

Sentiment during recessions

Diego García*

December 4, 2012

ABSTRACT

This paper studies the effect of sentiment on asset prices during the 20th century (1905 to 2005). As a proxy for sentiment, we use the fraction of positive and negative words in two columns of financial news from the New York Times. The main contribution of the paper is to show that, controlling for other well-known time-series patterns, the predictability of stock returns using news' content is concentrated in recessions. A one standard deviation shock to our news measure during recessions predicts a change in the conditional average return on the DJIA of 12 basis points over one day.

JEL classification: G01, G14.

Keywords: media content, stock returns, sentiment, recessions.

*Kenan-Flagler Business School, University of North Carolina at Chapel Hill. I would like to thank Malcolm Baker, Craig Carroll, Riccardo Colacito, Jennifer Conrad, Alex Edmans, Joey Engelberg, Francesca Gino, Mark Grinblatt, Michael Halling, Pab Jotikasthira, Patrick Kelly, Lisa Kramer, Øyvind Norli, Chris Parsons, Chris Roush, Francesco Sangiorgi, Günter Strobl, and Paul Tetlock for comments on an early draft, as well as seminar participants at UNC at Chapel Hill, the First International Moscow Finance Conference, the European Winter Finance Summit (2010), the 2011 FIRS conference, the Chinese University at Hong Kong, and the Central Bank of Serbia. The suggestions provided by the editor (Cam Harvey), an anonymous associate editor, and the referee were critical in shaping the paper. I would also like to thank David Ernsthausen for superb research support. Special thanks to the CDLA staff, Kirill Fesenko, Fred Stipe, and particularly Rita Van Duinen, for their help in the optical character recognition stage of this project.

1 Introduction

Shiller (2000) argues that the news media plays an important role in setting the stage for market moves and provoking them. His conjecture is that investors follow the printed word even though much of it is pure hype, suggesting that market sentiment is driven by news' content. More formally, Tetlock (2007) shows that the number of negative words in the "Abreast of the Market" column of the Wall Street Journal predicts stock returns at the daily frequency from 1984 to 1999.

This paper revisits both Shiller's conjecture and Tetlock's evidence by studying financial market news from the New York Times over the 1905 to 2005 period. The literature from psychology and economics, which we review below, suggests that investors' sensitivity to news is most pronounced when they are going through hard times. Our main result shows that the link between media content and Dow Jones Industrial Average (DJIA) returns is indeed concentrated in times of hardship, proxied empirically by NBER recessions. A one standard deviation change in our sentiment measure predicts a change in the daily average of the DJIA of 12 basis points during recessions, while the effect in expansions is only 3.5 basis points. The effect we uncover is robust to different normalizations of the media factors, volatility adjustments, outliers analysis, and controls for autocorrelation and other known determinants of stock returns.

We construct our proxy for market sentiment by counting the number of positive and negative words from two financial columns from the New York Times.¹ Both were published daily and covered general financial news — from stock market performance to industry news and macroeconomic events. Thus, they are natural candidates as gauges of the excitement and agitation in U.S. stock markets during the 20th century. It is important to emphasize that the columns we study contain the type of news that Shiller (2000) and Tetlock (2007) have in mind. Annual statements (Kogan et al., 2009; Loughran and McDonald, 2011), earnings announcements (Engelberg, 2008), or other firm-specific media coverage (Tetlock et al., 2008) are likely to contain new information. On the other hand, a daily description of the stock market is almost an opinion piece, speculating about why the market acted as it did over the recent past, and about what it may do in the near future.

For the majority of our sample the supply of news was much more concentrated than today. The two main media sources with regular coverage of business news were the Wall Street Journal and the New York Times. Thus, most investors would have read the columns that we study by the time the market opened. By studying a longer time series than Tetlock (2007), we not only have more statistical power, but we can also take advantage of the variation in the business

¹The columns were titled "Financial Markets" and "Topics in Wall Street" for roughly half of our sample period. The "Financial Markets" column was essentially a recount of major stock market moves the previous day. The "Topics in Wall Street" column (which used other names, such as "News, Comment and Incident in the Stock Exchange," "Financial and Business Sidelights of the Day," and eventually "Market Place") had a broader coverage of economic and financial news.

cycle and the content of a media outlet that was in most investors' hands every morning.

It is plausible that part of the effect we document is related to the arrival of new information. We find that the effect of news partially reverses over the following few trading days, which argues for a non-informational impact. The reversal is quantitatively large, in the sense that more than half of the initial drift disappears over four days. But one cannot underestimate the role of a newspaper, such as the New York Times, as an important channel of financial news during most of our sample period. While the time-series data we use do not allow us to directly disentangle the information versus sentiment hypotheses, we conduct indirect tests that seem to give more bite to the sentiment interpretation of our data.²

We first study whether financial reporting changes along the business cycle. The news in our sample have a “tag along” flavor, much as Shiller (2000) describes. They essentially report on the previous day's events, giving ex-post explanations about past asset price movements. In our data set, this is reflected in a strong correlation between media content on a given afternoon and stock returns on that day. Financial reporting does not appear to be related to the business cycle, so how journalists describe previous market movements does not drive our results.

We also document that the effect is stronger during weekends. News written on Saturdays and Sundays have a significant impact on Monday's stock returns. For a large part of our sample virtually all businesses, including the New York Times, were closed on Sundays.³ It is not clear what type of information journalists could have gathered over the weekend that would significantly move stock prices on Monday. On the other hand, circulation of newspapers is much higher on Sundays. Following DellaVigna and Pollet (2009), we interpret this evidence as supporting the hypothesis that traders tune out of their investments at the end of the week, catching up with the markets on Mondays.

To rule out the existence of an informational channel, we look at intraday data on the DJIA. If the New York Times columns contain information that did not make it into the previous day's closing prices, we should not be surprised about the predictability using close-to-close prices. On the other hand, if markets were weak-form efficient and quickly incorporated new information into prices, we should not expect any predictability after the open of the NYSE. Our analysis shows that our media measures can predict stock returns well after the NYSE opening (11am to close), which rules out alternative hypotheses based on new information in the newspapers that quickly makes it into prices.

Finally, we study the relationship between trading volume and media content. A behavioral story suggests that naive traders will react to extremely positive or negative media content. On

²To be able to separate the two, one must have some instrument that affects sentiment and not fundamentals. See Engelberg and Parsons (2011) and Dougal et al. (2012) for causal evidence on the link between media and asset pricing variables.

³Neilson (1973) discusses at length the state of business journalism in the first half of the 20th century. It was not until halfway through our sample period that the Monday edition of the Wall Street Journal was written on Sundays.

the other hand, if traders are fully rational and symmetrically informed, a public signal, such as that provided by a newspaper, should not affect trading volume. Providing further support for Tetlock’s (2007) evidence, we find that media content can indeed predict trading volume. Days with particularly optimistic or pessimistic media content are associated with high trading volume, which is at its lowest when media content is at its unconditional average.

While our analysis supports a behavioral story, it should be noted that some type of “rational expectations” explanation can also fit the data. One could interpret our evidence, for example, as supporting beauty contest theories (Morris and Shin, 2000), where coordination is achieved via the content of news stories. The difference in response to news can also be dependent on the state of the business cycle, as in Veronesi (1999). The fact that we find price reversals, and existing evidence on the media and stock returns from Shiller (2000) and Tetlock (2007) to Engelberg and Parsons (2011), makes it more natural to present our results using a behavioral alternative to the null.

Our analysis is predicated on the assumption that economic recessions correspond with times of heightened sensitivity to news. The psychology literature forcefully argues that emotions affect decision making, and information processing in particular. For example, Tiedens and Linton (2001) show that reliance on heuristic versus systematic processing varies with emotions. The literature also finds that anxiety, hope, and sadness are associated with a greater sense of uncertainty (Smith and Ellsworth, 1985; Ortony et al., 1988). Gino et al. (2009) show that anxiety makes agents more receptive to advice, even if this advice is bad.⁴ This literature shows that priming subjects into negative mood states changes their decision-making abilities. One can reasonably argue that in periods of expansion investors feel happy and optimistic, whereas during recessions they feel fearful and anxious. The job losses and uncertainty over the future that investors experience during recessions put the population at large in negative mood states.⁵ This evidence suggests that investors will use different decision-making rules in recessions than in expansions, being particularly sensitive to news in economic downturns.

Although the above experimental studies from psychology establish that moods and emotions affect individuals’ behavior, it is the behavioral economics literature that shows the extent to which sentiment can move aggregate quantities. For example, Hirshleifer and Shumway (2003) show that stock returns are affected by the weather across the world, and Edmans et al. (2007) associate the outcomes of sporting events, such as the World Cup, to drops in the stock market when the country loses a game (see Hirshleifer, 2001, for a survey on these topics). In discussing confidence and the Michigan Consumer Sentiment Index, Akerlof and Shiller (2009)

⁴See Forgas (1991) for an excellent survey of the earlier literature, and Bless et al. (1996), Forgas (1998), Park and Banaji (2000), and Lerner and Keltner (2000, 2001) for more recent work.

⁵It is worth noting that our classification of recessions, following NBER, was backdated for virtually our entire sample. While investors may not have gotten a formal NBER press release, clearly our classification seems the most reasonable one: NBER chose such dates because they were times of economic hardship.

state that “we conceive of the link between changes in confidence and changes in income as being especially large and critical when economies are going into a downturn, but not so important at other times.” Shiller (2000) highlights the potential role of the media in creating asset bubbles and triggering market crashes. Thus, both the existing empirical findings in finance and the experimental evidence from psychology suggest that human behavior is significantly different in times of anxiety and fear versus periods of prosperity and tranquility.

The paper contributes to the growing literature on the role and content of the media and its impact on asset prices and investor behavior. In the journalism literature, Bow (1980) argues that there were no predictive signs in the media prior to the 1929 stock market crash. Griggs (1963) gives a similar account in the context of the 1957 to 1958 recession. Norris and Bockelmann (2000) and Roush (2006) both provide extensive discussions on the role of the media since the beginning of the 20th century. Shiller (2000) discusses the 1929 and 1987 crashes in more detail.

The closest paper to ours is Tetlock (2007), who studies a column from the Wall Street Journal from 1984 to 1999. We show how the predictability he uncovers is much stronger during economic downturns. We also find that positive words help predict stock returns, whereas in his research only negative word counts have predictive power. Further, we show that the effect is particularly important over the weekend, when investors have a chance to read the news. Since Tetlock (2007), the literature has examined the cross-section of stock returns, other types of investor behavior, and news originating from sources other than media outlets.⁶ Our large sample provides an excellent laboratory in which to examine the predictability of a liquid market index using media content.

The paper is also related to the literature on investor sentiment (see Baker and Wurgler, 2007, for a survey). Our study is the first to show that sentiment can play a more important role during recessionary periods. Most of the sentiment indexes suggested by other authors have data requirements that restrict their implementation to the last 40 years.⁷ As a consequence, they are less likely to be able to take advantage of the frequency and severity of the business cycle that the U.S. experienced during the first half of the 20th century. Finally, another important advantage of our media-based measures is that they are available daily, and as such they can be used in high frequency studies.

The rest of the paper is structured as follows. Section 2 constructs the sentiment measures

⁶Standard databases of news start after 1980, so there is not much room for time-series research outside the 1980 to 2008 window. The only other pure time-series analysis on a wide index is Larkin and Ryan (2008), who study the effect of news feeds in intraday returns for the S&P500 index. For other related papers see Cutler et al. (1989), Huberman and Regev (2001), Klibanoff et al. (1998), Chan (2003), Kaniel et al. (2007), Schmitz (2007), Barber and Odean (2008), Engelberg (2008), Gaa (2008), Tetlock et al. (2008), Yuan (2008), Bhattacharya et al. (2009), Fang and Peress (2009), Griffin et al. (2011), Jegadeesh and Wu (2011), Tetlock (2011), and Solomon (2012).

⁷See Brown and Cliff (2004, 2005), Qiu and Welch (2006), and Lemmon and Portniaguina (2006) for recent research on sentiment measures and stock returns.

that we use in our study. Section 3 formally studies the relationship between daily returns and sentiment for the DJIA, and goes over robustness checks. Section 4 looks at the effect of the news on different days of the week, and studies intraday data and trading volume. Section 5 concludes.

2 The Data

The paper uses several sources of data. The first is stock return information. We collect the total return index for the DJIA from Williamson (2008).⁸ The Dow Jones Industrial index goes back to the turn of the 20th century, and thus allows us to have a metric of U.S. stock returns prior to the coverage in the more standard Center for Research in Security Prices (CRSP), which started in 1926. We note that during our sample period the Dow Jones Industrial index consisted of as few as 12 securities in 1905, and increased to 30 starting in 1928. We let R_t denote the log return on the DJIA on date t . We obtain business cycle information from the NBER. In Section 4 we use intraday data on the DJIA, provided by Global Financial Services, which are available as of 1933. Our last source of data relates to the novel measures of media content used in this paper, which we describe next.

The media content measures are constructed starting from the New York Times Article Archive, which goes back to the origins of the newspaper in 1851.⁹ This data set was built by scanning the full content of the New York Times newspaper. It is available to any subscriber of the New York Times, as well as to users of other media providers (i.e., ProQuest). To have a consistent set of articles that cover financial news, we focus on two columns that were published daily during this period: the “Financial Markets” column, and the “Topics in Wall Street” column. The “Topics in Wall Street” column ran daily under different titles (i.e., “Sidelights from Wall Street”, “Financial and Business Sidelights of the Day,” and “Market Place”) until the end of our sample period.¹⁰ The “Financial Markets” column stopped being published under this heading in the 1950s, although the New York Times continued to publish a column with the financial news for the day, which we use in our analysis. In this paper we study a total of 55,307 pdf files from the Article Archive that were associated with either of these two columns from January 1, 1905 through December 31, 2005.

Both columns under study were essentially summaries of the events on Wall Street during the

⁸Historical data are available free of charge from <http://www.djaverages.com/>, including the total return for the DJIA, but as of the end of 2009 this source did not include Saturday data. For this reason, we use the DJIA data from Williamson (2008); see <http://www.measuringworth.org/DJA/>. Exclusion of the Saturday data does not affect our results.

⁹The time-series of the data used in the paper are available on the journal’s website, at <http://www.afajof.org/supplements.asp>. See the Internet Appendix for further details.

¹⁰The format of the business news of the New York Times is very stable until the 1990s, when the proliferation of TV (CNBC) and the Internet made standard news sources change their formatting more frequently.

previous trading day. The “Financial Markets” column was somewhat shorter, around 700 words per day, versus 900 words for a typical “Topics in Wall Street.” The latter would typically be subdivided into multiple sections that described anything from particular companies or industries to commodities and general market conditions. The themes discussed in both columns were of a business nature, with a focus on financial matters. As such, they are good candidates to measure the content of financial news in the U.S. during the 20th century. Figures 1 and 2 present a sample of each of the columns.

To construct the media content measures, we transform the scanned images available from the New York Times Article Archive into text documents using “optical character recognition” (OCR). In particular, we use ABBYY software, the leading package in OCR processing, to convert the 55,307 image files into text files. A sample of the OCR output for the two columns in Figures 1 and 2 is included in the Appendix. Although the quality of the transcription is high, it is important to notice that the accuracy of OCR processing may be low for some files. The quality of the scanned images in the New York Times Article Archive is particularly low prior to 1905, thus our choice of starting date. The text samples in the Appendix contain a few typographical errors, all stemming from problems in the original scanned image. This adds random noise to our media content measures, and thus does not bias our conclusions.

To quantify the content of the New York Times articles, this paper takes a “dictionary approach.”¹¹ For each column i written on date t , we count the number of positive words, g_{it} , and negative words, b_{it} , using the word dictionaries provided by Bill McDonald.¹² As argued in Loughran and McDonald (2011), standard dictionaries fail to account for the nuances of finance jargon, and thus the categorization we use has particular merits for processing articles on financial events. We let w_{it} denote the total number of words in an article. We construct the media measures dating them to the day t in which they were written, with the understanding that they are published in the morning of day $t + 1$. The rationale is that the information contained in these columns clearly belongs to date t . The writing process for each article started at 2 to 3pm, typically just as the market was about to close, and the final copy was turned in for editing and typesetting at around 5 to 6pm.

We aggregate the media content measures to create a time series that matches the Dow Jones index return data available. The ultimate goal is to combine all the news printed before the market opened, and then examine whether the content of such news, our proxy for sentiment, can predict the following day’s stock returns. In essence, we are trying to measure the content of the financial news on investors’ desks prior to the opening of the market. The two columns of the

¹¹Non-dictionary approaches have gained much popularity in recent research on text content analysis, in which not just the words but their order and their role in a sentence are taken into account (i.e., the Diction software used in Demers and Vega, 2008). Given the OCR processing issues discussed above, these types of language processing algorithms are not appropriate for our study.

¹²See http://www.nd.edu/~mcdonald/Word_Lists.html for details.

New York Times would typically be published the day after the market closed, and they would discuss financial events related to that day. For example, the Sunday edition would discuss Wall Street events from Saturday. On some occasions, the New York Times would not print these columns on Sunday, but on Monday.

To aggregate the news, and not miss columns that appeared while the market was not open, we average the measures of positive and negative content from articles written between the market close and the next market open. When the market is open on consecutive days, t and $t + 1$, we define our daily measure of positive media content as $G_t = \sum_i g_{it} / \sum_i w_{it}$, where the summation is over all articles written on date t (given our news selection, there are two such articles for the majority of days in our sample). Similarly, we construct our daily measure of negative media content as $B_t = \sum_i b_{it} / \sum_i w_{it}$. In essence, we count the number of positive and negative words in the financial news under consideration, and normalize these counts by the total number of words. For nonconsecutive market days we follow a similar approach, including all articles published from close to open. To be more precise, consider two trading days t and $t+h+1$ such that $h > 0$, where the market was closed h days, from $t+1$ through $t+h$. We define the positive media content measure as $G_t = \sum_{i,s=t}^{s=t+h} g_{is} / \sum_{i,s=t}^{s=t+h} w_{sh}$. We proceed analogously for the negative media content variable to define $B_t = \sum_{i,s=t}^{s=t+h} b_{is} / \sum_{i,s=t}^{s=t+h} w_{sh}$. We define the pessimism factor as the difference between the negative and positive media content measures, that is, $P_t = B_t - G_t$.

For consecutive trading dates, our media measures G_t and B_t are constructed using information available as of the end of date t when the market is open on date $t + 1$ (the bulk of our sample). It is less clear whether market prices on date t reflected the information available to the journalists writing the columns, as the deadline for turning the article in to the editor was not until roughly 5 to 6pm, while the NYSE closed at 3 to 4pm. We further remark that for nonconsecutive trading dates, we use articles that may have been written on days after date t but prior to the market opening (i.e., in the case of holidays). The New York Times measures we use could contain information that the market would not have as of the close of trade. Furthermore, as highlighted in Tetlock (2007), it is not even clear to what extent our sentiment measures may lead/lag stock returns (see Figure 1 in his paper, and its discussion). We work along the most natural time dimension: we stamp the writing after market closure but before market opening and look at returns the following trading day using close-to-close, and post-open-to-close.¹³

Table 1 presents sample statistics on our media measures. Panel A shows that, over the entire sample period, the average number of positive words in an article was 1.20% of the total. Thus, given a typical “Financial Markets” column with 700 words, there were on average 8.4 positive words in each article. The standard deviation of the positive words measure is 0.42%.

¹³In Section 4 we study returns from 11am to close. Using open-to-close returns yields virtually identical results to using close-to-close returns.

The average number of negative words over the entire sample is 2.06%, with a standard deviation of 0.67%. The pessimism measure, as expected given the numbers just discussed, has a positive mean, that is, a typical article has about 0.86% more negative than positive words.

Panels B and C of Table 1 present the sample statistics broken down by business cycle. The average positive measure is slightly higher during expansions, by six basis points. On the other hand, the average negative measure is four basis points higher during recessions. The boxplots in Figure 3, which graphically illustrate the content of the two bottom panels of Table 1, show that our negative and pessimism media measures are different during recessions and booms, as one would expect. More importantly for our purposes, there is a large amount of variation within business cycles: the volatility of the measures is an order of magnitude larger than the mean differences across the business cycle.

In the rest of the paper we normalize our sentiment measures so they have zero mean and unit variance. This allows us to interpret the regression coefficients in terms of one standard deviation shocks to the sentiment measures, thus making it easier to gauge the economic magnitude of our results.

Table 2 gives summary statistics on DJIA returns. Panel A shows that the mean return on the DJIA was two basis points per day over the 1905 to 2005 period, with a daily volatility of 107 basis points. During recessions, which comprise 6,467 days out of the total 27,449 trading days in our sample, the return was -5.3 basis points, whereas during expansions it was 4.2. Equally important is the difference in the volatility across the business cycle. During recessions the daily volatility was 141 basis points, whereas it was only 94 basis points during expansions.

Panel B of Table 2 presents the estimates of a parsimonious time-series model, which we augment with our media variables in the next sections. We estimate the following model of stock returns:

$$R_t = (1 - D_t)\gamma_1\mathcal{L}_s(R_t) + D_t\gamma_2\mathcal{L}_s(R_t) + \boldsymbol{\eta}X_t + \epsilon_t, \quad (1)$$

where \mathcal{L}_s denotes an s -lag operator,¹⁴ D_t is a dummy variable that takes on the value one if and only if date t is during a recession, the vector X_t denotes a set of exogenous variables, and ϵ_t is a zero-mean error term with possibly time-varying volatility. The set of exogenous variables X_t includes a constant term, day-of-the-week dummies, and a dummy for whether date t belongs to a recession or an expansion. We estimate the specification in (1) letting the lag operators have $s = 5$. We report White (1980) heteroskedasticity-robust standard errors.

Starting with the first five rows of Panel B, we see that there is a statistically significant positive autocorrelation in the returns of the DJIA during expansions, but not during recessions. Thus, if autocorrelation is a concern for our results, the evidence from Panel B indicates that it should be a more serious concern during expansions than during recessions. The last three rows

¹⁴For an arbitrary random process Y_t , $\mathcal{L}_s(Y_t) = \{Y_{t-1}, \dots, Y_{t-s}\}$.

of Panel B show that there is a strong Monday effect. Returns on Monday are on the order of 13 to 19 basis points lower than the returns on other days of the week.

In some of our analysis we attempt to deal with the time variation in volatility that is present in our data. To do so, we fit a GARCH(1,1) model to the DJIA returns. In particular, we estimate a model with a constant mean, $R_t = \mu + \epsilon_t$, and time-varying volatility $\sigma_{t+1}^2 = \omega + \alpha_1 \epsilon_t^2 + \beta_1 \sigma_t^2$, where $\sigma_t^2 \equiv \text{var}(\epsilon_t)$. The estimated coefficients for the variance equation are given in Panel C of Table 2. As expected, there is strong evidence of time-varying volatility.

3 Sentiment and DJIA Returns

To formally analyze the relationship between stock returns and news measures, we postulate the following model for stock returns:

$$R_t = \beta \mathcal{L}_s(M_t) + \gamma \mathcal{L}_s(R_t) + \psi \mathcal{L}_s(R_t^2) + \eta X_t + \epsilon_t, \quad (2)$$

where M_t denotes one of our media measures, that is, $M_t = G_t$ in the case of positive news, $M_t = B_t$ in the case of negative news, and $M_t = B_t - G_t$ in the case of our pessimism factor. Throughout the paper, we set $s = 5$ for our lag operators.¹⁵ As the set of exogenous variables X_t we include a constant term, day-of-the-week dummies, as well as a dummy for whether date t belongs to a recession or an expansion, D_t .

Table 3 presents the point estimates β from the specification in (2), with White (1980) t -statistics. From the column titled “Positive,” we see that the fraction of positive words in the morning’s media helps predict stock returns on a given day. The effect is statistically different from zero at standard confidence levels. Turning to the column titled “Negative,” we see a similar pattern: more negative words predict lower stock returns. The magnitude of the effect is along the lines of what Tetlock (2007) reports — a one standard deviation shock to the “Pessimism” metric, for example, moves the conditional average return on the DJIA by 5.5 basis points (3.9 and -4.3 for the positive and negative media metrics).

The coefficients on the lagged media variable, from $t - 2$ through $t - 5$, give an indication as to whether the shock to stock returns caused by the media content written on date $t - 1$ is permanent or temporary. For example, in the case of the pessimism factor, we see that the loadings on all four lags, $t - 2$ through $t - 5$, are positive, suggesting at least part of the drop in asset prices associated with M_{t-1} reverses over the following four days. Panel B formally conducts a test of whether the sum of these coefficients is different than zero. Although the statistical power is not as high as the test on a single coefficient, the p -values in Table 3 suggest that there is a reversal.

¹⁵The choice of lags and controls in specification (2) does not affect any of the conclusions of the paper.

The regression in (2) is essentially the same model as estimated by Tetlock (2007), so the results in this section serve as an independent confirmation of his findings using more powerful tests. In particular, our evidence corroborates the fact that media content can indeed predict the returns on a market index, and that the effect of such predictability partially reverses over the following four days. We note that the finding that the positive word counts help predict stock returns is a novel result: Tetlock (2007) focuses on negative word counts for most of his analysis due to the lack of predictability using positive words (and other word categories). Following his lead, the literature has mostly ignored dictionaries of positive words. The evidence in Table 3 suggests that positive words can be as important as negative words.

To differentiate the effect of news content on stock returns along the business cycle, the focus of our paper, we estimate the following model:

$$R_t = (1 - D_t) (\beta_1 \mathcal{L}_s(M_t) + \gamma_1 \mathcal{L}_s(R_t) + \psi_1 \mathcal{L}_s(R_t^2)) + D_t (\beta_2 \mathcal{L}_s(M_t) + \gamma_2 \mathcal{L}_s(R_t) + \psi_2 \mathcal{L}_s(R_t^2)) + \eta X_t + \epsilon_t. \quad (3)$$

In the above specification, M_t denotes one of our media measures, and D_t is a dummy variable that takes on the value one if and only if date t is during a recession. The vector X_t denotes a set of exogenous variables, and ϵ_t is a zero-mean error term with possibly time-varying volatility. As the set of exogenous variables X_t we use a constant term, day-of-the-week dummies, and a dummy for whether date t belongs to a recession or an expansion.

Table 4 presents the estimates of the coefficients on the media variable M_t from (3). Panel A includes the coefficient estimates β_1 , which measure the effect of media content on stock returns during expansionary periods, roughly two-thirds of our sample. There is evidence of predictability for both the positive and the negative word counts, and as a consequence for the pessimism factor as well. The magnitude of the effect is nonetheless small — a one standard deviation change in the pessimism factor moves the DJIA by 3.5 basis points.

Panel B presents the point estimates of β_2 , which measure the effect of our news measures on stock returns during recessions. The point estimate on the positive news measure is 0.085, with a t -statistic of 3.4. Thus, a one standard deviation change in the counts of positive words from the two columns of the New York Times increases the returns in the DJIA by 8.5 basis points during recessions. The effect is both statistically and economically meaningful: positive word counts have a much more important effect during recessions than during expansions. A similar pattern obtains for the negative and pessimism media measures: while a one standard deviation shock to the pessimism factor (negative news metric) moves the DJIA by 3.5 (2.8) basis points during expansions, the effect is 12 (8.7) basis points during recessions. The first row of Panel C formally tests the differences in the coefficients in Panels A and B, and concludes that they are statistically different. As importantly, the economic magnitude of the differences is large: the

point estimates in Panel B are anywhere from three to four times bigger than in Panel A.

By any measure, our sentiment proxies help predict stock returns the following day, with similar magnitudes as those reported by Tetlock (2007) during expansions, and significantly larger magnitudes during recessions. There is some evidence of return reversals in Panels A and B, as the sum of the coefficients on lags $t - 2$ through $t - 5$ partially swamp the initial effect. The second and third rows in Panel C conduct formal F -tests. The second row shows that the sum of these four lags is different than zero during expansions (p -value 6.6% for the pessimism factor), but the statistical power, given the evidence from Table 3, is small. The F -test in the last row of Panel C formalizes the fact that we cannot reject the null hypothesis of no reversal during recessions. While the sum of the effects of lags $t - 2$ through $t - 5$ of the pessimism factor is about six basis points (half of the coefficient at $t - 1$), the F -test only yields a p -value of 13.2%.

We next consider the robustness of the results from Table 4. One potential issue is the fact that we do not explicitly model time-varying volatility in (3). Although our standard errors are adjusted, one could be concerned that periods of high volatility affect our results. To address this concern we use the estimated GARCH(1,1) model from Table 2. If we let the estimated daily volatility from the GARCH model be $\hat{\sigma}_t$, we normalize the returns of the DJIA by replacing R_t in (3) with $R_t/\hat{\sigma}_t$. This normalization essentially constructs a time series of stock returns with their volatility normalized to unity.

Panel A in Table 5 presents the estimates of the leading terms from the model in (3) using the unit-variance DJIA index returns. We observe a similar pattern as in Table 4. The effect of the positive metric is 5.1 basis points during recessions, but only 2.2 basis points during expansions. The point estimates for the negative metric are -7.0 during recessions and -2.5 during expansions. By scaling down the DJIA returns during recessions we do find slightly smaller coefficients than in Table 4, but the conclusions stand. In particular, the tests using the pessimism factor conclude that the coefficients in recessions and expansions are different, and that the effect only partially reverses.

The skeptical reader may also be concerned with the autocorrelation controls, particularly given the relationship between the content of a column written on date t and the stock returns on date t reported in Tetlock (2007). To address this issue, we replace our media measure M_t with the residuals from the following model:

$$\begin{aligned}
 M_t = & (1 - D_t) (\lambda_1 R_t + \beta_1 \mathcal{L}_s(R_t) + \gamma_1 \mathcal{L}_s(M_t)) \\
 & + D_t (\lambda_2 R_t + \beta_2 \mathcal{L}_s(R_t) + \gamma_2 \mathcal{L}_s(M_t)) + \boldsymbol{\eta} X_t + v_t.
 \end{aligned} \tag{4}$$

Thus, we strip the media content metrics of any linear relationship with returns, day-of-the-week effects, or their own lags. Panel B in Table 5 presents the leading terms after estimating (3)

replacing M_t with this orthogonalized media measure. As with the GARCH adjustment, we find somewhat smaller coefficients on our media variables, but the magnitudes of the point estimates during recessions are two to three times as large as during expansionary periods.

One further concern is the possibility that outliers may be driving our results, given the fat tails that characterize stock returns. Panel C in Table 5 presents the estimates of the model in (3) fitted by robust regression using the M estimator of Huber (1981). This technique, standard in statistics to deal with outliers, essentially caps the influence that single observations can have on the point estimates.¹⁶ The values reported in Panel C show that the effect reported in Table 4 is indeed robust. The magnitude of the point estimates is very similar to that reported in Table 4, and the statistical significance is actually even higher, which suggests that if anything the outliers are working against our findings.

We conclude this section by presenting two more robustness tests.¹⁷ The first addresses the concern of running time-series regressions on a data set that covers over 100 years. Although we adjust for heteroskedasticity in our standard errors, one could conjecture that regression coefficients change significantly, as U.S. markets changed dramatically during the 1905 to 2005 period. We estimate (2) separately during each business cycle, dropping the D_t variable and its interactions. This boils down to estimating (2) with business cycle fixed effects and interactions of these business cycle fixed effects with all the independent variables.

Figure 4 plots the estimates of the leading coefficients from our media variables for each business cycle in our sample. In particular, for each recession and for each expansion, we estimate (2) and plot the leading coefficient on the media variable M_{t-1} . Estimates from the 20 different recessionary periods are marked with an x. Estimates for the 21 expansionary periods are marked with crosses. The dashed lines give the time-series averages of the recessionary and expansionary coefficients. Focusing on the right figure in the graph, we see that in 18 out of the 20 recessions the pessimism factor loads with a negative sign. Further, the largest five coefficients all occur during recessions. It is also important to note that the effect does not seem to vary over time — there is significant predictability throughout the whole 20th century. Four of the largest five point estimates for the pessimism factor fall in the second half of the sample, indicating that the effect has not dissipated with the advent of new types of media.¹⁸

We also study nonparametrically the relationship between the DJIA index returns and our media measures. Figure 5 plots a nonlinear estimate of the relationship between stock returns and our media measures. In particular, the graphs are lowess (locally weighted scatterplot smoothing) plots of the residuals of a time series regression as in (3) (dropping the media variables) against each of our media variables. The solid line presents the estimates during

¹⁶For the technically inclined reader, we use the routine `r1m` from the CRAN depositories for the estimation.

¹⁷For the interested reader, the Internet Appendix includes further analysis.

¹⁸Broadcast television started in the U.S. in 1928. It reached a wide audience by the 1950s, halfway through our sample period.

expansions, whereas the dashed line corresponds to recessions. The nonparametric estimates also show that the effect of the media variables on stock returns is concentrated during recessions. We note that the linear approximation is a very good one for all media metrics during recessions. It is interesting to see not just the difference in magnitude of the slopes, but the fact that during expansions the relationship is not that robust: for the pessimism metric, for example, the slope is negative only for half of the range in Figure 5.

Overall, the data strongly support the OLS evidence from Table 4. Normalizing the stock returns and media measures in different ways yields similar conclusions. Moreover, the results are even stronger if we control for outliers. Finally, the estimates of the coefficients for each business cycle and our nonparametric analysis both show that the effect is significantly stronger in recessions.

4 Information Versus Sentiment

The asymmetric response of stock returns to financial news across the business cycle is the main finding of the paper. This result is consistent with a story in which media content proxies for investor sentiment (i.e., noise traders), with this sentiment more salient during recessions. The psychology literature discussed in the introduction suggests that reactions to news will be more pronounced during periods of anxiety and fear, that is, during economic downturns. While we believe that one of the key advantages of the media measures we construct is that they are unlikely to be related to fundamental information not possessed by traders, as argued in Shiller (2000) and Tetlock (2007), a skeptical reader may interpret the counts of positive and negative words from the New York Times as new information. The question of whether sentiment or information is behind our results is the focus of this section.

The alternative explanation for our results starts by assuming that journalists had information that was not impounded in closing prices. Given that the market closed a few hours before they had to turn in their story, it is possible that New York Times reporters would be ahead of traders. Our findings across the business cycle would then be consistent with information production by financial intermediaries, the New York Times journalists, generating more precise signals during recessions.

We first study variation in the media content itself along the business cycle. Differences in reporting style would be an indication of media processing information asymmetrically during recessions and expansions. In particular, we study the effect of the DJIA returns on the sentiment measures we constructed estimating the model in (4). We note that the system in (3) and (4) is not a VAR in a strict sense, since we postulate a contemporaneous term in (4).

Panel A from Table 6 presents the estimated coefficients, which measure the reaction of news content to the raw stock returns of the DJIA. As expected, stock returns are indeed important

predictors of the media content variables, both the positive and the negative measures. Positive returns increase the number of positive words and decrease the number of negative words, and as a consequence decrease the pessimism measure. Given the daily standard deviation of the DJIA over our entire sample period is 107 basis points, a one standard deviation increase in stock returns increases the percentage of positive words on the articles written that day during expansions by 0.36 standard deviations, and decreases the percentage of negative words by a similar amount.¹⁹ The pessimism factor decreases by 0.44 standard deviations for a one standard deviation increase in the DJIA returns during expansions. The effect is also persistent, as the second lag of returns also has significant predictive power on the media content variables.

The last five rows of Panel A show that the feedback effect is smaller during recessions. The leading coefficients are some 50% smaller than in the first five rows. Thus, it appears that journalists use more “signed words,” or “tag along” more, when writing about the day’s events on Wall Street during expansions than during recessions. As the last row in Panel A shows, a formal test of equality of the coefficients can easily reject the null that the reaction of news to stock returns is equal along the business cycle ($\lambda_1 = \lambda_2$). This last conclusion should be conditioned, however, by the fact that the volatility of stock returns is significantly higher during recessions than during expansions. If we use the standard deviations from Table 2, we see that the effect of a one standard deviation increase in DJIA returns during expansions scales down the unconditional point estimates by almost 15%, whereas a similar movement during recessions scales up the point estimates by 30%. Thus, the effect of a one standard deviation movement on stock returns, conditional on the state of the business cycle, appears to be similar in magnitude.

To formalize this statement, we let the estimated daily volatility from the GARCH model from Table 2, Panel C, be $\hat{\sigma}_t$. We then normalize the returns of the DJIA by replacing R_t in (4) with $R_t/\hat{\sigma}_t$. This normalization essentially constructs a time series of stock returns with their volatility normalized to unity. Table 6 gives the point estimates when the returns used in (4) are homoskedastic. The effect on the media variables of a one standard deviation movement in stock returns, reported in Panel B of Table 6, is identical during recessions and expansions for both the positive and the negative word counts. Journalists “tag along” quite a bit, in the sense that their words are quite predictable given the previous days’ stock returns. More importantly for our purposes, there does not seem to be variation in this reporting along the business cycle. Thus, to the extent that information reflected in the writing of the New York Times columnists is a function of market conditions, such information does not seem to vary in recessions and

¹⁹The marginal effects we find in Table 6 are an order of magnitude larger than those reported in Tetlock (2007), as we include a contemporaneous term in (4). The inclusion of such a term has to do with the nature of the data: the columns in the New York Times were finished once the market was closed, so it is natural to think that the market return on that day would have an effect on the news content. Our empirical results show that this is indeed the case — although lags one through four have an effect on our media measures, as in Tetlock (2007), the return on the day the columns were written is the biggest determinant of the tone used by journalists.

expansions.

If journalists were producing informative signals for traders, it is not clear why the precision of these signals would increase during recessions. During economic downturns the press is hit particularly hard, as both subscriptions and advertising revenues are highly procyclical. For example, during the Great Depression subscriptions to the Wall Street Journal dropped from 52,000 to 28,000 (see p. 60 in Roush, 2006). It is unlikely that better coverage of financial markets would accompany staff cuts. Even if journalists were producing higher quality signals during recessions, it is also hard to explain why there is virtually no predictability during expansions.

We next divide our sample based on whether the previous day was a trading day. Our goal is to study the effect of media on DJIA returns on different days of the week, in particular on Mondays and holidays versus on other weekdays. There are different motivations for this analysis. On the one hand, there are distinct levels of information production in weekdays versus weekends. During our sample period most businesses were closed on weekends, and hence it is natural to conjecture that information flow would be more intense during weekdays.²⁰ On the other hand, newspapers have significantly higher circulation on weekends, as investors have more time to read the news. We conjecture that investor attention to business news is concentrated on Mondays, as investors may be less attentive on Fridays (DellaVigna and Pollet, 2009).

More specifically, we define an indicator variable I_t that takes the value one if and only if the previous day is not a trading day (i.e., Mondays or days after holidays). We then estimate the model

$$\begin{aligned}
 R_t = & (1 - I_t) [(1 - D_t) (\beta_1 \mathcal{L}_s(M_t) + \gamma_1 \mathcal{L}_s(R_t) + \psi_1 \mathcal{L}_s(R_t^2)) + D_t (\beta_3 \mathcal{L}_s(M_t) + \gamma_3 \mathcal{L}_s(R_t) + \psi_3 \mathcal{L}_s(R_t^2))] \\
 & + I_t [(1 - D_t) (\beta_2 \mathcal{L}_s(M_t) + \gamma_2 \mathcal{L}_s(R_t) + \psi_2 \mathcal{L}_s(R_t^2)) + D_t (\beta_4 \mathcal{L}_s(M_t) + \gamma_4 \mathcal{L}_s(R_t) + \psi_4 \mathcal{L}_s(R_t^2))] \\
 & + \eta X_t + \epsilon_t,
 \end{aligned} \tag{5}$$

which boils down to (3) with all the coefficients interacted with the first trading day after a holiday indicator I_t . Table 7 gives the estimates of the leading coefficients on the media variables. During expansions we find that the predictability from Table 4 is concentrated on Mondays or days after holidays. The point estimates on such days are 5.6, -6.2 , and -7.9 basis points for the positive, negative, and pessimism factors, all highly significant. On weekdays the effect is much smaller; a formal test of differences of coefficients in the fifth row of Table 4 confirms this. During recessions we find an even larger effect for Mondays or days after holidays: the point estimate for the pessimism factor is -26.7 basis points, a very large effect given that we are studying daily stock returns. During weekdays, this effect is also statistically significant, with a point estimate of -8.5 basis points for the pessimism factor, still large in economic terms.

²⁰It should be noted that business writing on Sundays, for Monday's edition, did not start until halfway through our sample period at the New York Times (Neilson, 1973). Casual empiricism suggests that both businessmen and journalists work less hours during holidays.

Our previous conclusions are still unaltered — the effect of the media on the stock market is significantly larger during recessions than expansions, with the point estimates conditional on the day of the week anywhere from two to four times larger during recessionary periods. Although part of the effect could be due to information released during the weekend, the asymmetry of the effect along the business cycle again suggests a sentiment story. We interpret the stronger effect on Mondays via higher attention paid by investor to news that come out on Sundays and Mondays. But we note that our main results also show up on weekdays: there is virtually no predictability during expansions (-2.4 basis point coefficient on the pessimism factor), whereas the point estimates during recessions are both statistically significant and economically large (-8.5 basis points).

If liquidity varied along the business cycle, an informational story would yield predictions that are consistent with our main findings. Microstructure frictions, such as lower market depth during recessions, could potentially be driving part of our results. But since we are studying the most liquid sliver of the U.S. stock market, it is unclear why liquidity would exhibit the strong day-of-the-week effect we document in Table 7.

The more skeptical reader can view the tests on weekends as a robustness check. As mentioned previously, to the extent that there may be new information that journalists can include in their columns while the market was closed, it is possible that our study is contaminated by such inclusion. Holidays aside, this may be important for articles written on Saturdays, as the market closed at noon during the first half of our sample. It is possible that information that arrived in the afternoon of Saturday or during holidays could drive our results. Confirming that the effect is not just concentrated on weekends can be viewed as a robustness check on our main results.

We tackle this potential informational channel directly in our next set of tests. In particular, we look at the returns of the DJIA after the opening of the NYSE. If information is driving our results, and markets process this information quickly, we should find no predictability when looking at returns after the open. Table 8 presents the point estimates of a specification such as (3) where we replace the close-to-close DJIA returns by the returns from 11am to close.²¹

Table 8 presents the point estimates. We first see that the positive word counts have no predictive power at all during expansions, whereas they do help predict DJIA post-opening returns during recessions. The point estimate on positive words is smaller than in Table 4, 5.7 basis points versus 8.5, but still economically and statistically large. Similar conclusions arise from the negative word counts and as a consequence the pessimism factor. While this analysis

²¹Tetlock (2007) looks at returns from 10am to close, since during his sample period the NYSE opened at 9:30am. Trading hours during our sample period were 10am to 3pm weekdays and on Saturdays 10am to noon (1905–1952), 10am to 3.30pm (1952–1974), 10am to 4pm (1974–1985), 9.30am to 4pm (1985–2005); see http://www.nyse.com/pdfs/historical_trading_hours.pdf. The data provided by Global Financial Data have DJIA levels at hourly intervals during our sample period.

does not rule out an alternative hypothesis where information slowly diffuses into prices, it does dismiss informational theories where prices fully and quickly adjust to new data.

We conclude our analysis by studying the relationship between trading volume and our media metrics. The information hypothesis does not make obvious predictions about trading volume, whereas most versions of the sentiment hypothesis would predict that noise traders and rational traders disagree when the tone of news is extremely positive or negative. Disagreement in standard models drives volume, and thus a behavioral story in which noise traders “follow the printed word” naturally generates more volume. On the other hand, if we treat the content of news as a public signal, it should not generate disagreement among rational traders, and thus trading volume should not be affected by the content of the media. This is to be expected in models where all agents are symmetrically informed.²² Our next set of tests study NYSE aggregate trading volume in order to potentially rule out some of these theories.

For each year, we start by estimating the following model of trading volume:

$$V_t = \beta \mathcal{L}_s(V_t) + \gamma X_t + \epsilon_t, \quad (6)$$

where V_t denotes the natural logarithm of aggregate NYSE trading volume, and X_t includes a constant, a set of monthly dummies (one for each month but January), and a set of day-of-the-week dummies. We let the lag operator have $s = 5$ in (6). We normalize the trading volume by letting $\hat{V}_t = \hat{\epsilon}_t/s_\epsilon$ for each year, where the $\hat{\epsilon}_t$ are the estimate residuals from (6), and s_ϵ denotes their standard deviation. This procedure generates a white noise process (zero mean, unit variance) that removes low frequency components, as well as potential calendar effects, from the NYSE volume time series, which we do not attempt to explain with our media variables. Our procedure essentially follows Gallant et al. (1992), taking care of the time trends in trading volume via year fixed effects.

Following the nonlinear relationship between volume and media content reported in Tetlock (2007), we next estimate the following model:

$$\begin{aligned} \hat{V}_t = & (1 - D_t) (\beta_{1+} M_t^+ + \beta_{1-} M_t^- + \psi_1 R_t + \gamma_1 \mathcal{L}_s(R_t) + \eta_1 \mathcal{L}_s(R_t^2)) \\ & + D_t (\beta_{2+} M_t^+ + \beta_{2-} M_t^- + \psi_2 R_t + \gamma_2 \mathcal{L}_s(R_t) + \eta_2 \mathcal{L}_s(R_t^2)) + \epsilon_t. \end{aligned} \quad (7)$$

Table 9 presents the point estimates of (7). The first two rows present the effect of media content on volume during expansions, whereas rows 3 and 4 present the point estimates during recessions. Focusing on the columns that belong to the pessimism factor, we see that media content can indeed predict trading volume, in a rather nonlinear fashion. For example, dur-

²²Clearly the interpretation of the media content as a public signal uncorrelated with other economic shocks is important in this argument. If traders’ hedging needs are correlated with public signals, then there will be a nontrivial relationship between news and volume (i.e., Schneider, 2009). With asymmetrically informed agents, one should also expect a nonlinear relationship between trading volume and media content (Wang, 1993).

ing expansions, a one standard deviation increase in the pessimism factor, given pessimism is negative, moves trading volume by -0.12 standard deviations. The corresponding effect during expansions is -0.11 standard deviations. The effects are statistically strong throughout the different specifications presented in Table 9. More positive words and more negative words in financial news affect market-wide trading in a nonlinear way: when the columns are particularly positive or particularly negative trading volume peaks, with the lowest trading volume occurring on days with average media content.

Figure 6 presents a nonlinear (lowess) fit of the specification in (7), following the same procedure as in Figure 5. The evidence that we report in Table 9 is not dependent on the particular nonlinear specification that we use. The fit estimated in Figure 6 suggests that the effect we find is stronger during recessions than expansions. On the other hand, as the last two rows in Table 9 show, the fitted curves are not statistically different from each other, that is, the effect of the media on trading volume is not statistically different in recessions than in expansions.

5 Conclusion

This paper constructs a measure of sentiment based on financial news from the New York Times during 1905 to 2005, and studies its relationship to stock returns. The long time series allows us to study a period that had a fairly concentrated media sector, as well as plenty of variability in the business cycle. Our main finding is that news content helps predict stock returns at the daily frequency, particularly during recessions: while a one standard deviation change in our pessimism factor moves the DJIA by 12 basis points in recessions, the marginal effect during expansions is only 3.5 basis points.

We show that this asymmetric predictability is not driven by differences in reporting along the business cycle. The effect is especially strong on Mondays and on days after holidays, when investors have time to read the news, and it persists into the afternoon of the trading day. We also find that the effect partially reverses over the following four trading days. We conclude that our results support the hypothesis that investor sentiment has a prominent effect during bad times.

The New York Times Article Archive we use in this paper opens the door for other research questions. While our paper clearly documents differences in the impact of sentiment along the business cycle, it does not speak to the mechanism that drives the predictability. Finally, whether there are lower-frequency components to our sentiment measures is a natural avenue to explore, especially in connection to economic growth figures and long-run stock returns.

References

- Akerlof, George A., and Robert J. Shiller, 2009, *Animal Spirits* (Princeton University Press, Princeton).
- Baker, Malcolm, and Jeffrey Wurgler, 2007, Investor sentiment in the stock market, *Journal of Economic Perspectives* 21, 129–151.
- Barber, Brad, and Terry Odean, 2008, All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors, *Review of Financial Studies* 21, 785–818.
- Bhattacharya, Utpal, Neal Galpin, Xiaoyun Yu, and Rina Ray, 2009, The role of the media in the internet IPO bubble, *Journal of Financial and Quantitative Analysis* 44, 657–682.
- Bless, Herbert, Gerald L. Clore, Norbert Schwarz, Verena Golisano, Christina Rabe, and Marcus Wölk, 1996, Mood and the use of scripts: Does a happy mood really lead to mindlessness?, *Journal of Personality and Social Psychology* 71, 665–679.
- Bow, James, 1980, The “Times’s” financial markets column in the period around the 1929 crash, *Journalism Quarterly* 57, 447–450.
- Brown, Gregory W., and Michael T. Cliff, 2004, Investor sentiment and the near-term stock market, *Journal of Empirical Finance* 11, 1–27.
- Brown, Gregory W., and Michael T. Cliff, 2005, Investor sentiment and asset valuation, *Journal of Business* 78, 405–440.
- Chan, Wesley S, 2003, Stock price reaction to news and no-news: Drift and reversal after headlines, *Journal of Financial Economics* 70, 223–260.
- Cutler, David M., James M. Poterba, and Lawrence H. Summers, 1989, What moves stock prices?, *Journal of Portfolio Management* 15, 4–12.
- DellaVigna, Stefano, and Joshua M. Pollet, 2009, Investor inattention and Friday earnings announcements, *Journal of Finance* 64, 709–749.
- Demers, Elizabeth A., and Clara Vega, 2008, Soft information in earnings announcements: News or noise?, Technical report, INSEAD.
- Dougal, Casey, Joseph Engelberg, Diego García, and Chris Parsons, 2012, Journalists and the stock market, *Review of Financial Studies* 25, 639–679.
- Edmans, Alex, Diego García, and Øyvind Norli, 2007, Sports sentiment and stock returns, *Journal of Finance* 62, 1967–1998.
- Engelberg, Joseph, 2008, Costly information processing: Evidence from earnings announcements, Technical report, University of North Carolina.
- Engelberg, Joseph, and Chris Parsons, 2011, The causal impact of media in financial markets, *Journal of Finance* 66, 67–97.

- Fang, Lily H., and Joel Peress, 2009, Media coverage and the cross-section of stock returns, *Journal of Finance* 64, 2023–2052.
- Forgas, Joseph P., 1991, *Emotion and Social Judgments* (Pergamon Press, Oxford).
- Forgas, Joseph P., 1998, On being happy and mistaken: mood effects on the fundamental attribution error, *Journal of Personality and Social Psychology* 75, 318–331.
- Gaa, Charles, 2008, Good news is no news: Asymmetric inattention and the neglected firm effect, Technical report, University of British Columbia.
- Gallant, A. Ronald, Peter E. Rossi, and George Tauchen, 1992, Stock prices and volume, *Review of Financial Studies* 5, 199–242.
- Gino, Francesca, Alison Wood, and Maurice E. Schweitzer, 2009, How anxiety increases advice-taking (even when the advice is bad), Technical report, Wharton School of the University of Pennsylvania.
- Griffin, John M., Nicholas H. Hirschey, and Patrick J. Kelly, 2011, How important is the financial media in global markets?, *Review of Financial Studies* 24, 3941–3992.
- Griggs, Harry H., 1963, Newspaper performance in recession coverage, *Journalism Quarterly* 40, 559–564.
- Hirshleifer, David, 2001, Investor psychology and asset pricing, *Journal of Finance* 61, 1533–1597.
- Hirshleifer, David, and Tyler Shumway, 2003, Good day sunshine: Stock returns and the weather, *Journal of Finance* 58, 1009–1032.
- Huber, Peter J., 1981, *Robust Statistics* (Wiley, New York).
- Huberman, Gur, and Tomer Regev, 2001, Contagious speculation and a cure for cancer: A non-event that made stock prices soar, *Journal of Finance* 56, 387–396.
- Jegadeesh, Narasimham, and Di Wu, 2011, Word power: A new approach for content analysis, Technical report, Emory University.
- Kaniel, Ron, Laura T. Starks, and Vasudha Vasudevan, 2007, Headlines and bottom lines: Attention and learning effects from media coverage of mutual funds, Technical report, University of Texas.
- Klibanoff, Peter, Owen Lamont, and Thierry Wizman, 1998, Investor reaction to salient news in closed-end country funds, *Journal of Finance* 53, 673–699.
- Kogan, Shimon, Bryan R. Routledge, Jacob S. Sagi, and Noah Smith, 2009, Information content of public firm disclosures and the sarbanes-oxley act, Technical report, University of Texas at Austin.
- Larkin, Fiacc, and Conor Ryan, 2008, Good news: Using news feeds with genetic programming to predict stock prices, *Lecture Notes in Computer Science* 4971, 49–60.

- Lemmon, Michael, and Evgenia Portniaguina, 2006, Consumer confidence and asset prices: Some empirical evidence, *Review of Financial Studies* 19, 1499–1529.
- Lerner, Jennifer S., and Dacher Keltner, 2000, Beyond valence: Toward a model of emotion-specific influences on judgement and choice, *Cognition and Emotion* 14, 473–493.
- Lerner, Jennifer S., and Dacher Keltner, 2001, Fear, anger, and risk, *Journal of Personality and Social Psychology* 81, 146–159.
- Loughran, Tim, and Bill McDonald, 2011, When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks, *Journal of Finance* 66, 35–65.
- Morris, Stephen, and Hyun Song Shin, 2000, Global games: Theory and applications, Technical report, Yale University.
- Neilson, Winthrop, 1973, *What's News — Dow Jones* (Clinton Book Company, Radnor, PA).
- Norris, Floyd, and Christine Bockelmann, 2000, *The New York Times — Century of Business* (McGraw-Hill, New York).
- Ortony, Andrew, Gerald L. Clore, and Allan Collins, 1988, *The Cognitive Structure of Emotions* (Cambridge University Press, Cambridge, England).
- Park, Jaihyun, and Mahzarin R. Banaji, 2000, Mood and heuristics: The influence of happy and sad states on sensitivity and bias in stereotyping, *Journal of Personality and Social Psychology* 78, 1005–1023.
- Qiu, Lily Xiaoli, and Ivo Welch, 2006, Investor sentiment measures, Technical report, Brown University.
- Roush, Chris, 2006, *Profits and Losses* (Marion Street Press, Oak Park, IL).
- Schmitz, Philipp, 2007, Market and individual investors reactions to corporate news in the media, Technical report, University of Mannheim.
- Schneider, Jan, 2009, A rational expectations equilibrium with informative trading volume, *Journal of Finance* 64, 2783–2805.
- Shiller, Robert J., 2000, *Irrational Exuberance* (Princeton University Press, Princeton).
- Smith, Craig A., and Phoebe C. Ellsworth, 1985, Patterns of cognitive appraisal in emotion, *Journal of Personality and Social Psychology* 48, 813–838.
- Solomon, David, 2012, Selective publicity and stock prices, *Journal of Finance* 67, 599–637.
- Tetlock, Paul C., 2007, Giving content to investor sentiment: The role of media in the stock market, *Journal of Finance* 62, 1139–1168.
- Tetlock, Paul C., 2011, All the news that's fit to reprint: Do investors react to stale information?, *Review of Financial Studies* 24, 1481–1512.

- Tetlock, Paul C., Maytal Saar-Tsechansky, and Sofus Macskassy, 2008, More than words: Quantifying language to measure firms' fundamentals, *Journal of Finance* 63, 1437–1467.
- Tiedens, Larissa Z., and Susan Linton, 2001, Judgement under emotional certainty and uncertainty: The effects of specific emotions on information processing, *Journal of Personality and Social Psychology* 81, 973–988.
- Veronesi, Pietro, 1999, Stock market overreaction to bad news in good times: A rational expectations equilibrium model, *Review of Financial Studies* 12, 975–1007.
- Wang, Jiang, 1993, A model of intertemporal asset prices under asymmetric information, *Review of Economic Studies* 60, 249–282.
- White, Halbert, 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 48, 817–838.
- Williamson, Samuel H., 2008, Daily closing values of the DJIA in the United States, 1885 to present, Technical report, MeasuringWorth.
- Yuan, Yu, 2008, Attention and trading, Technical report, University of Iowa.

Appendix

This Appendix presents the output of the optical character recognition on two columns from the New York Times. The first one is the “Financial Markets” column from October 12, 1915, displayed in Figure 1. The second is the “Topics in Wall Street” column from June 25, 1916, displayed in Figure 2. Positive words, using the Loughran and McDonald (2011) dictionaries, are marked in italics, whereas negative words are marked in bold.

FINANCIAL MARKETS

October 12, 1915; pg. 14

Another Day of *Great* Activity, with Big *Gains* for the Industrials. Kxpccotations that yesterday’s Stock Kxohan’ge session, sandwiched in be-hvirn two holidays, would see much less activity were quickly **disappointed**. The market opened very active and *strong*, anr] afu-r a small reaction in the forenoon resumed its upward trend with almost the same **violence** shown in the *excited* sessions of last week. The list was again **irregular**, but by far the larger number of stocks scored substantial *gains* and the upward movement of some of the war issues, which had j been checked by banks and brokers who ’foiosaw **trouble** if the advance were not i held under control, was resumed with a *great* deal of visor. The most striking *gain* among such issues was scored by Baldwin Locomotive, which, after hang-infor several days around 113, returned yesterday to 127M>. **closing** at 12(i, with a net advance of 11 points.

This secondary stage of activity for Baldwin was accompanied by fresh merger rumors, which do not appear to have any substantial basis In fact. Even more active and relatively as *strong* was Westirighthouse. of which. more than WO.Udti shares changed hands — up a range of .”’i points. It **closed** at J 13.*!. with a *gain* of points above Saturday’s close. The American Car & Foundry made n *good* recovery to X.V-S. and *gains* of from t. ” points were numerous. The motor issues returned to *popularity*, all three classes of Maxwell stock *advancing* on the expectation of pom©kind of an announcement Wednesday of a plan looking to the payment of the accumulated dividend on the first preferred. Studebaker advanced 2]< and General Motors 1 point.

The rails retained some of their momentum from last week, and most of j the leaders sold at new high prices for : the. year. Xev.s of the note being prepared for dispatch to t treat Britain was received too **late** to affect the market, if indeed ûch news can have any ef-i feet on the present temper of traders, and the list **closed** pretty clofc to the top. j Some uneasiness was caused yesterday I by a new development of **weakness** in the foreign exchange market. Demand sterling, sold down to ?4.<<7% compared with the low price of Splits1., on Saturday. The **failure** of the conclusion of the So”>”,0(10,000 Anglo-French loan to help foreign exchange rates gave special interest to an important meeting of bankers held yesterday afternoon, which was addressed by Hr l-Zdward Holden, one of the visiting Commissioners.

TOPICS IN WALL STREET

June 25, 1916; pg. E6

American munition Orders.

Until yesterday the stock market gave no indication that the war stocks derived a chance of profit from war with Mexico. To speculators in these shares it was in fact a matter of the keenest **disappointment** that they went down on war news. Over and over they have repeated the **question**: "What sort of a war stock is it that is **depressed** by a new war?" Yesterday an advance of 17 points in Bethlehem Steel held out a ray of hope and *advances* in most of the others on covering by professionals *strengthened* hopes that the next turn would be for the *better*. Officers of many of the munitions companies expected orders from the United States Government in the near future, but nowhere was it believed that these orders would be placed at terms permitting as *great* profits as those obtained in some of the contracts with the Allies. **

The Extent of the **Declines**.

From the high point of week before last to the low point of last week, which was the low point of Friday's market, the average price of fifty representative stocks **declined**- 55.33 a share. These stocks included many railroad shares in which the **declines** were small compared with **losses** in some of the speculative industrials. Reading, which **lost** 8% points in this period, and Norfolk & Western, with a **loss** of 5%, were the only rails to **decline** more than the 51-3-point average of the fifty. A score of industrials sustained *greater losses* and many of these **losses** ran into double figures, among them being: New York Air Brake, 11; Mexican Petroleum, 12%; Baldwin Locomotive, 13; United States Smelting, 13; Tennessee Copper, W/y, American Zinc, 14%; Willys-Overland, 16; Butte and *Superior*, 10%; United States Industrial Alcohol, 2<% On the Curb Chevrolet Motors **lost** 46 points.

Sow Up, Now Down.

It is interesting to note the change in sentiment that sweeps over the floor of the Stock Exchange after a pronounced rise, or sharp **decline**. Traders who have been bearish for weeks were turning bullish yesterday morning. They figured that the **break** which had been needed had been supplied, and that, therefore, stocks were a purchase again. ***

The Mexican Fuetor.

An old-time member eaid after the close that neither the Mexican war **danger**, nor the **inadequacy** of our war machinery, was really back of the slump which took place last week. Those **arguments** were advanced to support the **decline**, but in his opinion the **break** would have come had the Mexican situation continued unchanged. This man's theory is that the market had become badly congested with stocks, and had to be cleaned out by a return to lower prices. A number of pools were carrying large amounts of stock which they had not been *able* to market, and there were some large individual accounts that needed shaking out. The low prices made on Friday brought in a number of fresh buyers, and if this trader's theory works out the market will develop a much *better* tone this week, regardless of developments across the border. When the list grows stale' nothing but a sharp **setback** will attract new money. That this market had become stale was evidenced by its utter **disregard** of *good* news, such as new and increased dividends. <<**

No Extra Holiday.

When the brokers gave up their expected extra holiday before May 3p, they looked for an extra day preceding the Fourth. The uncertainty of | the political situation appears to have **destroyed** any chance of getting it. No petition has been circulated on the floor, and it is unlikely that the situation will clear in time to allow the drafting of one before the next meeting of Governors. *>>*

Bonds) Have **Idle** Week.

The bond market **suffered** along with stocks last week, but without registering substantial **declines**. Bonds' were effected more through a let-down of buying than from the **liquidation** of securities. Some of the banks and large dealers were reported as sellers of a considerable amount of bonds which they had been carrying for a month or more, and on which they had *good* profits. If this actually did take place the offerings were rather *easily* absorbed, and inquiries among bond men **failed** to show that there had been any **urgent** selling through **fear** that the Mexican situation might wipe out profits before they could be realized. The investment demand is believed to be widening, now that supplies from Europe have begun to fall away, leaving room for other offerings, and the bankers are inclined to think that business will pick up again with the coming of definite developments south of the border.

Table 1
Sample Statistics for Media Content Variables during Recessions and Expansions

The table reports sample statistics for the media content measures used in the paper. These measures are constructed from the columns “Financial Markets” and “Topics in Wall-Street” published in the New York Times in the period 1905 to 2005. We construct the “Positive” and “Negative” measures by counting the number of positive and negative words and normalizing these counts by the total number of words of each article, using the Loughran and McDonald (2011) dictionaries. The “Pessimism” variable is the difference between the “Negative” and “Positive” measures. All numbers are given in percentages. Panel A presents the sample statistics for the entire sample period, which comprises 27,449 trading days. Panels B and C break the sample down by the business cycle. Panel B, which contains all trading days during a recession in the 1905 to 2005 period, has 6,467 observations, whereas Panel C, which contains the expansionary dates, has 20,722.

Media measure	Mean	Median	25%-quant.	75%-quant.	Stand. dev.
Panel A. All dates					
Positive	1.20	1.16	0.90	1.46	0.42
Negative	2.06	1.99	1.59	2.45	0.67
Pessimism	0.86	0.81	0.26	1.40	0.88
Panel B. Recessions					
Positive	1.15	1.12	0.88	1.38	0.39
Negative	2.09	2.04	1.64	2.48	0.64
Pessimism	0.94	0.90	0.38	1.46	0.84
Panel C. Expansions					
Positive	1.21	1.17	0.91	1.48	0.43
Negative	2.05	1.98	1.57	2.45	0.68
Pessimism	0.84	0.78	0.23	1.38	0.89

Table 2
Sample Statistics for Daily DJIA Returns, 1905–2005

The table reports sample statistics for the DJIA returns used in the paper. Panel A gives unconditional sample statistics for the daily log-returns of the DJIA for the period 1905 to 2005. The first row presents the sample statistics; the following two rows break the sample period into NBER recessions and expansions. Panel B reports the estimated coefficients from the model

$$R_t = (1 - D_t)\gamma_1\mathcal{L}_s(R_t) + D_t\gamma_2\mathcal{L}_s(R_t) + \eta X_t + \epsilon_t,$$

where \mathcal{L}_s denotes an s -lag operator, namely, $\mathcal{L}_s(R_t) = \{R_{t-1}, \dots, R_{t-s}\}$, and D_t is a dummy variable that takes on the value one if and only if date t is during a recession. We set $s = 5$ throughout the paper. As the set of exogenous variables X_t we include a constant term, day-of-the-week dummies, as well as a dummy for whether date t belongs to a recession or an expansion, D_t . Panel C presents the estimates of a GARCH(1,1) model, where we assume the return equation has a constant mean, $R_t = \mu + \epsilon_t$, but we allow for time-varying volatility of the form $\sigma_{t+1}^2 = \omega + \alpha_1\epsilon_t^2 + \beta_1\sigma_t^2$, where $\sigma_t^2 \equiv \text{var}(\epsilon_t)$. The sample period comprises 27,449 trading days, of which 6,467 were during recessions. The t -stats reported are computed using White (1980) standard errors.

Panel A. Sample statistics

	Mean	Median	25%-quant.	75%-quant.	Stand. dev.
All dates	0.020	0.044	-0.450	0.526	1.071
Expansions	0.042	0.056	-0.410	0.517	0.943
Recessions	-0.053	-0.011	-0.637	0.565	1.408

Panel B. Time-series regression

Expansions	γ_1	t -stat	Recessions	γ_2	t -stat
$(1 - D_t) \times R_{t-1}$	0.052	3.1	$D_t \times R_{t-1}$	0.024	0.8
$(1 - D_t) \times R_{t-2}$	-0.045	-2.7	$D_t \times R_{t-2}$	-0.019	-0.7
$(1 - D_t) \times R_{t-3}$	0.004	0.3	$D_t \times R_{t-3}$	0.004	0.2
$(1 - D_t) \times R_{t-4}$	0.005	0.5	$D_t \times R_{t-4}$	0.062	2.6
$(1 - D_t) \times R_{t-5}$	0.011	0.7	$D_t \times R_{t-5}$	0.022	0.9
	η	t -stat		η	t -stat
I_{Tue}	0.140	6.3	I_{Fri}	0.167	7.4
I_{Wed}	0.153	6.7	I_{Sat}	0.189	7.5
I_{Thu}	0.125	5.6	D_t	-0.091	-5.0

Panel C. GARCH(1,1) estimates

	$\omega, \alpha_1, \beta_1$	t -stat
Constant, ω	0.011	11.8
Innovations term, α_1	0.081	24.9
Autoregressive term, β_1	0.910	255.2

Table 3
Feedback from News Content to the DJIA

The table reports the estimated coefficients β from the model

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

The dependent variable R_t is the log-return on the DJIA from 1905 to 2005. The variable M_t is one of our media measures, that is, $M_t = G_t$ in the case of positive words, $M_t = B_t$ in the case of negative words, and $M_t = B_t - G_t$ in the case of our pessimism factor. The media measures are constructed from the columns “Financial Markets” and “Topics in Wall-Street” published in the New York Times during the period 1905 to 2005. The media variables are normalized to have unit variance. As the set of exogenous variables X_t we include a constant term, day-of-the-week dummies, as well as a dummy for whether date t belongs to a recession or an expansion, D_t . The sample period comprises 27,449 trading days, of which 6,467 were during recessions. The t -stats reported are computed using White (1980) standard errors.

Panel A. Media variables

	Positive		Negative		Pessimism	
	β	t -stat	β	t -stat	β	t -stat
M_{t-1}	0.039	5.2	-0.043	-5.2	-0.055	-6.3
M_{t-2}	0.003	0.4	0.003	0.3	0.001	0.2
M_{t-3}	-0.008	-1.1	0.005	0.7	0.008	1.0
M_{t-4}	-0.013	-1.8	0.008	1.0	0.013	1.6
M_{t-5}	-0.005	-0.6	0.009	1.2	0.010	1.3

Panel B. Tests

	Positive		Negative		Pessimism	
	F -stat	p -value	F -stat	p -value	F -stat	p -value
$\beta_1 = 0$	26.9	0.000	26.8	0.000	40.0	0.000
$\sum_{j=2}^5 \beta_j = 0$	3.1	0.077	3.6	0.059	5.6	0.018

Table 4
Feedback from News Content to the DJIA along the Business Cycle

The table reports the estimated coefficients β from the model

$$R_t = (1 - D_t) (\beta_1 \mathcal{L}_s(M_t) + \gamma_1 \mathcal{L}_s(R_t) + \psi_1 \mathcal{L}_s(R_t^2)) + D_t (\beta_2 \mathcal{L}_s(M_t) + \gamma_2 \mathcal{L}_s(R_t) + \psi_2 \mathcal{L}_s(R_t^2)) + \eta X_t + \epsilon_t.$$

All variables are defined as in Table 3. The sample period comprises 27,449 trading days, of which 6,467 were during recessions. The t -stats reported are computed using White (1980) standard errors.

Panel A. Expansions (β_1)

	Positive		Negative		Pessimism	
	β	t -stat	β	t -stat	β	t -stat
$(1 - D_t) \times M_{t-1}$	0.024	3.3	-0.028	-3.5	-0.035	-4.2
$(1 - D_t) \times M_{t-2}$	0.004	0.6	0.004	0.5	0.001	0.1
$(1 - D_t) \times M_{t-3}$	-0.004	-0.6	0.005	0.7	0.006	0.8
$(1 - D_t) \times M_{t-4}$	-0.012	-1.7	0.006	0.8	0.011	1.5
$(1 - D_t) \times M_{t-5}$	-0.004	-0.6	0.006	0.8	0.007	0.9

Panel B. Recessions (β_2)

	Positive		Negative		Pessimism	
	β	t -stat	β	t -stat	β	t -stat
$D_t \times M_{t-1}$	0.085	3.9	-0.087	-3.4	-0.117	-4.4
$D_t \times M_{t-2}$	0.004	0.2	-0.005	-0.2	-0.004	-0.2
$D_t \times M_{t-3}$	-0.021	-1.0	0.010	0.4	0.020	0.8
$D_t \times M_{t-4}$	-0.009	-0.4	0.016	0.7	0.019	0.8
$D_t \times M_{t-5}$	-0.005	-0.2	0.028	1.2	0.026	1.1

Panel C. Tests

	Positive		Negative		Pessimism	
	F -stat	p -value	F -stat	p -value	F -stat	p -value
$\beta_{11} = \beta_{21}$	7.2	0.007	5.0	0.025	8.6	0.003
$\sum_{j=2}^5 \beta_{1j} = 0$	1.6	0.205	2.6	0.109	3.4	0.066
$\sum_{j=2}^5 \beta_{2j} = 0$	0.7	0.403	1.6	0.212	2.3	0.132

Table 5
Volatility Adjustments, Orthogonal Media Content and Robust Regressions

The table reports the estimated coefficients β from the model

$$R_t = (1 - D_t) (\beta_1 \mathcal{L}_s(M_t) + \gamma_1 \mathcal{L}_s(R_t) + \psi_1 \mathcal{L}_s(R_t^2)) + D_t (\beta_2 \mathcal{L}_s(M_t) + \gamma_2 \mathcal{L}_s(R_t) + \psi_2 \mathcal{L}_s(R_t^2)) + \eta X_t + \epsilon_t.$$

All independent variables are as in Table 3. In Panel A, the dependent variable R_t denotes the normalized log-returns on the DJIA. These are constructed by taking the raw log-returns on the DJIA and dividing them by the estimates $\hat{\sigma}_t$ from the GARCH(1,1) model from Panel C of Table 2. In Panels B and C, R_t denotes the log-return on the DJIA average. In Panels A and C, the variable M_t denotes one of our media measures, as described in Table 3. In Panel B, the variable M_t is the residual from the model estimated in Table 6. Estimation in Panels A and B is via OLS. The estimation in Panel C is done using robust linear regression based on Huber's (1981) M -estimator. The sample period comprises 27,449 trading days, of which 6,467 were during recessions. The t -stats reported are computed using White (1980) standard errors.

	Positive		Negative		Pessimism	
Panel A. GARCH-adjusted returns						
	β	t -stat	β	t -stat	β	t -stat
Expansions, $(1 - D_t) \times M_{t-1}$	0.022	3.1	-0.025	-3.3	-0.033	-4.1
Recessions, $D_t \times M_{t-1}$	0.051	3.5	-0.070	-4.3	-0.087	-5.1
	F -stat	p -value	F -stat	p -value	F -stat	p -value
$\beta_{11} = \beta_{21}$	3.1	0.079	6.1	0.014	8.2	0.004
$\sum_{j=2}^5 \beta_{1j} = 0$	1.3	0.254	1.9	0.170	3.0	0.085
$\sum_{j=2}^5 \beta_{2j} = 0$	0.0	0.868	2.7	0.103	2.1	0.148
Panel B. Orthogonal media measures						
	β	t -stat	β	t -stat	β	t -stat
Expansions, $(1 - D_t) \times M_{t-1}$	0.022	3.3	-0.027	-3.8	-0.032	-4.4
Recessions, $D_t \times M_{t-1}$	0.078	3.8	-0.068	-3.1	-0.094	-4.1
	F -stat	p -value	F -stat	p -value	F -stat	p -value
$\beta_{11} = \beta_{21}$	6.9	0.008	3.1	0.080	6.6	0.010
$\sum_{j=2}^5 \beta_{1j} = 0$	0.1	0.815	0.2	0.667	0.2	0.688
$\sum_{j=2}^5 \beta_{2j} = 0$	0.0	0.898	0.0	0.867	0.0	0.917
Panel C. Robust regression						
	β	t -stat	β	t -stat	β	t -stat
Expansions, $(1 - D_t) \times M_{t-1}$	0.024	4.0	-0.025	-4.0	-0.034	-5.3
Recessions, $D_t \times M_{t-1}$	0.055	4.6	-0.086	-7.0	-0.101	-7.9
	F -stat	p -value	F -stat	p -value	F -stat	p -value
$\beta_{11} = \beta_{21}$	5.4	0.020	19.7	0.000	22.2	0.000
$\sum_{j=2}^5 \beta_{1j} = 0$	3.2	0.075	1.3	0.258	3.4	0.065
$\sum_{j=2}^5 \beta_{2j} = 0$	1.2	0.282	7.7	0.005	3.5	0.063

Table 6
Feedback from Stock Prices to News Content

The table reports the estimated coefficients λ and β from the model

$$M_t = (1 - D_t) (\lambda_1 R_t + \beta_1 \mathcal{L}_s(R_t) + \gamma_1 \mathcal{L}_s(M_t)) + D_t (\lambda_2 R_t + \beta_2 \mathcal{L}_s(R_t) + \gamma_2 \mathcal{L}_s(M_t)) + \eta X_t + v_t.$$

The variable M_t denotes one of our media measures, as described in Table 3. The set of exogenous variables X_t includes those in the specification of Table 3. In Panel A the variable R_t denotes the log-return on the DJIA. In Panel B the variable R_t denotes the normalized log-returns on the DJIA, constructed as in Panel A of Table 5. The sample period comprises 27,449 trading days, of which 6,467 were during recessions. The t -stats reported are computed using White (1980) standard errors.

	Positive		Negative		Pessimism	
	λ, β	t -stat	λ, β	t -stat	λ, β	t -stat
Panel A. Using raw returns ($\lambda_1, \beta_1, \lambda_2, \beta_2$)						
$(1 - D_t) \times R_t$	0.335	32.3	-0.332	-30.7	-0.414	-34.2
$(1 - D_t) \times R_{t-1}$	0.046	6.1	-0.056	-7.6	-0.059	-7.8
$(1 - D_t) \times R_{t-2}$	-0.008	-1.1	-0.012	-1.6	-0.005	-0.6
$(1 - D_t) \times R_{t-3}$	0.008	1.0	-0.001	-0.1	-0.004	-0.5
$(1 - D_t) \times R_{t-4}$	-0.011	-1.6	0.020	2.7	0.024	3.3
$D_t \times R_t$	0.197	16.8	-0.221	-19.5	-0.263	-20.3
$D_t \times R_{t-1}$	0.048	5.4	-0.045	-5.4	-0.052	-5.7
$D_t \times R_{t-2}$	0.011	1.3	-0.007	-0.8	-0.009	-1.0
$D_t \times R_{t-3}$	0.003	0.3	-0.019	-2.7	-0.012	-1.6
$D_t \times R_{t-4}$	-0.007	-0.9	0.022	2.4	0.023	2.5
	F -stat	p -value	F -stat	p -value	F -stat	p -value
Test $\lambda_1 = \lambda_2$	28.5	0.000	20.2	0.000	28.6	0.000
Panel B. Returns normalized by GARCH(1,1) ($\lambda_1, \beta_1, \lambda_2, \beta_2$)						
$(1 - D_t) \times R_t$	0.345	48.7	-0.342	-50.2	-0.427	-61.7
$(1 - D_t) \times R_{t-1}$	0.040	5.6	-0.044	-6.4	-0.047	-6.8
$(1 - D_t) \times R_{t-2}$	-0.017	-2.4	0.001	0.2	0.010	1.5
$(1 - D_t) \times R_{t-3}$	0.005	0.7	0.002	0.3	-0.000	-0.0
$(1 - D_t) \times R_{t-4}$	-0.007	-1.0	0.014	2.0	0.017	2.5
$D_t \times R_t$	0.324	26.3	-0.338	-31.4	-0.413	-37.0
$D_t \times R_{t-1}$	0.065	5.5	-0.050	-4.7	-0.063	-5.6
$D_t \times R_{t-2}$	-0.003	-0.3	0.003	0.3	0.008	0.8
$D_t \times R_{t-3}$	-0.007	-0.7	-0.016	-1.5	-0.001	-0.1
$D_t \times R_{t-4}$	-0.002	-0.1	0.026	2.4	0.024	2.2
	F -stat	p -value	F -stat	p -value	F -stat	p -value
Test $\lambda_1 = \lambda_2$	1.9	0.086	0.9	0.474	0.8	0.564

Table 7
Mondays, Post-Holiday Returns and Media

The table reports the estimated coefficients β from the model

$$R_t = (1 - I_t) [(1 - D_t) (\beta_1 \mathcal{L}_s(M_t) + \gamma_1 \mathcal{L}_s(R_t) + \psi_1 \mathcal{L}_s(R_t^2)) + D_t (\beta_3 \mathcal{L}_s(M_t) + \gamma_3 \mathcal{L}_s(R_t) + \psi_3 \mathcal{L}_s(R_t^2))] \\ + I_t [(1 - D_t) (\beta_2 \mathcal{L}_s(M_t) + \gamma_2 \mathcal{L}_s(R_t) + \psi_2 \mathcal{L}_s(R_t^2)) + D_t (\beta_4 \mathcal{L}_s(M_t) + \gamma_4 \mathcal{L}_s(R_t) + \psi_4 \mathcal{L}_s(R_t^2))] \\ + \eta X_t + \epsilon_t,$$

where the dummy variable I_t takes on the value one if and only if date $t - 1$ was not a trading date, and all other independent variables are defined as in Table 3. The dependent variable R_t is the log-return on the DJIA index from 1905 to 2005. The sample period comprises 27,449 trading days, of which 6,467 were during recessions. The t -stats reported are computed using White (1980) standard errors.

	Positive		Negative		Pessimism	
	β	t -stat	β	t -stat	β	t -stat
Expansion, Monday/holidays, β_{11}	0.056	3.3	-0.062	-3.5	-0.079	-4.1
Expansion, back-to-back weekdays, β_{21}	0.015	1.9	-0.019	-2.1	-0.024	-2.6
Recession, Monday/holidays, β_{31}	0.188	4.0	-0.208	-3.8	-0.267	-4.9
Recession, back-to-back weekdays, β_{41}	0.062	2.8	-0.061	-2.3	-0.085	-3.1
	F -stat	p -value	F -stat	p -value	F -stat	p -value
$\beta_{11} = \beta_{21}$	4.7	0.031	4.8	0.029	6.6	0.010
$\beta_{31} = \beta_{41}$	6.2	0.012	6.0	0.014	10.0	0.002
$\beta_{11} = \beta_{31}$	6.9	0.008	6.2	0.013	10.5	0.001
$\beta_{21} = \beta_{41}$	4.0	0.046	2.3	0.128	4.5	0.034

Table 8
Media Content and Returns after the NYSE Opening

The table reports the estimated coefficients β from the model

$$R_t = (1 - D_t) (\beta_1 \mathcal{L}_s(M_t) + \gamma_1 \mathcal{L}_s(R_t) + \psi_1 \mathcal{L}_s(R_t^2)) + D_t (\beta_2 \mathcal{L}_s(M_t) + \gamma_2 \mathcal{L}_s(R_t) + \psi_2 \mathcal{L}_s(R_t^2)) + \eta X_t + \epsilon_t,$$

where the dependent variable R_t is the log-return on the DJIA from 11am until close for the period 1933 to 2005. The set of independent variables is described in Table 3. The sample period comprises 19,184 trading days, of which 2,762 were during recessions. The t -stats reported are computed using White (1980) standard errors.

	Positive		Negative		Pessimism	
	β	t -stat	β	t -stat	β	t -stat
$(1 - D_t) \times M_{t-1}$	-0.001	-0.1	-0.015	-2.3	-0.012	-1.8
$D_t \times M_{t-1}$	0.057	3.1	-0.047	-2.6	-0.065	-3.5
	F -stat	p -value	F -stat	p -value	F -stat	p -value
Test $\beta_{11} = \beta_{21}$	8.9	0.003	2.9	0.088	7.3	0.007

Table 9
Media Content and Abnormal NYSE Trading Volume

The table reports the estimated coefficients β from the model

$$\begin{aligned} \hat{V}_t = & (1 - D_t) (\beta_{1+} M_t^+ + \beta_{1-} M_t^- + \psi_1 R_t + \gamma_1 \mathcal{L}_s(R_t) + \eta_1 \mathcal{L}_s(R_t^2)) \\ & + D_t (\beta_{2+} M_t^+ + \beta_{2-} M_t^- + \psi_2 R_t + \gamma_2 \mathcal{L}_s(R_t) + \eta_2 \mathcal{L}_s(R_t^2)) + \epsilon_t, \end{aligned}$$

where \hat{V}_t denotes abnormal normalized aggregate NYSE trading volume, the logarithm of total trading volume stripped of calendar effects and time trends, standardized to have a zero mean and unit variance. We let $M_t^+ = \max(M_t, 0)$, and $M_t^- = \min(M_t, 0)$. The time period in the estimation is 1929 to 2005, which includes 19,397 trading dates. The t -stats reported are computed using White (1980) standard errors.

	Positive		Negative		Pessimism	
	β	t -stat	β	t -stat	β	t -stat
$(1 - D_t) \times M_t^+$	0.037	3.1	0.014	1.2	0.024	1.9
$(1 - D_t) \times M_t^-$	0.005	0.3	-0.111	-7.4	-0.121	-8.8
$D_t \times M_t^+$	0.001	0.0	0.054	2.4	0.081	3.5
$D_t \times M_t^-$	-0.050	-1.8	-0.105	-3.3	-0.110	-3.7
	F -stat	p -value	F -stat	p -value	F -stat	p -value
$\beta_{1+} = \beta_{2+}$	2.0	0.152	3.1	0.080	5.9	0.016
$\beta_{1-} = \beta_{2-}$	3.8	0.051	0.1	0.848	0.1	0.717

FINANCIAL MARKETS

Another Day of Great Activity, with Big Gains for the Industrials.

Expectations that yesterday's Stock Exchange session, sandwiched in between two holidays, would see much less activity were quickly disappointed. The market opened very active and strong, and after a small reaction in the forenoon resumed its upward trend with almost the same violence shown in the excited sessions of last week. The list was again irregular, but by far the larger number of stocks scored substantial gains and the upward movement of some of the war issues, which had been checked by banks and brokers who foresaw trouble if the advance were not held under control, was resumed with a great deal of vigor. The most striking gain among such issues was scored by Baldwin Locomotive, which, after hanging for several days around 115, returned yesterday to 127½, closing at 126, with a net advance of 11 points. This secondary stage of activity for Baldwin was accompanied by fresh merger rumors, which do not appear to have any substantial basis in fact. Even more active and relatively as strong was Westinghouse, of which more than 100,000 shares changed hands on a range of 5½ points. It closed at

138, with a gain of 4½ points above Saturday's close. The American Car & Foundry made a good recovery to 85¼, and gains of from 2 to 5 points were numerous. The motor issues returned to popularity, all three classes of Maxwell stock advancing on the expectation of some kind of an announcement Wednesday of a plan looking to the payment of the accumulated dividend on the first preferred. Studebaker advanced 2¼ and General Motors 1 point.

The rails retained some of their momentum from last week, and most of the leaders sold at new high prices for the year. News of the note being prepared for dispatch to Great Britain was received too late to affect the market, if indeed such news can have any effect on the present temper of traders, and the list closed pretty close to the top.

Some uneasiness was caused yesterday by a new development of weakness in the foreign exchange market. Demand sterling sold down to \$4.67¾ compared with the low price of \$4.68¼ on Saturday. The failure of the conclusion of the \$500,000,000 Anglo-French loan to help foreign exchange rates gave special interest to an important meeting of bankers held yesterday afternoon, which was addressed by Sir Edward Holden, one of the visiting Commissioners.

Figure 1. Financial Markets column. Published in the New York Times on October 12, 1915.

TOPICS IN WALL STREET.

American Munition Orders.

Until yesterday the stock market gave no indication that the war stocks derived a chance of profit from war with Mexico. To speculators in these shares it was in fact a matter of the keenest disappointment that they went down on war news. Over and over they have repeated the question: "What sort of a war stock is it that is depressed by a new war?" Yesterday an advance of 17 points in Bethlehem Steel held out a ray of hope and advances in most of the others on covering by professionals strengthened hopes that the next turn would be for the better. Officers of many of the munitions companies expected orders from the United States Government in the near future, but nowhere was it believed that these orders would be placed at terms permitting as great profits as those obtained in some of the contracts with the Allies.

The Extent of the Declines.

From the high point of week before last to the low point of last week, which was the low point of Friday's market, the average price of fifty representative stocks declined \$5.33 a share. These stocks included many railroad shares in which the declines were small compared with losses in some of the speculative industrials. Reading, which lost 8½ points in this period, and Norfolk & Western, with a loss of 5½, were the only rails to decline more than the 5-1-3-point average of the fifty. A score of industrials sustained greater losses and many of these losses ran into double figures, among them being: New York Air Brake, 11; Mexican Petroleum, 12½; Baldwin Locomotive, 13; United States Smelting, 13; Tennessee Copper, 14½; American Zinc, 14½; Willys-Overland, 16; Butte and Superior, 10½; United States Industrial Alcohol, 26½. On the Curb Chevrolet Motors lost 46 points.

Now Up, Now Down.

It is interesting to note the change in sentiment that sweeps over the floor of the Stock Exchange after a pronounced rise, or sharp decline. Traders who have been bearish for weeks were turning bullish yesterday morning. They figured that the break which had been needed had been supplied, and that, therefore, stocks were a purchase again.

The Mexican Factor.

An old-time member said after the close that neither the Mexican war danger, nor the inadequacy of our war machinery, was really back of the slump which took place last week. Those arguments were advanced to support the decline, but in his opinion the break would have come had the Mexican situation continued unchanged. This man's theory is that the market had become badly congested with stocks, and had to be cleaned out by a return to lower prices. A number of pools were carrying large amounts of stock which they had not been able to market, and there were some large individual accounts that needed shaking out. The low prices made on Friday brought in a number of fresh buyers, and if this trader's theory works out the market will develop a much better tone this week, regardless of developments across the border. When the list grows stale nothing but a sharp setback will attract new money. That this market had become stale was evidenced by its utter disregard of good news, such as new and increased dividends.

No Extra Holiday.

When the brokers gave up their expected extra holiday before May 30, they looked for an extra day preceding the Fourth. The uncertainty of the political situation appears to have destroyed any chance of getting it. No petition has been circulated on the floor, and it is unlikely that the situation will clear in time to allow the drafting of one before the next meeting of Governors.

Bonds Have Idle Week.

The bond market suffered along with stocks last week, but without registering substantial declines. Bonds were effected more through a let-down of buying than from the liquidation of securities. Some of the banks and large dealers were reported as sellers of a considerable amount of bonds which they had been carrying for a month or more, and on which they had good profits. If this actually did take place the offerings were rather easily absorbed, and inquiries among bond men failed to show that there had been any urgent selling through fear that the Mexican situation might wipe out profits before they could be realized. The investment demand is believed to be widening, now that supplies from Europe have begun to fall away, leaving room for other offerings, and the bankers are inclined to think that business will pick up again with the coming of definite developments south of the border.

Figure 2. Topics in Wall Street column. Published in the New York Times on June 25, 1916.

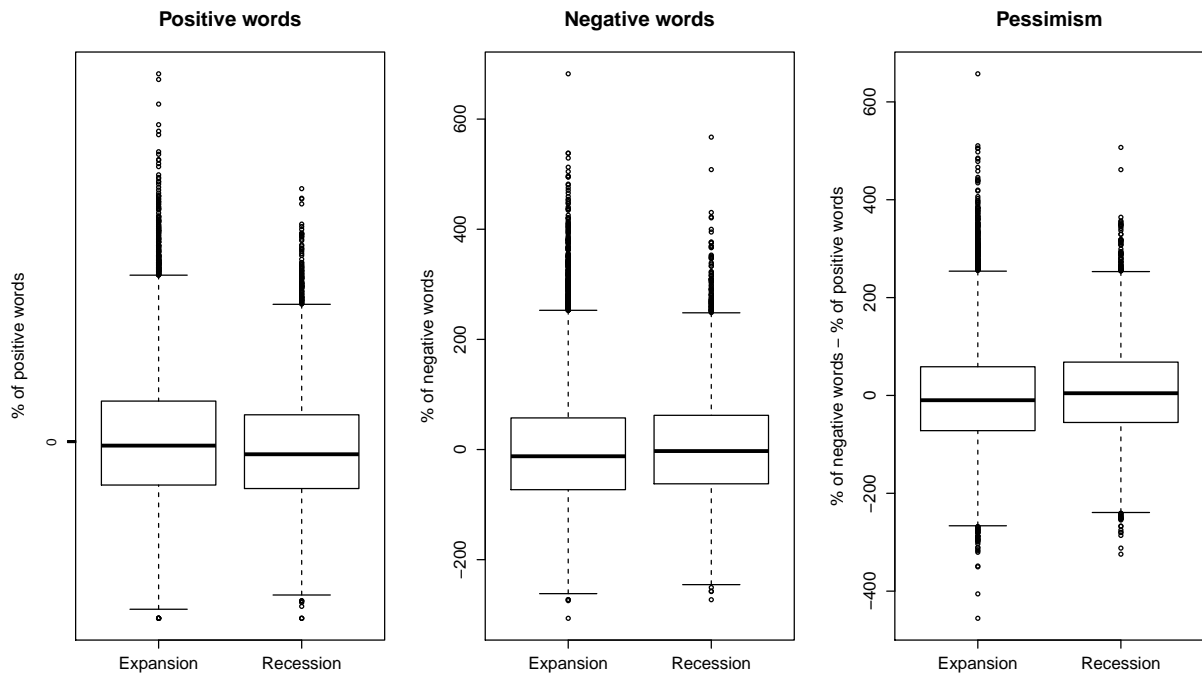


Figure 3. Boxplots of the media measures. Each graph presents boxplots of each of the three media measures (positive, negative, pessimism) used in the paper, as a function of the business cycle. These measures are constructed from the columns “Financial Markets” and “Topics in Wall-Street” published in the New York Times in the period 1905 to 2005.

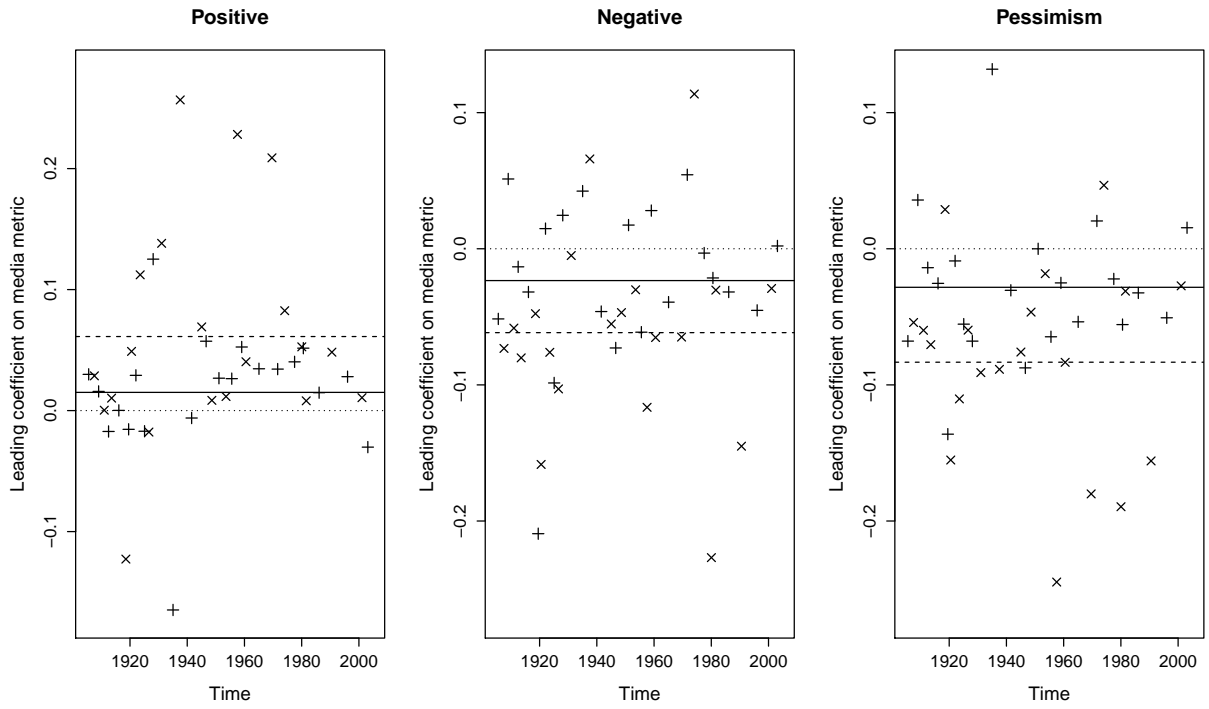


Figure 4. Plot of leading coefficients. Each graph plots the leading coefficients on the media measures as predictors of stock returns for each business cycle during our sample period (1905–2005). The crosses correspond to expansionary periods, according to NBER, whereas the x’s correspond to recessions. The dashed line is the time-series average of the estimates during recessions, whereas the solid line corresponds to expansions.

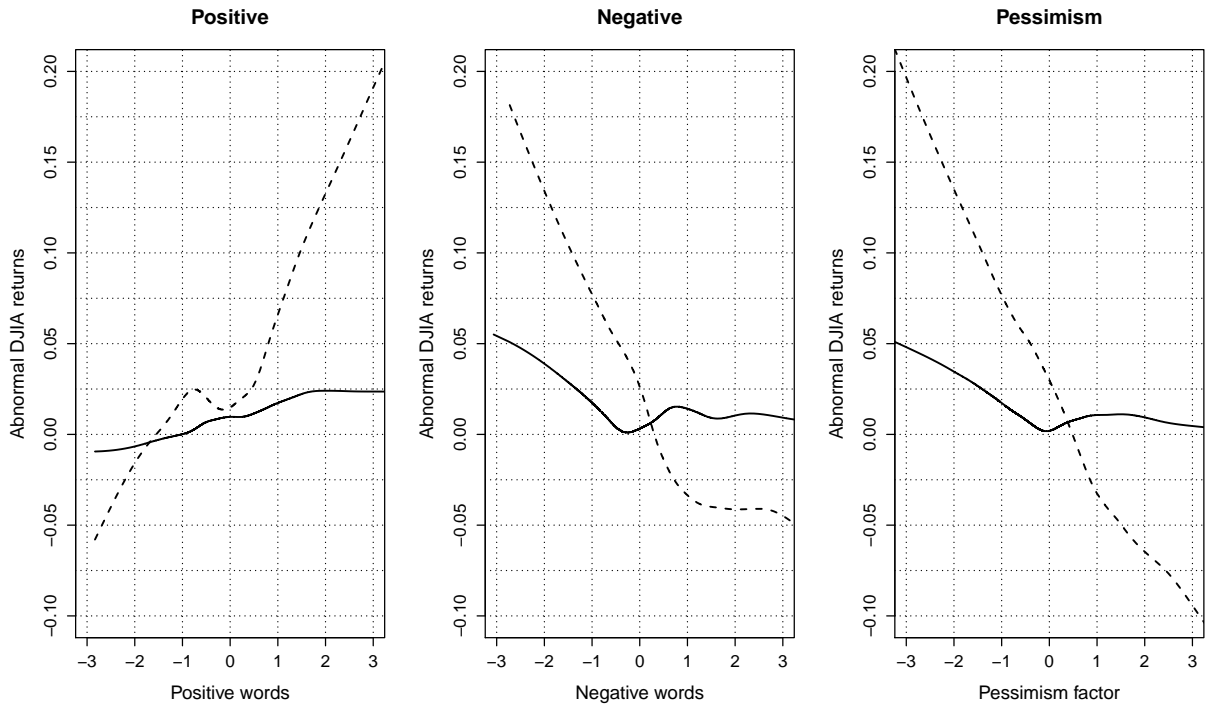


Figure 5. Non-parametric estimates of DJIA returns and media content. Non-parametric estimates of the conditional average DJIA returns as a function of the media measures constructed in the paper. The solid line is the estimate during expansionary periods, whereas the dashed line corresponds to NBER recessions.




Figure 6. Non-parametric estimates of abnormal NYSE trading volume and media content. Non-parametric estimates of the conditional abnormal NYSE trading volume as a function of the media measures constructed in the paper. The solid line is the estimate during expansionary periods, whereas the dashed line corresponds to NBER recessions.