrieve the daily historical data of the S&P 500 (at market close) from Bloomberg. In total, there are 3776 trading days from 2006 to 2020. The average S&P 500 log return for this period is 0.029, with a median of 0.072 and a standard deviation of 1.279.

3. Investor sentiment

In their seminal papers, Baker and Wurgler (2006, 2007) stress the importance of market or investor sentiment in explaining stock market behavior. Following (Baker and Wurgler, 2007), we define sentiment as the investors' beliefs about future market conditions that cannot be justified by facts. As a consequence, investment decisions, and thus market prices, are to some extent irrationally influenced by such investor mood. Furthermore, Shiller (2016) suggests that the media plays a vital part in capturing and propagating such market sentiment. In this section, we explain the methodology of transforming the raw media data into various indices aimed at capturing market or investor sentiment.

3.1. Methodology

3.1.1. Identification strategy

Articles appearing on social media and traditional print media often convey both rational information and irrational investor mood. To isolate the effect of investor sentiment on stock returns from the effect of financial market news, it is necessary to make some adjustments depending on the publication time of the articles (Garcia, 2013).

Following the efficient market hypothesis (Fama, 1970), new information should be immediately impounded in prices. Hence, for investor sentiment to be independent of news effects, it has to be measured after new information is released. The idea is that investors who read a summary of market movements along with potential explanations are influenced by its sentiment for their subsequent investment decisions — and not by the news as such. Put differently, our identification strategy is that news affects the market immediately, while sentiment affects the stock market in the subsequent trading days. Shiller (2016) suggests that such "tag-along news" can influence trading behavior even days after the actual event.

The daily articles appearing in the *Markets* column of the WSJ are usually published on the WSJ website within a few hours after the market close (normally at 4:00 pm ET).⁷ By contrast, as a pure online medium, SA articles as well as comments can be

⁶ All main results are qualitatively and quantitatively similar when measuring market returns using the CRSP index.

⁷ More recently, the WSJ introduced also a live coverage of the market. This feature was not available during the time period we analyze (2006–2020).



Fig. 2. Timeline of articles published on the weekend.

published at any time, seven days a week. While most publication activity is during and just after U.S. trading hours, SA authors and readers living all over the world mean a fairly continuous publication of articles and comments.

To separate news from sentiment effects, we use a timing convention similar to Garcia (2013). Articles published in the 24-h time window after market closing are assumed to be related to that trading day. For example, an article published after 4:01 pm ET on Monday (day t), but before Tuesday 4:00 pm ET, is considered a "Monday publication", since it essentially refers to market movements on Monday. The sentiment reflected in these articles is then expected to predict stock returns on the next day's trading session, i.e., on Tuesday (t + 1). Fig. 1 illustrates how we treat articles published on weekdays. Similarly, the content of articles published between Friday 4:01 pm ET and Monday 4:00 pm ET refers to the Friday trading session, and are hence considered a "Friday publication". Hence, their sentiment is expected to predict market prices until the market closure on Monday (see Fig. 2).

This timing convention creates a short overlap for SA articles published during regular U.S. trading hours. Since these articles are published on the same day as stock returns are measured, the clear separation between news and sentiment effects might be violated. As discussed in Section 2.1.1, it is important to remember that SA articles do not refer to current news but provide analyses about events that happened on past trading days, or even in the future. Thus, articles appearing during trading hours should convey investor sentiment only. The two sample articles reproduced in Figs. 3 and 5, for example, are written on the same day after the market close. While the WSJ article reports a Fed announcement earlier that day, the SA article provides some investment ideas for the month to come, not mentioning the Fed announcement at all. Nevertheless, in Section 6.1, we further examine the potential effects of news contained in SA articles appearing during trading hours.

The timing adjustment of reader comments in SA is similar to those of the articles. Yet, we only use comments published within the first 24 h after the publication of the SA article they respond to, similar to Chen et al. (2014). The idea is that comments posted after the 24-h window are generally little related to the original article. Panel B of Table 1 shows that approximately 68% of reader comments appear within the first 24 h after the publication time of the article they respond to, implying that SA users mainly comment about recently published articles.

3.1.2. Data cleaning

Since our objective is to analyze the relation between market sentiment and broad stock market returns, we only retain SA articles and comments that contain the keywords "market" or "markets" in the sample, similar to the methodology used in Tetlock et al. (2008). Since the *Markets* column of the WSJ contains only articles that cover general market topics, this step is omitted for the WSJ sample.

The final step in data cleaning involves removing hyperlinks and other HTML tags, punctuation, stop words, and finally the top-ten of unique words, which represent spelling errors, usernames, etc. This is a standard procedure when working with textual data, which allows us to remove noise and is applied to both SA and WSJ data.

'My View Of The S&P 500 Index - November 2017'

November 1, 2017, 10:16PM ET

Author: Walter Zelzniak Jr

Another month and another gain in the stock market. This bull market keeps on going higher and higher. The market, as measured by the S&P 500 index, was up again in October and closed 2.22% higher. The SPDR S&P 500 ETF (SPY) gained 2.36% for the month. As for my pension plan assets, I too gained for October. Consequently, my first investment goal, preservation of capital, was achieved. Unfortunately, I did not beat the returns of the S&P 500 index as measured by SPY. So only half of my investment objectives were met for the month of October. Table 1 below shows my returns for the month and Table 2 below shows my returns for the past 12 months. To review the purpose of this series of articles, my retirement account only allows me to buy the following four ETFs: iShares Core U.S. Aggregate Bond ETF (AGG), SPDR S&P 500 ETF, iShares Russell 2000 ETF (IWM), and iShares MSCI EAFE ETF (EFA). I can also have my money in cash. The question is how to decide where and when to allocate money to these various ETFs. I use my moving average crossover system combined with relative strength charts to determine how to allocate my pension plan assets. My moving average crossover system uses the 6-month and the 10-month exponential moving averages to identify which of the four ETFs are in a position to be bought. If the 6-month moving average is above the 10-month moving average, then the ETF is a buy. I call this setup being in bullish alignment. When the 6-month moving average is below the 10-month moving average, the setup is referred to as a bearish alignment. When a bearish alignment happens I don't want to hold that asset. See Chart 1 below for a long-term look at the S&P 500 index using my moving average crossover system. You can see that the moving average crossover system provided some excellent long-term buy and sell signals that would have allowed investors to capture long duration moves in the index, while avoiding costly drawdowns. Avoiding these costly drawdowns allows me to meet the objective of capital preservation. To me, the last place you want to experience a large drawdown is in your pension plan. During the 2008-2009 market crash, many people didn't even look at their retirement statements because they were afraid of what they would find. I submit that if those people would have used a market strategy similar to what I outline in this series of articles, they would have been able to avoid much of the decline during the bear market and consequently would have had less emotional stress during that time period. I find this investment strategy to be particularly useful for managing the assets in my pension plan. If your pension plan is anything like mine, then you just have a choice of ETFs or mutual funds to select from. Looking at standard fundamental criteria such as P/E ratios, free cash flow, dividend yield, and the like is not easily applied to stock indices and mutual funds. What is easily applied to stock indices and mutual funds is trend following technical analysis as I show in this article. The following charts show the current status of the ETFs that I am allowed to buy in my retirement account. As mentioned in the opening paragraph, SPY was up for the month of October gaining 2.36%. Looking at Chart 2, you can see the strong uptrend that is in place since the last buy signal was given. The trend has been easy to follow since the buy signal occurred when the 6-month moving average crossed above the 10-month moving average and that is the idea. Trend following is meant to be easy and profitable. Chart 3 shows that IWM gained 0.73% for the month and closed at a new high. In last month's article, I said I would allocate some money to IWM and that is what I did. Since the buy signal was generated in August of 2016 IWM has gained over 24%. Chart 4 shows the relative strength of IWM compared to SPY. When the signal line is trending higher it means IWM is outperforming SPY. When the signal line is trending lower it means the opposite; IWM is underperforming SPY. The downward sloping trend line from early 2014 to early 2016 shows that investors would have made more money investing in SPY versus investing in IWM. For October, IWM underperformed SPY by 1.59%, giving back some of the outperformance IWM had in the month of September. If the ratio can take out the highs of last December, then the downtrend in the ratio that was discussed in a previous article will have reversed. The ratio is still above its recent low in June. I will continue to maintain a small position in IWM and monitor this ratio. EFA had another good month and closed up 1.68%. EFA closed at a new high again in October and EFA remains in bullish alignment. The trend is growing stronger as the distance or the whitespace between the 6-month moving average and the 10-month moving average is getting wider. Chart 6 shows that EFA, like IWM mentioned previously, was outperformed by SPY. Despite the underperformance, the ratio remains in bullish alignment. The relative strength ratio in Chart 6 is one to watch closely. Since 2013, US stocks have outperformed international stocks until early 2017. Readers can see this as the signal line falls in Chart 6 from the upper left corner to the lower right corner. However, I see that a change is taking place as the signal line is now starting to trend higher since early 2017, and the 6-and 10-month moving averages have gone into bullish alignment. As I mentioned last month, the EFA:IWM ratio could be rolling over. EFA did outperform IWM so the ratio closed higher. Two months ago the 6- and 10-month moving averages moved into a bullish alignment and they remain in bullish alignment, but both moving averages are sloping downward. I will continue to monitor this ratio. AGG had a minimal gain and remains in bullish alignment. Bonds investors are still making money but investors in bonds are lagging the returns of investors in equities as can be seen in Chart 9. AGG continues to be outperformed by SPY and the ratio closed at a new low in October. Investors simply prefer stocks over bonds at this time. This simple ratio confirms that equities continue to be the place to put your money. None of my retirement assets are allocated to AGG. For the month of October, I was allocated 50% EFA, 25% IWM and 25% SPY, and I will maintain that allocation for the seasonally strong month of November. The ETFs for EFA, IWM, and SPY are in bullish alignment and the current trends are strong as depicted in Charts 5, 3, and 2, respectively. I have used the relative strength charts to allot my assets and I favor the indication in Chart 6, so I have 50% of my allotment to EFA as I expect it to outperform SPY in the near future. All in all, I think that my moving average crossover system has me properly aligned and I will continue to monitor the charts.

Fig. 3. Sample article of the Today's Market section on Seeking Alpha, appearing on the November 1, 2017 (a Fed announcement date). Thirty-two positive words (green, bold, and underlined), eighteen negative words (red, bold), 778 total words. Positive sentiment: 0.041, negative sentiment: 0.023, pessimism index: -0.017.

1. Commenter: mrjustice

November 2, 2017, 10:06AM ET

'I like the approach and use one very similar to it in managing my stable of individual stocks and sector ETF's. SOXL, for example, provided me a buy signal in Aug 2016 at \$30 and has remained in <u>strong</u> bullish alignment, without a single **threat** of any whipsaw, right up thru \$150. I like **forcing** the stock to do the work required to stay in the game. If and when it decides to give up and quits trying, off the field it goes...'

2. Commenter: blanco bill

November 3, 2017, 0627AM ET

'Excellent strategy simple and effective'

Fig. 4. Comments written in response to the Today's Market article on Seeking Alpha, see Fig. 3. three positive words (green, bold, and underlined), two negative words (red, bold), 88 total words (the two comments together). Positive sentiment: 0.034, negative sentiment: 0.022, pessimism index: -0.011.

Panel C of Table 1 presents the final data set after all pre-processing and cleaning steps. There are 3776 WSJ articles with 656 words per article on average, and about 76,100 SA articles with an average of 574 words. The final number of comments is around 146,200, with an average length of 60 words.

3.2. Sentiment indices

3.2.1. Index construction

Similar to Garcia (2013), we use the (Loughran and McDonald, 2011) word dictionary to construct the sentiment indices. This lexicon contains lists of 2355 negative and 354 positive words that have been identified as appropriate for assessing the tone or sentiment of finance-related texts. The examples in Figs. 3, 4, and 5 highlight negative words in bold, and positive words in bold and underlined.

The sentiment indices obtained from the WSJ are based on the daily *Markets* column, consisting of one article a day. By contrast, the daily SA sentiment index aggregates all articles appearing in a 24-h window. Positive and negative sentiment indices are defined as the fraction of positive (negative) words out of total words. For every article and comment *i* on day *t*, we count the number of positive g_{it} , negative b_{it} , and total words, w_{it} . Then we aggregate these numbers to construct the daily positive sentiment index, $G_t = \sum_i g_{it} / \sum_i w_{it}$, and the daily negative sentiment index, $B_t = \sum_i b_{it} / \sum_i w_{it}$. This means that the daily SA sentiment index is a weighted average of the sentiment scores of all SA articles on a given day, where the weight depends on the number of words in each article (longer SA articles have a greater weight in the daily SA sentiment index).⁸ Using the positive and negative sentiment indices, G_t and B_t , we also calculate a pessimism index as the difference between negative and positive sentiment indices, i.e., $P_t = B_t - G_t$.⁹

3.2.2. Summary statistics of sentiment indices

Table 2 presents the descriptive statistics of the sentiment indices extracted from the SA and WSJ data. In all media sources, there are on average about twice as many negative words than positive words, reflecting the uneven distribution of negative relative to positive words in the (Loughran and McDonald, 2011) dictionary. The WSJ has the lowest fraction of positive and negative words, with averages of 1.3% and 2.4%, respectively. SA articles and comments have on average considerably larger fractions of dictionary words. About 1.8% of the words in SA articles are positive words, and around 3.6% are negative words. The sentiment indices derived from the SA comments reveal an even slightly higher fraction of dictionary words.

Table 3 presents the correlation statistics between the various sentiment indices. Most of the pairwise correlations are substantial, and highly significant. This suggests that there is a co-movement between the sentiment obtained from the different media sources, which is especially strong for the pessimism indices. Only the correlation between positive sentiment indices of WSJ articles and SA comments fails to be statistically significant, suggesting that these two indices capture different facets of positive investor sentiment.

⁸ The forecast combination literature (Timmermann, 2006; Elliott and Timmermann, 2016; Wang et al., 2023) suggests that an equally weighted aggregation results in even better forecasts. In untabulated robustness checks, we thus calculated an equally weighted SA sentiment index, where $G_t = \frac{1}{N} \sum_i (g_{it}/w_{it})$ and $B_t = \frac{1}{N} \sum_i (b_{it}/w_{it})$. In this approach, all articles have the same weight, regardless of their length. The results are very similar.

⁹ Negations before dictionary words can reverse their meaning, and hence tone. However, when taking into account words that express negation, the sentiment indices do not change significantly.

'Stocks Bounce Higher to Start November'

By Amrith Ramkumar and Riva Gold

November 1, 2017, 05:15PM ET

The Dow Jones Industrial Average and S&P 500 climbed to start November, after major indexes posted a flurry of records last month. Investors and analysts have said they expect earnings growth around the world to keep supporting major indexes, which have continued to hover near all-time highs in recent sessions. "You're getting a pretty good fundamental backdrop on the earnings front-that's the primary focus," said Nathan Thooft, senior managing director of global asset allocation at Manulife Asset Management. The Dow industrials rose 57.77 points, or 0.2%, to 23435.01 after the blue-chip index posted its seventh straight monthly gain Tuesday, its longest streak since April 2012. The S&P 500 climbed 4.10 points, or 0.2%, to 2579.36, and the techheavy Nasdag Composite swung between small gains and losses and closed down 11.14 points, or 0.2%, at 6716.53. With nearly 70% of S&P 500 companies having reported as of the market close Wednesday, roughly three-guarters of those firms have exceeded earnings expectations, compared with the five-year average of 69%, according to FactSet. Estée Lauder was among the biggest gainers Wednesday. Shares climbed \$10.31, or 9.2%, to \$122.12 after the beauty company reported an increase in organic sales and raised its forecast for fiscal year 2018. Buoyant oil prices supported shares of energy firms, with Devon Energy and Marathon Oil also among the best performers in the S&P 500. The index's energy sector rose 1.1%, with U.S. crude oil prices near their highest level of the year amid signs that production cuts by major suppliers are helping to rebalance the market. Mr. Thooft said his firm is more positive on the energy sector now than it has been for much of the year, given the possibility of higher oil prices in the next six months. Stocks inched slightly higher after the Fed left rates unchanged as expected but signaled it would consider lifting them before year's end amid signs the economy was gaining momentum. The dollar was also slightly higher after the Fed's statement, which investors and analysts said offered few surprises. "They've really laid the groundwork to raise rates in December," said Brad McMillan, chief investment officer at Commonwealth Financial Network. "This is a Fed that's fairly confident about the economy," he said. The WSJ Dollar Index, which tracks the U.S. currency against a basket of 16 others, was up 0.2% Wednesday after its best month since November 2016. Some investors and analysts have said formal details on the next chair of the Fed or on potential changes to taxes will also be critical for understanding the future path of U.S. monetary policy. The Wall Street Journal reported after the market closed that the White House has notified Fed governor Jerome Powell that President Donald Trump intends to nominate him for the role. Elsewhere, the Stoxx Europe 600 advanced 0.4%. Clothing retailer Next fell 9.1% after the firm's earnings, while shares of Standard Chartered declined 6% after it reported higher costs and lower revenue than expected. Gains in technology shares boosted major Asian bourses. Japan's Nikkei Stock Average closed up 1.9% in its biggest daily point gain since May, while Hong Kong's Hang Seng Index rose 1.2%.

Fig. 5. Sample article of the Markets column in the Wall Street Journal, appearing on the November 1, 2017 (a Fed announcement date). twelve positive words (green, bold, and underlined), seven negative words (red, bold), 525 total words. Positive sentiment: 0.022, negative sentiment: 0.013, pessimism index: -0.009.

4. Investor sentiment and stock returns

In this section, we revisit the hypothesis by Shiller (2016) that investor sentiment as reflected in the media can influence stock market returns, using the measures of investor sentiment presented in the previous section. Similar to Garcia (2013), we employ the following model:

$$R_t = \beta_{SA} \mathcal{L}_s(SA_t) + \beta_C \mathcal{L}_s(C_t) + \beta_W \mathcal{L}_s(W_t) + \gamma \mathcal{L}_s(R_t) + \psi \mathcal{L}_s(R_t^2) + \eta X_t + u_t,$$
(1)

where R_t denotes the log return of the S&P 500 on day t. The set of explanatory variables consists of the sentiment indices, lagged returns, and some additional exogenous variables. SA_t denotes the sentiment index extracted from the SA articles, C_t denotes the sentiment index based on the SA comments, and W_t denotes the sentiment index extracted from the WSJ articles. Depending on the analysis, these can represent either the positive sentiment index G_t of each media source, the negative sentiment index B_t , or the pessimism index P_t . Prior to the regression analysis, we normalize the sentiment indices to have zero mean and unit variance. This allows for an easier interpretation of the coefficient estimates in terms of one standard deviation changes to investor sentiment.

The \mathcal{L}_s denotes an *s*-lag operator, such that, for example, $\mathcal{L}_s(W_t) = \{W_{t-1}, \dots, W_{t-s}\}$. We set s = 5, i.e., Eq. (1) includes up to five lags. This allows testing, for example, whether investor sentiment as reflected in the media on a Monday can predict stock

I. Lachana and D. Schröder

Table 2

Descriptive statistics of sentiment indices.

Sentiment measure	Mean	Median	25%-quant.	75%-quant.	Std. dev.
Panel A: WSJ articles					
Positive sentiment	1.261	1.196	0.821	1.638	0.610
Negative sentiment	2.386	2.321	1.754	2.958	0.873
Pessimism index	1.126	1.094	0.282	1.945	1.242
Panel B: SA articles					
Positive sentiment	1.814	1.808	1.632	1.993	0.302
Negative sentiment	3.643	3.624	3.240	4.074	0.652
Pessimism index	1.829	1.805	1.347	2.305	0.730
Panel C: SA comments					
Positive sentiment	2.051	2.028	1.715	2.366	0.739
Negative sentiment	4.073	4.012	3.491	4.598	1.154
Pessimism index	2 022	1 972	1 331	2 654	1 361

This table reports sample statistics of sentiment measures obtained from the *Wall Street Journal* articles, the Seeking Alpha articles, and the Seeking Alpha comments. The positive sentiment measure (G_i) is the fraction of positive words out of total words, using the (Loughran and McDonald, 2011) dictionary. The negative sentiment measure (B_i) is the fraction of negative words out of total words. The pessimism index is the difference between the negative and the positive sentiment measures ($P_i = B_i - G_i$). Panel A presents the summary statistics for the WSJ, Panel B and C present the summary statistics for the SA articles and comments, respectively.

Table 3

Correlation statistics of sentiment indices.

Panel A: Negative sentiment indices			
	WSJ articles	SA articles	SA comments
WSJ articles	1.00	0.24***	0.13***
SA articles	0.24***	1.00	0.26***
SA comments	0.11****	0.21***	1.00
Panel B: Positive sentiment indices			
	WSJ articles	SA articles	SA comments
WSJ articles	1.00	0.13***	0.02
SA articles	0.12***	1.00	0.14***
SA Comments	0.01	0.10***	1.00
Panel C: Pessimism indices			
	WSJ articles	SA articles	SA comments
WSJ articles	1.00	0.25***	0.10***
SA articles	0.25***	1.00	0.22***
SA comments	0.09***	0.20***	1.00

This table reports the correlation coefficients of the three sentiment indices obtained from the *Wall Street Journal* (WSJ) articles, Seeking Alpha (SA) articles, and SA comments written in response to the SA articles. Panel A shows the correlation coefficients of the positive sentiment indices (G_i), and B the correlation coefficients of the negative sentiment indices (G_i), and panel C the correlation coefficients of the parel shows the linear correlation coefficients ($P_i = B_i - G_i$). The sample period is from 2006 to 2020. The lower triangle of each panel shows the linear correlation coefficients (Pearson's r). The upper triangle of each panel shows the rank correlation coefficients (Spearman's ρ). ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

market returns up to the following Monday (five trading days). To ensure that the sentiment effects are not driven by short-term autocorrelation (Lo and MacKinlay, 1990) or reversals in stock returns (Cox and Peterson, 1994), the model also includes lagged stock market returns, $\mathcal{L}_s(R_t)$, and lagged squared returns, $\mathcal{L}_s(R_t^2)$, as control variables. The set of additional exogenous variables X_t consists of days-of-the-week dummies to control for day-effects in stock returns, as documented by Conrad et al. (1997), and a constant. Finally, u_t denotes the error term, with standard errors adjusted for heteroskedasticity following (White, 1980).

4.1. Social media versus traditional media

We next analyze and compare the relation between investor sentiment and subsequent stock returns using the sentiment indices obtained from SA and the WSJ.

First, we analyze the relation between investor sentiment and stock market returns for the two sentiment measures separately. Panel A of Table 4 presents the coefficient estimates $\hat{\beta}_{SA}$ when explaining stock market returns by the investor sentiment obtained from SA articles only.

Panel A shows that the sentiment indices obtained from the SA articles help to predict stock returns. For example, the higher the fraction of positive words in SA articles, the higher subsequent stock returns. That is, a positive investor sentiment as reflected in the social media platform translates in higher asset prices. The relation is statistically different from zero at the 1% significance

Table 4

Sentiment indices and stock market returns - WSJ and SA articles.

	Positive sentime	ent	Negative sentime	ent	Pessimism index	
	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.
SA_{t-1}	0.052***	2.6	-0.082***	-3.0	-0.092***	-3.6
SA_{t-2}	-0.006	-0.3	0.037	1.3	0.036	1.4
SA_{t-3}	-0.021	-1.1	-0.007	-0.3	0.010	0.4
SA_{t-4}	-0.012	-0.6	0.033	1.2	0.037	1.5
SA_{t-5}	0.037*	1.9	0.023	0.8	-0.003	-0.1
Panel B: WSJ	articles					
	Positive sentime	ent	Negative sentime	ent	Pessimism index	
	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.
W_{t-1}	0.028	1.2	-0.040	-1.6	-0.047*	-1.7
W_{t-2}	-0.008	-0.3	-0.010	-0.3	-0.003	-0.1
W_{t-3}	-0.022	-1.0	0.008	0.4	0.019	0.8
W_{t-4}	0.036	1.5	0.017	0.7	-0.007	-0.2
W_{t-5}	-0.030	-1.2	0.012	0.5	0.026	0.9
Panel C: SA a	and WSJ articles					
	Positive sentim	ent	Negative sentime	ent	Pessimism index	
	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.
SA_{t-1}	0.051***	2.5	-0.078***	-2.9	-0.088***	-3.6
SA_{t-2}	-0.004	-0.2	0.040	1.4	0.037	1.4
SA_{t-3}	-0.021	-1.1	-0.007	-0.3	0.009	0.4
SA_{t-4}	-0.011	-0.6	0.032	1.2	0.036	1.4
SA_{t-5}	0.038*	1.9	0.022	0.8	-0.005	-0.2
W_{t-1}	0.021	0.9	-0.037	-1.5	-0.039	-1.4
W_{t-2}	-0.011	-0.4	-0.012	-0.4	-0.002	-0.0
W_{t-3}	-0.022	-1.0	0.006	0.2	0.016	0.6
W_{t-4}	0.035	1.4	0.015	0.6	-0.009	-0.3
W_{t-5}	-0.034	-1.4	0.010	0.4	0.026	0.9

This table reports the estimated coefficients $\hat{\beta}_{SA}$ and $\hat{\beta}_{W}$ of following:

 $R_t = \beta_{SA} \mathcal{L}_s(SA_t) + \beta_W \mathcal{L}_s(W_t) + \gamma \mathcal{L}_s(R_t) + \psi \mathcal{L}_s(R_t^2) + \eta X_t + u_t.$

The dependent variable R_i is the daily log return of the S&P 500 from 2006 to 2020. The variables SA_i and W_i denote the media sentiment indices obtained from the Seeking Alpha (SA) and *Wall Street Journal* (WSJ) articles, respectively. We consider the positive sentiment index (G_i), the negative sentiment index (B_i), and the pessimism index ($P_i = B_i - G_i$). \mathcal{L}_s is the lag operator with s = 1...5. The sentiment indices are normalized to have zero mean and unit variance. X_i represents the set of exogenous variables that consists of day-of-the-week dummies and a constant, and u_i is the error term. The *i*-statistics are computed using (White, 1980) standard errors. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

level, and means that a one standard deviation shock to positive investor sentiment moves average stock returns by 5.2 bps. The results are even stronger for negative sentiment and the pessimism index.

While the results show that investor sentiment is moving stock markets, the effect is short-lived as the coefficient estimates are generally significant for one trading day only. The estimated coefficients of lagged sentiment indices, however, allow for an assessment of whether the sentiment effect on stock markets is temporary or permanent. In fact, for all three sentiment indices, the sum of the coefficient estimates for lags 2 to 5 ($\sum_{i=2}^{5} \hat{\beta}_{SA}$) has the opposite sign than the coefficient estimates of the first lag. For negative sentiment, for example, the coefficients of lags 2 to 5 add up to 0.086, which compares to -0.082 of the first lag. An *F*-test shows that this sum is significantly different from zero (*p*-value < 0.05), while the sum of all 5 coefficients is not significantly different from 0. The pattern is similar for the other indices, albeit at lower significance levels.

Thus, consistent with Tetlock (2007) and Garcia (2013), these findings show that the effect of investor mood is not only shortlived, but also (partly) reversed in the following trading days. Taken together, this means that investor sentiment has only a transitory effect on stock markets, and can thus be considered driven by irrational mood rather than rational (fundamental) information.

Next, we repeat the analysis using the sentiment indices extracted from the WSJ articles instead. The results show that these sentiment indices are considerably less related to subsequent stock returns. The coefficient estimates of the three sentiment indices are about half the magnitude than those in Panel A. As a result, they lose much of their statistical significance. Only the pessimism index is marginally related to subsequent stock returns, at a confidence level of 10%. Similar to Panel A, the sentiment indices lagged by more than one trading day are not significantly related to stock returns, but nevertheless suggest some reversal of the effect on subsequent trading days. Taken together, the results indicate a weak relationship between traditional media sentiment and market returns.

These results provide support for (Tetlock, 2007) and Garcia (2013) who document a relation between positive, negative, and pessimism sentiment indices obtained from traditional print media sources, and broad stock market returns. Yet, while in Tetlock (2007) and Garcia (2013) these association are highly statistically significant, the results in Panel B are weaker and sometimes fail to be significant.

This difference can have a variety of reasons. Most of all, we analyze a more recent time period, covering the years from 2006 to 2020. By contrast, Tetlock (2007) considers the years from 1984 to 1999, and Garcia (2013) even extends back to the year 1905. Given the tremendous changes in the media ecosystems, the results might not be directly comparable. In addition, we rely on the S&P 500 to proxy for stock market returns which covers a broader set of companies, while (Tetlock, 2007) and Garcia (2013) use the DJIA that includes a smaller set of companies. Finally, Garcia (2013) uses media data of the *New York Times*, while the traditional print media data we use are retrieved from the *Wall Street Journal*.

Appendix A presents some further analyses that explore these differences using the original data set provided by Garcia (2013). The results show that the relation between sentiment measures obtained from traditional media and the stock markets decays through time. In the years after 1990, there is only little relation between sentiment indices based on traditional media and the financial markets. This means that the weak results we document in Panel B are effectively consistent with (Tetlock, 2007) and Garcia (2013).

In a last test, Panel C of Table 4 reports the results using both sentiment indices from SA and WSJ articles to explain stock market returns. Combining the two sentiment measures in a joint regression does not change our previous results. The relation between the SA sentiment indices is slightly lower, while remaining strongly and significantly related to stock returns. Similarly, the no-result of the WSJ sentiment index persists; the negative sentiment index is now no longer related to the returns of the S&P 500 index. This result also suggests that, despite a significant and substantial correlation between both measures (see Table 3), they capture noticeably different dimensions of investor sentiment.

Overall, these results support the (Shiller, 2016) hypothesis that investor sentiment as reflected in the media is important in explaining short-term stock market behavior. Yet, during the period from 2006 to 2020, it seems that sentiment indices obtained from traditional print media lost their predictive power for stock market returns. By contrast, investor sentiment as measured from social media platforms seems a strong predictor for stock market returns.

4.2. The role of social media comments

(Chen et al., 2014) show that sentiment measures based on reader comments posted in response to SA articles have significant relation to firm-specific stock returns. In this section, we investigate whether a similar pattern holds true for the entire stock market. More specifically, we analyze the relation between sentiment indices obtained from SA comments and broad stock market returns.

We first fit Eq. (1) using the sentiment index extracted from SA comments as the only measure of investor mood. The results are in Panel A of Table 5. The fraction of positive words in SA comments helps to predict subsequent stock returns. The relation is statistically significant, and different from zero at the 10% level. The magnitude of the effect is similar to the results obtained from traditional media: a one standard deviation change in positive investor sentiment, for example, results in a return on the S&P 500 of 2.9 bps, on average. Yet, in comparison with the positive sentiment of SA articles, the relation is much smaller. A similar pattern applies to the negative sentiment index obtained from the SA comments. Negative sentiment results in lower stock returns, as expected. This relation is statistically significant at the 5% level, and again of similar size compared to the effect of negative investor sentiment extracted from traditional media. Again, the effect of negative sentiment of SA comments is considerably lower than the effect of SA articles. When looking at the pessimism index, the relation between investor sentiment of the S&P 500 of 5.4 bps, on average. This relation is highly significant at the 1% level. Similar to the results obtained from SA and WSJ articles, the effect of SA comments ceases after one day, again suggesting that investor sentiment is quickly incorporated in the stock market without any long-term effects. Rather, there is some reversal of the effect of sentiment extracted from SA reader comments on subsequent trading days, especially for the negative sentiment index.

The remainder of Table 5 presents the results when including the investor sentiment based on SA articles (Panel B), and the investor sentiment based on WSJ articles (Panel C). When adding investor sentiment obtained from SA articles as control variables, the regression coefficient of positive sentiment (SA comments) loses some of its magnitude, such that it is no longer statistically different from zero. Yet, the estimates for negative sentiment and the pessimism index remain substantial and significant. As before, the effect of investor sentiment derived from SA articles is highly significant. Adding investor sentiment obtained from WSJ articles does not change the results, similar to the analysis presented in Section 4.1. Again, investor sentiment as captured by this traditional media source is not significantly related to stock returns.

Overall, this section highlights that not only sentiment indices based on SA articles have a strong relation to subsequent stock returns, but user comments written in response to these articles also exhibit important and useful investor sentiment. These results strengthen the perception that investor mood reflected on social media platforms is useful when trying to understand market behavior, going beyond standard traditional print media.

4.3. Robustness

In this subsection, we discuss some robustness checks of the main results. A possible concern with Eq. (1) is that it does not explicitly allow for a time-varying volatility of stock market returns. Hence, the results might be partially driven by periods of higher-than-average stock market volatility. To rule out this possibility, we re-estimate Eq. (1) using normalized stock returns. The normalization divides stock returns by daily estimates of stock market volatility, i.e., $R_t/\hat{\sigma}_t$. We estimate market volatility using a GARCH(1,1) model. This adjustment effectively creates a time series of stock returns with constant volatility.

The results in Table 6 show that the overall pattern does not change after correcting stock returns for time-varying stock market volatility. While the coefficient estimates are slightly smaller than in Tables 4 and 5, the general conclusions remain unchanged. In untabulated tests, we also examine whether there is a difference when including sentiment indices from both SA articles and comments in the regression model, and we also add the sentiment indices extracted from the WSJ, as control variable. Again, the main results do not change.

Table 5

Sentiment indices and stock market returns - SA comments.

Panel A: SA	comments					
	Positive sentime	ent	Negative sentime	ent	Pessimism index	
	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.
C_{t-1}	0.029*	1.7	-0.044**	-2.2	-0.054***	-2.8
C_{t-2}	0.018	1.1	0.001	0.1	-0.009	-0.4
C_{t-3}	0.005	0.3	-0.009	-0.4	-0.011	-0.5
C_{t-4}	-0.004	-0.2	0.011	0.5	0.014	0.7
C_{t-5}	-0.008	-0.5	0.015	0.7	0.019	0.9

Panel B: SA comments and articles

	Positive sentiment		Negative sentiment		Pessimism index	
	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.
C_{t-1}	0.024	1.4	-0.042**	-2.1	-0.051***	-2.8
$C_{t=2}$	0.018	1.1	0.000	0.0	-0.007	-0.3
$C_{t=3}$	0.004	0.2	-0.010	-0.5	-0.010	-0.5
C_{t-4}	-0.005	-0.3	0.010	0.5	0.013	0.6
C_{t-5}	-0.012	-0.7	0.011	0.6	0.018	0.9
SA_{t-1}	0.072***	3.0	-0.086***	-2.8	-0.105***	-3.6
SA_{t-2}	-0.023	-0.9	0.039	1.2	0.045	1.5
SA_{t-3}	-0.025	-1.1	0.004	0.1	0.019	0.7
SA_{t-4}	-0.013	-0.5	0.021	0.7	0.025	0.9
SA_{t-5}	0.032	1.3	0.035	1.1	0.013	0.4

Panel C: SA comments and articles, and WSJ articles

	Positive sentime	ent	Negative sentime		Pessimism index	
	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.
C_{t-1}	0.023	1.4	-0.042**	-2.1	-0.050***	-2.6
C_{t-2}	0.018	1.1	0.001	0.0	-0.007	-0.3
C_{t-3}	0.005	0.3	-0.010	-0.5	-0.010	-0.5
C_{t-4}	-0.007	-0.4	0.012	0.6	0.014	0.7
C_{t-5}	-0.012	-0.7	0.011	0.5	0.017	0.9
SA_{t-1}	0.071***	3.0	-0.082***	-2.8	-0.102***	-2.6
SA_{t-2}	-0.021	-0.8	0.039	1.2	0.044	1.5
SA_{t-3}	-0.025	-1.1	0.002	0.1	0.017	0.6
SA_{t-4}	-0.012	-0.5	0.018	0.6	0.024	0.8
SA_{t-5}	0.033	1.4	0.033	1.0	0.010	0.3
W_{t-1}	0.022	0.9	-0.039	-1.4	-0.038	-1.3
W_{t-2}	-0.012	-0.4	-0.004	-0.1	0.005	0.1
W_{t-3}	-0.023	-1.0	0.020	0.8	0.028	1.1
W_{t-4}	0.035	1.4	0.015	0.5	-0.009	-0.3
W_{t-5}	-0.030	-1.1	0.011	0.4	0.023	0.7

This table reports the estimated coefficients β_C , β_{SA} , and β_W of the model:

 $R_t = \beta_C \mathcal{L}_s(C_t) + \beta_{SA} \mathcal{L}_s(SA_t) + \beta_W \mathcal{L}_s(W_t) + \gamma \mathcal{L}_s(R_t) + \psi \mathcal{L}_s(R_t^2) + \eta X_t + u_t.$

The dependent variable R_t is the daily log return of the S&P 500 from 2006 to 2020. The variables C_t , SA_t and W_t denote the media sentiment indices obtained from the Seeking Alpha (SA) comments, the SA articles, and *Wall Street Journal* (WSJ) articles, respectively. For more details, see Table 4.

5. The advantage of social media

We established that investor sentiment as captured in social media is a better predictor of stock returns than sentiment extracted from traditional media sources, we explore various potential underlying reasons for this finding.

Although articles appearing in SA and the WSJ both discuss developments of the financial markets, there are some significant differences. Given its nature as a social media platform, the number of articles published on SA vastly exceeds the number of articles published in the WSJ, along with the number of authors. In addition, language, timing, and length of the articles appearing in both media sources differ considerably. In this section, we explore whether these differences can explain why SA is a better alternative than the WSJ to capture investor sentiment.

5.1. Volume of articles

While the daily *Markets* column of the WSJ consists of just one article, there is no limit for the number of articles appearing in the Today's Market section of SA, as discussed before. During our sample period, a maximum of 85 articles appeared in that section within a 24-h window (8/24/2009). As a result, the sentiment indices based on SA articles can draw on a much larger data set compared to the WSJ. Table 1 shows that the SA sample is more than 20 times larger than the WSJ sample. Thus, the better

 $C_{t=5}$

Table 6 Robustness analysis

0.8

Panel A: SA	articles					
	Positive senti	ment	Negative sentin	nent	Pessimism inde	x
	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.
SA_{t-1}	0.033*	1.9	-0.074***	-3.4	-0.075***	-3.7
SA_{t-2}	0.009	0.5	0.026	1.2	0.019	0.9
SA_{t-3}	-0.026	-1.4	-0.009	-0.4	0.010	0.5
SA_{t-4}	-0.002	-0.1	0.029	1.3	0.026	1.2
SA_{t-5}	0.034*	2.0	0.026	1.3	0.002	0.1
Panel B: SA	comments					
	Positive senti	ment	Negative sentin	nent	Pessimism inde	x
	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.
C_{t-1}	0.033*	1.9	-0.031*	-1.9	-0.047***	-2.8
C_{t-2}	0.023	1.5	0.012	0.7	-0.003	-0.2
C_{t-3}	0.005	0.3	-0.024	-1.3	-0.023	-1.3
C_{t-4}	-0.011	-0.6	0.009	0.5	0.016	0.9

This table reports the estimated coefficients β_{SA} and β_C of the model:

-0.009

 $R_t = \beta_{SA} \mathcal{L}_s(SA_t) + \beta_C \mathcal{L}_s(C_t) + \gamma \mathcal{L}_s(R_t) + \psi \mathcal{L}_s(R_t^2) + \eta X_t + u_t.$

-0.5

The dependent variable R_i is the normalized daily log return of the S&P 500. The normalization divides the log returns by the estimated volatility $\hat{\sigma}_i$ using a GARCH(1,1) model:

0.009

0.5

0.014

$$R_t = \mu + \epsilon_t$$

$$\sigma_{t+1}^2 = \omega + \alpha_1 \epsilon_t^2 + \alpha_2 \sigma_t^2,$$

where $\sigma_t^2 \equiv var(\epsilon_t)$. The variables SA_t and C_t denote the media sentiment indices obtained from the Seeking Alpha (SA) articles and SA comments, respectively. We consider the positive sentiment index (G_t) , the negative sentiment index (B_t) , and the pessimism index $(P_t = B_t - G_t)$. \mathcal{L}_s is the lag operator with s = 1...5. The sentiment indices are normalized to have zero mean and unit variance. X_t represents the set of exogenous variables that consists of day-of-the-week dummies and a constant, and u_t is the error term. The *t*-statistics are computed using (White, 1980) standard errors. ***, ***, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

predictive power of investor sentiment indices extracted from social media sources might be explained by the data availability, i.e., the number of articles published.

In fact, a well-known finding of the forecast combination literature (Wang et al., 2023) is that a simple average of many individual forecasts often beats any single forecast. In the context of this study, this means – if interpreting the sentiment scores extracted from each article as forecast indicators – that the SA sentiment index can be considered an average of many individual forecasts. Thus, the superiority of the SA sentiment index to predict the market might not be surprising.

To test this hypothesis, we adopt a two-step test procedure. In the first step, we create 25 sub-samples of SA articles that consists of one randomly drawn SA article per day. This gives 25 SA data sets that have approximately the same number of articles as the entire WSJ sample, each containing 3759 SA articles (see Panel A of Table 7).¹⁰ We then construct the pessimism index for each of these 25 sub-samples, following the methodology described in Section 3.2.1. Then we estimate Eq. (1) separately for each of the sub-samples:

 $R_t = \beta_{SA} \mathcal{L}_s(SA_t) + \gamma \mathcal{L}_s(R_t) + \psi \mathcal{L}_s(R_t^2) + \eta X_t + u_t.$

In the second step, we use a *t*-test to assess whether the average of the 25 coefficient estimates of $\hat{\beta}_{SA}$ of the first-pass regressions is significantly different from zero.¹¹ If the better explanatory power of the SA sentiment indices is only due to the larger data set they can draw on, the advantage should disappear in this setting and result in insignificant coefficient estimates.

Panel B of Table 7 shows that the relation between the first lag of the pessimism index and log stock market returns is still statistically different from zero at the 1% significance level. This means that using just one single SA article per day to construct the pessimism index (instead of the entire data set) does not change the ability of SA sentiment in predicting the market. To conclude, the significantly larger volume of articles published on the SA platform cannot explain the higher predictive ability of investor sentiment of social media over traditional media.

¹⁰ In the first months of 2006, the *Today's Markets* section of SA did not publish an article every day. Hence, the SA sample is slightly smaller than the WSJ sample.

¹¹ For brevity, this section only presents the results for the pessimism index. The results for positive and negative sentiment are qualitatively similar. Including the pessimism index of the WSJ articles as a control variable in the regression does not change the results (see the Online Appendix).

Table 7		
Matching the number of articles.		
Panel A: Sample statistics		
	Random sampli	ıg
Mean number of authors	979	
Mean number of days in sample	3,759	
Mean number of articles	3,759	
Panel B: Pessimism index		
	Random sampli	ıg
	Coeff.	t-stat.
SA _{t-1}	-0.028***	-5.7
SA_{t-2}	-0.000	0.1
SA_{t-3}	-0.008^{*}	-1.8
SA_{t-4}	0.004	1.1
SA_{t-5}	-0.005	-1.2

The table analyzes the relation between stock returns and the pessimism index obtained from Seeking Alpha (SA) articles when using smaller sub-samples of SA articles that correspond to the number of articles in the *Wall Street Journal* (WSJ) sample. We first select 25 random samples of one SA article per day. Then we calculate the average coefficient estimates obtained from the 25 random samples. Panel A shows the average summary statistics of the 25 sub-samples, while Panel B shows the average coefficient estimates. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

5.2. The wisdom of crowds

Another potential explanation for the higher explanatory power of SA sentiment measures is the large pool of authors and contributors on social media platforms. This idea has been coined "wisdom of crowds" by Surowiecki (2005). He suggests that the aggregation of multiple opinions within a group can outperform individuals, including experts, in prediction tasks (Budescu and Chen, 2014). Following this idea, Chen et al. (2014) partly attribute the larger informational content extracted from SA articles and comments for individual stock returns to the large diversity of contributors, which allows aggregating a much larger information set than a few professional equity analysts could ever process. A similar pattern might hold true when it comes to investor sentiment. While there are in total only 152 authors who have written the 3776 articles in the WSJ sample, there are more than 4000 different authors contributing to the articles of the SA sample. Hence, in the this subsection we examine whether the number of authors who write and publish articles could explain the differences in predictive power of SA and WSJ.

To examine if the larger number of SA authors (and thus variety) is a potential explanation for the superior ability of social media sentiment in explaining market returns, we perform two tests. In the first test, we form a subsample of SA articles written by the 152 authors that have contributed most articles to the *Today's Markets* section of SA. Thus, we match the number of authors of SA to those of the WSJ. Then, using this sub-sample of articles, we construct the pessimism index and repeat the analysis as before.

In the second test, we adopt again the two-step random sampling procedure, as explained in Section 5.1. More precisely, we create 25 sub-samples of SA articles written by a random selection of 152 SA authors, thereby matching the number of SA authors to those of the WSJ. Then we test whether the average coefficient estimates of $\hat{\rho}_{SA}$ of the 25 first-pass regressions are significantly different from zero. If the higher explanatory power of SA sentiment indices is due to the larger pool of authors, this advantage (i.e., significance) should disappear in both tests.¹²

Table 8 reports the results. When restricting the SA sample to the 152 authors who have written most articles, the sample size is still substantial, reaching more than 46,000 individual articles (see Panel A). This indicates that most articles appearing on SA are written by a rather small group of contributors. This is a common feature of online platforms. For example, Kittur et al. (2007) show that in the early days of Wikipedia, most articles were written by only a few users. Using this sub-sample to construct the pessimism index, Panel B shows that investor sentiment extracted from SA articles still predicts market returns. The effect is highly statistically different from zero at the 1% level.

The second test based on 25 random samples of 152 authors confirms these findings. A random sample of 152 authors publishes on average a much smaller number of articles, reaching just over 2700 articles. As a result, there are many days without any articles in the Today's Market section. Looking at the "Random sampling" columns, the first lag of the pessimism index is, however, again statistically significant at the 1% level.

 $^{^{12}}$ In case SA authors read each other's articles before writing their own contribution, the investor mood conveyed in their article is not independent of previous SA articles. In this case, SA articles capture the investor mood of all previous SA articles, and a randomly selected SA article contains the "wisdom of crowds". However, untabulated robustness checks show that there is significant daily cross-sectional dispersion of the pessimism scores extracted from SA articles. Thus, we can conclude that the investor mood conveyed in individual SA articles is at least to some extent independent of previous SA articles, and thus does not contain the "wisdom of crowds".

Table 8

Matching the number of authors.						
Panel A: Sample statistics						
	152 authors w	ith most articles	Random sampl	ing		
Number of authors	152		umber of authors 152		152	
Number of days in sample	3,742		1,702			
Number of articles	46,566		2,733			
Panel B: Pessimism index						
	152 authors w	ith most articles	Random sampl	ing		
	Coeff.	t-stat.	Coeff.	<i>t</i> -stat		
SA _{t-1}	-0.092***	-3.7	-0.028***	-3.6		
SA_{t-2}	0.018	0.7	-0.006	-0.8		
SA_{t-3}	0.026	1.2	0.009	1.3		
SA_{t-4}	-0.013	-0.5	-0.002	-0.2		
SA ₁₋₅	0.037	1.4	0.001	0.1		

The table provides the results of the analyses the relation between stock returns and the pessimism index obtained from Seeking Alpha (SA) articles when using smaller sub-samples of SA articles written by 152 authors, matching the number of authors of the *Wall Street Journal* (WSJ) data set. In the first test (on the left), we present the coefficient estimates using the sub-sample of 152 SA authors that have written most articles during the sample period. In the second test (on the right), we first select 25 random samples of 152 SA authors. Then we calculate the average coefficient estimates obtained from the 25 random samples. Panel A shows the summary statistics of the samples. In the random sampling test, the average summary statistics of the 25 random samples is presented. Panel B shows the (average) coefficient estimates. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

The results imply that the larger pool of SA authors cannot explain why social media sentiment on SA is better than traditional media sentiment on the WSJ in predicting the market. Social media sentiment remains statistically significant, even if using a subsample of SA authors that matches the total number of WSJ authors. Put differently, the "wisdom of crowds" hypothesis cannot explain the outperformance of social media sentiment on SA.

5.3. Writing style and language

The writing style and language of social media platforms is different from traditional media. Articles appearing in the traditional media are usually more structured, adhering to some specific formats, which contrasts with the rather non-structured user content published on social media. Even standard newspapers often adopt a different and richer language style on their social media accounts relative to their print editions (Hagvar, 2019). Besides using emojis, they tend to express more emotions and convey more subjective options online.

Differences in writing style and language could be important in this context, since investor sentiment is measured as the fraction of positive and negative words as specified by the (Loughran and McDonald, 2011) dictionary. Thus, there might be a direct impact of the writing style on the usefulness of sentiment indices estimated from the different media sources. In fact, Table 2 shows that the fraction of both positive and negative dictionary words is around 50% larger in the sample of SA articles than in the WSJ sample. Put differently, SA articles might convey more investor sentiment, and are thus better in predicting stock market movements.

We test this conjecture by forming another subsample of SA articles, this time based on the fraction of dictionary words. More precisely, we create a sub-sample of SA articles such that their average dictionary ratio matches the dictionary ratio of the WSJ, which is at approximately 3.90%. To achieve this lower ratio in the SA sample, we drop articles from the original SA sample that have a high fraction of dictionary words, until the ratio matches the average of the WSJ sample. Panel A of Table 9 shows that the resulting sub-sample is significantly smaller than the full sample, but still substantial.

Panel B of Table 9 reports the regression results when using this sub-sample of SA articles. The results show that none of the coefficient estimates of the first five lags of the pessimism index is statistically different from zero. This suggests that the pessimism index has lost its explanatory power for stock returns when restricting the SA sample to articles with a low fraction of dictionary words.

Overall, this finding confirms the hypothesis that the richer language style of SA articles provides explanatory power for why sentiment captured from social media is better than sentiment extracted from traditional media in predicting the market. In addition, this finding reinforces the view that the success of the lexicon approach hinges on a good match between the lexicon and the text data to be analyzed.

5.4. Publication time

In this subsection, we examine the impact of the publication time of WSJ and SA articles on their ability to detect investor mood. As a print medium, the WSJ articles of the *Markets* column are traditionally written after the market close. While in most of the 20th century, the WSJ was not available until the next morning, these articles can now be accessed on the internet as soon as they are written, also during the evening hours. Articles on SA are published, in contrast, around the clock, including during trading hours.

Matching the fraction of dictionary words.			
Panel A: Sample statistics			
Number of authors	2,650		
Number of days in sample	3,727		
Number of articles	32,065		
Fraction of dictionary words	3.90%		
Panel B: Pessimism index			
	Articles with lov	w fraction of dictionary words	
	Articles with low	w fraction of dictionary words t-stat.	
<i>SA</i> _{<i>i</i>-1}	Articles with low Coeff. -0.016	w fraction of dictionary words t-stat. -0.8	
SA ₁₋₁ SA ₁₋₂	Articles with lov Coeff. -0.016 0.003	w fraction of dictionary words t-stat. -0.8 0.1	
SA_{t-1} SA_{t-2} SA_{t-3}	Articles with lov Coeff. -0.016 0.003 -0.008	x fraction of dictionary words t-stat. -0.8 0.1 -0.3	
SA_{t-1} SA_{t-2} SA_{t-3} SA_{t-4}	Articles with lov Coeff. -0.016 0.003 -0.008 -0.007	t-stat. -0.8 0.1 -0.3	

The table analyzes the relation between stock returns and the pessimism index obtained from Seeking Alpha (SA) articles when using a subset of articles with a lower fraction of dictionary words, matching the fraction of dictionary words of the *Wall Street Journal* (WSJ) data set. Panel A shows the summary statistics of the sample, while Panel B shows the regression results. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Recent empirical evidence on irrational investor sentiment by Sun et al. (2016) and Renault (2017) show that investor mood is short-lived and can considerably fluctuate over short periods of time. Hence, timing differences between WSJ and SA might play an important role for the effectiveness of the sentiment indices obtained from both media sources to predict the market. For example, if there is a shift in investor mood in the morning hours, the tone of SA articles might be affected, while the tone of the WSJ article is not as it was likely written the evening before.¹³ In this sense, SA is timelier in capturing investor mood than the WSJ, and the sentiment reflected in SA articles might be therefore more relevant for asset prices.

To test this hypothesis, we divide the sample of SA articles into two subsamples depending on their publication time: during trading hours or outside trading hours. If investor sentiment is indeed short-lived, the sentiment indices from SA articles should have substantial explanatory power for stock returns during trading hours, but lose their predictive power if they are published outside trading hours, similar to the WSJ sentiment.

Panel A of Table 10 shows that roughly two-thirds of SA articles are published outside trading hours in the U.S., while the rest are published during these trading hours. Panel B shows a strong relation between investor sentiment captured from SA articles appearing during trading hours on stock returns. More important, the relation between investor sentiment and stock returns during trading hours is highly significant at the 1% level. By contrast, there is no significant relation between sentiment extracted from articles appearing outside trading hours and stock returns — despite the larger sample size.

These results confirm that the predictive power of investor sentiment on stock returns is indeed short-lived. Thus, the immediate and continuous flow of articles on SA has an advantage over WSJ print media to capture investor sentiment. A possible explanation for this pattern is that the higher frequency of SA articles allows for greater iteration and thus amplification (Hirshleifer, 2020). At the same time, the results in this section also further suggest the reason why the sentiment of traditional print media like the WSJ has lost is explanatory power for stock returns over the last century (see Appendix A).

5.5. Article length

As discussed in Section 5.5, SA articles used to extract investor sentiment are, on average, slightly shorter than WSJ articles. Yet, their length varies considerably more. This difference in dispersion is a direct consequence of differences in the media format. If the tone conveyed in articles varies with their length and complexity, some of the stronger association between SA sentiment and stock returns might be explained by the higher variability in article length. For example, if SA authors publish longer articles with more in-depth analyses on days with a more pronounced investor mood (either positive or negative), SA articles might be more suitable to detect investor sentiment relative to the WSJ that has a more uniform article length.

To examine whether differences in article length can explain why the sentiment extracted from SA is better in explaining shortterm stock returns than the WSJ, we again use a subsample analysis. We form a subset of SA articles that matches the mean and standard deviation of the article length of the WSJ sample. Table 11 shows, however, that when using such a subsample, the sentiment extracted from SA is still significantly related to short-term stock returns. Thus, differences in article length caused by the different publication formats do not seem to have an impact on the results.

¹³ WSJ articles appearing online can be updated. Yet, there are only a few of such revisions. In addition, these updates happen very quickly after the initial publication of the articles.

Table 10

Analysis	by	publication	time

Panel A: Sample statistics					
	During trading	hours	Outside trading hours		
Number of authors	2,715		3,329		
Number of days in sample	3,568		3,757		
Number of articles	23,414		52,721		
Panel B: Pessimism index					
	During trading	hours	Outside trading hours		
	Coeff.	t-stat.	Coeff.	t-stat.	
SA_{t-1}	-0.125***	-5.2	-0.030	-1.3	
SA_{t-2}	0.020	0.9	0.015	0.6	
SA_{t-3}	0.015	0.7	0.005	0.2	
SA_{t-4}	0.014	0.7	0.012	0.5	
SA. 5	0.006	0.3	0.017	0.7	

The table provides the results of analyses of the relation between stock returns and the pessimism index obtained from Seeking Alpha (SA) articles for two sub-samples. The first sample (on the left) consists of articles published during U.S. trading hours from 9:30 am ET to 4:00 pm ET; the second sample (on the right) consists of articles published outside trading hours. Panel A shows the summary statistics of the sample, while panel B shows the regression results. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 11		
Matching the article length.		
Panel A: Sample statistics		
Number of words	656	
Standard deviation of words	168	
Number of articles	35,201	
Panel B: Pessimism index		
	Coeff.	t-stat.
SA_{t-1}	-0.063***	-2.6
SA_{t-2}	0.019	0.8
SA_{t-3}	-0.000	-0.0
SA_{t-4}	0.014	0.6
SA_{t-5}	0.019	0.8

The table analyzes the relation between stock returns and the pessimism index obtained from Seeking Alpha (SA) articles for a sub-sample that matches the mean and standard deviation of the article length of the *Wall Street Journal* (WSJ) sample. Panel A shows the summary statistics of the sample, while Panel B shows the regression results. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

5.6. The choice to publish

While the *Markets* column of the WSJ is published every day after the market close (regardless of the market environment), SA authors can chose whether or not to publish articles. This difference might create some endogeneity problem: If SA authors chose to publish articles only on days with more stock market activity, then the higher predictive power of investor mood extracted from SA articles might be explained by such a self-selection bias.

There might be two reasons for a selective publication strategy. SA authors might genuinely be interested in sharing their views, opinions, or interpretations on days with higher market activity. Alternatively, more people are likely to consult SA for opinions and analyses on such days, thereby increasing the monetary compensation for SA authors.

Yet, empirically, for all SA authors combined, the choice to publish is effectively made almost every day. The raw data set contains at least one SA article on 3761 trading days, i.e., there are only 15 days without any SA article. The exception are some days at the beginning of the sample period, when the number of articles in the Today's Market section was low. Thus, regardless of the market environment, there is at least one article in the SA sample. This means that there are almost equally many observations of investor sentiment of SA and the WSJ. Collectively, there is no endogeneity problem.

Nevertheless, the choice to publish might not be evenly distributed across the sample period. For example, SA authors might be more inclined to publish articles on days with some important macroeconomic announcements. Untabulated analyses show that for some announcements days, there are indeed significantly more SA articles relative to non-announcement days. Further analyses in Section 6.2 show, however, that after controlling stock marker returns for such announcement effects, the sentiment extracted from SA articles is still significantly related to stock returns.

6. Further analyses

We show that transitory irrational investor mood extracted from media sources significantly affects stock returns. Furthermore, sentiment obtained from SA is better in explaining short-term stock returns than sentiment extracted from the WSJ.

This interpretation of our findings rests upon the assumption that the investor sentiment measures are unrelated to any fundamental news. If the sentiment measures are not fully independent of any news component, then this interpretation might be wrong. More important, if SA articles have more informational content than WSJ articles, we would expect sentiment extracted from SA articles to have a greater ability to predict stock returns — not because of irrational investor mood, but because if new rational information.

Indeed, the timing adjustments following (Garcia, 2013) described in Section 3.1.1 to separate information from sentiment might not be ideally suited for social media. The WSJ articles of our sample are published after the market is closed, discussing events that took place earlier in the day. Thus, they refer to past information, which has already been impounded in market prices. If the numbers of positive and negative words of these articles are predicting next day's stock returns, this effect cannot be explained by new information, but investor sentiment only. By contrast, although SA articles tend to cover past events, a SA article appearing, say during Tuesday trading hours, might reflect current events. In this case, new information released on Tuesday afternoon influences both SA articles (and hence the sentiment index) and stock returns of that day. As a consequence, the higher predictive power of the SA sentiment index over the WSJ sentiment index might not be an indication of SA carrying more investor sentiment, but just some new information. The finding in Section 5.4, showing that sentiment extracted from SA articles appearing during trading hours is driving the main results, also might point in this direction. Therefore, in this section, we revisit the problem of separating information content from sentiment by removing new information from the SA articles.

6.1.1. Opinion versus facts

The first approach aims at eliminating any news components from SA articles published during trading hours. To this end, we classify the sentences of all SA articles as facts or opinions/analyses. Then we drop the sentences containing facts from all SA articles appearing during trading hours. The remaining SA data set should therefore not contain any new information, at least not during trading hours, and is thus more comparable to WSJ articles.

We define facts as sentences that contain at least one of the following: numbers, numerical words, or words that refer to amounts such as "billions" or "thousands". In turn, all other sentences are flagged as opinions or analyses. To perform this task, we use the named entity recognition (NER) approach, a popular method in natural language processing (NLP).¹⁴

Fig. 6 shows exemplary the classification of sentences into facts and opinions for the SA article presented in Fig. 3. The majority of the article is classified as opinions, as very few sentences contain numbers, and are thus labeled as facts.

We acknowledge that this classification algorithm is not perfect. Although most relevant financial news is expressed in numbers (e.g., earnings announcements, share prices, exchange rates), some important news does not contain any numbers (e.g., political events). In turn, some opinions and analyses contain historical facts, which should not be classified as news (e.g., some year dates, as in Fig. 6). Nevertheless, we believe that this classification is still informative for our analysis.

Panel A of Table 12 presents the fraction of sentences that are classified as facts versus opinions for the WSJ and SA data sets. On average, 81% of WSJ articles are classified as facts, while SA tends to offer more opinions than facts. This result shows that the WSJ articles are more concerned with reporting facts, while the SA articles include more discussions, analyses, and opinions, as conjectured earlier.

Next, we drop all sentences classified as opinions from SA articles appearing during trading hours, and build the pessimism index as before. Then we re-estimate Eq. (1). Panel B shows that SA articles remain statistically significant at the first lag, similar as when we use the entire SA data set. This indicates that even if facts published during trading hours are removed from SA articles, the sentiment index extracted from SA is still able to explain market returns.

6.1.2. Excluding concurrent events

In the second approach, we aim to eliminating concurrent information contained in SA articles published during trading hours by removing entire articles based on their content. While most SA articles refer to past events (like the example presented in Fig. 6), some articles do cover current events. To identify these articles, we randomly select 100 articles from the SA articles published during trading hours, and read them carefully to identify whether they discuss concurrent events. Out of this sample of 100 articles, 21 articles discuss concurrent events. Many of these articles discuss Asian or European markets that open before the U.S. markets, along with other events that took place earlier in the day.

Then we identify the key expressions that serve as an indicator for this classification (such as *this morning*). Using this set of expressions, we go back to all SA articles published during U.S. trading hours and eliminate any article that contains at least one of these expressions. Out of 23,414 articles published during these trading hours, we identify 3826 as referring to concurrent events. Appendix B provides the full list of expressions.

¹⁴ NER identifies certain entities within a text and then classifies them into a pre-specified category (Repke and Krestel, 2021). In our study, we use NER to assign all words of the text into one of the pre-specified categories. Then we assume that words assigned to one of the CARDINAL, PERCENT or MONEY categories to represent facts.

'My View Of The S&P 500 Index - November 2017'

November 1, 2017, 10:16PM ET

Author: Walter Zelzniak Jr

Another month and another gain in the stock market. This bull market keeps on going higher and higher. The market, as measured by the S&P 500 index, was up again in October and closed 2.22% higher. The SPDR S&P 500 ETF (SPY) gained 2.36% for the month. As for my pension plan assets, I too gained for October. Consequently, my first investment goal, preservation of capital, was achieved. Unfortunately, I did not beat the returns of the S&P 500 index as measured by SPY. So only half of my investment objectives were met for the month of October. Table 1 below shows my returns for the month and Table 2 below shows my returns for the past 12 months. To review the purpose of this series of articles, my retirement account only allows me to buy the following four ETFs: iShares Core U.S. Aggregate Bond ETF (AGG), SPDR S&P 500 ETF, iShares Russell 2000 ETF (IWM), and iShares MSCI EAFE ETF (EFA). I can also have my money in cash. The guestion is how to decide where and when to allocate money to these various ETFs. I use my moving average crossover system combined with relative strength charts to determine how to allocate my pension plan assets. My moving average crossover system uses the 6-month and the 10-month exponential moving averages to identify which of the four ETFs are in a position to be bought. If the 6-month moving average is above the 10-month moving average, then the ETF is a buy. I call this setup being in bullish alignment. When the 6-month moving average is below the 10-month moving average, the setup is referred to as a bearish alignment. When a bearish alignment happens I don't want to hold that asset. See Chart 1 below for a long-term look at the S&P 500 index using my moving average crossover system. You can see that the moving average crossover system provided some excellent long-term buy and sell signals that would have allowed investors to capture long duration moves in the index, while avoiding costly drawdowns. Avoiding these costly drawdowns allows me to meet the objective of capital preservation. To me, the last place you want to experience a large drawdown is in your pension plan. During the 2008-2009 market crash, many people didn't even look at their retirement statements because they were afraid of what they would find. I submit that if those people would have used a market strategy similar to what I outline in this series of articles, they would have been able to avoid much of the decline during the bear market and consequently would have had less emotional stress during that time period. I find this investment strategy to be particularly useful for managing the assets in my pension plan. If your pension plan is anything like mine, then you just have a choice of ETFs or mutual funds to select from. Looking at standard fundamental criteria such as P/E ratios, free cash flow, dividend yield, and the like is not easily applied to stock indices and mutual funds. What is easily applied to stock indices and mutual funds is trend following technical analysis as I show in this article. The following charts show the current status of the ETFs that I am allowed to buy in my retirement account. As mentioned in the opening paragraph, SPY was up for the month of October gaining 2.36%. Looking at Chart 2, you can see the strong uptrend that is in place since the last buy signal was given. The trend has been easy to follow since the buy signal occurred when the 6-month moving average crossed above the 10-month moving average and that is the idea. Trend following is meant to be easy and profitable. Chart 3 shows that IWM gained 0.73% for the month and closed at a new high. In last month's article, I said I would allocate some money to IWM and that is what I did. Since the buy signal was generated in August of 2016 IWM has gained over 24%. Chart 4 shows the relative strength of IWM compared to SPY. When the signal line is trending higher it means IWM is outperforming SPY. When the signal line is trending lower it means the opposite: IWM is underperforming SPY. The downward sloping trend line from early 2014 to early 2016 shows that investors would have made more money investing in SPY versus investing in IWM. For October, IWM underperformed SPY by 1.59%, giving back some of the outperformance IWM had in the month of September. If the ratio can take out the highs of last December, then the downtrend in the ratio that was discussed in a previous article will have reversed. The ratio is still above its recent low in June. I will continue to maintain a small position in IWM and monitor this ratio. EFA had another good month and closed up 1.68%. EFA closed at a new high again in October and EFA remains in bullish alignment. The trend is growing stronger as the distance or the whitespace between the 6-month moving average and the 10-month moving average is getting wider. Chart 6 shows that EFA, like IWM mentioned previously, was outperformed by SPY. Despite the underperformance, the ratio remains in bullish alignment. The relative strength ratio in Chart 6 is one to watch closely. Since 2013, US stocks have outperformed international stocks until early 2017. Readers can see this as the signal line falls in Chart 6 from the upper left corner to the lower right corner. However, I see that a change is taking place as the signal line is now starting to trend higher since early 2017, and the 6-and 10-month moving averages have gone into bullish alignment. As I mentioned last month, the EFA:IWM ratio could be rolling over. EFA did outperform IWM so the ratio closed higher. Two months ago the 6- and 10-month moving averages moved into a bullish alignment and they remain in bullish alignment, but both moving averages are sloping downward. I will continue to monitor this ratio. AGG had a minimal gain and remains in bullish alignment. Bonds investors are still making money but investors in bonds are lagging the returns of investors in equities as can be seen in Chart 9. AGG continues to be outperformed by SPY and the ratio closed at a new low in October. Investors simply prefer stocks over bonds at this time. This simple ratio confirms that equities continue to be the place to put your money. None of my retirement assets are allocated to AGG. For the month of October, I was allocated 50% EFA, 25% IWM and 25% SPY, and I will maintain that allocation for the seasonally strong month of November. The ETFs for EFA, IWM, and SPY are in bullish alignment and the current trends are strong as depicted in Charts 5, 3, and 2, respectively. I have used the relative strength charts to allot my assets and I favor the indication in Chart 6, so I have 50% of my allotment to EFA as I expect it to outperform SPY in the near future. All in all, I think that my moving average crossover system has me properly aligned and I will continue to monitor the charts.

Fig. 6. Sample article of the Today's Market section on Seeking Alpha. The highlighted text indicates sentences classified as facts, based on NER. All other sentences are classified as opinions or analyses.

Table 12 Opinions versus	facts.		
Panel A: Classi	fication of sentences		
	Opinions	Facts	
WSJ	19%	81%	
SA	65%	35%	
Panel B: Pessir	nism index		
	Coeff.	t-stat.	
SA_{t-1}	-0.066*	-2.6	
SA_{t-2}	0.031	1.2	
SA_{t-3}	0.010	0.4	
SA_{t-4}	0.019	0.7	
SA_{t-5}	0.000	0.0	

Panel A presents the faction of sentences that are classified as facts versus opinions for the *Wall Street Journal* (WSJ) and Seeking Alpha (SA) datasets. Sentences are classified as facts if they contain some numerical value or some word that refer to numbers or amounts, while all other sentences are classified as opinions. Panel B shows the regressions results when explaining stock returns using the pessimism index obtained from SA articles for the sub-sample which excludes all sentences classified as facts appearing during trading hours. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 13		
Excluding concu	rrent events.	
Pessimism inde	х	
	Coeff.	t-stat.
SA_{t-1}	-0.078**	-3.1
SA.	0.039	1.5

0.011

0.028

-0.007

0.5

1.1

-0.3

 SA_{t-3}

 SA_{t-4}

 SA_{t-5}

The table shows the regressions results when explaining stock returns using the pessimism index obtained from Seeking Alpha (SA) articles for the sub-sample which excludes all articles published during trading hours that refer to concurrent events. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Again, this approach might lead to a noisy exclusion of SA articles, as there might be articles referring to concurrent events that do not contain any of the 15 expressions identified. Yet, the 3826 articles referring to concurrent events correspond to approximately 16% of the within-trading-hours population, which is close to the sample (21%). Furthermore, most concurrent events are clearly identified using a few key expressions only.

Table 13 shows that excluding articles published during trading hours that refer to concurrent events before constructing the pessimism index does not change the baseline results. Again, the SA pessimism index coefficient remains statistically significant at the first lag.

Taken together, these analyses support the view that the predictive power of the sentiment index extracted from SA articles for stock returns is not a result of information contained in SA articles, but is likely to reflect the effects of irrational investor mood.

6.2. Reverse causality

The continuous publication process of SA articles during trading hours may raise yet another concern, known as reverse causality: the sentiment of SA articles may reflect the reaction of investors to what had happened earlier in the trading day, even if the article discusses some events in the distant past. For example, a SA article written and published during falling markets might have a more negative tone than the same story would have in a positive market environment. Hence, instead of sentiment predicting stock returns, stock returns influence investor sentiment. The higher ability of investor sentiment extracted from SA articles to predict stock returns could therefore be explained by reverse causality, which is stronger for social media than for print media.

To address this concern, we add additional control variables to capture news affecting the stock market. As control variables, we use a set of macroeconomic announcements that have been used in the literature to analyze the impact of news on stock

Table 14			
Controlling	for	macro-economic	announcements.

	Coeff.	t-stat.	Coeff.	t-stat.
Panel A: Ann	ouncement dummies			
GDP	0.196*	2.1	0.193**	2.1
PPI	-0.175	-1.5	-0.170	-1.5
CPI	0.037	0.4	0.032	0.4
UR	0.068	0.7	0.063	0.7
RS	0.065	0.7	0.061	0.7
TB	0.121	1.3	0.122	1.4
Panel B: Pess	imism index			
SA_{t-1}			-0.092***	-3.6
SA_{t-2}			0.036	1.4
SA_{t-3}			0.011	0.5
SA_{t-4}			0.035	1.4
SA_{t-5}			-0.001	0.0

The table analyzes the relation between stock returns and the pessimism index obtained from Seeking Alpha (SA) articles, when controlling stock returns for macro-economics announcements. For Panel A, the column of the left shows the coefficients of the macro-economic dummy variables when excluding the pessimism index from the regression model (2) while the column on the right presents the coefficients of the macro-economic dummy variables when including the pessimism index. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

markets (Andersen and Bollerslev, 1998). Such news is likely to be reflected in the sentiment of the articles published the same day. We resort to six important macro-variables: Gross Domestic Product (GDP), Production-Price Index (PPI), Consumer Price Index (CPI), Unemployment Rate (UR), Retail Sales Growth (RS), and Trade Balance (TB).¹⁵

For each announcement, we create a dummy variable taking the value of 1 on an announcement day, and 0 otherwise (Savor and Wilson, 2013; Lucca and Moench, 2015). Then we include these dummies as follows,

$$R_t = \beta_{SA} \mathcal{L}_s(SA_t) + \gamma \mathcal{L}_s(R_t) + \delta \sum_{m=1}^6 D_{mt} + \psi \mathcal{L}_s(R_t^2) + \eta X_t + u_t.$$
(2)

Table 14 shows that most macro-economic announcements have no effect on stock returns. Only GDP announcements have a statistically and positive effect. When adding the sentiment index to the equation, the results do not change. The first lag of the SA sentiment index remains statistically significant even if we control for macro announcements.¹⁶

7. Conclusion

Using a large data set of daily articles and reader comments from 2006 to 2020, this study shows that investor sentiment extracted from the social media platform Seeking Alpha is better in predicting daily stock market returns than investor sentiment obtained from the *Wall Street Journal*, a traditional print medium. The relation between investor sentiment and stock market returns is however highly transitory, in line with (Baker and Wurgler, 2007), who consider that investment decisions driven by investor mood largely irrational.

Seeking Alpha is more suitable to extract investor sentiment because of the richer language and the timeliness of online media. Contrary to conventional wisdom, the power of sentiment extracted from Seeking Alpha in predicting market returns cannot be explained by differences in the volume and length of articles published on social media platforms, or the greater diversity of its contributors.

The mechanism behind these results is simple. Sentiment indices based on the lexicon approach convert the fraction of positive and negative words into measures of investor mood. Since social media exhibits a significantly richer and more expressive language compared to traditional media, it has a clear advantage to extract a valuable measure of investor sentiment. Furthermore, given that investor mood is highly transitory, the continuous feed of articles gives social media a clear timing advantage to extract investor sentiment.

The results of this study highlight the importance of considering novel media sources to explain market behavior driven by investor sentiment. At the same time, this study contributes to a better understanding of the mechanisms of lexicon-based sentiment indices. For the lexicon approach to be beneficial, it has to match the language of the media analyzed.

CRediT authorship contribution statement

Ioanna Lachana: Writing – original draft, Software, Investigation, Formal analysis, Data curation. **David Schröder:** Writing – review & editing, Supervision, Conceptualization.

¹⁵ For more details, see Appendix C.

¹⁶ We also test whether dropping all announcement days from the model changes the results. The SA sentiment coefficient remains statistically significant in this case as well.

Declaration of competing interest

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Appendix A. Traditional print media and stock returns

In this appendix, we further explore the relation between sentiment indices obtained from traditional print media and stock market returns. More precisely, we aim to reconcile the difference in results reported in Panel B of Table 4 (no statistical significance) with those obtained by Garcia (2013).

We replicate some of the analyses by Garcia (2013). First, we re-estimate model (1) using the sentiment indices by Garcia (2013) based on *New York Times* (NYT) articles, using the S&P 500 index instead of the DJIA. Second, we analyze whether the effect of investor sentiment on stock returns has changed over the last century.

Since the S&P 500 is not available since 1905, the starting point of the analysis in Garcia (2013), we limit the analysis to the period from 1930 to 2005. Panel A of Table A.1 shows the descriptive statistics of the NYT sentiment indices. The distribution of the NYT sentiment indices are very similar to those of the WSJ (see Panel A of Table 2), suggesting that the fraction of positive and negative words that appear in both newspapers are not very different.

Next, Table A.2 replicates the analysis of Garcia (2013) for the period from 1930 to 2005. In this specification, the first lag for all three sentiment indices is statistically significant at the 1% significance level. Thus, we are able to replicate the results of Garcia (2013), despite using a different time horizon and stock market index.

Finally, in Table A.3, we divide the entire time period into five 15-year intervals. The results indicate that the relation between sentiment indices obtained from traditional print media sources and stock market returns decays through time. While in the first two time intervals 1930–1944 and 1945–1959, the sentiment indices have a strong predictive power for stock market returns, this effect is hardly present in later sub-samples. In the period from 1990 to 2005, only the negative sentiment index is related to stock returns, similar to the results reported in Panel B of Table 4 for the 2006–2020 period using the WSJ sentiment index.

These findings show that investor sentiment obtained from traditional media lost its predictive power for stock market returns over the last century. This result also mirrors evidence by Gan et al. (2020), showing that the return responses to newspaper-based investment sentiment halved in the period from 2012 to 2017. Further analysis shows that this decline is not due to a change in newspapers language over time (see Section 5.3), as the fraction of dictionary words is fairly stable around 3.3% in the NYT sample.

Table A.1

builling statistics.					
	Mean	Median	25%-quant.	75%-quant.	Std. dev.
Panel A: New York Times	articles				
Positive sentiment	1.19	1.15	0.88	1.47	0.44
Negative sentiment	2.12	2.05	1.63	2.53	0.70
Pessimism index	0.93	0.87	0.30	1.50	0.92
Panel B: S&P 500 index					
Log returns	0.021	0.040	-0.463	0.529	1.165

Panel A reports sample statistics of sentiment measures obtained from the *New York Times* articles as provided by Garcia (2013) for the time period from 1930 to 2005. The positive sentiment measure (G_i) is the fraction of positive words out of total words; the negative sentiment measure (B_i) is the fraction of negative words out of total words. The pessimism index is the difference between the negative and the positive sentiment measures ($P_i = B_i - G_i$). Panel B presents the summary statistics of daily log returns of the S&P 500 index from 1930 to 2005.

Table A.2

Sentiment	indices	and	stock	market	returns ·	— Ne	w York	Times.
-----------	---------	-----	-------	--------	-----------	------	--------	--------

	Positive sentiment		Negative sentir	nent	Pessimism index		
	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.	
M_{t-1}	0.060***	6.6	-0.066***	-6.6	-0.084***	-7.9	
M_{t-2}	-0.016*	-1.9	0.017*	1.8	0.023**	2.4	
M_{t-3}	-0.014	-1.6	-0.004	-0.5	0.004	0.4	
M_{t-4}	-0.008	-1.0	0.012	1.3	0.014	1.5	
M_{t-5}	0.008	1.0	0.018**	2.0	0.011	1.1	

This table reports the estimated coefficients β_M of the model:

 $R_t = \beta_M \mathcal{L}_s(M_t) + \gamma \mathcal{L}_s(R_t) + \psi \mathcal{L}_s(R_t^2) + \eta X_t + u_t.$

The dependent variable R_i is the daily log return of the S&P 500 from 1930 to 2005. The variable M_i denotes the media sentiment index obtained from the *New York Times* articles. We consider the positive sentiment index (G_i), the negative sentiment index (B_i), and the pessimism index ($P_i = B_i - G_i$). \mathcal{L}_s is the lag operator with s = 1...5. The sentiment indices are normalized to have zero mean and unit variance. X_i represents the set of exogenous variables that consists of day-of-the-week dummies and a constant, and u_i is the error term. The *t*-statistics are computed using (White, 1980) standard errors. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table A.3

Sentiment indices and stock market returns over time - New York Times.

Panel A:	Positive sentime	ent								
	1930–1944		1945–1959		1960–1974	1960–1974			1990-2005	
	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.
M_{t-1}	0.117***	3.2	0.069***	4.9	0.039***	2.9	0.016	1.1	0.001	0.1
M_{t-2}	-0.011	-0.3	-0.031**	-2.0	0.005	0.3	-0.012	-0.8	0.004	-0.2
M_{t-3}	-0.055	-1.6	0.002	0.1	0.015	1.1	-0.012	-0.8	0.001	0.1
M_{t-4}	-0.026	-0.7	-0.012	-0.9	-0.004	-0.3	-0.013	-1.0	0.004	0.3
M_{t-5}	0.016	0.4	-0.005	-0.3	0.024*	1.9	0.011	0.7	0.033**	2.0
Panel B:	Negative sentim	ent								
	1930–1944		1945–1959		1960–1974		1975–1989		1990-2005	
	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.
M_{t-1}	-0.197***	-4.7	-0.068***	-3.5	0.002	0.1	-0.012	-0.8	-0.028**	-1.9
M_{t-2}	0.013	0.3	0.005	0.2	0.033**	2.2	-0.003	-0.2	-0.002	-0.1
M_{t-3}	-0.062	-1.5	0.000	0.0	-0.001	0.0	-0.002	-0.2	-0.001	-0.1
M_{t-4}	0.075*	1.9	0.002	0.1	0.019	1.2	0.008	0.5	-0.012	-0.9
M_{t-5}	0.031	0.8	0.035**	2.0	-0.015	-1.0	0.016	1.1	0.006	0.4
Panel C:	Pessimism index	C C								
	1930–1944		1945–1959		1960–1974		1975–1989		1990-2005	
	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.	Coeff.	t-stat.
M_{t-1}	-0.223***	-5.1	-0.097***	-4.9	-0.023	-1.3	-0.018	-1.2	-0.025	-1.6
M_{t-2}	0.020	0.5	0.021	1.1	0.025	1.6	0.004	0.3	0.000	0.0
M_{t-3}	-0.022	-0.5	-0.001	-0.1	-0.009	-0.6	0.005	0.3	-0.001	0.0
M_{t-4}	0.070*	1.8	0.009	0.5	0.019	1.2	0.013	0.9	-0.012	-0.8
M_{t-5}	0.017	0.4	0.031	1.8*	-0.027*	-1.8	0.006	0.4	-0.008	-0.5
		1 60 1			-		-			

This table reports the estimated coefficients β_M of the model:

 $R_t = \beta_M \mathcal{L}_s(M_t) + \gamma \mathcal{L}_s(R_t) + \psi \mathcal{L}_s(R_t^2) + \eta X_t + u_t.$

for 5 different time sub-samples. For more details, see Table A.2. ***, ***, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

It seems therefore likely that the increased competition of more timely media sources, such as TV and the internet, is contributing to this decline (for more discussion, see Section 5.4).

Appendix B. Expressions to identify concurrent events

In Table B.1, we present the expressions we use to identify concurrent events in the SA articles published during U.S. trading hours, along with their occurrence in the random sample of 100 articles, as well as in the total population. We also present in Fig. C.1 an example of a SA article published during U.S. trading hours which we identify to cover concurrent events.

Expression	Sample	Population	Fraction
This morning	9	2,863	12.23%
Pre-market	3	999	4.27%
Pre-game	2	279	1.19%
The European morning	1	10	0.04%
Early afternoon trading	1	8	0.03%
In overnight trading	1	107	0.46%
Intra-day	1	198	0.85%
Premarket	1	433	1.85%
Earlier today	1	97	0.41%
Finished higher today	1	28	0.12%
At the moment	1	983	4.20%
At the time of the writing	1	3	0.01%
Closed higher today	1	2	0.01%
Finished lower today	1	21	0.09%
Closed lower today	1	4	0.02%
Total	100	23,414	

The table presents the expressions used to identify concurrent events, along with their occurrence in the random sample of 100 Seeking Alpha (SA) articles, as well as in the total SA population.

'Daily State of the Markets: Taking Another Look?'

March 11, 2011, 10:19AM ET

Author: David Moening

Good morning. A powerful earthquake in Japan overnight has triggered a massive tsunami and a wave of selling in overseas markets (the earthquake hit just 14 minutes before the close in Japan). While tragic from a human perspective, forces of Mother Nature usually don't have long-lasting effects on the economies of the world. As such, it is important to recognize that the current selling being done appears to be a continuation of traders taking another look at the premise for recent rally.

By now, everybody knows that up until February 22, the stock market had been a one-way street based on the concept that the U.S. economy had finally reached a growth rate that appeared to be sustainable. The thinking was that with the economic data coming in largely better-than-expected, job creation (as well as all the good things that accompany job growth) couldn't be too far behind. As such, traders discounted blue skies ahead on a daily basis and bought each and every dip.

Along the way, the market became very overbought and investor sentiment reached extremely positive levels – something that traditionally leads to a corrective or consolidation phase. Thus, the events in the MENA (Middle East, North Africa) regions and the resulting spike in oil became the trigger for the current pullback.

However, up until yesterday, the tape action seen in the market seemed to indicate that this was a garden variety consolidation pattern and that the chances of a meaningful correction (which we'll define as a decline in excess of -5%) seemed remote. That is, right up until the Chinese reported an upsetting balance of trade number and the U.S. trade deficit spiked yesterday.

In short, the fact that exports only grew by 2% in China (largely due to the increasing cost of oil) got traders' attention -and suddenly the worry about the Chinese economy was back on the front burner. We had already seen that the tightening measures employed by the People's Bank of China were starting to have an impact. And up until yesterday, the fears that the Chinese might go too far had seemed to largely fade into the background. But with this data, the worry that the world's leading economy might be starting to slow caused stock traders to revisit the issue and to switch back into "risk off" mode. It's easy to see why traders might have instantly reacted to the China data. However, understanding why anybody trading stocks cares about the U.S. Balance of Trade from a stock market perspective might be a different story. The key here is that the trade data is one of the inputs in GDP. And if the deficit continues to spike higher, Q1 GDP could be affected. In fact, Ned Davis Research was out with a note Thursday saying that some of the recent indicators suggest a GDP rate closer to 2.1%, which is a far cry from the current rate of 2.7%.

So, given that the situation in Libya might stick around a while -- which could easily lead to more unrest in the region, as well as to oil prices remaining "sticky" above \$100 -- the issue of European sovereign debt rearing its ugly head again, and reports of a crowd being fired upon (which may have been a bit inflammatory) in Saudi Arabia, the concept of traders taking another look at the premise yesterday certainly makes some sense.

From a chart perspective, the bulls may (key word) have a problem here. The bottom line is that the charts of the major indices broke down yesterday on a closing basis -- although not by much, and I can argue that the support zone on the Nasdaq remains intact. But if the market can rebound quickly -- as it has done over the past two years -- then all will be forgiven and the buyers will likely return post haste. However, should the bears find a way to push the indices lower in the next few days, traders may also be taking another look at their commitment levels in the stock market.

Turning to <u>this morning</u>... The big story is the earthquake and the resulting 10-meter tsunami that destroyed some of the Pacific islands and may reach Hawaii or the U.S. Northwest coast. As a result, Asian stock markets were down hard. However, with oil (currently trading under \$100), copper, silver, etc. in decline, traders in the U.S. have yet to hit the panic button. Oh, and the Saudi "day of rage" hasn't materialized (yet?).

On the economic front ... The Commerce Department reported that retail sales rose in the month of February by 1.0%. This was below the consensus for +1.1%. When you strip out the sales of autos, sales were up 0.7%, which was also a tenth below the consensus for an increase of 0.8%.

We will also get the University of Michigan Consumer Sentiment Report later this morning at 9:55 am Eastern. Thought for the day: Best of luck on this Friday and be sure to enjoy the weekend!

Pre-Game Indicators

Here are the Pre-Market indicators we review each morning before the opening bell ...

Fig. C.1. Sample article of the Today's Market section on Seeking Alpha. The underlined text indicates expressions used to identify articles that refer to concurrent events.

Appendix C. Macro-economic announcements

In this appendix, we present an overview of the macroeconomic announcements that we use when assessing the impact of reverse causality on our results. Table C.1 provides information on the six U.S. macroeconomic variables, including data source, frequency, and total number of observations from 2005 to 2020. While GDP data are quarterly, GDP estimates are updated on a monthly basis.

Table C.1

Summary of U.S. macroeconomic announcements (2006-2020).

	GDP	PPI	CPI	UR	RS	TB
Panel A: Information						
Source	BEA	BLS	BLS	BLS	Census	BEA
Frequency	Quarterly	Monthly	Monthly	Monthly	Monthly	Monthly
Number of announcements	177	180	180	180	180	177
Panel B: Number of announcements (by day of the week)						
Monday	1	0	2	4	16	0
Tuesday	15	55	31	1	39	34
Wednesday	33	42	62	0	38	39
Thursday	72	45	41	5	34	55
Friday	56	38	44	170	53	47

Panel A provides an overview of the macro-economic announcements, including the data source, frequency, and total number of observations from 2005 to 2020. The six announcements are Gross Domestic Product (GDP), Producers Price Index (PPI), Consumer Price Index (CPI), Unemployment Rate (UR), Retail Sales Growth (RS), and Trade Balance (TB). BEA is the U.S. Bureau of Economic Analysis, BLS is the Bureau of Labor Statistics, and Census is the U.S. Census Bureau. We only consider announcements on trading days in the U.S. (a few announcements have been made on the weekend). Panel B provides the number of announcements made by the day of the week.

Most macro-economics announcements, such as PPI and RS, are well distributed throughout the week. The unemployment rate is typically announced on Friday. Also, macroeconomic announcements are rare on Mondays.

Appendix D. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.finmar.2025.100970.

References

- Andersen, Torben G., Bollerslev, Tim, 1998. Deutsche mark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies. J. Financ. 53 (1), 219–265.
- Baker, Malcolm, Wurgler, Jeffrey, 2006. Investor sentiment and the cross-section of stock returns. J. Financ. 61 (4), 1645-1680.

Baker, Malcolm, Wurgler, Jeffrey, 2007. Investor sentiment in the stock market. J. Econ. Perspect. 21 (2), 129-151.

- Bartov, Eli, Faurel, Lucile, Mohanram, Partha S., 2018. Can twitter help prdict firm-level earnings and stock returns? Account. Rev. 93 (3), 25–57.
- Budescu, David V., Chen, Eva, 2014. Identifying expertise to extract the wisdom of crowds. Manag. Sci. 61 (2), 267-280.
- Chen, Hailiang, De, Prabuddha, Hu, Yu, Hwang, Byoung-Hyoun, 2014. Wisdom of crowds: The value of stock opinions transmitted through social media. Rev. Financ. Stud. 27 (5), 1367–1403.
- Conrad, Jennifer, Gultekin, Mustafa N., Kaul, Gautam, 1997. Profitability of short-term contrarian strategies: Implications for market efficiency. J. Bus. Econom. Statist. 15 (3), 379–386.
- Cox, Don R., Peterson, David R., 1994. Stock returns following large one-day declines: Evidence on short-term reversals and longer-term performance. J. Financ. 49 (1), 255–267.
- Da, Zhi, Engelberg, Joseph, Gao, Pengjie, 2015. The sum of all FEARS: Investor sentiment and asset prices. Rev. Financ. Stud. 28 (1), 1-32.

Edmans, Alex, Garcia, Diego, Norli, Oyvind, 2007. Sports sentiment and stock returns. J. Financ. 62 (4), 1967–1998.

- Elliott, Graham, Timmermann, Allan, 2016. Economic Forecasting. Princeton University Press.
- Fama, Eugene, 1970. Efficient capital markets: A review of theory and empirical work. J. Financ. 25 (2), 383-417.
- Frankel, Richard, Jennings, Jared, Lee, Joshua, 2022. Disclosure sentiment: Machine learning vs. Dictionary methods. Manag. Sci. 68 (7), 5514–5532.
- Gan, Baoqing, Alexeev, Vitali, Bird, Ron, Yeung, Danny, 2020. Sensitivity to sentiment: News vs social media. Int. Rev. Financ. Anal. 67, 1139-1168.
- Garcia, Diego, 2013. Sentiment during recessions. J. Financ. 68 (3), 1267-1300.
- Hagvar, Yngve Benestad, 2019. News media's rhetoric on facebook. Journal. Pr. 13 (7), 853-872.
- Hanna, Alan J., Turner, John D., Walker, Clive B., 2020. News media and investor sentiment during bull and bear markets. Eur. J. Financ. 26 (14), 1377–1395. Hillert, Alexander, Jacobs, Heiko, Müller, Sebastian, 2018. Journalist disagreement. J. Financ. Mark. 41, 57–76.
- Hirshleifer, David, 2020. More than words: Quantifying language to measure firms' fundamentals. J. Financ. 74 (4), 1779-1832.
- Hirshleifer, David, Shumway, Tyler, 2003. Good day sunshine: Stock returns and the weather. J. Financ. 58 (3), 1009-1032.
- Jiao, Peiran, Veiga, Andre, Walther, Ansgar, 2020. Social media, news media and the stock market. J. Econ. Behav. Organ. 176, 63-90.

Kim, Soon-Ho, Kim, Dongcheol, 2014. Investor sentiment from internet message postings and the predictability of stock returns. J. Econ. Behav. Organ. 107 (1), 708–729.

- Kittur, Aniket, Pendleton, Bryan A., Suh, Bongwon, Mytkowicz, Todd, 2007. Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. WWW: World Wide Web 1 (2), 19–28.
- Kraussl, Roman, Mirgorodskay, Elizaveta, 2017. Media, sentiment and market performance in the long run. Eur. J. Financ. 23 (11), 1059-1082.
- Leung, Henry, Ton, Thai, 2015. The impact of internet stock message boards on cross-sectional returns of small-capitalization stocks. J. Bank. Financ. 55 (7–8), 37–55.
- Li, Ting, van Dalen, Jan, van Rees, Pieter, 2018. More than just noise? Examining the information content of stock microblogs on financial markets. J. Inf. Technol. 33 (2), 50–69.
- Lo, Andrew W., MacKinlay, Craig A., 1990. An econometric analysis of nonsynchronous trading. J. Econometrics 45 (1-2), 181-211.
- Loughran, Tim, McDonald, Bill, 2011. When is a liability not a liability? Textual analysis, dictionaries, and 10-ks. J. Financ. 66 (1), 35-65.

Lucca, David, Moench, Emanuel, 2015. The pre-FOMC announcement drift. J. Financ. 70 (1), 329-371.

- Renault, Thomas, 2017. Intraday online investor sentiment and return patterns in the U.S. stock market. J. Bank. Financ. 84 (3), 25-40.
- Repke, Tim, Krestel, Ralf, 2021. Extraction and representation of financial entities from text. In: Consoli, Sergio, Reforgiato Recupero, Diego, Saisana, Michaela (Eds.), Data Science for Economics and Finance: Methodologies and Applications. Springer International Publishing, Cham, pp. 241–263.

Savor, Pavel, Wilson, Mungo, 2013. How much do investors care about macroeconomic risk? Evidence from scheduled economic announcements. J. Financ. Quant. Anal. 48 (2), 343-375.

Shiller, Robert J., 2016. Irrational Exuberance, third ed. Princeton University Press, Princeton, NJ.

Siganos, Antonios, Vagenas-Nanos, Evangelos, Verwijmeren, Patrick, 2017. Divergence of sentiment and stock market trading. J. Bank. Financ. 78 (3), 130–141.
Sprenger, Timm O., Sandner, Philipp G., Tumasjan, Andranik, Welpe, Isabell M., 2014a. News or noise? Using Twitter to identify and understand company-specific news flow. J. Bus. Financ. Account. 41 (7–8), 791–830.

Sprenger, Timm O., Tumasjan, Andranik, Sandner, Philipp G., Welpe, Isabell M., 2014b. Tweets and trades: the information content of stock microblogs. Eur. Financ. Manag. 20 (5), 926–957.

Sul, Hong Kee, Dennis, Alan R., Yuan, Lingyao (Ivy), 2017. Trading on Twitter: Using social media sentiment to predict stock returns. Decis. Sci. 48 (3), 454–488.
Sun, Licheng, Najand, Mohammad, Shen, Jiancheng, 2016. Stock return predictability and investor sentiment: A high-frequency perspective. J. Bank. Financ. 73, 147–164.

Surowiecki, James, 2005. The Wisdom Of Crowds: Why the Many are Smarter than the Few and How Collective Wisdom Shapes Business, Economics, Society and Nations. Knopf Doubleday Publishing Group, New York, NY.

Tetlock, Paul C., 2007. Giving content to investor sentiment: The role of media in the stock market. J. Financ. 62 (3), 1139-1168.

Tetlock, Paul C., Saar-Tsechansky, Maytal, Macskassy, Sofus, 2008. More than words: Quantifying language to measure firms' fundamentals. J. Financ. 63 (3), 1437–1467.

Timmermann, Allan, 2006. Chapter 4 forecast combinations. In: Elliott, G., Granger, C.W.J., Timmermann, A. (Eds.), In: Handbook of Economic Forecasting, vol. 1, Elsevier, pp. 135–196, URL https://www.sciencedirect.com/science/article/pii/S1574070605010049.

Wang, Xiaoqian, Hyndman, Rob J., Li, Feng, Kang, Yanfei, 2023. Forecast combinations: An over 50-year review. Int. J. Forecast. 39, 1518–1547.

White, Halbert, 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. Econometrica 48 (3), 817-838.