

A Friendly Turn: Advertising Bias in the News Media*

Florens Focke, Alexandra Niessen-Ruenzi, and Stefan Ruenzi

First Version: February 2014, This Version: November 2015

Abstract

This paper investigates whether newspapers report more favorably about advertising corporate clients than about other firms. Our identification strategy based on high-dimensional fixed effects and high frequency advertising data shows that advertising leads to more positive press coverage. This advertising bias in reporting is found to be mitigated but not eliminated by newspapers' reputational concerns. Advertising bias manifests particularly in less negative reporting of a newspaper about bad news events of its advertising clients as compared to firms not advertising in this newspaper. Our findings cast doubt on the independence of the press from corporate pressure.

JEL-Classification Codes: M37, M38

Keywords: Media, Advertising Bias, Newspapers, Commercial Bias, Computer Linguistics

*All authors are at the University of Mannheim, L9, 1-2, 68131 Mannheim, florens.focke@gess.uni-mannheim.de; niessen@bwl.uni-mannheim.de; ruenzi@bwl.uni-mannheim.de. We thank Patrick Kelly, Tim Loughran, David Oesch, Felix von Meyerinck, Jakub Redlicki, seminar participants at the University of Mannheim and the Schmalenbach-Gesellschaft as well as participants at the BEROCC 2014, German Economic Association 2014, the German Finance Association 2014, the Swiss Finance Association 2015, the LAGV 2015 conferences, and the Media Economics Workshop 2015 for valuable comments. We also thank the Julius Paul Stiegler Memorial Foundation and the German Research Foundation for financial support. All errors are our own.

I headed the "Wall Street Journal" for 16 years. We made money from advertising and we constantly wrote about our advertising clients. Sometimes we criticized them harshly, sometimes they threatened us to stop advertising, sometimes they stopped advertising. However, we didn't allow them to bias our reports.

– Paul Steiger

1 Introduction

Independence of the news press is one of the pillars of a functioning democracy. Ideally, newspapers and other media outlets should report truthfully and objectively about news items of interest to their readers, thus allowing them to make rational and unbiased decisions based on the information reported. This is important because although consumers, voters or investors may understand the incentives of firms or politicians to manipulate information, the media is a third party that might appear independent. In line with this, sizable persuasion effects of the media have been found, whereas the direct effects of communication devices such as advertising appear to be small (see [DellaVigna and Gentzkow \(2010\)](#)). In this paper, we examine whether advertising influences news coverage. Ours is the first paper to provide direct evidence that newspapers report about their advertising clients in a positively biased way.

Most industrialized countries established freedom of press laws in their constitutions to ensure that the media are not captured by politicians or powerful business groups.¹ The literature argues that the rise of advertising in the nineteenth-century and the associated revenues for newspapers created a press largely independent from political influence (e.g., [Petrova \(2011\)](#)).² But newspapers that receive advertising revenues act in a classical two-sided market ([Rochet and Tirole \(2003\)](#)) and might be subject to corporate influence: on the one hand, they provide news and typically charge a price for each issue from their readers. Among other things, these readers may value accuracy and entertainment ([DellaVigna and](#)

¹See 2013 World Press Freedom Index, <http://en.rsf.org/press-freedom-index-2013,1054.html>

²However, [Di Tella and Franceschelli \(2011\)](#) find that government advertising can reintroduce a dependence of the media on the government.

[La Ferrara \(2015\)](#)). On the other hand, newspapers sell advertising space to firms that are interested in reaching as many readers (i.e. potential customers) as possible. More important in our context, these firms also are interested in (too) positive coverage because many of their customers and other stakeholders read newspapers.³ The conflicting preferences of advertisers and readers create a conflict of interest for newspapers (and other advertising-financed media outlets like television stations or magazines). As suggested by Paul Steiger's quote above, at least some media outlets are aware of this conflict and find it important to uphold a reputation for accuracy. Whether such reputational concerns are sufficient to eliminate bias in favor of advertisers depends on the amount of pressure either group exerts on the newspaper. It is therefore an empirical question.

According to the Newspaper Association of America, around 70% of newspapers' revenues are obtained from advertising, while only about 30% are obtained from circulation.⁴ Thus, advertising is the dominant source of revenue for most newspapers, providing a first indication that advertisers may be able to exert sufficient direct or indirect pressure on newspapers to distort coverage. In addition, there is some anecdotal evidence suggesting that newspapers might tilt reporting in order to please their advertisers. For example, magazines that had a high part of their revenues coming from tobacco companies did report less about smoking causing lung cancer ([Weis and Burke \(1986\)](#)). Furthermore, [Nyilasy and Reid \(2011\)](#) present survey results according to which 70 to 90% of US newspaper editors had experienced pressure by advertisers. However, the authors further find that the overwhelming majority of editors indicate that they do not bend to such pressure. This may be because no bias exists or because reputational concerns play an important role even in this anonymous survey.

³An additional reason why newspapers might be inclined to positively bias reports about their advertising clients is provided in a theoretical paper by [Germano and Meier \(2013\)](#). They show that newspapers have incentives to suppress news that would negatively affect the sales of their advertising clients (even if advertisers do not directly monitor or try to influence what newspapers write), because this helps their advertisers to be able to afford advertising in the future, too.

⁴<http://www.naa.org/Trends-and-Numbers/Newspaper-Revenue/Newspaper-Media-Industry-Revenue-Profile-2012.aspx>

We investigate whether advertising has an impact on how newspapers report about their advertising corporate clients. The sample spans from 1999 to 2012 and covers advertising expenditures and press coverage on firms in a comprehensive sample of local and national US newspapers. Based on weekly firm- and newspaper- level advertising data, we show that newspapers are more likely to cover their advertising client firms and that they tend to write longer articles about firms that advertise in them. More importantly, we document that articles on a firm are indeed positively biased if the firm spends more on advertising in this newspaper. To objectively measure the tone of media reports, we focus on the negative word list developed for financial text by [Loughran and McDonald \(2011\)](#). In robustness tests, we additionally consider a positive word list, but note that it is less reliable because of the difficulty of detecting negations.

Due to the high frequency and high dimensionality (firm, time, and newspaper dimension) of our advertising and press coverage data, we are able to make a crucial step towards establishing causality by including firm-week as well as firm-newspaper fixed effects in our analysis. This allows us to control for all time-varying information on firms as well as any preferences certain newspapers might have for certain firms for other reasons than their advertising. Thus, our approach alleviates endogeneity concerns as our main result can only be driven by time-varying advertisement expenditures within a firm-newspaper pair.

We also implement a panel vector autoregression (VAR) model to analyze the temporal dynamics of the relation between advertising and news tone. Our results support the notion that it is past advertising that has an impact on the tone of future press coverage, while there is no significant evidence suggesting a link between the past tone of press coverage and future advertising. This pattern is consistent with the view that newspapers report favorably about their existing advertising clients, but there is no evidence for ‘anticipatory obedience’ in the sense of newspapers positively biasing reports about potential future advertising clients.

To investigate the impact of reputational concerns on advertising-induced bias, we then study the relationship between newspaper audience characteristics and the strength of the bias. The readers of different newspapers likely differ in their demand for accuracy as opposed to, for example, entertainment. The higher the value placed on accuracy by the readers of a newspaper, the weaker we expect the advertising bias to be. Using the complexity of newspapers’ language and whether they cater to a local or a national audience, we find that advertising bias is indeed significantly weaker in newspapers with a more complex language and those with a national audience. Importantly, however, these newspaper characteristics only mitigate but do not fully eliminate the bias.

To shed some light on the channels via which advertising bias manifests, we analyze the newspapers articles following earnings announcements and following extreme stock market returns. These two events offer the advantage that we observe the exact timing of the event and that they can be objectively classified as good or bad events. We find that advertising bias is strongest in articles on firms with bottom quintile earnings surprises and with extremely bad stock market returns. These results show that newspapers mainly bias coverage of bad news, while there is no evidence of positively biased articles on good news.

Overall, our results suggest that economic incentives have an adverse impact on the independence of the news media from the corporate world. The need to generate profits based on advertising makes news outlets dependent on their corporate clients and creates an incentive to report about these clients in a more positive way. This pattern has important implications, as biased information printed in newspapers leads to suboptimal behavior of newspaper readers—who base many important decisions like purchases of consumption goods or investment decisions on this information—and eventually to adverse welfare consequences.

Furthermore, to the extent that investment decisions on financial markets are also affected by the newspaper articles potential investors read ([Engelberg and Parsons \(2011\)](#)), asset prices might be biased (at least if investors are not able to fully de-bias the content of

the reports they read), that is, informational efficiency is reduced and resource allocation in the economy would be less efficient.

Our paper contributes to various strands of the literature. There are several theoretical models that investigate the impact of advertising on news tone, i.e. advertising bias of newspapers. [Blasco, Pin, and Sobbrío \(2014\)](#) set up a model where advertisers may pay a media outlet to conceal negative information about the quality of their products. In their model, advertisers' success in influencing media outlets depends on the degree of competition between advertisers. In a similar vein, [Stroemberg \(2004\)](#) models the incentives of the media to deliver news to different groups. The author argues that advertising financing of media firms induces them to provide more news to groups that are valuable to advertisers. Finally, [Ellman and Germano \(2009\)](#) develop a game-theoretical model in which advertisers are also concerned about the content of articles surrounding their advertisements. We contribute to this theoretical literature by showing empirically that advertisers indeed have an impact on media reporting and are successful in mitigating negative news coverage.

There is also a small number of papers investigating the relation between advertising and a firm's press coverage empirically. Focusing on the mutual fund industry, [Reuter and Zitzewitz \(2006\)](#) show that funds are more likely to be recommended in personal finance magazines in which they advertise.⁵ However, they find no evidence for biased coverage in national newspapers. In contrast to these studies, we are able to study articles on firms in all industries, which increases the representativeness of our findings. Moreover, our comprehensive sample of newspapers allows more direct tests of the impact of reputational concerns on bias. [Gambaro and Puglisi \(2015\)](#) show that the frequency of newspaper coverage of a firm is positively related to the amount of advertisements that the firm commissioned at that newspaper. However, their analysis is based on a very small sample of only 13 Italian firms and 6 different newspapers in 2006 and 2007. A similar result is shown in [de Smet and Vanormelingen \(2012\)](#) based on 57 Belgian companies and 8 newspapers between 2001 and 2005. The setting of both papers, [Gambaro and Puglisi \(2015\)](#) and [de Smet and Vanormelin-](#)

⁵A similar result is obtained by [Reuter \(2009\)](#) for wine ratings in a wine magazine.

gen (2012), does not allow for the inclusion of the same set of high-dimensional fixed effects for identification that we can use. Furthermore, while these papers only look at coverage per se, our investigation mainly focuses on the tone of newspaper coverage based on a linguistic analysis. The paper that is probably most closely related to ours is Gurun and Butler (2012). They show that local media outlets report more favorable about local firms than about distant firms and argue that this might be due to local firms advertising in local media outlets. However, due to data limitations they can only provide indirect evidence on a potential impact of advertising on the news tone in local (but not in national) newspapers, because they do not have disaggregate information on how much a firm spends on advertising in which individual newspaper. In contrast, the granularity of our data allows us make a big step towards establishing causality in the relationship between advertising and media coverage (for local as well as national newspapers) for the first time.

The paper is organized as follows. Section 2 describes the different data sets used in our analysis and presents univariate statistics. In Section 3, we investigate the impact of advertising on media coverage of corporate advertising clients. The impact of specifically defined news events on advertising bias is analyzed in Section 4. Section 6 concludes.

2 Data and summary statistics

2.1 Advertising data

We obtain advertising data from Kantar Media, a subsidiary of WPP plc. Kantar Media collects all advertisements that are published in a core set of 155 US newspapers. Their proprietary database, “Kantar Media Strategy”, contains the advertisement’s publication date, the news outlet in which it is published, and the firm that commissioned the advertisement. In addition, each advertisement is associated with a cost using “rate cards” that indicate advertising prices depending on size, product categories, and days of week or sections. These data would allow us to run the analysis on the firm-newspaper-day level.

However, as there are typically only a very small number of news articles on a firm per day, we will later mainly use aggregate data on the firm-newspaper-week level. Advertising data are available from 1999 to 2012. Over this time period, Kantar tracks advertising expenditures totaling 322bn USD.⁶ Figure 1 shows the evolution of aggregate advertising spending and changes in GDP over our sample period. The graph reveals several characteristics of the newspaper advertising data. First, it follows a seasonal pattern with advertising being about 20 percent higher in November and December than in January. Second, aggregate newspaper advertising was closely related to changes in GDP until about 2009. Even as the economy recovered after the recession, newspaper advertising continued to shrink, although at a lower rate than before. These strong patterns highlight the importance of controlling for time fixed effects, which we do in all our regressions.

— Please insert FIGURE 1 approximately here —

Based on the Kantar Media data, we create a list of all firms that commissioned advertising in one of the 155 US newspapers in our sample period. We then link these firms to publicly listed firms between 1999 and 2012 that are included in the CRSP database (share codes 10 or 11). We focus on publicly listed firms as the news events we will later analyze (see Section 4) can be clearly defined for those firms only. Furthermore, these firms typically are the firms with the largest advertising budgets and should thus be most relevant in our context.⁷ We match firms based on their names (corrected for abbreviations), industry and main products. All matches are manually checked and where in doubt, we have excluded a potential match. Using this approach, we successfully matched 5,290 companies.

⁶The Newspaper Association of America reports advertising revenues totaling 533bn USD for all covered US newspapers. Thus, our data cover roughly 60% of all print advertisements in US newspapers.

⁷The average amount spent on newspaper advertising per year is 5,557,752 USD for the public companies in our analysis and 74,518 USD for private companies.

2.2 Press coverage of firms

We retrieve newspaper articles on our sample firms from various news providers (e.g., Factiva and LexisNexis). Out of the 155 US newspapers that are covered by Kantar Media, 41 newspapers are also covered by these providers. The full list of newspapers that are covered by both, Kantar and the news providers, and are thus included in our analysis, is provided in Table 1. Unlike previous studies, our sample includes all big national newspapers, i.e. the New York Times, USA Today, the Wall Street Journal, and the Washington Post as well as most major local newspapers.

— Please insert TABLE 1 approximately here —

From the news providers, we download all articles on publicly listed firms between 1999 and 2012. Then, we merge these articles to the list of firms that we extract from the Kantar Media database, using the CRSP company identifier (permco). Thus, the sample includes all CRSP firms that appear at least once as an advertiser in the Kantar data *and* are at least covered once in one of the 41 newspapers.

We exclude articles with less than 20 words and conduct cross checks to make sure that the newspaper articles are indeed mainly about the respective firm.⁸ Overall, we obtain 1,567,629 articles on 4,009 different companies. Aggregated to the weekly level, this results in 998,481 firm-newspaper-week observations.

In the next step, we use the bag-of-words approach from computer linguistics to analyze the content of newspaper articles. Specifically, we compare the words in each article against a pre-defined dictionary. To classify articles as positive or negative, we follow the approach of prior papers concerned with textual analysis and use the Loughran and McDonald (2011) dictionary (LMD) of positive and negative words.⁹ Following Loughran and McDonald

⁸See Appendix A1 in the Supplementary Material for further details on the data cleansing process.

⁹The LMD dictionaries are obtained from Bill McDonald’s webpage. They are designed to capture the tone of text in a business context.

(2011), we account for negations that could bias the results of word lists designed to measure positive tone. Simple negation is taken to be observations of one of six words (no, not, none, neither, never, nobody) occurring within three words preceding a positive word. Despite this attempt to control for negation, prior literature has generally found weaker results using positive words, which has been attributed to the failure to fully capture negation patterns. Unlike positive words, negative words are negated very rarely in the English language. Thus, we would not expect negations to bias the results of word lists that measure negative tone phrases and do not consider them for the negative word dictionary. Based on the LMD dictionaries, we calculate a positive and a negative tone measure, LMD^+ and LMD^- , by dividing the number of (non-negated) positive and negative words, respectively, from the word lists by the total number of words in an article. If there is more than one article on a firm in one newspaper in a given week, we use the mean across articles.

— Please insert FIGURE 2 approximately here —

In Figure 2 we plot the monthly average of LMD^- (across all firms) and changes in GDP. We observe that the two variables are quite highly negatively correlated (correlation coefficient is -0.4833): In times of poor economic conditions, newspapers use more negative words in articles on companies, and vice versa. The exceptionally low fraction of negative words during the Dotcom Bubble as well as the strong spike after the bankruptcy of Lehman Brothers are particularly noteworthy in this regard. This illustrates that our measure of tone, LMD^- , indeed captures variation in newspaper reports on companies. Together with Figure 1, it further highlights the importance of controlling for time effects.

Summary statistics on our advertising and press coverage variables are presented in Panel A of Table 2. Data are aggregated on the weekly level to ensure that a sufficient number of observations for each point in time is available.

— Please insert TABLE 2 approximately here —

There are two ways of constructing our main sample. In Panel A.1, we only include observations of firms that are mentioned in at least one newspaper article in a given week. Observations of firms that are not mentioned in the newspapers in our sample in a given week are dropped from the sample. In Panel A.2, we set the respective values for the latter equal to zero and thus also include these observations for the calculation of our summary statistics. We keep firms in the sample in between their first and last observation with positive advertising expenditure or newspaper article count.

In the first line we present average advertising expenditures over the week before a newspaper article is published on an advertising firm. We find that the dollar amount of advertising expenditures is much higher one week before a newspaper writes an article on a given firm (21,200 USD in Panel A) as compared to all weeks in the sample (2,440 USD in Panel B). A similar pattern is observed in the second line, where we compare the sum of advertising expenditures over a time period of four weeks before a newspaper article on a firm is published (84,810 USD in Panel A) to all four week periods in the sample (9,770 USD in Panel B). These patterns are expected, as larger firms are more important and of more interest to the general reader—and are thus covered more frequently by newspapers—and at the same time typically have larger advertising budgets.

Regarding our press coverage variables, we find that there are on average 1.57 articles on a firm in a given week in Panel A.1. The average article comprises 800 words. The fraction of negative words per article according to the LMD measure is 1.86%, while the fraction of positive words is 0.7%. Given that the negative word list comprises more words than the positive word list, the higher fraction of negative words than positive words in a newspaper article is not surprising.

While the tone measure based on the LMD dictionary has the appeal that it is straightforward and has a simple interpretation, it has two disadvantages: First, it does not take into account the length of an article, i.e. a very short article with 5% negative words receives the same score as a very long article with 5% negative words. Second, a realization of zero

for the negative (positive) measure indicates that there are no negative (positive) words in an article at all, that is, that the article is extremely positive (negative). Thus, zero can not be interpreted as a neutral realization of this measure and consequently it is not sensible to assign zero to cases where there are no articles on a specific firm in a newspaper in a certain week. Because some of our later empirical analysis requires uninterrupted time-series of our tone measure per company and newspaper, we also compute a media content measure that allows us to include these cases. Specifically, this is computed as minus one times the demeaned tone measure multiplied by the logarithm of an article’s number of words.¹⁰

$$(1) \quad MC = -1 \cdot (LMD^- - \overline{LMD^-}) \cdot \log(\text{word } \#).$$

This measure has the appealing features that (i) it is always decreasing in negativity and increasing in article length for above average tone articles and vice versa for below-average tone articles, (ii) that firm/newspaper/week/-combinations without any article can be assigned a neutral value of zero. This allows us to differentiate between the impact of, for example, a very short negative article and a very long negative article. Longer articles are presumably favorable only for good news, whereas a short negative article might be preferred to a long one.

2.3 Univariate differences in press coverage conditional on advertising

To test whether there are any significant differences in news coverage of firms that advertise in a given newspaper in a specific week and firms that do not advertise in that newspaper-week, we conduct a univariate analysis based on our press coverage variables. The sample

¹⁰The construction of all variables is described in more detail in Appendix A2 in the Supplementary Material.

includes those observations with at least one article about a firm within a newspaper-week combination.

We define a firm as an “Advertiser” at a given newspaper if it has positive advertisement expenditures in the previous one (or, alternatively, four) weeks at this newspaper. If the firm has not spent any money on advertising at a given newspaper over the past one (or, alternatively, four) weeks, we define it as a “Non-Advertiser”. Then, we analyze differences in press coverage and in our news tone measures between advertisers and non-advertisers. Results are presented in Panels B.1 and B.2 of Table 2.

Consistent with the differences found between Panels A.1 and A.2, results in Panel B.1 of Table 2 show that the number of articles published about a firm in a given newspaper is significantly larger if the firm has commissioned advertisement in this newspaper during the previous week. This difference is economically large; while there are on average 1.47 articles on a firm that did not advertise in a given newspaper, there are on average 1.93 articles if the firm did advertise over the past week. We observe an equally large difference if we look at the number of words in a given article (where the number of words is calculated conditional on at least one article on the firm appeared in the newspaper). While an article comprises about 1,040 words for firms classified as “Advertiser”, it only comprises about 740 words for firms classified as “Non-Advertiser”.

With respect to the tone of the newspaper articles, results based on the LMD^- tone measures indicate that articles on advertisers are significantly less negative than articles on non-advertisers. At the same time, newspaper articles on advertisers are significantly more positive than articles on non-advertisers. Our media content measure that takes into account the length of an article also portray a consistent picture by showing that advertisers receive more positive and less negative news coverage than non-advertisers.

We obtain very similar results if we compare differences in advertising expenditures over the past four weeks instead of one week before a newspaper article on a particular firm is published. Results are displayed in Panel B.2. We find that media coverage is higher for

advertisers than for non-advertisers. In addition, the fraction of positive words in an article is significantly higher for advertisers than for non-advertisers, while the opposite pattern is observed for the fraction of negative words. Results based on the media coverage measure confirm this finding.

Taken together, results in Table 2 provide first indicative evidence that a firm’s press coverage is correlated with its advertising expenditures at the corresponding newspaper. In the next step, we turn to a multivariate regression analysis including high dimensional fixed effects as an attempt to establish a causal link between lagged advertising expenditures and future media coverage.

3 The impact of advertising on press coverage and news tone

3.1 Main results

To investigate whether a newspaper reports differently about its advertising clients than about firms that do not advertise in this newspaper, we relate various proxies of a firm’s media coverage to its advertising expenditures in a given newspaper. The analysis is conducted on the firm-newspaper-week level. Therefore, we can include various combinations of fixed effects in our regressions, which will help us to identify a causal link between a firm’s advertising expenditures and its media coverage. We always include firm-week fixed effects to control for general firm characteristics (like its size) as well as for variations of the relevant information environment that impacts news reports, such as, for example, product launches, earnings announcements, corporate scandals, or other newsworthy events. This also controls for the effect on media coverage of retaining an investor relations firm documented in Solomon (2012). We argue that including fixed effects is preferable to explicitly controlling for firm characteristics and news events, as we would always face a severe omitted variables problem in the latter case. In contrast, firm-week fixed effects control for all fixed and time-varying firm-related variables that might have an impact on the quantity

and tone of news coverage. These interacted fixed effects subsume the time fixed effects that were shown to be important in our setting in Sections 2.1 and 2.2. We combine these firm-week fixed effects with two sets of additional fixed effects. In our main specification, we additionally include a dummy for each newspaper. Our regression model therefore is:

$$(2) \quad Media_{ijt} = \alpha + \beta_1 \cdot \text{Log}(4 - \text{week Ads}_{ijt}) + u_{it} + v_j + \varepsilon_{ijt}$$

, where i is the firm, j the newspaper and t the week. $Media_{ijt}$ is either the extent of media coverage measured by the number of articles written about a firm, the tone of coverage measured by LMD^- or media content measured by MC . $\text{Log}(4 - \text{week Ads})$ is the log of the sum of advertising over the previous four weeks. u_{it} is a firm-week interacted fixed effect and v_j is a newspaper fixed effect. The newspaper fixed effects account for different writing styles and tastes across newspapers. For example, some newspapers might in general be more critical of companies than others. Table 1 shows that there is indeed substantial heterogeneity across newspapers. This specification already controls for time-varying firm information as well as firm-invariant newspaper characteristics. In an even stricter specification, we include interacted firm-newspaper fixed effects (in addition to firm-week fixed effects). Hence, we substitute v_{ij} for v_j in Equation 2. This controls for any preferences certain newspapers might have for specific individual firms that might not be related to advertising. In this extremely restrictive specification, only time-varying variables in the firm-newspaper relation can drive our findings. Any results we obtain using this specification are a lower bound for the actual impact of advertising on news coverage and news tone as they are exclusively driven by within firm-newspaper pair time-series variation and none of the effect is driven by cross-sectional variations. Results from the two alternative specifications are presented in Table 3.

— Please insert TABLE 3 approximately here —

In columns (1) to (3), we present results from our main specification with media coverage (column (1)), media tone (column (2)) or media content (column (3)) as the dependent variables. We find that there are significantly more articles written on a firm in a given newspaper the higher the firm’s advertising expenditures in this newspaper over a time period of four weeks before the publication date are.¹¹ This impact is economically significant since our results suggest that spending 100,000 USD on advertising within four weeks implies an increase of 11.1% in the number of articles published on the firm in the following week. One might argue that all news create publicity for a firm and can thus be considered as good news. Nevertheless, we think it is more plausible that firms are mainly interested in positive coverage and try to avoid negative coverage. Thus, we turn to an analysis of the impact of advertising on article tone in column (2). We find that high advertising expenditures lead to a significantly lower fraction of negative words in articles subsequently published about the firm. In economic terms, the estimated coefficient implies that spending 100,000 USD on advertising within four weeks before the publication date of an article results in a reduction of negative words of 3.07%. For comparison, the effect of a one standard deviation change in GDP growth on aggregate newspaper article tone on firms is 3.78%, suggesting that the effect of advertising is also economically. In column (3), we consider the extent and the tone of coverage jointly by using our media content measure (MC) as the dependent variable. Again, we find a highly significant positive effect of advertising. This implies that articles on a given firm are longer and more positive after high advertising expenditures at a given newspaper.

In columns (4) to (6), we present results from our more restrictive specification using interacted firm-newspaper fixed effects. On top of the 328,208 dummies that are effectively included for the firm-week fixed effects, this specification adds another 47,672 dummies for the firm-newspaper pairs. Nevertheless, we still find a highly significant effect of advertising on media coverage, tone and content. Note that these estimates only serve as a lower

¹¹Alternatively, we compute the sum of a firm’s advertising expenditures in a given newspaper one week instead of four weeks before the newspaper publishes an article on the firm. Our results (not reported) are stable and even slightly stronger in economic terms.

bound of the overall effect due to the restrictiveness of the specification. The inclusion of firm-newspaper fixed effects forces identification to come exclusively from time variation in advertising and tone within a firm-newspaper pair. Therefore, if a firm spends a similar amount on advertising in a given paper throughout the sample period, the estimate of the effect would be close to zero. In our cross-sectional tests we will thus rely on the specification using newspaper fixed effects instead.

Taken together, results in Table 3 provide support for the view that firms can use advertising to influence the frequency and tone of articles that a newspaper publishes on the firm. In the next step, we explore the temporal dynamics of our main effect.

3.2 Temporal dynamics

There are two - not mutually exclusive, but equally worrisome - possible explanations for the advertising bias we document. First, newspapers might portray certain companies (too) positively in order to attract advertisements from those companies. This explanation suggests some kind of “anticipatory obedience” of newspapers. Second, newspapers might draw a positively biased picture of the firms that already advertise in them because advertising firms exert pressure on editors to get positive coverage and newspapers aim to keep these firms as advertising clients. Thus, the effect would be driven by direct “advertiser pressure”. To shed more light on these possibilities, we analyze the temporal dynamics of the relationship between advertising and news coverage. We implement a panel vector autoregressive (VAR) model to simultaneously relate advertising expenditures and news tone on leads and lags of the same as well as the other variable, respectively. If the “anticipatory obedience” channel is active, we would mainly expect an impact of lagged news tone on future advertising expenditures, if the “advertiser pressure” channel is active, we would mainly expect an impact of lagged advertising expenditures on future news tone.

Our VAR model follows the extension of standard vector autoregression models to the panel context by [Holtz-Eakin, Newey, and Rosen \(1988\)](#). It controls for firm-newspaper

fixed effects and corrects the dynamic panel bias analyzed in [Nickell \(1981\)](#).¹² Results are presented in Table 4.

— Please insert TABLE 4 approximately here —

Results in columns (1) and (2) of Table 4 show that lagged advertising has a significantly positive impact on the current number of articles on a given firm. At the same time, the lagged number of articles on a firm has a significantly positive impact on current advertising dollars spent by this firm. These findings suggest that causality goes from past coverage to future advertising and at the same time from past advertising to future coverage. However, joint coefficient tests show that the impact of lagged advertising on current media coverage is much stronger (χ^2 -stat=434.40) than the impact of lagged media coverage on current advertising expenditures (χ^2 -stat=156.12). The latter effect could be driven by firms realizing that a paper started to cover them in their articles more frequently and because of this eventually direct more advertising money to them to try to induce positive coverage.

With respect to media tone, results in columns (3) and (4) of Table 4 show that only past advertising expenditures have a significant impact on the tone of future newspaper reports (χ^2 -stat=22.02), while lagged media tone has no significant impact on future advertising expenditures (χ^2 -stat=5.16). This pattern is consistent with the notion that high advertising expenditures lead to a less negative tone of subsequent newspaper articles on the firm, but a less negative tone does not attract higher advertising expenditures in the future. Results in columns (5) and (6) portray a similar picture. While lagged advertising expenditures have a significant impact on future tone-adjusted media coverage (χ^2 -stat=36.06) we find no significant impact of lagged tone-adjusted media coverage on future advertising expenditures (χ^2 -stat=4.14).

¹²[Nickell \(1981\)](#) shows that the dynamic panel bias is inversely proportional to T. The length of our panel is T=732 weeks, so that one could argue that the bias should be small in our case. However, using a panel vector autoregressive model allows us to estimate both equations simultaneously for all firms. Therefore, we use this model even though the dynamic panel bias may be small.

Taken together, the results in this section suggest that the “advertiser pressure” channel is much stronger than the “anticipatory obedience” channel for media tone and tone-adjusted media coverage.¹³ That is, past advertising expenditures have a significant impact on how newspapers write about their advertising clients in the future. With respect to the frequency of media coverage, we can not entirely rule out “anticipatory obedience” as an additional explanation. However, the impact of lagged advertising on future media coverage is still stronger than vice versa.

Overall, we think that the results from the temporal dynamics analysis together with the high dimensional fixed-effects analysis in Section 3.1 address any omitted variables and reverse causality concerns to a large extent and allow us to make at least a big step towards identifying a causal relationship between advertising and newspaper tone.

3.3 Advertising bias and newspaper audience

In this section, we turn to the investigation of cross-sectional differences between newspapers. Newspapers operate in a two-sided market in which they sell news to readers and advertising space to advertisers (e.g. [Anderson and Gabszewicz \(2006\)](#), [Ellman and Germano \(2009\)](#), and [Germano and Meier \(2013\)](#)). Readers are likely to value accuracy, but might also be interested in entertainment. In a recent survey, [DellaVigna and La Ferrara \(2015\)](#) conclude that entertainment is the most important determinant of consumer demand for media like general interest television or radio outlets. For print publications like newspapers, demand for objective information might be relatively more important, and can vary across newspapers. For instance, [George and Waldfogel \(2003\)](#) and [George and Waldfogel \(2006\)](#) show that newspapers adapt their content to the demographic or educational characteristics of their potential readers, which vary across newspaper markets. Advertisers

¹³An observationally equivalent explanation to the “anticipatory obedience” explanation would be that journalists completely ignore the amount firms might spend on advertising, but firms direct advertising money to those newspapers that recently wrote in a more positive tone about them than others. However, given that we find no evidence for the anticipatory obedience explanation to start with, we also do not try to distinguish it further from this alternative story.

value reaching readers with their advertisements, but might also be interested in receiving favorable coverage by the newspaper (which is another, possibly even more effective, form of positive publicity). This market structure creates competing incentives for newspapers. On the one hand, newspaper content must be interesting enough to attract readers. On the other hand, newspapers have an incentive to please their advertisers by (positive) coverage to maximize advertising revenue. The outcome of the resulting trade-off depends on the value that the readers of a newspaper place on accurate and detailed reports as compared to, for example, entertainment. The evidence in Section 3.1 shows that newspapers on average solve the trade-off by biasing their reports in favor of advertisers to some extent. However, the readership of the newspapers in our sample likely differ in their preferences for accurate information. For example, accurate information on companies is most likely perceived to be even more crucial by readers of a business publication like the Wall Street Journal than to readers of a local newspaper like the New York Daily News. If the readers of a newspaper place a high value on objectivity, then maintaining a reputation for high reporting standards is crucial for a newspaper to retain its readership. We therefore test whether the strength of the advertising bias depends on newspapers’ audience.

As our first measure of newspapers’ audience, we use the complexity of articles in a newspaper. We hypothesize that the more complex the language used in a newspaper is, the more interested its readers are in detailed and objective information. In contrast, readers of newspapers that use simpler language might be relatively more interested in entertainment. We measure article complexity by the Fog-Index, a well-known measure of readability originally developed by Gunning (1952). It has been used extensively for financial text (see e.g. Li (2008), Dougal, Engelberg, Garcia, and Parsons (2012) or De Franco, Hope, Vyas, and Zhou (2014)).¹⁴ The index is calculated as:

¹⁴Loughran and McDonald (2014) argue that the fraction of complex words component of the Fog-Index can be problematic when applied to 10-K statements and suggest using the total file size (in megabytes) of the 10-K instead. However, Bonsall, Leone, and Miller (2015) find that the significance of the file size is mostly driven by non-text components of the 10-K document like XML code or pdfs. Such components are not applicable to newspaper articles. Hence, we use the more conventional Fog-Index here to focus on the readability of the text itself. Using just the number of words per sentence, which Loughran and McDonald (2014) do not find to be misspecified for 10-Ks, we obtain highly similar results, however.

$$(3) \quad Fog - Index = 0.4 * (Words\ per\ sentence + \frac{complex\ words\#}{word\#})$$

Complex words are those with more than two syllables. The index is scaled such that its value roughly corresponds to the years of formal education that a person needs to understand the text at a first reading. Hence, a value of twelve would indicate that someone roughly needs a U.S. high school degree to read the text. The writing style of a newspaper is presumably highly persistent. To ensure that this is indeed the case, we randomly draw 500 articles in each newspaper-year to compute the Fog-Index. We find that article complexity is highly persistent over time (autocorrelation coefficient of 0.8283) and therefore take the time series average for each newspaper. Table 1 shows that there is substantial dispersion in the Fog-Index across newspapers. The newspapers with the lowest values of the index are the New York Daily News and the Palm Beach Post, with a Fog-Index of 11.12 and 11.25, respectively. The Wall Street Journal has the highest Fog-Index with a value of 13.94.

To investigate whether newspapers differ in their reports on companies depending on the complexity of their articles' language, we first regress our measures of media reporting on lagged advertising and the Fog-Index. We include firm-week fixed effects, but omit newspaper fixed effects because these would subsume the coefficient on the Fog-Index. In the next step, we include an interaction between lagged advertising and the Fog-Index to test the hypothesis that advertising bias is weaker in newspapers with more complex language. Since we are interested in an interaction term in this specification, we can include newspaper fixed effects.

— Please insert TABLE 5 approximately here —

Results are presented in Panel A of Table 5. In columns (1) to (3), we report results on the baseline effect of the Fog-Index on media coverage (column (1)), tone (column (2))

and content (column (3)). In column (1), we find that newspapers with a higher Fog-Index write significantly more articles about companies. The coefficient estimate implies that the four newspapers with the highest Fog-Index write about 9.99% more business articles than those four newspapers with the lowest Fog-Index. This indicates that their readers are indeed more interested in business news. Moreover, column (2) shows that newspapers with a higher Fog-Index also write significantly more negatively about companies. Column (3) confirms these findings based on the media content measure. In columns (4) to (6), we present the results on the interaction term between lagged advertising and the Fog-Index. We find that the coefficients all point in the expected direction: the impact of advertising is smaller for newspapers with a higher Fog-Index. The difference is significant for LMD^- and MC , but not for the number of articles. This could imply that biased media coverage is perhaps less visible or less costly to readers. The estimated coefficient in the regressions using LMD^- and MC imply that advertising bias is about 60% smaller for those newspapers with the highest Fog-Index compared to those with the lowest index values. However, the sum of the coefficient on the baseline effect of lagged advertising and its interaction with the Fog-Index is significantly different from zero in all models, showing that even those newspapers with the most complex language exhibit advertising bias. Hence, although reader demand for detailed and objective information seems to counteract the pressure from advertisers for positive coverage, it is not sufficient to eliminate the bias.

Our second measure of newspapers' audience is whether they cater to a local or national readership. [Reuter and Zitzewitz \(2006\)](#) find that mutual fund recommendations are correlated with past advertising in three personal finance journals, but in none of the two national newspapers they analyze (New York Times and Wall Street Journal). They hypothesize that this might be due to higher reputational concerns of these newspapers compared to personal finance magazines. However, the authors also note that their data on newspaper coverage of mutual funds in national newspapers may be insufficient to detect the bias. Specifically, they only consider one column per newspaper and in the case of the Wall Street Journal

they are unable to differentiate positive and negative fund mentions.¹⁵ The structure of our data allows us to re-examine differences between local and national newspapers in more detail. This is because our data include all firm-related articles in a large number of local as well as all four large national newspapers. Moreover, we can relate firm-level advertising expenditures at individual newspapers to the extent and tone of media coverage of firms in this newspaper.

To investigate which individual newspapers are subject to advertising bias, we rerun our main regression for sub-samples of local and national newspapers. National newspapers are defined in a standard way as those with the highest circulation and that are distributed nationwide (i.e. New York Times, Wall Street Journal, USA Today, and Washington Post). All other newspapers in our sample are defined as local newspapers (see Table 1). Results are presented in Panel B of Table 5. The regression includes newspaper fixed effects as well as firm-week fixed effects.

Results in Panel B show that lagged advertising expenditures have a significant impact on advertising bias in both local as well as national newspapers. This result holds for the level of media coverage, the tone of the articles, as well as the tone-adjusted media coverage measure. We find that they are stronger for local as compared to national newspapers: A 100,000 USD increase in advertising spending leads to a 13.41% (5.53%) increase in the number of articles in local (national) newspapers. Furthermore, a 100,000 USD increase in advertising spending leads to a 4.83% (2.42%) reduction in the fraction of negative words in articles published in local (national) newspapers.

National newspapers might not be a homogeneous group, however. For instance, while the New York Times and the Wall Street Journal are both in the tenth decile of the Fog-Index distribution, the USA Today and the Washington Post are in the first and fifth decile,

¹⁵Similarly, [Gurun and Butler \(2012\)](#) provide some indirect evidence that local newspapers report about local firms in a biased way and argue that this might be due to advertising of local firms. They find no evidence for an impact of firm advertising budgets in national newspapers on the news tone about those firms in the lead paragraph of articles in the Wall Street Journal.

respectively. In addition, [George and Waldfogel \(2006\)](#) finds that the circulation of the New York Times is much more strongly positively correlated with average educational levels in a region than circulation of the USA Today. In the next step, we therefore estimate our main regressions (columns (1) to (3) in Table 3), adding an interaction of advertising expenditures with a vector of 41 dummy variables indicating each newspapers in our sample. This allows us to test for the impact of lagged advertising on future media coverage of advertising corporate clients for each of the newspapers individually. We report results on the combined coefficients of advertising and the interaction terms of the four national newspapers in Panel C. We find that lagged advertising expenditures have a significant impact on the number of articles published on firms for all national newspapers in our sample. With respect to the tone of the article and our tone-adjusted media coverage measures, we find a significant impact for the New York Times and the USA Today. Advertising bias is strongest for USA Today. We do not find a significant impact for the Wall Street Journal and the Washington Post.

Of course, each interaction term is estimated on a relatively small number of observations. We therefore conduct joint F-tests on the national newspapers' combined impact on the relation between past advertising and future media coverage. Results are reported in Panel D. Consistent with the results from Panel C, we find that the number of articles on an advertising corporate client is significantly related to this client's past advertising expenditures for all combinations of national newspapers. Regarding the tone of an article, we do not find a significant joint impact of the Wall Street Journal and Washington Post, while all possible combinations of three and four national newspapers yield significant results. The same result is observed for our tone-adjusted media coverage measure.

Taken together, the results in this section indicate that there is substantial heterogeneity in the degree of advertising bias across newspapers. This heterogeneity seems to be linked with newspapers' audience, indicating that reputational concerns drive the differences. How-

ever, our results suggest that even some of the largest national newspapers are not immune to advertising pressure from their corporate clients.

4 Advertising bias around corporate news events

In the final step, we try to better understand the channels through which advertising bias manifests. Specifically, we want to analyze whether tone is particularly biased after good or after bad news and whether bias manifests itself via less coverage of bad news events (‘selective omission’). To investigate these channels in more detail, we define two news events based on exogenous sources, that is, not based on press reports in newspapers: earnings announcements (Section 4.1) and days with extreme stock returns (Section 4.2). Earnings announcements are news events themselves, while extreme stock market returns might be newsworthy in themselves, but—more importantly—they usually happen when something else and newsworthy happened to the firm that heavily impacts its value.

4.1 Earnings announcements

Earnings announcements are scheduled news events that are easily observable and whose objective information content can easily be measured. Thus, they are well suited to further investigate how exactly newspapers bias reports about their corporate advertising clients around these events, particularly whether advertising bias differs after good or bad news. Earnings announcements are capital market events that generally trigger a lot of attention among investors ([Aharony and Swary \(1980\)](#)). There is a vast literature showing that stock markets react strongly upon the non-anticipated component of earnings announcements (for an overview, see [Kothari \(2001\)](#)). Thus, earnings announcements constitute an information

event that is particularly important for a firm and media coverage of earnings announcements can largely amplify a positive or negative investor reaction to the announcement.¹⁶

We use I/B/E/S analyst forecasts to compute quarterly earnings surprises for the firms in our sample. Specifically, we follow the previous literature on earnings announcements (DellaVigna and Pollet (2009)) and subtract the median I/B/E/S analyst forecast in the 30 days prior to the announcement from a firm’s announced earnings and divide this difference by the stock price 5 days prior to the announcement. Then, we sort firms into quintiles according to their earnings surprise (defined as standardized unexpected earnings, SUE). Summary statistics on our media coverage and advertising variables for the bottom and top earnings surprise quintile, respectively, are presented in Panel A of Table 6.

— Please insert TABLE 6 approximately here —

We focus on the most extreme quintiles as these are the cases when earnings strongly fell short of or strongly beat expectations, that is, these are cases with significant news content. Slightly beating or falling slightly short of expectations probably provides less of a newsworthy event. While our media coverage variables are computed over two days after an earnings announcement, advertising expenditures refer to a time period of 30 days before an earnings announcement. This choice of the time structure also ensures that earnings themselves are not endogenous with respect to advertising expenditures, because the earnings announced on a given earnings announcement day typically refer to a reporting period that ends earlier than 30 days prior to the announcement.

Results in Panel A show that there are slightly fewer articles on a firm if it belongs to the bottom quintile of earnings surprises in the preceding calendar month (1.37 articles) than if it belongs to the top quintile of earnings surprises (1.40 articles). As expected, there are more negative words in an article covering a firm in the bottom earnings surprise quintile (2.57%),

¹⁶Peress (2008) shows that announcements with more media coverage generate a stronger price and trading volume reaction at the announcement date and less subsequent drift.

than in an article covering a firm in the top earnings quintile (2.26%). A similar result obtains for our tone-adjusted media coverage measure. Furthermore, advertising expenditures 30 days before an earnings announcement are higher if a firm belongs to the top earnings quintile than if it belongs to the bottom earnings quintile. This is also not an unexpected pattern and shows that firms in a good economic situation generally advertise more.

More important in our context, we next investigate whether newspapers cover earnings announcements differently depending on whether the announcement firm belongs to its recent corporate advertising clients or not. We split our data into subsamples conditional on earnings surprise quintiles and re-run our main regression (column (1) in Table 3) with one of the media coverage variables as dependent variable. Results are presented in Table 7.

— Please insert TABLE 7 approximately here —

We find a significant impact of past advertising expenditures on the number of articles written about firms in the top as well as in the bottom earnings surprise quintile. One might have expected that newspapers just refrain from reporting about bad news events of its advertising clients. At least in the context of earnings announcements, we find no evidence of such “selective omission”.

Regarding the tone of an article, we find that articles are significantly less negative for firms in the bottom earnings surprise quintile the more this firm advertised in a given newspaper 30 days before its earnings announcement. There is no significant impact for firms with very good earnings news, i.e. firms in the top SUE quintile.¹⁷ We also observe a significant impact of past advertising on firms in the bottom SUE quintile for our tone-adjusted media coverage variable.¹⁸ No significant effect is found in the top SUE quintile.

¹⁷Results (not reported) for firms belonging to quintiles 2 to 4 are not statistically significant.

¹⁸We obtain similar results if we replace newspaper fixed effects by newspaper-industry fixed effects, while results are not significant anymore if we include newspaper-firm fixed effects. This is due to lack of statistical power, because the number of observations is extremely reduced in this case as we only focus on the two-day windows after the firms’ quarterly earnings announcements.

Results in Table 7 show that advertising bias is strongest for firms in the bottom earnings surprise quintile. Thus, newspapers mainly bias their coverage of bad news in favor of their corporate advertising clients, while reports on good news seem to be unbiased.

4.2 Extreme stock returns

An alternative way to define corporate news events based on daily stock returns is suggested by Barber and Odean (2008). This approach circumvents the challenge of defining specific news events and measuring their importance for a given firm. We follow this approach and sort firms according to their excess stock returns relative to the CRSP value-weighted index in the previous calendar month. We then classify all observations at the bottom 1% and the top 99% of the excess return distribution as potentially important information events. Summary statistics on these firms are reported in Panel B of Table 6. We find that there are on average 1.39 articles on firms that belong to the bottom percentile of the excess return distribution, while there are 1.34 articles on firms that belong to the top percentile of the excess return distribution. Not surprisingly, the fraction of negative words is higher for firms in the bottom percentile of the return distribution (2.50) than for firms in the top percentile (2.02).

In the next step, we run regressions based on subsamples of firms in the bottom and top return distribution percentile. Results are reported in Panel B of Table 7.

We find that newspapers write significantly more articles about a corporate advertising clients in the top but also in the bottom return distribution percentile. This pattern confirms our findings from above that there is no evidence for selective omission.

The tone of an article is significantly less negative for firms in the bottom return distribution percentile, a result that also holds for our tone-adjusted media coverage variable. At the same time, there is only a substantially smaller impact of past advertising on news tone for firms in the top return distribution percentile.

Taken together, the results in Table 7 portray a consistent picture independent of whether we define a corporate news event based on earnings announcements (Panel A) or extreme stock returns (Panel B): Advertising bias is strongest after negative news events of a firm. Advertising thus seems to allow firms to hedge against negative news coverage around bad corporate events and prevent an amplification effect of bad news due to negative media coverage.

5 Robustness

We investigate the robustness of our results to a variety of alternative specifications. The results of these tests are reported in Tables A1, A2, A3 and A4 in the Supplementary Material. Specifically, we conduct our analysis of the relationship between advertising and media coverage on the entire sample, i.e. we do not restrict our analysis to observations with at least one article in the firm-newspaper-week. We further use the number of words and the fraction of positive words as alternative measures of media coverage and tone. Moreover, we consider further alternatives to the inclusion of newspaper fixed effects that control for newspaper-specific industry preferences or potentially time-varying regional activity of firms.

Additionally, we consider alternative advertising horizons and use dummies instead of dollar expenditures. We also estimate a poisson model for the regression with the number of articles as the dependent variable and a tobit model for the regressions with media tone measures as the dependent variables. Finally, we estimate models using two-way clustering of the standard errors by firm and time, firm and newspaper, and newspaper and time, respectively. We find our results to be highly robust across the various specifications.

6 Conclusion

This paper investigates whether an advertising bias exists in the US newspaper industry. Our results provide support for the view that economic incentives arising from newspapers' revenue generating process indeed bias their reports. Newspapers report more frequently and less critically about their corporate advertising clients. This effect is mitigated but not eliminated by reputational concerns of newspapers. The effect is also more pronounced after bad corporate news events than after good corporate news events. Newspaper advertising thus seems to offer additional benefits for a firm beyond attracting new customers for their products: It can hedge a firm against bad media coverage if it has to announce bad news.

Overall, our results cast serious doubts on the independence of the press from the corporate world. While articles on advertising firms in local newspapers are particularly biased, we also find an advertising bias at least among some national newspapers.

Previous papers have argued that the rise of advertising in the nineteenth-century successfully created a press that is independent from political influence, because profits could then be generated from advertising revenues ([Petrova \(2011\)](#)). Our paper implies that regulatory policies might be needed to fully establish an independent press that conveys unbiased and accurate information to its readers. Specifically, our results suggest that it might make sense to require a stricter separation between the advertising department and the editorial department of newspapers.

References


- Aharony, Joseph and Itzhak Swary. 1980. "Quarterly Dividend and Earnings Announcements and Stockholders Returns: An Empirical Analysis." *Journal of Finance* 35 (1):1–12.
- Anderson, Simon P. and Jean J. Gabszewicz. 2006. "The Media and Advertising: A Tale of Two-Sided Markets." *Handbook of the Economics of Art and Culture* 1 (06):567–614.
- Barber, Brad and Terrance 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 (2):785–818.
- Blasco, Andrea, Paolo Pin, and Francesco Sobbrío. 2014. "Paying Positive to Go Negative: Advertisers' Competition and Media Reports." Working Paper.
- Bonsall, Samuel B V, Andrew J Leone, and Brian P Miller. 2015. "A Plain English Measure of Financial Reporting Readability." Working Paper.
- De Franco, Gus, Ole Kristian Hope, Dushyantkumar Vyas, and Yibin Zhou. 2014. "Analyst report readability." *Contemporary Accounting Research* 32 (1):76–104.
- de Smet, Dries and Stijn Vanormelingen. 2012. "The Advertiser is Mentioned Twice. Media Bias in Belgian Newspapers." Working Paper.
- DellaVigna, S and JM Pollet. 2009. "Investor inattention and Friday earnings announcements." *Journal of Finance* 64 (2):709–749.
- DellaVigna, Stefano and Matthew Gentzkow. 2010. "Persuasion: Empirical Evidence." *Annual Review of Economics* 2:643–669.
- DellaVigna, Stefano and Eliana La Ferrara. 2015. "Economic and Social Impact of the Media." In *Handbook of Media Economics*, edited by Simon P. Anderson, David Strömberg, and J. Waldfogel. Amsterdam: North-Holland, 1st ed.

- Di Tella, Rafael and Ignacio Franceschelli. 2011. "Government Advertising and Media Coverage of Corruption Scandals." *American Economic Journal: Applied Economics* 3 (October):119–151.
- Dougal, Casey, Joseph Engelberg, Diego Garcia, and Christopher A. Parsons. 2012. "Journalists and the Stock Market." *Review of Financial Studies* 25 (3):639–679.
- Ellman, Andrew and Fabrizio Germano. 2009. "What do the Papers Sell? A Model of Advertising and Media Bias." *Economic Journal* 119 (537):680–704.
- Engelberg, Joseph E. and Christopher A. Parsons. 2011. "The Causal Impact of Media in Financial Markets." *Journal of Finance* 66:67–97.
- Fang, Lily and Joel Peress. 2009. "Media Coverage and the Cross-section of Stock Returns." *Journal of Finance* 64 (5):2023–2052.
- Gambaro, Marco and Riccardo Puglisi. 2015. "What do ads buy? Daily coverage of listed companies on the Italian press." *European Journal of Political Economy* 39 (September):41–57.
- George, Lisa and Joel Waldfogel. 2003. "Who Affects Whom in Daily Newspaper Markets?" *Journal of Political Economy* 111 (4):765–784.
- . 2006. "The New York Times and the Market for Local Newspapers." *American Economic Review* 96 (1):435–447.
- Germano, Fabrizio and Martin Meier. 2013. "Concentration and self-censorship in commercial media." *Journal of Public Economics* 97:117–130.
- Gunning, Robert. 1952. *The Technique of Clear Writing*. New York, NY: McGraw-Hill.
- Gurun, Umit G. and Alexander W. Butler. 2012. "Don't believe the Hype: Local Media Slant, Local Advertising, and Firm Value." *Journal of Finance* 67 (2):561–597.

- Holtz-Eakin, D, W Newey, and HS Rosen. 1988. "Estimating vector autoregressions with panel data." *Econometrica* 56 (6):1371–1395.
- Honoré, BE. 1992. "Trimmed LAD and least squares estimation of truncated and censored regression models with fixed effects." *Econometrica* 60 (3):533–565.
- Kothari, S. P. 2001. "Capital markets research in accounting." *Journal of Accounting and Economics* 31 (1):105–231.
- Li, Feng. 2008. "Annual report readability, current earnings, and earnings persistence." *Journal of Accounting and Economics* 45 (2-3):221–247.
- Loughran, Tim and Bill McDonald. 2011. "When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks." *Journal of Finance* 66 (1):35–65.
- . 2014. "Measuring Readability in Financial Disclosures." *The Journal of Finance* 69 (4):1643–1671.
- Love, Inessa and Lea Zicchino. 2006. "Financial development and dynamic investment behavior: Evidence from panel VAR." *The Quarterly Review of Economics and Finance* 46 (2):190–210.
- Nickell, Stephen. 1981. "Biases in dynamic models with fixed effects." *Econometrica* 49 (6):1417–1426.
- Nyilasy, Gergely and Leonard N. Reid. 2011. "Advertiser Pressure and the Personal Ethical Norms of Newspaper Editors and Ad Directors." *Journal of Advertising Research* 51 (3):538–551.
- Peress, Joel. 2008. "Media coverage and investors' attention to earnings announcements." Working Paper.
- Petrova, Maria. 2011. "Newspapers and Parties: How Advertising Revenues Created an Independent Press." *American Political Science Review* 105 (4):790–808.


- Reuter, Jonathan. 2009. “Does Advertising Bias Product Reviews? An Analysis of Wine Ratings.” *Journal of Wine Economics* 4 (2):125–151.
- Reuter, Jonathan and Eric Zitzewitz. 2006. “Do Ads Influence Editors? Advertising and Bias in the Financial Media.” *Quarterly Journal of Economics* 121:197–227.
- Rochet, Jean-Charles and Jean Tirole. 2003. “Platform Competition in Two-Sided Markets.” *Journal of the European Economic Association* 1 (4):990–1029.
- Solomon, David H. 2012. “Selective Publicity and Stock Prices.” *Journal of Finance* 67 (2):599–638.
- Stroemberg, David. 2004. “Mass Media Competition, Political Competition, and Public Policy.” *Review of Economic Studies* 71 (1):265–284.
- Weis, William L. and Chauncey Burke. 1986. “Media Content and Tobacco Advertising: An Unhealthy Addiction.” *Journal of Communication* 36 (4):59–69.

Figure 1: Aggregate Advertising and GDP changes



This figure displays the evolution of total advertising spending and changes in GDP from 1999 until 2012. Monthly data is used to provide better readability.

Figure 2: Aggregate LMD^- and GDP changes



This figure displays the evolution of the cross-sectional average newspaper tone, measured by LMD^- , and changes in GDP from 1999 until 2012. LMD^- is measured by Loughran and McDonald's negative word list:

$$LMD^- = 100 * \frac{\text{negative word\#}}{\text{word\#}}$$
Monthly data is used to provide better readability.

Table 1: Summary statistics on newspapers

This table provides summary statistics on the newspapers used in our analysis. Article # is the total number of articles. LMD^- is measured by Loughran and McDonald’s negative word list: $LMD^- = 100 * \frac{\text{negative word\#}}{\text{word\#}}$. Media Content (MC) is calculated as $-1 * (LMD^- - \overline{LMD^-}) * \log(\text{word\#})$, where $\overline{LMD^-}$ is the sample mean. LMD^- and MC are sample means. Newspaper article complexity is measured by the Fog-Index, which is defined as $0.4 * (\text{Words per sentence} + \frac{\text{complex words\#}}{\text{word\#}})$. The decile rank according to the Fog-Index of a newspaper is given in parentheses.

Newspaper	Article #	LMD^-	MC	Fog-Index (Dec.)	Start	End
Arkansas Democrat-Gazette	22,051	1.60	1.45	13.05 (7)	1999	2010
Atlanta Journal-Constitution	53,322	1.77	0.20	12.47 (2)	1999	2012
Austin American-Statesman	23,414	1.65	1.10	13.00 (6)	1999	2012
Bergen Record	48,849	1.99	-1.12	12.91 (5)	1999	2012
Birmingham News	17,166	1.61	1.17	12.51 (3)	1999	2010
Boston Herald	24,076	1.95	-0.70	13.50 (9)	1999	2010
Buffalo News	28,365	1.59	1.28	12.55 (3)	1999	2012
Chicago Sun-Times	59,843	1.94	-0.68	13.12 (7)	1999	2010
Cleveland Plain Dealer	33,990	1.63	1.17	12.93 (5)	1999	2010
Dallas Morning News	59,392	1.47	2.11	12.90 (5)	1999	2011
Dayton Daily News	18,071	1.55	1.62	12.78 (4)	1999	2012
Denver Post	24,623	1.72	0.65	12.45 (2)	1999	2012
Fresno Bee	6,620	1.32	3.01	11.60 (1)	1999	2010
Houston Chronicle	65,288	2.05	-1.52	13.68 (10)	1999	2010
Las Vegas Review-Journal	11,998	1.71	0.69	13.00 (6)	1999	2012
Minneapolis Star-Tribune	25,375	1.69	0.80	13.14 (7)	1999	2012
New Orleans Times-Picayune	19,018	1.54	1.66	12.15 (2)	1999	2010
New York Daily News	21,371	1.87	-0.24	11.11 (1)	1999	2012
New York Post	37,317	2.16	-1.97	12.71 (3)	1999	2012
New York Times	142,528	1.89	-0.52	13.76 (10)	1999	2012
Newark Star-Ledger	33,899	1.92	-0.77	13.17 (8)	1999	2010
Norfolk Virginian Pilot	21,755	1.93	-0.69	12.99 (6)	1999	2012
Oklahoma Oklahoman	12,604	1.35	2.88	13.20 (8)	2003	2010
Palm Beach Post	21,114	1.61	1.32	11.25 (1)	1999	2012
Philadelphia Inquirer	51,510	1.96	-0.92	12.86 (4)	1999	2012
Pittsburgh Post-Gazette	42,135	2.00	-1.14	12.02 (2)	1999	2011
Portland Oregonian	11,623	1.74	0.38	13.31 (8)	1999	2010
Providence Journal	30,684	1.99	-1.01	13.06 (7)	1999	2010
Richmond Times-Dispatch	26,351	1.75	0.48	12.85 (4)	1999	2012
Sacramento Bee	12,796	1.69	0.69	12.68 (3)	2001	2010
Salt Lake Tribune	25,336	2.00	-1.21	13.44 (8)	1999	2012

continued on next page

continued from previous page

San Antonio Express News	27,629	1.65	1.05	13.66 (10)	1999	2010
San Francisco Chronicle	30,004	1.89	-0.42	13.49 (9)	1999	2010
San Jose Mercury News	76,436	1.94	-0.79	13.56 (9)	1999	2012
Seattle Post-Intelligencer	33,165	1.81	0.09	13.59 (9)	1999	2009
St. Louis Post-Dispatch	65,626	2.07	-1.65	12.83 (4)	1999	2012
Tampa Tribune	17,951	1.58	1.58	11.81 (1)	1999	2010
Tulsa World	57,258	1.81	-0.00	13.03 (6)	1999	2012
USA Today	34,431	1.87	-0.35	11.97 (1)	1999	2012
Wall Street Journal	109,135	2.24	-2.62	13.94 (10)	1999	2012
Washington Post	83,510	1.92	-0.73	12.96 (5)	1999	2012
Total	1,567,629	1.79	0.15	12.85	1999	2012

Table 2: Summary statistics

This table gives summary statistics and conditional mean comparisons. Panel A reports summary statistics. The sample in Panel A.1 is conditional on at least one article being published within a firm-newspaper-week, while Panel A.2 includes all observations. Panel B reports results from mean comparison tests over advertising status. A firm is considered an advertiser in a given newspaper if its lagged 1 (4)-week advertising is larger than zero in Panel B.1 (B.2). Article # is the number of articles published on a firm in a newspaper-week. I_{coverage} is a dummy equal to one if the number of articles written about a firm in a newspaper-week is larger than zero. Word # is the number of words in these articles. LMD^- and LMD^+ are measured by Loughran and McDonald's negative and positive word lists, respectively: $LMD^i = 100 * \frac{\text{word in list}\#^i}{\text{word}\#}$, where $i = -, +$. Media Content (MC) is calculated as $-1 * (LMD^- - \overline{LMD^-}) * \log(\text{word}\#)$, where $\overline{LMD^-}$ is the sample mean. 1 (4)-week Ads is the sum of advertising in the previous 1 (4) weeks. Further variable definitions can be found in Appendix A2. The unit of observation is the firm-newspaper-week. Standard errors in Panel B are clustered by firm.

Panel A: Summary statistics						
	Mean	Median	SD	75%ile	25%ile	N
Panel A.1: Conditional on at least one article						
1-week Ads (\$'000)	21.20	0	91.24	0	0	997,014
4-week Ads (\$'000)	84.81	0	335.57	10	0	992,075
Article #	1.57	1	1.44	2	1	998,481
Word # ('000)	0.80	1	1.04	1	0	998,481
Tone (LMD^-)	1.86	2	1.58	3	1	998,481
Tone (LMD^+)	0.70	1	0.62	1	0	998,481
Media content (MC)	-0.31	2	9.32	7	-5	998,481
Panel A.2: All observations						
1-week Ads (\$'000)	2.44	0	26.37	0	0	31,677,344
4-week Ads (\$'000)	9.77	0	95.64	0	0	31,526,652
Article #	0.05	0	0.37	0	0	31,727,448
I_{coverage}	0.03	0	0.17	0	0	31,727,448
Word # ('000)	0.03	0	0.23	0	0	31,727,448
Media content (MC)	-0.01	0	1.65	0	0	31,727,448

continued on next page

continued from previous page

Panel B: Conditional mean comparisons				
	Advertiser	Non-Advertiser	Difference	t-stat
Panel B.1: 1-week Advertising				
Article #	1.93	1.47	0.45	5.39
Word #	1.04	0.74	0.30	5.92
Tone (LMD ⁻)	1.78	1.89	-0.11	-3.18
Tone (LMD ⁺)	0.73	0.69	0.04	2.97
Media Content	0.27	-0.47	0.74	3.22
Panel B.2: 4-week Advertising				
Article #	1.90	1.44	0.46	6.46
Word #	1.02	0.72	0.30	6.96
Tone (LMD ⁻)	1.80	1.89	-0.09	-2.94
Tone (LMD ⁺)	0.73	0.69	0.04	3.17
Media Content	0.12	-0.50	0.62	2.98

Table 3: Media coverage, tone and content

This table reports the result of ordinary least squares regressions of media coverage, tone and content on lagged 4-week advertising. $\text{Log}(\text{Art. \#})$ is the log of the number of articles published on a firm in a newspaper-week. LMD^- is measured by Loughran and McDonald's negative word list: $\text{LMD}^- = 100 * \frac{\text{negative word\#}}{\text{word\#}}$. Media Content (MC) is calculated as $-1 * (\text{LMD}^- - \overline{\text{LMD}^-}) * \log(\text{word\#})$, where $\overline{\text{LMD}^-}$ is the sample mean. Firm-week and newspaper-firm are interacted fixed effects. Further variable definitions can be found in Appendix A2. The unit of observation is the firm-newspaper-week, conditional on at least one article being published. t statistics are provided in parentheses. Standard errors are clustered by firm. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

	Log(Art. #) (1)	LMD ⁻ (2)	MC (3)	Log(Art. #) (4)	LMD ⁻ (5)	MC (6)
Log(4-week Ads)	0.0229*** (8.17)	-0.0124*** (-4.38)	0.0822*** (4.93)	0.0034*** (3.17)	-0.0060*** (-2.63)	0.0360*** (2.67)
Firm-week FE	Yes	Yes	Yes	Yes	Yes	Yes
Newspaper FE	Yes	Yes	Yes	—	—	—
Newspaper-firm FE	—	—	—	Yes	Yes	Yes
No. obs.	992,075	992,075	992,075	992,075	992,075	992,075
R ²	0.393	0.534	0.533	0.562	0.569	0.568

Table 4: Temporal dynamics

This table reports the results of a panel VAR with advertising dollars and media measures as the components. The components of the VAR are weekly advertising as well as media coverage (VAR 1), tone (VAR 2) and content (VAR 3). $\text{Log}(\text{Art. \#})$ is the log of the number of articles published on a firm in a newspaper-week. LMD^- is measured by Loughran and McDonald's negative word list and demeaned: $\text{LMD}^- = 100 * \frac{\text{negativeword\#}}{\text{word\#}} - \overline{\text{LMD}^-}$, where $\overline{\text{LMD}^-}$ is the overall mean and observations with no articles are set to 0. Media Content (MC) 2 is calculated as $-1 * (\text{LMD}^- - \overline{\text{LMD}^-}) * \log(\text{word\#})$. The model is estimated following the extension of standard vector autoregression to the panel context by [Holtz-Eakin, Newey, and Rosen \(1988\)](#) and using the program developed by [Love and Zicchino \(2006\)](#). The given χ^2 -statistics and p-values are those from a joint test that the given coefficients are equal to 0 and have six degrees of freedom. Media refers to $\text{Log}(\text{Art. \#})$ in columns 1 and 2, LMD^- in columns 3 and 4 and MC in columns 5 and 6. Further variable definitions can be found in [Appendix A2](#). The unit of observation is the firm-newspaper-week, including observations with no articles being published. Standard errors are adjusted for heteroskedasticity. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

	VAR1: Media coverage		VAR2: Media tone		VAR3: Media content	
	$\text{Log}(\text{Art. \#})$	$\text{Log}(\text{Ads})$	LMD^-	$\text{Log}(\text{Ads})$	MC	$\text{Log}(\text{Ads})$
	(1)	(2)	(3)	(4)	(5)	(6)
Media						
Lag 1	0.1243***	0.0100***	0.0391***	-0.0000	0.0615***	0.0000
Lag 2	0.0870***	0.0037***	0.0263***	0.0005	0.0433***	0.0000
Lag 3	0.0743***	0.0054***	0.0228***	0.0003	0.0375***	0.0001
Lag 4	0.0695***	0.0043***	0.0205***	0.0011*	0.0344***	-0.0001
Lag 5	0.0655***	0.0041***	0.0205***	0.0006	0.0339***	-0.0001
Lag 6	0.0641***	0.0008	0.0176***	-0.0003	0.0305***	-0.0000
Ads						
Lag 1	0.0013***	0.2812***	-0.0006**	0.2813***	0.0046***	0.2813***
Lag 2	0.0000	0.1874***	-0.0007**	0.1874***	0.0049***	0.1874***
Lag 3	0.0002	0.0668***	-0.0003	0.0669***	0.0015	0.0669***
Lag 4	0.0006***	0.1054***	0.0006*	0.1055***	-0.0030*	0.1055***
Lag 5	0.0008***	0.0551***	0.0001	0.0551***	0.0001	0.0551***
Lag 6	0.0007***	0.0535***	-0.0004	0.0535***	0.0029	0.0535***
N	31,106,955	31,106,955	31,106,955	31,106,955	31,106,955	31,106,955
χ^2 -stat (Media=0)	129,876.00	156.12	6,986.34	5.16	20,662.98	4.14
p-value (Media=0)	0.00	0.00	0.00	0.52	0.00	0.66
χ^2 -stat (Ads=0)	434.40	810,738.06	22.02	809,797.14	36.06	809,753.04
p-value (Ads=0)	0.00	0.00	0.00	0.00	0.00	0.00

Table 5: Advertising bias and newspaper audience

This table examines the relationship between newspaper audience and advertising bias. In Panel A, we proxy for audience by the complexity of articles in a newspaper. Newspaper article complexity is measured by the Fog-Index, which is defined as $0.4 * (Words\ per\ sentence + \frac{complex\ words\#}{word\#})$. Fog-Index^R is the decile rank of each newspaper, scaled to range from 0 to 1. Log(Art. #) is the log of the number of articles published on a firm in a newspaper-week. LMD⁻ is measured by Loughran and McDonald's negative word list: $LMD^- = 100 * \frac{negative\ word\#}{word\#}$. Media Content (MC) is calculated as $-1 * (LMD^- - \overline{LMD^-}) * \log(word\#)$, where $\overline{LMD^-}$ is the sample mean. In Panels B to D, we separately consider newspapers with a local and national audience. Panel B reports the results from regressions of media coverage, tone and content on lagged 4-week advertising separately for local and national newspapers. National newspapers are Wall Street Journal, New York Times, USA Today and the Washington Post. Panel C reports the combined coefficients on the four national newspapers from the following regression (run using all 41 newspapers): $Media = \alpha_0 + \alpha_1 * Ads + \beta * Ads * Paper + Paper$. **Paper** is a 1x41 vector of newspaper dummies. *Media* is either Log (Art. #), LMD⁻ or MC. Panel D reports F-Tests on the national newspapers' combined coefficients ($\alpha_1 + \beta$) from this regression. Firm-week is an interacted fixed effect. Further variable definitions can be found in Appendix A2. The unit of observation is the firm-newspaper-week, conditional on at least one article being published. t statistics are provided in parentheses. Standard errors are clustered by firm. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Article complexity						
	Log(Art. #) (1)	LMD ⁻ (2)	MC (3)	Log(Art. #) (4)	LMD ⁻ (5)	MC (6)
Log(4-week Ads)	0.0289*** (9.17)	-0.0052* (-1.81)	0.0348* (1.96)	0.0249*** (5.32)	-0.0210*** (-5.91)	0.1392*** (6.14)
Fog-Index ^R	0.1123*** (5.65)	0.2259*** (14.60)	-1.3628*** (-15.58)	— —	— —	— —
Fog-Index ^R × Ads				-0.0032 (-0.50)	0.0138*** (2.87)	-0.0915*** (-3.22)
Firm-week FEs	Yes	Yes	Yes	Yes	Yes	Yes
Newspaper FEs	—	—	—	Yes	Yes	Yes
No. obs.	992,075	992,075	992,075	992,075	992,075	992,075
R ²	0.376	0.528	0.527	0.393	0.534	0.534
Panel B: Newspaper subsamples						
	Log(Art. #)		LMD ⁻		MC	
	Local	National	Local	National	Local	National
Log(4-week Ads)	0.0273*** (6.26)	0.0117*** (4.21)	-0.0191*** (-5.80)	-0.0106** (-2.03)	0.1225*** (5.50)	0.0876*** (3.15)
Firm-week FEs	Yes	Yes	Yes	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes
No. obs.	777,736	214,339	777,736	214,339	777,736	214,339
R ²	0.393	0.724	0.553	0.789	0.551	0.782

continued on next page

continued from previous page

Panel C: Advertising-newspaper interactions, combined coefficients						
	Log(Art. #)		LMD ⁻		MC	
	$\alpha_1 + \beta$	p-value	$\alpha_1 + \beta$	p-value	$\alpha_1 + \beta$	p-value
New York Times	0.0314	0.00	-0.0116	0.00	0.0702	0.01
USA Today	0.0167	0.00	-0.0172	0.00	0.1068	0.00
Wall Street Journal	0.0201	0.00	-0.0027	0.71	0.0176	0.60
Washington Post	0.0277	0.00	-0.0097	0.16	0.0657	0.16

Panel D: Advertising-newspaper interactions, F-Tests						
	Log(Art. #)		LMD ⁻		MC	
	F-Stat	p-value	F-Stat	p-value	F-Stat	p-value
All nationals:	14.54	0.00	5.14	0.00	4.92	0.00
NYT, WSJ, WP:	16.06	0.00	4.46	0.00	4.13	0.01
USAT, WSJ, WP:	15.79	0.00	4.18	0.01	3.98	0.01
NYT, USAT, WP:	19.14	0.00	6.75	0.00	6.45	0.00
NYT, USAT, WSJ:	15.16	0.00	5.38	0.00	4.86	0.00
WSJ and WP:	18.93	0.00	1.12	0.33	1.18	0.31

Table 6: Media and corporate events - Summary statistics

This table gives summary statistics by earnings surprise quintiles in Panel A and extreme stock returns in Panel B. In Panel A, quintile cutoffs are based on the distribution of earnings surprises in the preceding calendar month. Standardized earnings surprise (SUE) is defined as $\frac{A_q - E_q}{P_q}$, where A_q is the actual earnings, E_q the median analyst forecast in the 30 days prior to the announcement from I/B/E/S, and P_q the stock price 5 days prior to the announcement from CRSP [DellaVigna and Pollet \(2009\)](#). Article # is the number of articles published on a firm in a newspaper-week. LMD^- is measured by Loughran and McDonald's negative word list: $LMD^- = 100 * \frac{\text{negative word\#}}{\text{word\#}}$. Media Content (MC) is calculated as $-1 * (LMD^- - \overline{LMD^-}) * \log(\text{word\#})$, where $\overline{LMD^-}$ is the sample mean. Article #, LMD^- and MC2 are based on days zero to two relative to an earnings announcement or extreme return event. 30-day Ads is the sum of advertising over the 30 days prior to the event. In Panel B, we consider an excess return (in excess of CRSP value-weighted index), denoted by Ex. Return, to be extreme if it is in the 5th or 95th percentile of the excess return distribution of the preceding calendar month. Further variable definitions can be found in [Appendix A2](#). The unit of observation is the firm-newspaper-earnings surprise day in Panel A and the firm-newspaper-extreme return day in Panel B, conditional on at least one article being published.

Panel A: Earnings surprises						
	Mean	Median	SD	25%ile	75%ile	N
1 st quintile						
SUE	-3.72	-0	28.37	-1	-0	12,420
Article #	1.37	1	1.00	1	1	12,420
LMD^-	2.57	2	1.98	1	3	12,420
MC	-3.16	-2	9.75	-9	5	12,420
30-day Ads	71.26	0	299.87	0	5	12,420
5 th quintile						
SUE	2.32	1	23.55	0	1	12,245
Article #	1.40	1	1.14	1	1	12,245
LMD^-	2.26	2	1.84	1	3	12,245
MC	-1.43	0	9.09	-7	6	12,245
30-day Ads	99.14	0	385.13	0	15	12,245

continued on next page

continued from previous page

Panel B: Extreme stock returns						
	Mean	Median	SD	25%ile	75%ile	N
5 th percentile						
Ex. Return [%]	-11.09	-9	8.70	-13	-6	67,165
LMD ⁻	2.50	2	1.92	1	3	67,165
Article #	1.39	1	1.14	1	1	67,165
MC	-1.21	0	9.88	-7	7	67,165
30-day Ads	58.46	0	264.25	0	0	67,165
95 th percentile						
Ex. Return [%]	12.10	9	11.46	6	14	70,739
LMD ⁻	2.02	2	1.72	1	3	70,739
Article #	1.34	1	1.01	1	1	70,739
MC	1.37	3	9.07	-4	8	70,739
30-day Ads	57.09	0	275.03	0	0	70,739

Table 7: Media and corporate events - Fixed effects analysis

This table reports the results of ordinary least squares regressions of media coverage, tone and content on lagged 30-day advertising within subsamples based on financial market events. Panel A considers earnings announcements. Quintile cutoffs are based on the distribution of earnings surprises in the preceding calendar month. Panel B considers extreme stock returns, where extreme returns are defined using the 5th and 95th percentiles from the preceding calendar month. Log(Art. #) is the log of the number of articles published on a firm in a newspaper-week. LMD^- is measured by Loughran and McDonald's negative word list: $LMD^- = 100 * \frac{negative\ word\#}{word\#}$. Media Content (MC) is calculated as $-1 * (LMD^- - \overline{LMD^-}) * \log(word\#)$, where $\overline{LMD^-}$ is the sample mean. Log(Art. #), LMD^- and MC2 are based on days zero to two relative to an earnings announcement or extreme return event. Log(30-day Ads) is the log of the sum of advertising over the 30 days prior to the event. Further variable definitions can be found in Appendix A2. The unit of observation is the firm-newspaper-earnings surprise day in Panel A and the firm-newspaper-extreme return day in Panel B, conditional on at least one article being published. t statistics are provided in parentheses. Standard errors are clustered by firm. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Earnings surprise quintiles						
	Log(Art. #)		LMD^-		MC	
	1 st	5 th	1 st	5 th	1 st	5 th
Log(30-day Ads)	0.0190*** (6.56)	0.0190*** (6.47)	-0.0442*** (-2.88)	0.0110 (0.63)	0.1985*** (3.02)	-0.0618 (-0.78)
Firm-event FEs	Yes	Yes	Yes	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes
No. obs.	12,420	12,245	12,420	12,245	12,420	12,245
R ²	0.371	0.362	0.543	0.517	0.545	0.517
Panel B: Extreme stock returns						
	Log(Art. #)		LMD^-		MC	
	5 th	95 th	5 th	95 th	5 th	95 th
Log(30-day Ads)	0.0240*** (8.22)	0.0196*** (6.49)	-0.0291*** (-4.53)	-0.0103 (-1.62)	0.1288*** (4.23)	0.0757** (2.43)
Firm-event FEs	Yes	Yes	Yes	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes
No. obs.	67,165	70,739	67,165	70,739	67,165	70,739
R ²	0.418	0.400	0.618	0.617	0.623	0.622

Appendix

This appendix provides supplementary information to the paper "A Friendly Turn: Advertising Bias in the News Media" by Florens Focke, Alexandra Niessen-Ruenzi and Stefan Ruenzi.

A1 Description of data cleansing process

To merge our advertising data with newspaper articles from the data providers we use (LexisNexis and Factiva), we follow the procedure applied by Fang and Peress (2009). We exemplify the procedure for LexisNexis. LexisNexis provides company identifiers for each article. In addition, a "relevance score" indicates how strong a given article is related to a specific firm. This score is based on criteria such as the keyword's frequency, and its weight and location within the document. According to LexisNexis, an article is classified as a "major reference" with respect to a given firm, if its relevance score is at least 85%.¹⁹ We manually checked 1,000 articles and find that a relevance score of 80% still correctly classifies articles that are mainly about a given firm. To maintain a reasonably large sample size for identification purposes, we thus keep all articles on a firm that have a relevance score of at least 80% in our sample.

In the next step, we drop all sponsored articles. These articles are explicitly flagged in LexisNexis. In addition, all duplicate articles are deleted from the sample as well. We identify duplicates as follows. First, articles that are identical with respect to the company covered, newspaper, day of the week and number of words are classified as duplicates. Out of two duplicates, we keep the one that was first added to the database (the variable "Load Date" allows us to identify the exact date when an article was added to the database). Second, articles that are identical with respect to the company covered, newspaper, day of the week, and headline are classified as duplicated. These duplicates frequently consist in

¹⁹See http://wiki.lexisnexis.com/academic/index.php?title=Relevance_Score

an (often shorter) online version and a print version of the same article. We drop online versions of these duplicates from our sample if there is also a print version of the same article. If there are still duplicates left with respect to company covered, newspaper, day of the week, and headline, we keep the longest article with the earliest load date.

Finally, we aggregate articles on the firm-newspaper-week level by computing the mean tone and the sum of all articles.

A2 Brief definitions and sources of main variables

Advertising Advertising spent in a specific newspaper by a firm. Data is obtained from Kantar Media Strategy. $\text{Log}(1\text{ (4)-week Ads})$ is the natural log of the sum of advertising expenditures within a firm-newspaper pair over the preceding 1 (4) weeks. $\text{Log}(30\text{-day Ads})$ is the natural log of the sum of advertising expenditures within a firm-newspaper pair over the preceding 30 days.

Article # The number of articles written about a company in a newspaper-week.

Excess Return The excess return of firm i at time t is defined as its stock return in excess of the CRSP value-weighted index: $ExcessReturn_{i,t} = Return_{i,t} - CRSP_{VW}$.

Fog-Index This readability measure is calculated as $0.4 * (\text{Words per sentence} + \frac{\text{complex words\#}}{\text{word\#}})$. Complex words are those that contain more than two syllables.

Fog-Index^R Fog-Index^R is the decile rank of each newspaper, scaled to range from 0 to 1.

I_{coverage} Dummy equal to one if the number of articles written about a firm in a newspaper-week is larger than zero.

LMD⁻ Negative article tone as measured by the negative word list developed in [Loughran and McDonald \(2011\)](#): $LMD^- = 100 * \frac{\text{negativeword\#}}{\text{totalword\#}}$.

LMD⁺ Positive article tone as measured by the positive word list developed in [Loughran and McDonald \(2011\)](#): $LMD^+ = 100 * \frac{\text{positiveword\#}}{\text{totalword\#}}$.

Log(Art. #) The log of the number of articles written about a company in a newspaper-week plus one.

Log(Word #) The log of the number of words written about a company in a newspaper-week plus one.

Media Content Media Content (MC) is calculated as $-1 * (LMD - \overline{LMD}) * \log(word\#)$, where $\overline{LMD}$ is the overall mean of LMD. MC is decreasing in negativity and increasing in article length only for less negative than average news, otherwise it is decreasing. Higher values indicate more favorable media coverage.

Standardized earnings surprise (SUE) SUE is defined as $\frac{A_q - E_q}{P_q}$, where A_q is the actual earnings, E_q the median analyst forecast in the 30 days prior to the announcement from I/B/E/S, and P_q the stock price 5 days prior to the announcement from CRSP [DellaVigna and Pollet \(2009\)](#).

Word # The number of words written about a company in a newspaper.

Table A1: Robustness - Further results

This table provides additional results. In Panel B, we use the entire sample, i.e. we do not restrict our analysis to observations with at least one article in the firm-newspaper-week. For summary statistics on this sample, see Table 2, Panel A.2. Our dependent variables are $\text{Log}(\text{Art. \#})$, which is the logarithm of the number of articles written about a firm in a newspaper-week, and I_{coverage} , which is a dummy equal to one if the number of articles written about a firm in a newspaper-week is larger than zero. In Panel B, we use two alternative dependent variables: Word #, which is the number of words written about a firm in a newspaper-week, and LMD^+ , which is measured by Loughran and McDonald's positive word list: $\text{LMD}^+ = 100 * \frac{\text{positive word\#}}{\text{word\#}}$. The unit of observation is the firm-newspaper-week in Panel A and B, conditional on at least one article being published in Panel B only. Further variable definitions can be found in Appendix A2. t statistics are provided in parentheses. Standard errors are clustered by firm. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Entire sample				
	Log(Art. #) (1)	I_{coverage} (2)	Log(Art. #) (3)	I_{coverage} (4)
Log(4-week Ads)	0.0164*** (11.02)	0.0148*** (12.97)	0.0016*** (3.97)	0.0010*** (2.86)
Firm-week FE	Yes	Yes	Yes	Yes
Newspaper FE	Yes	Yes	–	–
Newspaper-firm FE	–	–	Yes	Yes
No. obs.	31,526,652	31,526,652	31,526,652	31,526,652
R ²	0.337	0.308	0.482	0.407
Panel B: Alternative dependent variables				
	Log(Word #) (1)	LMD^+ (2)	Log(Word #) (3)	LMD^+ (4)
Log(4-week Ads)	0.0431*** (8.83)	0.0026*** (2.64)	0.0042** (2.38)	-0.0005 (-0.68)
Firm-week FE	Yes	Yes	Yes	Yes
Newspaper FE	Yes	Yes	–	–
Newspaper-firm FE	–	–	Yes	Yes
No. obs.	992,075	992,075	992,075	992,075
R ²	0.649	0.457	0.721	0.494

Table A2: Robustness - Alternative fixed effects specifications

This table provides results from alternative fixed effects specifications. In Panel A, we present results from regressions including interacted newspaper-industry fixed effects, where we use the Fama-French 12 industries in columns (1) to (3) and SIC 2-digit industries in columns (4) to (6). In Panel B, we present results from regressions using firm-newspaper location-year fixed effects, where location is measured at the state level in columns (1) to (3) and at the city level in columns (4) to (6). Firm-week, firm-state-year and firm-city-year are interacted fixed effects. LMD^- is measured by Loughran and McDonald's negative word list: $LMD^- = 100 * \frac{\text{negative word\#}}{\text{word\#}}$. Media Content (MC) is calculated as $-1 * (LMD^- - \overline{LMD^-}) * \log(\text{word\#})$, where $\overline{LMD^-}$ is the sample mean. Further variable definitions can be found in Appendix A2. The unit of observation is the firm-newspaper-week, conditional on at least one article being published. t statistics are provided in parentheses. Standard errors are clustered by firm. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Newspaper-specific industry preferences						
	Log(Art. #) (1)	LMD^- (2)	MC (3)	Log(Art. #) (4)	LMD^- (5)	MC (6)
Log(4-week Ads)	0.0192*** (7.51)	-0.0125*** (-4.46)	0.0763*** (4.46)	0.0188*** (7.99)	-0.0121*** (-4.22)	0.0730*** (4.15)
Firm-week FE	Yes	Yes	Yes	Yes	Yes	Yes
Newspaper-FF12 FE	Yes	Yes	Yes	—	—	—
Newspaper-SIC2 FE	—	—	—	Yes	Yes	Yes
No. obs.	842,472	842,472	842,472	842,472	842,472	842,472
R ²	0.447	0.529	0.529	0.462	0.533	0.533
Panel B: Time-varying regional activity of firms						
	Log(Art. #) (1)	LMD^- (2)	MC (3)	Log(Art. #) (4)	LMD^- (5)	MC (6)
Log(4-week Ads)	0.0145*** (6.35)	-0.0100*** (-3.23)	0.0731*** (3.94)	0.0139*** (5.46)	-0.0120*** (-3.26)	0.0770*** (3.45)
Firm-week FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-state-year FEs	Yes	Yes	Yes	—	—	—
Firm-city-year FEs	—	—	—	Yes	Yes	Yes
No. obs.	992,075	992,075	992,075	992,075	992,075	992,075
R ²	0.558	0.606	0.635	0.613	0.634	0.605

Table A3: Robustness - Alternative advertising definitions

This table examines the robustness of our results to variations in the definition of advertising. In Panel A, we use alternative time horizons to measure advertising. In Panel B, we use dummy variables equal to one if the firm advertised in the newspaper in the previous 4 (8) weeks. Log(4 (8)-week Ads) is the log of the advertising expenditure over the previous 4 (8) weeks. Log(5-8-week Ads) is the log of advertising expenditures between weeks 5 and 8. LMD^- is measured by Loughran and McDonald's negative word list: $LMD^- = 100 * \frac{negative\ word\#}{word\#}$. Media Content (MC) is calculated as $-1 * (LMD^- - \overline{LMD^-}) * \log(word\#)$, where $\overline{LMD^-}$ is the sample mean. Firm-week is an interacted fixed effect. Further variable definitions can be found in Appendix A2. The unit of observation is the firm-newspaper-week, conditional on at least one article being published. t statistics are provided in parentheses. Standard errors are clustered by firm. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Alternative advertising horizons						
	Log(Art. #) (1)	LMD^- (2)	MC (3)	Log(Art. #) (4)	LMD^- (5)	MC (6)
Log(8-week Ads)	0.0231*** (8.56)	-0.0118*** (-4.55)	0.0764*** (4.92)			
Log(4-week Ads)				0.0139*** (8.50)	-0.0089*** (-3.69)	0.0596*** (4.44)
Log(5-8-week Ads)				0.0128*** (7.46)	-0.0050** (-2.39)	0.0321** (2.55)
Firm-week FEs	Yes	Yes	Yes	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes
No. obs.	992,075	992,075	992,075	992,075	992,075	992,075
R ²	0.395	0.534	0.533	0.394	0.534	0.533
Panel B: Advertising dummies						
	Log(Art. #) (1)	LMD^- (2)	MC (3)	Log(Art. #) (4)	LMD^- (5)	MC (6)
I _{4-week Ads > 0}	0.0847*** (6.99)	-0.0460*** (-4.03)	0.3116*** (4.48)			
I _{8-week Ads > 0}				0.0926*** (7.71)	-0.0465*** (-4.37)	0.3011*** (4.64)
Firm-week FEs	Yes	Yes	Yes	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes
No. obs.	992,075	992,075	992,075	992,075	992,075	992,075
R ²	0.391	0.534	0.533	0.392	0.534	0.533

Table A4: Robustness - Alternative econometric models

This table examines the robustness of our results to alternative econometric models. In Panel A, we consider a panel poisson regression on the number of articles in column (1), and a panel tobit regression as in [Honoré \(1992\)](#) in column (2) and (3). To make the coefficient from the poisson regression comparable to the ordinary least squares results, the table displays $e^\beta - 1$, i.e. the incident rate ratio minus one. In Panel B, we cluster by firm and time, in Panel C, we cluster by firm and newspaper and in Panel D, we cluster by newspaper and time. LMD^- is measured by Loughran and McDonald's negative word list: $LMD^- = 100 * \frac{\text{negative word\#}}{\text{word\#}}$. Media Content (MC) is calculated as $-1 * (LMD^- - \overline{LMD^-}) * \log(\text{word\#})$, where $\overline{LMD^-}$ is the sample mean. Firm-week is an interacted fixed effect. Further variable definitions can be found in [Appendix A2](#). The unit of observation is the firm-newspaper-week, conditional on at least one article being published. t statistics are provided in parentheses. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Non-linear models			
	Article # (1)	LMD^- (2)	LMD^+ (3)
Log(4-week Ads)	0.0477*** (8.04)	-0.0048*** (-4.26)	0.0062*** (12.34)
Firm-week FEs	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes
No. obs.	813,177	992,075	992,075
Panel B: Clustering by firm and time			
	Log(Art. #) (1)	LMD^- (2)	MC (3)
Log(4-week Ads)	0.0229*** (6.68)	-0.0124*** (-3.55)	0.0822*** (3.99)
Firm-week FEs	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes
No. obs.	992,075	992,075	992,075
R ²	0.393	0.534	0.533
Panel C: Clustering by firm and newspaper			
	Log(Art. #) (1)	LMD^- (2)	MC (3)
Log(4-week Ads)	0.0229*** (5.40)	-0.0124*** (-2.75)	0.0822*** (3.13)
Firm-week FEs	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes
No. obs.	992,075	992,075	992,075
R ²	0.393	0.534	0.533

continued on next page

continued from previous page

Panel D: Clustering by newspaper and time			
	Log(Art. #) (1)	LMD ⁻ (2)	MC (3)
Log(4-week Ads)	0.0229*** (6.66)	-0.0124*** (-2.99)	0.0822*** (3.39)
Firm-week FEs	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes
No. obs.	992,075	992,075	992,075
R ²	0.393	0.534	0.533