

Journalists and the Stock Market *

Casey Dougal, Joseph Engelberg, Diego García, and Christopher A. Parsons[†]

August 23, 2011

Abstract

We use exogenous scheduling of Wall Street Journal columnists to identify a causal relation between financial reporting and stock market performance. To measure the media's unconditional effect, we add columnist fixed effects to a daily regression of excess Dow Jones Industrial Average returns. Relative to standard control variables, these fixed effects increase the R^2 by about 35 percent, indicating each columnist's average persistent "bullishness" or "bearishness." To measure the media's conditional effect, we interact columnist fixed effects with lagged returns. This increases explanatory power by yet another one-third, and identifies amplification or attenuation of prevailing sentiment as a tool used by financial journalists.

Keywords: *Media Bias, Market Efficiency, Financial Journalism.*

JEL: H53, I38, J31, J33.

*We thank Mark Grinblatt, Alexandra Niessen-Runzi (discussant), Stephan Siegel, Denis Sosyura (discussant), Paul Tetlock, Sheridan Titman, and seminar participants at Boston College, INSEAD, University of Arizona, UNC, University of Toronto (Economics), SFS Cavalcade, Napa Conference on Financial Markets Research, and Helsinki Finance Summit for helpful feedback and suggestions. The usual disclaimer applies.

[†]All authors are from the Kenan-Flagler Business School, University of North Carolina. Casey Dougal, (Email) casey_dougal@kenan-flagler.unc.edu, (Tel) 208-284-3117; Joseph Engelberg, (Email) joseph_engelberg@kenan-flagler.unc.edu, (Tel) 919-962-6889; Diego García, (Email) diego_garcia@unc.edu, (Tel) 919-962-8404; Christopher Parsons, (Email) chris_parsons@kenan-flagler.unc.edu, (Tel) 919-962-4132.

1 Introduction

The media is often modeled as a faceless institution, but its main product – news content – is generated by people. This is important because unlike, say, making tires or processing paper, writing is a fiercely individualistic craft that allows an author’s style, persuasion, views, or bias to be injected into the finished product. In this paper, we present direct evidence that the writing of specific journalists has a *casual* effect on aggregate market outcomes.

This is surprising because, at any point in time, individual columnists are unlikely to possess information relative to the market as a whole, let alone consistently over a period of several years. Thus, any persistent return predictability related to specific authors must arise from their “sentiment” or spin of public events. From 1970 to 2007, we find that the short-term returns on the Dow Jones Industrial Average (DJIA) can be predicted using only the author of a widely read market summary article, the *Wall Street Journal’s* “Abreast of the Market” (AOTM) column.

Ordinarily we would be concerned about the endogenous nature of news coverage in an article summarizing market events. As Tetlock (2007, p. 1139) notes, “It is unclear whether the financial news media induces, amplifies, or simply reflects investors’ interpretations of stock market performance.” Making a distinction between a reflective and a causal role for financial media thus requires exogenous variation in news content, or reporting uncorrelated with underlying events.

Our setting is particularly useful in this regard. Over the nearly four decades we study, columnists rotate frequently – three different journalists write the AOTM column in the typical month – and often according to regular schedules. Moreover, journalists differ markedly in their writing styles such as sentence structure, complexity, article length, and even pessimism or optimism about market conditions. Our empirical strategy exploits exogenous rotation and these content differences across journalists to identify a causal effect on investor behavior.

In our main tests, the dependent variable in a linear regression is the daily excess return on the DJIA Index. The control variables include several lags of returns, day of the week dummies, time effects, lagged volume, and lagged volatility. Our primary interest is the twenty-five vectors of journalist indicators, one for each financial columnist writing for the *WSJ* during our sample period. These fixed effects are statistical stand-ins for both observable and unobservable content differences that persist between financial columnists.

We find that journalist fixed effects are significant predictors of future DJIA returns. Specifying only the name of the financial columnist writing for the *WSJ* on a given day increases predictive power of the regression by 30-40% relative to the other control variables. In joint linear restriction tests, journalists are significant at predicting returns on both the day the journalist’s article is published, as well as the day immediately following. Because we are examining signed rather than absolute market returns, we can interpret the columnist coefficients as capturing their average bullishness or bearishness. When a bullish (bearish) columnist writes, the market inches upward (downward) a few basis points.

However, this specification masks what is potentially a more important question: can journalists exert a unilateral influence on investor behavior, or must certain conditions be met for them to have an effect? Answering this speaks directly to the mechanism underlying any media effects we observe. If the role of financial journalists is ultimately to provide color and interpretation to market events, we would expect for their effects to be highest around news events and volatile returns. On the other hand, evidence on limited attention might suggest that investors are least persuadable during these busy periods, and therefore, that financial journalism might matter in “quieter” times.¹

To address the issue, we augment our benchmark specification by interacting each of the columnist fixed effects with lagged stock returns. Positive interaction coefficients identify journalists that contribute to positive serial correlation, effectively amplifying whatever investor sentiment may exist. A negative coefficient suggests the opposite – a contrarian writer who tempers enthusiasm, dampening the market response. We find that richer specifications including these interactions increase the explanatory power of the regression by yet another one-third, with roughly the same number of journalists being significant as in the unconditional case. Together a coherent story emerges: financial journalism appears to causally influence stock returns, even more so during times of extreme market sentiment.

One challenge to a causal interpretation is that the selection of journalists may not be orthogonal to future market returns. For example, one might worry about an editor assigning a certain writer after steep declines – unless we can perfectly control for any continuation or reversal effects,² future

¹See DellaVigna and Pollet (2009) and Hirshliefer, Lim, and Teoh (2009).

²Note that we already include several days of lagged returns, but the underlying relationship may be more complicated than this linear specification.

returns might be spuriously correlated with the presence of certain journalists. Fortunately, the fact that there is considerable predictability in columnist scheduling allows us to deal directly with this possibility. Rather than explaining stock returns using which journalist’s article was published that day, we instrument for the author using his past scheduling information.

For example, a common arrangement is for one journalist to write on Monday to Thursday for a few weeks, and for a different one to spell him on Fridays. These and other scheduling patterns make past writing activity a valid instrument for future activity, but importantly, not in a way that can be plausibly related to future returns. Although somewhat weaker than the benchmark regressions, the instrumental variable specifications yield jointly significant coefficients for both the journalist indicators and their interactions with returns. Because we are using only information exogenous to returns to predict journalist arrival, this specification represents perhaps the strongest evidence for a causal relation.

A second reason our results could be spurious is data mining. Although generally hard to refute, we conduct a number of robustness checks that, combined with the main results, should make a successful mining expedition less likely. For example, while our main regressions use DJIA close-to-close prices, the results are similar if we use DJIA open-close prices, or if we analyze other series such as the CRSP value-weighted or S&P 500 Index. Also, our results do not appear to be driven by outliers, for either returns or journalists. If we use GARCH-adjusted or winsorized returns as our dependent variable, the results hold. Similarly, if we include only the ten most frequently credited authors, the relations we document in the main analysis remain.

The primary, and to our knowledge novel, contribution is to identify a causal link between media reporting and aggregate stock prices. A number of studies have documented the media’s ability to shift public opinion, particularly with regard to voting patterns (DellaVigna and Kaplan (2007), and Gerber, Karlan, and Bergan (2009)). What makes the present result so surprising is the strong theoretical assumption that the media should not, apart from information effects, be able to influence prices, certainly at the aggregate level.³ Whereas there are models to explain why consumers might be susceptible to biased reporting – and by extension, why media outlets may

³An information story would require a few financial columnists to have persistent information advantages over the entire market, a claim that seems implausible in the short term, even more so over many years. Second, recalling that columnists are affiliated with return patterns of a particular sign, these information advantages would need to be both journalist- and sign-specific. For example, columnist John Smith would need to consistently receive private, *positive* signals about future returns.

then have an incentive to misreport to them (Gentzkow and Shapiro (2006)) – prices of financial securities should not reflect bias due to informed traders.

A secondary contribution is methodological, and is related to how we proxy for exogenous content differences. Almost without exception, research in this area has followed Tetlock’s (2007) seminal work and characterized written articles using computerized algorithms that count, for example, “negative” or “positive” words using financial dictionaries (e.g., Loughran and McDonald (2009)). While this procedure has the advantage of identifying specific stylistic elements to investor behavior, a shortcoming is that automated programs may neither completely, nor accurately, summarize how a *human* audience interprets the written word.

To give a specific example, it is clear that Charles Dickens and William Faulkner employ different themes and rhetorical techniques in their writings, and indeed, such differences have occupied the attention of literary critics for decades. But it seems equally obvious that what truly makes Dickens ‘Dickens’ or Faulkner ‘Faulkner’ cannot be easily quantified, regardless of how sophisticated the analysis may be. Like performing a violin concerto or preparing a fine meal, summarizing the nuances of a written article may be impossible, relative to simply specifying its creator.⁴

Ultimately, this argument highlights both the main strength and weakness of our empirical strategy. By focusing on journalist fixed effects, we implicitly capture *any* persistent differences in stylistic or thematic choices, no matter how difficult to directly measure. Moreover, because these differences are orthogonal to returns, we can make causal inferences about the media on stock prices. At the same time, such a reduced form approach is agnostic about the tools (e.g., word tone, sentence length) each author employs to distinguish his writing from that of his peers. How important one views this distinction determines, to a large extent, the relevance of our results.

The paper is organized as follows. In Section 2, we describe the data and define our key variables. Section 3 presents evidence that specific journalists influence the aggregate market. Within this section, we also characterize whether the effects of financial journalism vary with market conditions. We deal with the potential endogeneity of journalist scheduling in Section 4, presenting the results when we instrument for journalist arrival using past scheduling information. Section 5 presents additional robustness tests, and Section 6 concludes.

⁴See Bertrand and Schoar (2003) for a similar argument and empirical strategy in the context of corporate decisions, and Chevalier and Ellison (1999) in the context of mutual funds.

2 Data

2.1 Market returns and news articles

Two main data sources are used in this paper: the “Abreast of the Market” (AOTM) column from the *Wall Street Journal*, and the Dow Jones Industrial Average (DJIA) price and dividend series. Our sample period spans January 1, 1970 to December 31, 2007. Data further back is available for both sources, but it was not until approximately 1970 that the AOTM column was consistently published with the accompanying author’s name.

Our main dependent variable is the excess daily return on the DJIA Index.⁵ From Yahoo! Finance, we extract a daily series of closing prices for the DJIA and then we add the price-weighted dividend yield for each of the index’s components because the DJIA is a price-weighted index.⁶ Defining r_t as the DJIA excess return and p_t as the level of the DJIA index at the close of day t , we have

$$r_{t+1} = \frac{p_{t+1} - p_t}{p_t} - r_{f,t+1} + dp_{t+1} \quad (1)$$

where $r_{f,t+1}$ is the one-month Treasury bill rate obtained from the Center for Research in Securities Prices (CRSP), and dp_{t+1} is the price-weighted average dividend yield for the stocks in the DJIA index defined as

$$dp_{t+1} = \frac{\sum_{i \in \text{DJIA}} d_{i,t+1}}{\sum_{i \in \text{DJIA}} p_{i,t}}. \quad (2)$$

One-day lagged prices are used in the calculation of the DJIA aggregate dividend yield. This is to avoid the price adjustments that occur following a dividend issue or stock split. Over our sample period, the total excess return of the DJIA averaged 2.6 basis points, equating to an annualized excess return of 6.5 percent. Daily excess return volatility is approximately 101 basis points, implying an annualized Sharpe ratio of 0.41, nearly identical to that found in other diversified return indices.⁷

Additionally, in the regressions that follow other variables are included to control for known

⁵We use DJIA returns as our dependent variable following Tetlock (2007), who argues that the AOTM column tends to disproportionately cover the blue-chip stocks of the DJIA. However, our main results are nearly identical when we use other aggregate return series, e.g., the S&P 500 Index or the CRSP Value-Weighted Index (see Table 8).

⁶Changes in the level of the DJIA ignore distributions to shareholders. See Sialm and Shoven (2000).

⁷For example, over this same time period the CRSP Value-Weighted Index annualized Sharpe ratio was also 0.41.

sources of return predictability such as day-of-the-week or liquidity effects and microstructure effects such as bid-ask bounce or non-synchronous trading. We construct the *Controls* vector which includes five lags of detrended daily log volume from the New York Stock Exchange (NYSE) obtained from CRSP, five lags of detrended squared DJIA residuals which proxy for volatility, day-of-the-week dummies, and a dummy variable for the month of January. To further address potential calendar effects, the *Control* vector also includes year fixed-effects. Both log volume and squared DJIA residuals are detrended by subtracting their past 60-day moving average. DJIA residuals are demeaned DJIA returns. To control for heteroskedasticity or auto-correlation in regression residuals all regression standard errors are calculated using both White and Newey-West standard errors with five lags. Using these controls makes our regression specifications comparable to those in Tetlock (2007).

AOTM is one of the most widely read market summary columns in the United States. It provides analysis of prior market activity, describes some notable company-specific events, and sometimes offers predictions for the future.⁸ Electronic text copies of AOTM columns dated after 1984 are available from different sources. Data prior to 1984 is obtained from the historical *Wall Street Journal* archive, which stores the articles as scanned images. To convert these images to text files, we use ABBYY OCR software.⁹ Typically, this process yields a high quality of transcription. Any errors in this process are likely to be idiosyncratic, and will thus bias the coefficients of interest to zero.

During our sample period, the AOTM column was published Monday through Friday with a few exceptions on national holidays. Occasionally, there are two articles published before the next trading day. In this instance only the most recent article is used. Additionally, during this period there are 40 days when the stock market is open, but no AOTM column is available from our data sources. Overall, our sample includes 9,552 articles, over a period of 9,592 open market days. We restrict our attention to journalists which wrote at least fifty AOTM columns. For a small number of articles (76), we are either unable to identify the author, or the author wrote fewer than fifty

⁸See Tetlock (2007) for more discussion.

⁹OCR, or optical character recognition, is the electronic translation of scanned images of handwritten, typewritten, or printed text into machine-encoded text. The ABBYY OCR software we use performs OCR using intelligent character recognition (ICR). This type of OCR works by searching the scanned image for common elements such as open spaces, closed forms, lines, diagonals intersecting and so on to identify letters. Typically, the accuracy rates using ICR are very high.

articles. A similarly small number (516) were co-written, in which case authorship is credited to neither journalist.¹⁰ Overall, this results in a set of twenty-five authors, which account for over 80 percent of the articles in our sample period.

2.2 Journalist scheduling

Table 1 presents a number of statistics related to each journalist’s writing schedule. Moving across the table, we first list the journalist’s last name, the years he or she was active, and the total number of articles written. As seen, a few journalists are responsible for the majority of the articles, with Hillery (2,413), O’Brien (1,215), and Talley (915) being credited the most frequently. The median author, McLean, is associated with 103 articles.

A crucial feature of our identification strategy is that journalists tend to alternate or rotate with one another over the same time period. We show this graphically in Figure 1, which plots with X’s the dates each journalist wrote, separately for each columnist by row. We note that five authors were responsible for the bulk of the writing during 1970–1984, whereas in the late 1980s and early 1990s there was significantly more turnover. However, the more important observation is that at any point in time, there are multiple active authors. For example, at the year 1980 mark, we see frequent activity from four different columnists: Hillery, Elia, Marcial, and Metz. Inspection of other dates reveals a similar pattern. Without this overlap we would not be able to separately identify any impact journalists might have on investor behavior from simple time trends.

Returning to Table 1, we see also that journalists tend to write articles in relatively brief spells. Shown in the fourth column is *Number of Rotations*, which identifies the number of instances where, for each time a journalist wrote, a different columnist wrote the following day. For example, Marcial penned 625 articles, but had only only 364 *Rotations*. This means that for $625 - 364 = 261$ days, Marcial directly followed one of his own articles the next business day, e.g., writing on a consecutive Wednesday and Thursday. Because our empirical tests will ultimately compare market returns between days when, for example, Marcial’s articles were published to days when they were not, these transitions are important. For journalists that write less frequently (the bottom half on the table), the typical spell falls between 1 and 2 days, whereas for the more frequent authors, spells are two to three times longer on average. Garcia is a notable outlier, writing over 588 columns, but

¹⁰Our results are nearly identical if dual authorship is credited.

rotating only 70 times, for an average spell length of over 8 days.

The final five columns show the breakdown of each columnist’s writing by day of the week. From this it is clear that there are two distinct types of writers – those assigned for all or most of the business week, and those slated for one particular day. As an example of the former type, Garcia’s articles are distributed relatively equally across all days, with Monday (16%) being only slightly less common than the other four days. O’Brien, Raghavan, and Gonzalez are other examples. By contrast, Browning’s articles are almost always published on Mondays (90%), with Sease (69% on Mondays), Rosenberg (70% on Fridays), Ip (76% on Mondays), and Levingston (51% on Mondays), exhibiting similar concentration on a particular day.

Such strong day-of-the-week patterns across journalists suggests that at least part of what we observe arises from pre-determined, semi-regular schedules. Figure 2 gives some graphical intuition for this claim, plotting detailed schedules over three sample sub-periods: July–December 1972, July–December 1994, and July–December 2000. Each row corresponds to a different journalist. For example, in the top panel, the writings of Hillery are shown in the bottom row, Rosenberg in the row above, Dorfman above him, and an indicator for “No Author” in the top row. The X’s correspond to Mondays, circles to Fridays, and crosses to the other three weekdays.

Looking first at the top panel, note the remarkable regularity between the two dominant writers, Hillery and Rosenberg. Of the 26 Mondays, Hillery wrote the AOTM article for 21 of them, and of the 26 Fridays, Rosenberg was active for all but 4. The pattern for the other weekdays is even more pronounced. Beginning after Hillery’s first article (the second week shown), he missed only 13 Tuesdays, Wednesdays, or Thursdays. Six of these days were missed consecutively in the last two weeks of August, and three more the week leading up to Christmas – all almost certainly corresponding to vacation time. For nearly half the sample (12 weeks), a *completely* deterministic alternation between Hillery and Rosenberg is observed, with Rosenberg writing only on, but on every, Friday.

The second panel plots journalist schedules for the second sub-period, and although not as predictable as the first, nevertheless indicates considerable regularity. Kansas writes most Mondays, as well as some other days from time to time, while Pettit is the mode weekday writer over the period. Mollison appears to spell Pettit from his regular duties for two weeks in late August and early September respectively, but was otherwise active only sporadically. The remaining journalists

– Bauman, Granahan, Levingston, Frank, and Arvedlund – wrote only an article or two each, and at seemingly random times. The final graph shows a similar pattern, with O’Brien being the regular, weekday writer, and except on rare occasions (e.g., the likely vacation week seen yet again prior to Labor Day), being active every day except Mondays. Over this period, AOTM articles published on Monday were authored by five journalists, with no apparent pattern.

An important caveat is that although the evidence in Table 1, and Figure 2 seem to suggest some degree of scheduling predictability, this is not necessary (although it would be sufficient) for our later return regressions to be properly specified. Because we will be examining returns *after* a given columnist’s article is published, the main concern is that certain writer selection somehow depends on future market returns. Clearly, this will not be the case if journalist rotations are completely deterministic, as some periods in our sample appear to be, or completely random.

3 Journalists and market returns

3.1 Unconditional effects

We begin our main analysis with some simple univariate comparisons. In Table 2, for each columnist we list: 1) \bar{r}_{wrote} , the average excess return on the days his articles are published, 2) $\bar{r}_{day\ after}$, the average excess return the day after his article is published, and 3) \bar{r}_{other} , the average excess return on all other days over his writing tenure. For example, Table 1 indicates that Hillery authored 2,413 AOTM articles from 1970 to 1984. The average DJIA return on the days these articles were published was slightly less than 1 bp, and on the day afterward, 3 bp. By contrast, the average DJIA return from 1970 to 1984 on the complement set of days when another journalist wrote the AOTM column was -3.4 bp. This comparison thus holds constant the average returns over each journalist’s tenure, so that the fact that one journalist wrote mostly in the 1970s (when average returns were low), while another wrote in the late 1990s (when they were not), is not a concern.

We see that of the twenty-five *WSJ* columnists, over one-third are associated with significant abnormal returns, either the day their articles are published or the day afterward. These are split fairly evenly between positive and negative abnormal returns. In an absolute sense, a few journalists (e.g., Pasha and McGee) are associated with particularly striking abnormal returns, in the range of 40 bp per day. However, the more representative case, based on number of articles written, implies

magnitudes roughly half to a quarter as large (e.g., O’Brien 20 bp on day $t + 1$, Pettit -6 bp on day t).

To more formally characterize the univariate patterns seen in Table 2, we estimate the following linear regression:

$$r_t = c + \sum_{i=1}^{25} (\beta_{i,1} \cdot \mathit{Journalist}_{i,t-1} + \beta_{i,0} \cdot \mathit{Journalist}_{i,t}) + \eta \cdot \mathit{Controls}_t + \epsilon_t, \quad (3)$$

where r_t is the excess return of the DJIA index on day t , as defined in Equation (1). All returns are nominal and are reported in basis points.

The variables of interest, the *Journalist* fixed effects, correspond to each of the twenty-five *WSJ* columnists, i , active from 1970-2007. If columnist i is the author of the AOTM for publication the morning of day t , then $\mathit{Journalist}_{i,t}$ takes a value of one, and zero otherwise. Going backward in time, $\beta_{i,1}$ captures the effect of columnist i ’s article published yesterday on today’s returns, or alternatively, the effect of journalist i one day after his article is published. We consider the day after publication not only because it allows a columnist’s effect to last more than one day but also because it introduces “a significant time gap between the release of the column in the afternoon and the beginning of the event return window” (Tetlock 2007).

Only a few variables have been shown to be significant predictors of one- or two-day returns: lagged returns, trading volume, lagged volatility, day of the week, and a dummy for the month of January. Of these, perhaps the most important in our context is the day of the week, given that Figure 2 indicates strong intra-week patterns for AOTM columnists. Our specification controls for the effects of these variables on market returns. In particular, the *Controls* vector includes five lags of returns, five lags of detrended daily log NYSE volume, five lags of detrended squared DJIA residuals (i.e. lagged volatility), day-of-the-week dummies, a dummy variable for the month of January, and year fixed effects.

We begin with the first column of Table 3, which focuses on the effects of articles published the same day returns are measured (analogous to \bar{r}_{wrote} in Table 2). The introduction of the control variables reduces the statistical significance, but in general, the point estimates are similar to those found in the univariate comparisons. The point estimates have the same sign in 21 of the 25 cases, and for all 4 cases of disagreement, the coefficients are estimated imprecisely. Below the coefficient

estimates, we show the p -values for the joint test that $\beta_{i,0} = 0$ for all journalists i . Depending on how standard errors are calculated, the joint significance is between 1% and 4%.

Moving to the right, we see the impact on current returns of articles written for publication yesterday – i.e., we are explaining Wednesday’s excess returns as a function of which journalist authored AOTM on Tuesday. This corresponds to the $\bar{r}_{\text{day after}}$ column in Table 2. As we see, the evidence is even stronger in this column, suggesting that the impact of specific journalists lasts more than one day. We find that seven journalists are now significant at the 10% level, compared to only four in the previous column. O’Brien, Rosenberg, Ip, Pasha, Dorfman, and McGee are all associated with abnormal excess return at better than the 5% level. The p -values for the joint linear restriction test indicate strong statistical significance ($p\text{-value} < 0.001$).

At the bottom of the first column, we present the R^2 of the excess return regression, with and without the full set of journalist fixed effects. The low explanatory power for both the restricted ($\beta_{i,1} = \beta_{i,0} = 0$) and unrestricted cases is expected, given that we are examining high frequency returns. Still, percentage wise, the improvement is impressive. Journalist fixed effects explain more than an additional 35% of daily excess returns, relative to that explained by time effects, recent returns, volatility, and trading volume.

3.2 Conditional effects

The tests so far have considered the average marginal impact of each journalist, but have ignored whether their effects are dependent upon market conditions. It is not obvious what we should expect, and largely depends on whether we view financial journalists as mostly providing interpretation of underlying events, or as primarily creating *de novo* content. To draw an analogy with other branches of journalism, should we think of financial journalists as being like sportswriters (who requires a game on which to comment), or to an investigative reporter expected to dig up her own facts and write a groundbreaking story?

Table 4 addresses whether the effects of financial journalism are strongest after days of extreme returns, both positive or negative. We start with the specification in Table 3, but interact each journalist indicator (both on day $t - 1$ and day t) with the excess return the day before the article was published. That is, we interact the set of journalist that publish on day t with returns on day $t - 1$ and journalists that publish on day $t - 1$ with returns on day $t - 2$:

$$\begin{aligned}
r_t = & c + \sum_{i=1}^{25} \{ \beta_{i,1} \cdot Journalist_{i,t-1} + \beta_{i,0} \cdot Journalist_{i,t} + \gamma_{i,1} \cdot r_{t-2} Journalist_{i,t-1} \\
& + \gamma_{i,0} \cdot r_{t-1} Journalist_{i,t} \} + \eta \cdot Controls_t + \epsilon_t.
\end{aligned} \tag{4}$$

The goal is to gauge the extent to which return conditions on say, Monday, influence how a journalist’s whose article is published on Tuesday are perceived by the market, either on Tuesday or Wednesday.

There are two reasons why we might expect Monday’s return conditions (in our example) to influence how investors respond to financial journalism on Tuesday or Wednesday. First, returns proxy for how much information is released to the market. To the extent that financial columnists use this information to “set the stage” for their articles, we might expect stronger marginal effects. Second, extreme returns – particularly large negative returns – may proxy for swings in investor sentiment. Borrowing the methodology from Tetlock (2007), García (2010) shows that news-response coefficients in return regressions are stronger in recessions, and interprets this result as investors being more susceptible to slant in reported news during bad times.

Looking first at columns 1 and 2, we see that even in the presence of the interaction terms, the coefficients and statistical significance of the unconditional journalist indicators are similar to those in Table 3. However, note also the final two columns, which show how these slopes are affected by the return environment. In the third column, we find that roughly one-third of the journalists have a significant interaction coefficient, with half again as significant in the final column.

The diagnostic statistics at the bottom of the table formalize the importance of the return interactions in the return regressions. Recall that in Table 3, the inclusion of day t and $t - 1$ journalists increased the R^2 from 2.8% to 3.8%. Here, we see an even bigger improvement with the journalist-return interactions, to 5%. Recall that all specifications (even the baseline without journalists) include lagged returns – instead, it is the interaction with the journalist fixed effects that makes the difference. At the bottom of the table, we present p -values for the journalist indicators, the return-journalist indicators, and their union. Regardless of how standard errors are calculated, the joint hypothesis that our coefficients are zero is rejected at better than the 0.1% level.

Given that the journalist-return interactions are important predictors – perhaps, even more

important than just the journalist indicators alone – it is worth being explicit about their interpretation. While the unconditional columnist fixed effects can be interpreted as capturing each columnist’s average bullishness or bearishness, the interaction terms measure whether a given columnist contributes to, or detracts from, short-term return continuation. In other words, the interactions tell us whether a journalist *amplifies* the effects of past returns, or whether he plays an *attenuating* role.

Take as examples columnists Karen Talley and David McLean. For both journalists, the first two coefficients, β_1 and β_0 , indicate that unconditionally, their writings have neither a positive nor negative impact on future returns. However, conditioning on past returns, we see that Talley has a journalist-interaction coefficient $\gamma_0 = -0.211$, while McLean has a coefficient of $+0.319$. This indicates that following an excess return of +100 basis points (roughly the standard deviation of the DJIA), the market would be expected to decline 21 basis points on days when Talley’s articles are published (on the margin), but rise roughly 32 basis points on days when McLean writes. In other words, Talley appears to attenuate – i.e., detract from positive serial correlation – while McLean appears to amplify.

Taken together, the results in this section not only indicate a casual impact for financial journalism, but also paint a more complete picture of *how* it matters. If we think of journalists as actors with the goal of persuading an audience, the results in Tables 2-3 suggest that an actor’s identity *alone* tells us something about the performance he will give. However, Table 4 indicates that an actor’s performance also depends upon the stage he is given. When the stage is set for good news, some journalists make the good news sound even better, while others temper such enthusiasm. By contrast, when the stage is set for bad news, some journalists are somber while others look on the bright side. Interestingly, it is precisely in times with large price movements when the power of rhetoric has the greatest effect on investors.

4 Endogeneity of journalist arrival

The evidence so far indicates persistent correlation between certain journalists writing and subsequent market returns. However, the selection of columnists may not be exogenous with respect to future market returns – e.g., an editor choosing a certain journalist around releases of bad news.

If not, then the patterns we observe may be spurious, and thus, tell us nothing about the financial media's ability to influence investor behavior. This is clearly not a problem if the process by which journalists are selected is completely exogenous – preset schedules being a special case. However, although there appears to be considerable predictability in how journalists are chosen (see Figure 2), it is equally clear that the selection is not completely deterministic, leaving our specifications open to omitted variable bias.

The traditional solution for this problem is to look for an instrumental variable that is correlated with the potentially endogenous variable (here the vector of journalist dummies), but is otherwise unrelated to the dependent variable (future excess market returns). Intuitively, we want to project each journalist's actual schedule onto explanatory variables that we know cannot be systematically correlated with future returns, and use these projections in place of the potentially endogenous variables. Any observed relation thus results from the part of the endogenous variable that is explained solely by exogenous factors, and thus cannot be susceptible to the endogeneity critique.

We are fortunate to be afforded a nearly perfect instrument: each journalist's recent writing schedule. Because journalists tend to write in relatively short bursts, and often on the same day of the week, we can use past writing activity as an instrument for current writing activity. The key to this being exogenous is the time lag. For a given Tuesday in 1972, for example, we use as instruments which journalist wrote on Monday (yesterday), as well as which journalist wrote the Tuesday one week ago and day-of-the-week dummies for that year. Together, these instruments are powerful predictors of actual writing activity, but have no systematic relation to stock returns.

Table 5 shows the performance of our instrumental variable specification. We run a daily linear probability model for each journalist, over his respective tenure – i.e., for only the years 1970-1984 for Hillery, 1995-2002 for O'Brien, etc. In the second column, we report the R^2 using only year fixed effects; as seen, this generally produces a poor fit. The third column adds market variables which include five lags of recent returns, five lags of de-trended lagged volume, and five lags of squared DJIA residuals. In most cases, these controls lead to a small increase in explanatory power. In the fourth column, in addition to these variables, we include as regressors year fixed effects interacted with day of the week dummies, and two dummies for journalist i – one dummy indicating if journalist i wrote on the previous day and one indicating if journalist i wrote on the same day the previous week. In all cases, these additional controls substantially increase explanatory power,

mostly due to the day of the week interactions, which Figure 2 indicates are strong predictors for most journalists.

Column five presents a formal test of whether the market variables help predict journalist arrivals. For six of the twenty-five journalist there is some evidence that past returns influence their selection, including two (Pasha and McGee) that have significant return coefficients in Table 3, which argues for the IV analysis that we shall construct below. The final column shows the additional explanatory power afforded by our instrumental variables, the one-day and one-week lagged *Journalist* indicators, interacted with the relevant year. In all but one case – Gonzalez – specifying whether a columnist wrote yesterday, or that particular day the previous week dramatically improves fit. Increases in R^2 in the range of 30% are common, and even bigger improvements are seen in some cases. The p -values in the rightmost column show this formally, calculated against the null that the coefficients on the instruments are simultaneously zero. As indicated by their very low values (nearly all are below 0.001), a journalist’s past writing schedule is very valuable for predicting his near-term future activity.

It is important to note the difference in terms of fit that the market variables provide, relative to the rotation variables. Even though market variables help predict the arrival of Browning (p -value 6%), the increase in the R^2 of the specification from adding such variables is 5.6%. In contrast, adding the rotation variables increases the R^2 to 45.8% – clearly knowing who wrote last Thursday is more important for determining who writes this Thursday than what happened in the markets over the last week. This last source of variation is what drives identification in the IV second stage.

In Table 6 we use the fitted values of the LPMs from Table 5, instead of the journalist’s actual writing activity, in order to predict DJIA returns. For example, for a given day in 1972, if the journalist rotation model (Table 5) predicts that Hillery will write with 60% probability, we use this fitted value rather than a zero or one, as we did when referring to Hillery’s *actual* writing activity in Table 4. During the period when a journalist is not actively writing, these fitted values are automatically set to zero.

Comparing the second column of Table 6 with the second column in Table 4, the similarity is apparent. All seven columnists with return coefficients significant at the 10% level in Table 4 have point estimates of the same sign in Table 6, and four (Smith, Rosenberg, Dorfman, and McGee) retain similar statistical significance. The evidence is a bit weaker in the first column.

Although we observe similar magnitudes in the IV regression for the four significant journalists in the baseline specification (O’Brien -0.113 non-IV vs. -0.194 IV, Garcia 0.186 non-IV vs. 0.251 IV, Pasha 0.487 non-IV vs. 0.291 IV, and McGee 0.352 non-IV vs. 0.358 IV), the noise introduced by the instrumental variable specification reduces the statistical significance below conventional levels for these journalists. However, the joint significance of all unconditional columnist coefficients, β_1 and β_0 , is well below the 1% level, as seen by the joint linear restriction test directly below the first column.

The third and fourth columns of Table 6 show the coefficients on the return interactions under the IV specification. Like the previous columns, the instrumental variable specification produces qualitatively similar point estimates, but somewhat weaker statistical significance. In the third column, all seven of the significant author interactions from Table 4 have point estimates of the same sign in Table 6, however, the statistical significance of the estimates in column three of Table 6 is much weaker. The IV results in the fourth column are not particularly useful, mostly because the non-IV results (Table 4 column 4) are already relatively weak. The joint linear restriction test that all interactions are simultaneously zero ($\gamma_1 = \gamma_0 = 0$) is soundly rejected ($p < 0.001$ for all standard error assumptions), as is the joint test of all columnist indicators and their interactions. This latter result is unsurprising, given that the improvement in R^2 due to the columnist variables is comparable to that seen in Table 4, increasing from 0.028 to 0.054.¹¹

5 Extensions and robustness

In this section, we present additional evidence that our results are not spurious, and that in fact, they support a causal interpretation. We organize this discussion as follows. First, in subsection 5.1, we dig a little deeper into the journalist fixed effects, until now about which we have been completely agnostic. The point of this exercise is to reinforce the idea that the journalist fixed effects are in fact capturing persistent content differences, and that these differences are plausibly related to the return patterns we see. Then we address two types of model mis-specification issues: residuals that may not be normally and independently distributed (subsection 5.2) and data mining

¹¹The results are similar if, instead of using the raw probabilities (the outputs of Table 5) as inputs into the IV regression: 1) we take the maximum probability observed across journalists, and assign a value of one to that journalist, or 2) we scale all probabilities so that they sum to one across journalists for every date.

(subsection 5.3).

5.1 Do journalist fixed effects capture content differences?

As we have mentioned throughout, our identification strategy is not free: although journalist rotations give us plausibly exogenous variation in reporting, they do not pinpoint the specific rhetorical devices journalists use to create these differences. By using statistical stand-ins for each journalist, we may be picking up literally *any* persistent stylistic element that differ across journalists and yet influence investor behavior. How important this is largely comes down to the researcher’s objective. If the broad goal is to understand whether, and if so how much, financial journalism matters for stock returns, then this is of less importance. However, those interested in stylistic nuances might not agree – a linguist, for example, might find the specific writing elements of considerable interest.

Because we are more aligned with the general objective, we have to this point emphasized the exogeneity of the journalist fixed effects, and neglected the stylistic elements they are capturing. Here, we attempt to partially bridge this gap, although two reasons ensure that any such progress will be incomplete. First, there are a seemingly endless number of ways that one can measure content, and second, there is almost no theory linking any of these to stock returns. Are longer articles more effective? Should the “main point” be stated in the first couple of paragraphs? Are easily digestible article, with short, pithy phrases more likely to induce an investor response? Do negative-sounding words have to be juxtaposed with less inflammatory diction to retain their impact? Do investors become accustomed to the style of particular authors?

Because answering these or similar questions seems difficult if not impossible, we have more modest ambitions in this section. First, we wish to simply show that for at least a few easily measurable content metrics, differences between journalists are large and persistent. Second, we provide suggestive evidence that in at least one of these dimensions – an article’s “pessimism” – the results make intuitive sense, consistent with Tetlock’s (2007) original findings. That is, the most pessimistic journalists are associated with the most negative short term returns, and that these reverse soon thereafter.

Table 7 studies five content metrics we can easily quantify using computerized algorithms: *Syllables*, *Words Per Sentence (WPS)*, *Percentage complex words*, *Fog* and *Pessimism*. The *Pessimism*

variable is constructed counting the number of “positive” and “negative” words, as well as the total number of words used by each author daily, using Loughran and McDonald (2009) dictionaries.¹² *Pessimism* is equal to the difference between the percentage of negative and positive words, normalized across all authors to have zero mean and unit variance. The average number of syllables per word (*Syllables*), sentences per AOTM column (*Sentences*), and words per sentence (*WPS*). Complex words (*Complex*) are defined as having three or more syllables and are tabulated as a fraction of total words. The final metric, *Fog*, reports the Fog readability score, which indicates the number of years of formal education a reader of average intelligence would need to read and understand the article in one sitting. For example, a *Fog* score of 18 would be considered unreadable, a score of 12 appropriate for a high school graduate, and so on. As indicated by a mean of 11.1, the typical AOTM article would be appropriate for a high school senior.

In Table 7, we regress each of these content metrics on the vector of columnist fixed effects. Moving across the table, we find that columnist fixed effects are important for the number of *Syllables* (column one), *Words per sentence* (column 2), *Percentage complex words*, and *Fog*. The incremental increase in explanatory power ranges from a full ten percent (for *Syllables*) to five percent (for *Fog*), and in most cases, well over half the journalists are individually significant at conventional levels. Note also that differences in writing style do not appear correlated with writing activity; journalists at the top of the list (the most frequently credited authors) do not appear to write systematically more brief or complex than their more sporadic counterparts.

The final two columns show that journalists are important predictors of the positive-negative word balance in AOTM articles, both unconditionally (column five), as well as when interacted with past returns (column six). The interpretation of each is straightforward. The unconditional coefficients in column five capture the average incremental *Pessimism* of each journalist, irrespective of prevailing market conditions. Hillary (−0.17) for example, writes consistently more bullish-sounding articles than O’Brien (0.27). The sixth column allows these cross-journalists tonal differences to depend on recent returns, which, as in Table 4, is also important. In both columns, we observe *p*-values consistently less than 0.01%, regardless of we compute standard errors under the OLS, White (1980), or Newey-West assumptions.

¹²See http://www.nd.edu/~mcdonald/Word_Lists.html. We use these dictionaries in particular because they account for a number of the nuances related to financial language.

We wish to stress once more the limitations of Table 7. The content metrics we analyze here necessarily paint a necessarily incomplete picture of cross-journalist content differences – these are likely the tip of the iceberg. However, in at least one case, *Pessimism*, we have previous research with which to compare our results. Tetlock (2007) shows that articles with a higher percentage of negative words foreshadow low next day returns, and that this effect reverses over the following week. Here, the next question is obvious: are the most pessimistic journalists associated with the lowest next day returns, and do these reverse?

Panels A and B of Figure 3 make this comparison, plotting each journalist’s return coefficient against his *Pessimism* coefficient (Table 7, column 5). The day t return effect is shown in Panel A, and the day $t + 1$ effect in Panel B. The radius of each circle corresponds to the precision of the *Pessimism* coefficient, so that more precisely estimated journalists are given more weight. As clearly seen, the slope of the line in Panel A is negative, with almost all the significant journalists lying close to the line (Pasha is the notable exception), confirming a causal interpretation for Tetlock’s (2007) original finding. When we advance a day in Panel B, the pattern reverses. Journalists that write pessimistic articles on day t are associated with *positive* returns a day later, consistent with a reversal of the original day t effect. As before, the fit of this line is impressive using only 25 data points, with Pasha again as the lone outlier.

These figures represent our only attempt to explicitly link content and returns, and even here, we urge a cautious interpretation. While Figure 3 suggests that the use of pessimistic words might be one way that journalists temporarily influence stock prices, there are undoubtedly others, perhaps ones far more important.¹³ We present this correlation only to show that our findings are consistent with prior work, and to emphasize that to the extent that article pessimism is important, causation can likely be inferred.

5.2 Standard errors

Our inferences for the statistical significance of the journalist fixed effects rely on tests that assume independently and normally distributed residuals. This assumption can be violated for a number of reasons, including homoskedasticity and serial correlation, both of which we know apply to stock returns (Schwert (1989), Campbell, Grossman, and Wang (1993)). Moreover, one may also

¹³See section 5.3.5 for more analysis related to this issue.

be concerned about using asymptotic approximations for finite samples, particularly with a large number of fixed effects (50-100 depending on specification).

In response to such concerns, we adopt both parametric and non-parametric approaches. Tables 4 and 5 already show the former, where we report the p -values for F -tests of joint significance under three different assumptions for standard errors: OLS, White (1980), and Newey-West with five lags. The modeling choice does not seem to matter. We observe p -values less than 1:10,000 in each case except for the non-interacted journalist fixed effects on day t (Table 3, column 1), and even here, they are jointly significant at the 4.2% level in the worst case.

Non-parametric analysis yields similar conclusions. We first see this in Table 6, where we instrument for each journalist's writing activity using his past scheduling information. To account for the errors-in-variables problem introduced by using generated regressors, we bootstrap the second stage using 1,000 iterations. As a side benefit, this simultaneously addresses other potential violations of the spherical residuals assumption. That we still obtain p -values in the 1:10,000 range (or better) suggests that incorrect inference is unlikely to explain our results.

We conduct an even more detailed bootstrapping exercise in the Appendix. Conceptually, the experiment can be thought of as the following: given a series of $N=9,592$ daily excess return observations, what would the distribution of p -values for the F -tests look like if our 25 journalists were *randomly* assigned? This is similar to the approach taken by Fee, Hadlock and Pierce (2010) who are concerned about the validity of F -tests in the presence of serial correlation and many fixed effects. With perfectly specified statistical tests, we would expect such a distribution on a placebo data set to line up on the 45 degree line, e.g., to obtain 7% significance level in 7% of the simulations, 8% significance 8% of the time, and so on. Deviations from the 45 degree line might indicate finite sample bias, serially correlated residuals despite Newey-West corrections, heteroskedasticity despite White corrections, or other mis-specifications. However, as the plots in the Appendix show, the test statistics are well behaved. Although we observe somewhat fat tails, particularly with the Newey-West correction, these biases are modest, especially given the high level of statistical significance we observe (up to 10^{-7} in some cases). We refer the reader interested in more detail to the methodological description provided in Appendix.

5.3 Data mining

It is also possible that our results are a product of data mining. This can occur intentionally, when authors cycle through a large number of specifications, and report only those with the desired level of significance. It can also occur unintentionally – i.e., across authors – where the number of tested specifications gets large not because each author (or set of authors) mines for significance, but simply because lots of researchers are running similar regressions. Unlike the model misspecification issues discussed in section 5.2, there are no formal statistical tests that answer the critique of data mining. Instead, our strategy is to present a number of extensions and robustness checks, in hopes of making it more difficult for us to produce (if done intentionally) or less likely to stumble upon (if done unintentionally) spurious correlations.

5.3.1 Alternative return series

The first five columns of Table 8 presents the results of the Table 4 specification, but vary the return series. Panel A shows the p -values for the joint linear restriction test that all unconditional journalist coefficients (β_0 and β_1) are zero, and Panel B shows the corresponding p -values for the conditional effects (γ_0 and γ_1). Linear restriction tests for both unconditional and conditional effects are considered in Panel C.

In the first column, we consider open-to-close DJIA returns. The second and third column present the results when instead we use excess returns on the S&P 500 Index or CRSP value-weighted index as our dependent variable. As seen, the results are similar for each of these alternatives, compared to the DJIA results shown in Tables 3 and 4.

5.3.2 GARCH adjustments and winsorization

The fourth and fifth columns consider GARCH-adjusted and winsorized returns, respectively, which we conduct to address the concern that a few outliers may be responsible for the return patterns we document. In the regression results reported in the *Winsorized* column returns, volume, and all non-dummy variables are winsorized at the 5 percent level. GARCH-adjusted returns are defined as DJIA close-to-close returns divided by the estimated daily volatility from a GARCH(1,1) model estimated on the same return series. As seen, we observe high levels of statistical significance for

the unconditional effects, the conditional effects, and their union.

5.3.3 Restriction to most frequent authors

To rule out the possibility that a few very strong, but infrequently observed, journalists are responsible for the joint significance we observe in Tables 3 and 4, the sixth column of Table 8 repeats the analysis, but considers only the ten most frequently credited authors. As indicated in the summary stats (Table 1), this restriction implies that we are considering only authors with at least 157 written articles. Table 8 indicates that the unconditional effects are jointly significant at between the 0.2% and 0.5% level, and the interactions at below the 0.1% level.

5.3.4 Time-varying return effects

Because all our previous regressions include year dummies, the relevant variation in journalist-specific content is *intra*year.¹⁴ Consequently, that Raghavan wrote in the 1990s while Rosenberg wrote in the 1970s is irrelevant for identification purposes.

In the penultimate column of Table 8, we employ even finer controls for time effects. Rather than including a single fixed effect for each year (e.g., 1972, 1973, 1974, etc.), we include twelve fixed effects for each year (January-1972, February-1972, etc.). In aggregate, 456 additional fixed effects are added. This more stringent specification identifies journalist-return effects over small windows (21 or 22 trading days), and, with the other controls, leaves very few degrees of freedom. Furthermore, as part of this specification we include controls for a variety of calendar-time anomalies, including: a dummy variable that equals one if the trading-day falls within the first 15 days of the month (Ariel (1987), Lakonishok and Smidt (1988)), a dummy variable that equals one on trading days that fall before scheduled NYSE trading holidays (Ariel (1990)), and weekly dummies for the first four weeks and the last four weeks of the year as a finer control for the January effect. Nonetheless, all rows show statistical significance at the 10^{-4} level or better, suggesting that even at very high frequencies, content differences across journalists appear to influence returns.

The final column allows for time-varying differences in first order daily autocorrelations. For a variety of reasons, the microstructure of the DJIA may have changed over time, e.g., stale limit

¹⁴Even this specification appears to be overkill, given that year dummies are not statistically significant predictors of daily excess returns ($p = 0.84$).

orders that could not be easily filled at close before computerization, etc. Because some of our tests interact journalist fixed effects with the previous day’s return, r_{t-1} , and because we are comparing journalist interactions separated by decades, we include the same controls used the in the penultimate column and interactions of these variables with the previous day’s return, r_{t-1} . Additionally, we include the interaction of the previous day’s volume interacted with the previous day’s return as a control for return autocorrelation on days of intense trading similar to Campbell, Grossman, and Wang (1993). Overall, these controls add 933 additional regressors. As with our other robustness checks however, this makes little difference, affecting only the statistical significance of the journalist interactions (γ_0, γ_1) , and then, only slightly.

5.3.5 Controlling for computer-measured content

We have argued throughout that the vectors of journalist fixed effects are proxies for content generally – both measurable and immeasurable elements. As an example of the former, Figure 3 presents suggestive evidence that writers using more pessimistic words are associated with more negative next day returns, a pattern that reverses the next day. This finding is consistent with a causal interpretation of Tetlock (2007). A natural extension is to ask whether this is the full story – i.e., does authorship *only* proxy for pessimistic words, or do they also capture more subtle stylistic elements?

To shed some light on this issue, we augment the specification in Table 3 with the normalized negative-positive word count of the AOTM column, both the day returns are measured (t) and the previous day ($t - 1$). If either current or lagged *Pessimism* drives out the journalist fixed effects, then we will have identified both a causal relation for journalists as well as a specific journalistic mechanism responsible for the effect.

However, comparing the columnists coefficients in Panel A of Table 9 with the corresponding ones in Table 3 reveals only trivial differences. All four journalists originally significant during the reporting day (t) are still significant after controlling for pessimistic word counts, with nearly identical magnitudes. The same is true for the seven significant authors on the day afterward ($t - 1$). Thus, while at least part of the “explanation” for the journalist fixed effects appears to be negative-positive word mixes (as Figure 3), other stylistic elements are important too.

The final row shows that when journalist fixed effects are included *Pessimism* is only weakly

related to returns ($t = -1.2$), although with the expected, negative sign. Importantly, we note that this is neither inconsistent with Tetlock (2007), nor with Figure 3. We already know (Table 7) that there are substantial negative-positive word counts differences across journalists, so the journalist fixed-effects partially capture *Pessimism*. Indeed, when the journalists fixed-effects are omitted, the raw *Pessimism* variable becomes more significant, on par with previous studies.

If one takeaway from the first two rows is that modern day computers aren't "great readers", the last column suggests one way we might give them a head start. By interacting *Pessimism* with the set of journalist fixed effects, we hope to give some journalist-specific content to negative-sounding words, and in so doing, allow them to have differential impacts on stock returns. For example, imagine some journalists consistently use certain words in specific, predictable settings, such that interacting words and their authors might improve the information quality relative to using the words alone.

The final column suggests that at least in a few cases, this hypothesis appears plausible. In Panel B of Table 9, we report the coefficients on each journalist (on day t) interacted with *Pessimism*. Although this subsumes the raw, non-interacted *Pessimism* variable, we find different marginal *Pessimism* effects for Hillery and McLean, perhaps extending to Wilson and Pettit. In aggregate, the joint significance of the *Pessimism-Journalist* interactions is relatively weak, with a p -value of 6.7% (OLS). Thus, while computerized word counts appear to be slightly more informative when conditioned on author, the main source of content variation is, at least with current technology, not easily parametrized with automated programs. For even with the *Pessimism-Journalist* included, the statistical significance of the authors remains highly significant ($p < 0.0001$).

5.3.6 A falsification test: content published after returns

We end the analysis with a falsification test, and further study whether the effect of journalists' writing can extend over the two-day window that we have focused on this far. In particular, we augment (3) with one more lag and two more leads. The variable $\beta_{i,2}$ captures any residual effect from an article published two days prior. Going the other direction, $\beta_{i,-1}$ and $\beta_{i,-2}$ are estimated for falsification. They measure the effect of articles written for publication on future dates (e.g., an article for publication on Thursday influencing Wednesday's returns), and consequently, should have no effect.

Table 10 presents the point estimates of such a specification. The rightmost column indicates that generally, any observed journalist-return effects will show up within two business days. Only two journalists – Hillery and Ip – have significant coefficients, and the high p -values at the table’s bottom indicate that together, they add little explanatory power to the regression.

The first two columns test for return effects that, under a causal interpretation, should not produce significant results. Each vector β_{-1} and β_{-2} measures columnist who write on future days, after controlling for current and past authors. For example, if we are measuring Wednesday’s excess returns, columns 5, 4, and 3, respectively, tell us which AOTM author was published on Monday ($t - 2$), Tuesday ($t - 1$) and Wednesday (t), while columns 1 and 2 pick up who *will* be published on Thursday and Friday respectively. As expected, future authors have no apparent relation to current returns, with Newey-West p -values of 0.52 and 0.66 for $t + 1$ and $t + 2$, respectively.

6 Conclusion

There is widespread speculation that the news media has the power to influence financial markets, apart from simply reporting events. Yet, such claims are often based on anecdotal associations that make causal inferences difficult. For example, times of negative financial reporting frequently coincide with bad economic news, and vice versa. Stripping away the effects of only the *reporting* thus requires variation in news content that is unrelated to underlying fundamentals.

The identification strategy of this paper is based on two assumptions. The first is that authors of the *Wall Street Journal’s* “Abreast of the Market” column exhibit persistent stylistic differences, such that even for the same set of facts, article content will vary. The second is that the selection of journalists is not systematically related to future returns, an assumption relatively easy to justify given that we are examining returns on a nearly unpredictable market index.

Our results suggest that financial journalists have the potential to influence investor behavior, at least over short time horizons. Adding journalist fixed effects to a daily return regression significantly increases explanatory power, and when these fixed effects are interacted with recent returns, the implied return predictability is even stronger. Overall, our results suggest that the interpretation of public news is important, as the effects we uncover are strongest when journalists write about significant market moves.

An important caveat we have mentioned throughout is that although our empirical design permits a causal interpretation, our analysis does not shed light on the specific rhetorical tools that authors use to influence investor behavior. That is, we do not attempt to say whether longer articles, more complex words, or less pessimism leads to a predictable market response. This is not because we cannot quantify a number of content measures, but instead because we think that attempting to gauge a human audience's response may be difficult using computerized algorithms, relative to using statistical stand-ins for human authors. Clearly, we sacrifice the ability to pinpoint specific stylistic techniques, but we hopefully gain by capturing other unobservable elements that vary across journalists.

By documenting causal effects of the media on aggregate market prices, our findings paint a somewhat ominous picture of financial journalism. One recalls Shiller's (2000) less-than-veiled indictment: "The history of speculative bubbles begins roughly with the advent of newspapers" (p. 85). His implication is as clear as it is concerning – if financial journalists can manipulate investor beliefs apart from fundamentals, then their actions and incentives play a direct role in prices and allocations. The evidence in this paper, particularly as it applies to aggregate allocations, calls for a better understanding of these issues.


References

- [1] Ariel, Robert A., 1987. "A monthly effect in stock returns," *Journal of Financial Economics* 18, 161-174.
- [2] Ariel, Robert A., 1990. "High stock returns before holidays: existence and evidence on possible causes," *Journal of Finance* 45, 1611-1626.
- [3] Campbell, J., Grossman, S., Wang, J., 1993, Trading Volume and Serial Correlation in Stock Returns, *The Quarterly Journal of Economics* 108, Vol. 4, 905-939.
- [4] Chesher, A., Jewitt, I., 1987. "The bias of a heteroskedasticity consistent covariance matrix estimator," *Econometrica* 55 (5), 1217-1222.
- [5] Chevalier, J., Ellison, G., 1999. "Are Some Mutual Fund Managers Better Than Others? Cross-Sectional Patterns in Behavior and Performance," *Journal of Finance* 54 (3), 875-899.
- [6] DellaVigna, S., Kaplan, E., 2007. "The Fox News Effect: Media Bias and Voting," *The Quarterly Journal of Economics* 122 (3), 1187-1234.
- [7] DellaVigna, S., Pollet, J., 2009. "Investor Inattention and Friday Earnings Announcements," *Journal of Finance* 64 (2), 709-49.
- [8] Engelberg, J., 2008. "Costly Information Processing: Evidence from Earnings Announcements," *Working Paper*, University of North Carolina.
- [9] Engelberg, J., Parsons, C., 2010. "The Causal Impact of Media in Financial Markets," *Forthcoming at the Journal of Finance*.
- [10] Fama, E., 1970. "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance* 15 (2), 383-417.
- [11] Fee, C., Hadlock, C., and Pierce, J., 2010. "Managers Who Lack Style: Evidence from Exogenous CEO Changes," *Working Paper*.
- [12] French, K., 1980. "Stock Returns and the Weekend Effect," *Journal of Financial Economics* 8 (1), 55-69.

- [13] García, D., 2010. “Sentiment During Recessions,” *Working Paper*, University of North Carolina, Kenan-Flagler School of Business.
- [14] Gentzkow, M., and Shapiro, J., 2006. “Media Bias and Reputation,” *Journal of Political Economy* 114 (2), 280-316.
- [15] Gerber, A., Karlan, D., and Bergan, D., 2009. “Does the Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions,” *American Economic Journal: Applied Economics* 1 (2), 35-52.
- [16] Hirshleifer, D., Lim, S., Teoh, S., 2009. “Driven to Distraction: Extraneous Events and Underreaction to Earnings News,” *Journal of Finance* 64 (5), 2289-2325.
- [17] Huberman, G., Regev, T., 2001. “Contagious Speculation and a Cure for Cancer: A Nonevent that Made Stock Prices Soar,” *Journal of Finance* 56 (1), 387-396.
- [18] Lakonishok, J., Smidt, S., 1988. “Are Seasonal Anomalies Real? A Ninety Years Perspective,” *Review of Financial Studies* 1(4), 403-425.
- [19] Loughran, T., McDonald, B., 2009. “When is a Liability not a Liability?,” *Journal of Finance*, forthcoming.
- [20] Milton, J., 1644. “Areopagitica: A Speech of Mr. John Milton for the Liberty of Unlicensed Printing, to the Parliament of England.”
- [21] Mullainathan, S., Shleifer, A., 2005. “The Market for News,” *The American Economic Review* 95 (4), 1031-1053.
- [22] Newey, W., West, K. 1987. “A simple, positive definite, heteroscedastic and autocorrelation consistent covariance matrix,” *Econometrica* 55, 703-708.
- [23] Petersen, M., 2004. “Information: Hard and Soft,” *Working Paper*, Northwestern University.
- [24] Plumlee, M., 2003. “The Effect of Information Complexity on Analysts’ Use of that Information,” *The Accounting Review* 78 (1), 275-296.
- [25] Reuter, J., Zitzewitz, E., 2006. “Do Ads Influence Editors? Advertising and Bias in the Financial Media,” *Quarterly Journal of Economics* 121 (1), 197-227.

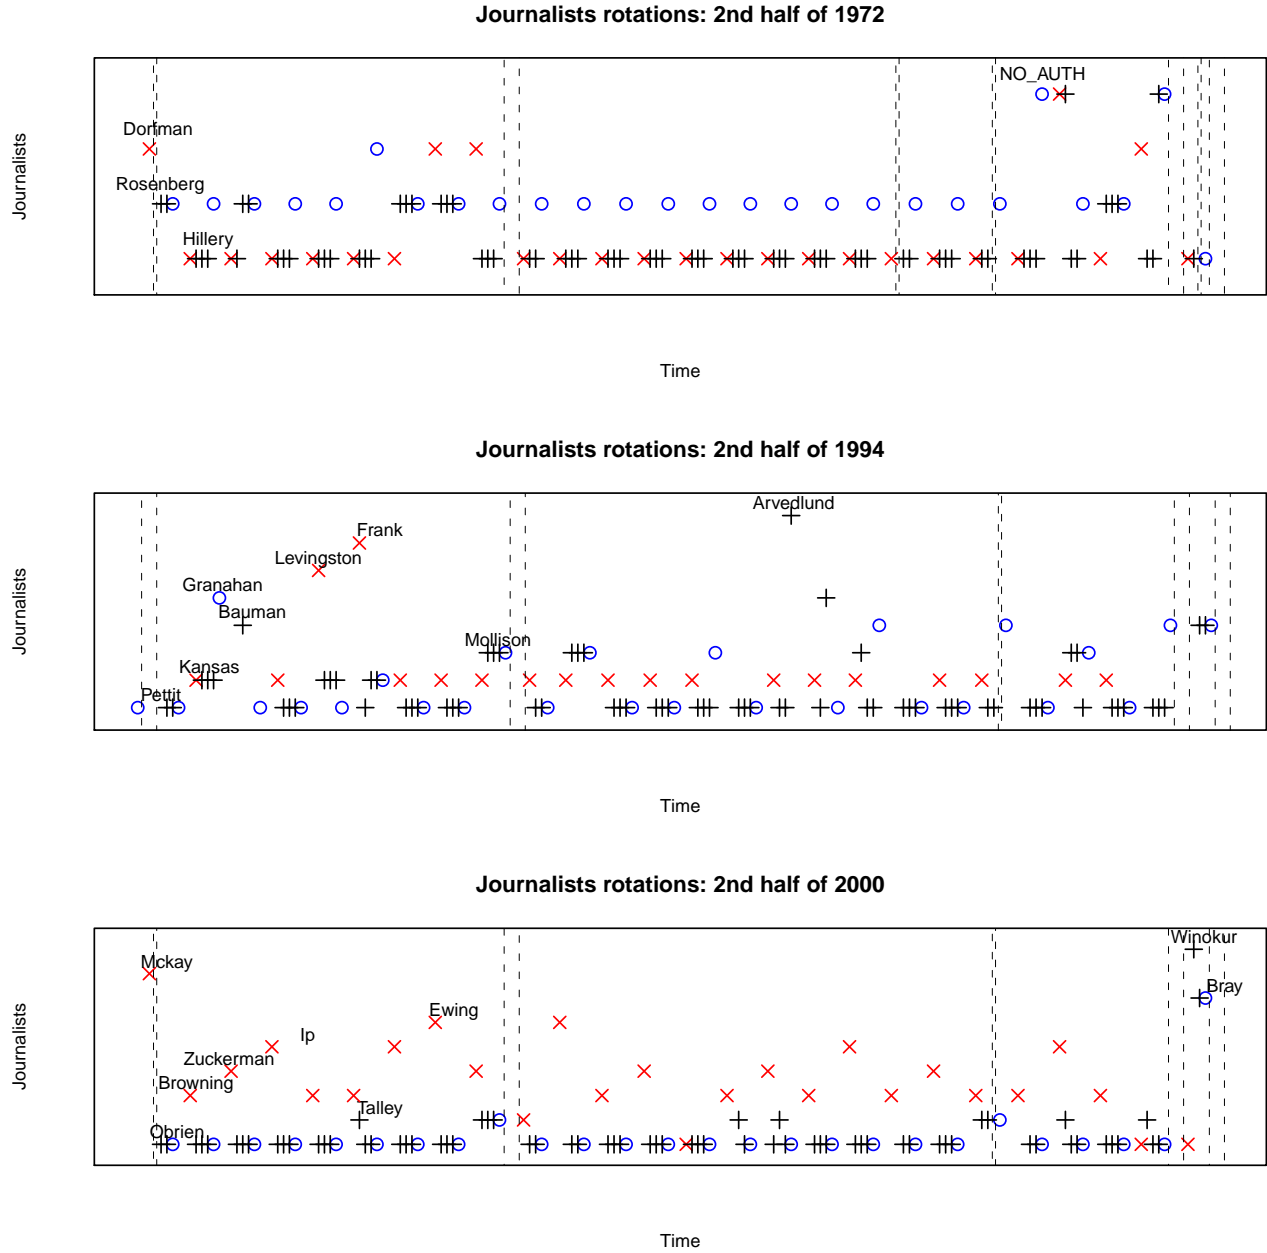
- [26] Schwert, W., 1989, Why Does Stock Volatility Change Over Time? *Journal of Finance* 44, 1115-1153.
- [27] Shiller, R., 2000. "Irrational Exuberance," *Princeton University Press, Princeton*.
- [28] Sialm, C., Shoven, J., 2000. "The Dow Jones Industrial Average: The Impact of Fixing Its Flaws," *The Journal of Wealth Management* 3 (3), 9-18.
- [29] Tetlock, P., 2007. "Giving Content to Investor Sentiment: The Role of Media in the Stock Market.," *Journal of Finance* 62 (3), 1139-1168.
- [30] Tetlock, P., 2009. "All the News That's Fit to Reprint: Do Investors React to Stale Information?," *Working Paper*, Columbia University.
- [31] Wooldridge, J., 2009. "Introductory Econometrics," *South-Western College Pub* (4), 269.

Figure 1: Timeline of AOTM journalists



This Figure documents the authorship of the *Wall Street Journal* "Abreast of the Market" column for our full sample time-period. Each point corresponds to an author writing the AOTM column on a given day.


Figure 2: Sample of journalist writing days



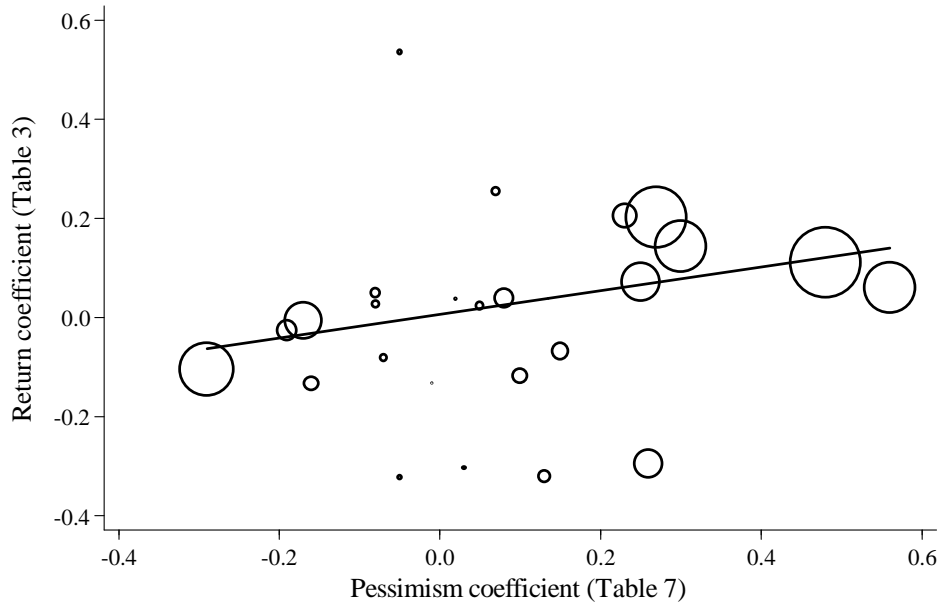
This Figure documents the authorship of the *Wall Street Journal* “Abreast of the Market” column for three 6-month periods. For every period, each row corresponds to a different author. For example, in the first panel the authors represented from top row to bottom row are NO_AUTH, Dorfman, Rosenberg, and Hillery. In each row, markings represent the days on which the row’s corresponding author wrote. Red x’s signify writing on Monday, blue circles writing on Friday, and crosses writing on Tuesday, Wednesday, or Thursday. Holidays are denoted by dashed vertical lines.

Figure 3: Linking journalist pessimism to returns

Panel A: Day of publication



Panel B: Day after publication



Panel A plots journalist's pessimism coefficients from Table 7 versus the journalist's day of publication return coefficients from Table 3. Panel B plots journalist's pessimism coefficients from Table 7 versus the journalist's day after publication return coefficients from Table 3. In both panels coefficients are weighted by their corresponding squared t-statistic from the pessimism regressions in Table 7, additionally both plots show a weighted least squares fit line.

Table 1: Statistics on journalists' tenure

This table presents statistics for each journalist who wrote more than fifty articles for the AOTM column. In particular, it lists the last names of the journalists, the years they were actively writing for the AOTM column (Years Active), the total number of articles they published (Articles), the total number of consecutive writing days for each journalist (Number of rotations), the average length of these rotations (Average rotations), and the percentage of articles each journalist published on each weekday.

Journalist	Years Active	Articles	Number of rotations	Average length	% Mon.	% Tue.	% Wed.	% Thu.	% Fri.
Hillery	1970 – 1984	2413	708	3.4	0.18	0.23	0.25	0.25	0.09
O'Brien	1995 – 2002	1215	415	2.9	0.01	0.24	0.26	0.25	0.24
Talley	2000 – 2007	915	289	3.2	0.01	0.23	0.26	0.25	0.25
Marcial	1974 – 1981	625	364	1.7	0.23	0.19	0.13	0.12	0.34
Garcia	1984 – 1988	588	70	8.4	0.16	0.21	0.21	0.20	0.21
Smith	1985 – 1993	302	140	2.2	0.02	0.19	0.24	0.25	0.30
Wilson	1988 – 1990	251	97	2.6	0.01	0.23	0.25	0.28	0.23
Browning	1996 – 2007	250	249	1.0	0.90	0.10	0.00	0.00	0.00
Pettit	1992 – 1995	222	109	2.0	0.00	0.21	0.22	0.29	0.28
Sease	1987 – 1993	157	115	1.4	0.69	0.13	0.06	0.06	0.06
Rosenberg	1972 – 1974	125	95	1.3	0.00	0.07	0.13	0.10	0.70
Kansas	1994 – 1996	104	77	1.4	0.61	0.15	0.10	0.09	0.06
McLean	1970 – 1971	103	69	1.5	0.00	0.14	0.12	0.17	0.58
Raghavan	1993 – 1994	93	58	1.6	0.30	0.22	0.22	0.16	0.11
Ip	1996 – 2005	90	80	1.1	0.76	0.12	0.03	0.06	0.03
Gonzalez	1995 – 1995	87	50	1.7	0.18	0.22	0.18	0.23	0.18
Metz	1978 – 1988	80	57	1.4	0.16	0.20	0.15	0.15	0.34
Levingston	1992 – 1994	77	50	1.5	0.51	0.18	0.12	0.10	0.09
Pasha	2001 – 2005	74	36	2.1	0.03	0.23	0.26	0.27	0.22
Rose	1985 – 1986	65	14	4.6	0.14	0.22	0.22	0.23	0.20
Steiner	1991 – 1996	63	33	1.9	0.10	0.37	0.30	0.16	0.08
Dorfman	1970 – 1973	62	56	1.1	0.21	0.06	0.05	0.06	0.61
Bauman	1994 – 2000	52	33	1.6	0.00	0.21	0.27	0.17	0.35
McGee	1996 – 2001	51	51	1.0	0.96	0.04	0.00	0.00	0.00
Elia	1973 – 1981	50	46	1.1	0.68	0.08	0.06	0.02	0.16

Table 2: Univariate return tests

For each journalist, this table presents the average daily excess return of the DJIA on the days they wrote, \bar{r}_{wrote} , the days after they wrote, $\bar{r}_{day\ after}$, and for all other days, \bar{r}_{other} , during the period they were actively writing for the AOTM column. Column four presents the t -statistic for a test of the difference $\bar{r}_{wrote} - \bar{r}_{other}$, and column six the t -statistic for $\bar{r}_{day\ after} - \bar{r}_{other}$. The t -statistics that are statistically significant at the 1% level are indicated by ***, at the 5% level by **, and at the 10% level by *.

	\bar{r}_{other}	\bar{r}_{wrote}	t -stat	$\bar{r}_{day\ after}$	t -stat
Hillery	-0.034	0.007	1.0	0.031	1.2
O'Brien	-0.036	-0.011	0.3	0.202***	2.6
Talley	-0.021	0.025	0.9	0.066	1.0
Marcial	0.046	0.036	-0.2	-0.055*	-1.8
Garcia	0.005	0.112	1.1	-0.249	-1.0
Smith	0.035	0.067	0.5	0.170	1.3
Wilson	0.107	-0.041*	-1.8	-0.020	-1.2
Browning	0.029	0.066	0.5	0.002	-0.4
Pettit	0.050	-0.057**	-2.0	0.084	0.5
Sease	0.026	0.053	0.3	0.133	1.2
Rosenberg	0.099	-0.095*	-1.9	-0.304***	-3.6
Kansas	0.076	0.119	0.6	-0.005	-1.0
McLean	0.103	0.069	-0.3	-0.035	-1.1
Raghavan	-0.024	0.200***	2.9	-0.021	0.0
Ip	0.013	0.195	1.4	0.118	0.8
Gonzalez	0.010	0.107	1.2	0.225**	2.3
Metz	0.029	0.031	0.0	0.041	0.1
Levingston	0.029	0.042	0.2	-0.031	-0.7
Pasha	0.011	0.275*	1.9	-0.431**	-2.3
Rose	0.123	0.034	-0.8	-0.084	-0.9
Steiner	0.057	-0.076	-1.6	0.142	0.7
Dorfman	0.027	0.028	0.0	-0.258***	-2.6
Bauman	0.073	0.001	-0.5	-0.064	-0.7
McGee	0.003	0.411**	2.4	0.406**	2.4
Elia	-0.005	-0.191	-1.3	0.027	0.2

Table 3: Multivariate return regressions

This table presents OLS coefficient estimates and their corresponding t -statistics for

$$r_t = c + \sum_{i=1}^{25} \{\beta_{i,1} \cdot \text{Journalist}_{i,t-1} + \beta_{i,0} \cdot \text{Journalist}_{i,t}\} + \eta \cdot \text{Controls}_t + \epsilon_t$$

where the *Controls* vector includes five lags of daily excess DJIA return, five lags of detrended daily log NYSE volume, five lags of detrended squared DIJA residuals, day-of-the-week dummies, a dummy variable for the month of January, and year fixed-effects. Also presented are the number of observations, the R -squared for an unreported regression with no journalist fixed effects, R_{noJFE}^2 , and the R -squared for the reported regression which includes the journalist fixed effects, R_{JFE}^2 . The p -values and F -statistics from F -tests with the following null hypotheses are also recorded: $\beta_{i,1} = 0, \forall i$; $\beta_{i,0} = 0, \forall i$; and $\beta_{i,1} = \beta_{i,0} = 0, \forall i$. For each test, results are reported using the OLS variance/covariance matrix, a heteroscedasticity robust variance/covariance matrix (WHITE), and a Newey-West variance/covariance matrix using five lags (NW5). t -statistics that are statistically significant at the 1% level are indicated by ***, at the 5% level by **, and at the 10% level by *.

	publication date		day after			
	β_0	t -stat	β_1	t -stat		
Hillery	0.038	0.6	-0.007	-0.1		
O'Brien	-0.121*	-1.9	0.202***	3.2		
Talley	0.010	0.2	0.039	0.6		
Marcial	0.078	1.0	-0.105	-1.4		
Garcia	0.204*	2.0	-0.118	-1.1		
Smith	-0.026	-0.3	0.144*	1.7		
Wilson	-0.032	-0.4	-0.133	-1.5		
Browning	0.064	0.8	0.024	0.3		
Pettit	-0.098	-1.0	0.111	1.1		
Sease	0.010	0.1	0.072	0.8		
Rosenberg	-0.047	-0.4	-0.304***	-2.6		
Kansas	0.030	0.3	0.050	0.4		
McLean	0.091	0.7	-0.026	-0.2		
Raghavan	0.175	1.3	0.037	0.3		
Ip	0.110	0.9	0.255**	2.1		
Gonzalez	-0.076	-0.5	0.206	1.4		
Metz	0.073	0.6	-0.081	-0.6		
Levingston	0.038	0.3	-0.068	-0.5		
Pasha	0.487***	3.4	-0.296**	-2.1		
Rose	0.285	1.3	-0.323	-1.5		
Steiner	-0.134	-0.9	0.027	0.2		
Dorfman	0.050	0.3	-0.321**	-2.2		
Bauman	-0.016	-0.1	-0.134	-0.8		
McGee	0.386**	2.5	0.536***	3.4		
Elia	-0.118	-0.8	0.060	0.4		
Observations	9592					
R_{noJFE}^2	0.028					
R_{JFE}^2	0.038					
H ₀ :	$\beta_1 = 0$		$\beta_0 = 0$		$\beta_1 = \beta_0 = 0$	
p -value/ F -stat OLS	0.000	2.4	0.011	1.8	0.000	1.9
p -value/ F -stat WHITE	0.000	2.4	0.042	1.8	0.000	1.9
p -value/ F -stat NW5	0.000	2.7	0.025	1.6	0.000	2.2

Table 4: Multivariate return regressions with interactions

This table presents OLS coefficient estimates and t -statistics for

$$r_t = c + \sum_{i=1}^{25} \{ \beta_{i,1} \cdot \text{Journalist}_{i,t-1} + \beta_{i,0} \cdot \text{Journalist}_{i,t} + \gamma_{i,1} \cdot r_{t-2} \times \text{Journalist}_{i,t-1} + \gamma_{i,0} \cdot r_{t-1} \times \text{Journalist}_{i,t} \} + \eta \cdot \text{Controls}_t + \epsilon_t.$$

The *Controls* vector is as in Table 3. This table also presents the number of observations, the R -squared for an unreported regression with no journalist fixed effects, R_{noJFE}^2 , and the R -squared for the reported regression that includes the journalist fixed effects, R_{JFE}^2 . Also recorded are p -values and F -statistics from F -tests testing the following null hypotheses: $\beta_{i,1} = \beta_{i,0} = 0, \forall i$; $\gamma_{i,1} = \gamma_{i,0} = 0, \forall i$; and $\beta_{i,1} = \beta_{i,0} = \gamma_{i,1} = \gamma_{i,0} = 0, \forall i$. For each test, results are reported using the OLS variance/covariance matrix, a heteroscedasticity robust variance/covariance matrix (WHITE), and a Newey-West variance/covariance matrix using five lags (NW5). The t -statistics that are statistically significant at the 1% level are indicated by ***, at the 5% level by **, and at the 10% level by *.

	publication date		day after		publication date		day after	
	β_0	t -stat	β_1	t -stat	γ_t	t -stat	γ_1	t -stat
Hillery	0.036	0.6	-0.007	-0.1	-0.011	-0.3	-0.031	-0.9
O'Brien	-0.113*	-1.8	0.205***	3.3	-0.149***	-4.2	-0.077**	-2.2
Talley	0.017	0.3	0.042	0.6	-0.211***	-4.6	-0.014	-0.3
Marcial	0.072	0.9	-0.105	-1.4	0.089*	1.7	-0.103*	-1.9
Garcia	0.186*	1.8	-0.117	-1.1	-0.061	-1.4	0.048	1.1
Smith	-0.046	-0.5	0.162*	1.9	-0.090	-1.2	-0.095	-1.3
Wilson	-0.028	-0.3	-0.130	-1.5	-0.082	-1.1	-0.040	-0.5
Browning	0.071	0.9	0.025	0.3	-0.128*	-1.9	0.116*	1.7
Pettit	-0.107	-1.1	0.122	1.2	-0.004	0.0	0.058	0.5
Sease	-0.013	-0.1	0.091	1.0	0.056	0.8	-0.215***	-2.9
Rosenberg	-0.055	-0.5	-0.295**	-2.5	-0.025	-0.2	-0.136	-1.3
Kansas	0.030	0.3	0.027	0.2	-0.041	-0.3	0.039	0.2
McLean	0.096	0.7	-0.011	-0.1	0.319***	2.7	0.089	0.8
Raghavan	0.191	1.4	0.050	0.4	-0.170	-0.8	-0.018	-0.1
Ip	0.106	0.9	0.252**	2.1	0.039	0.4	-0.036	-0.4
Gonzalez	-0.092	-0.6	0.225	1.6	0.004	0.0	-0.245	-1.2
Metz	0.072	0.5	-0.103	-0.8	-0.085	-0.7	0.063	0.5
Levingston	0.027	0.2	-0.057	-0.4	-0.056	-0.4	-0.021	-0.1
Pasha	0.487***	3.4	-0.283*	-2.0	-0.226**	-2.4	0.122	1.3
Rose	0.297	1.4	-0.335	-1.5	-0.277*	-1.7	-0.037	-0.2
Steiner	-0.160	-1.0	0.057	0.4	-0.038	-0.2	-0.178	-0.9
Dorfman	0.042	0.3	-0.326**	-2.2	0.099	0.7	-0.182	-1.3
Bauman	-0.050	-0.3	-0.118	-0.7	0.252	1.3	-0.147	-0.8
McGee	0.352**	2.3	0.537***	3.4	-0.224**	-2.0	-0.200*	-1.8
Elia	-0.149	-1.0	0.050	0.3	0.146	1.0	0.073	0.5
Observations	9592							
R_{noJFE}^2	0.028							
R_{JFE}^2	0.050							
H ₀ :	$\beta_1 = \beta_0 = 0$		$\gamma_1 = \gamma_0 = 0$		$\beta_1 = \beta_0 = \gamma_1 = \gamma_0 = 0$			
p -value/ F -stat OLS	0.000	1.9	0.000	2.4	0.000	2.2		
p -value/ F -stat WHITE	0.000	2.2	0.000	2.2	0.000	2.2		
p -value/ F -stat NW5	0.000	2.4	0.000	3.0	0.000	2.7		

Table 5: Linear probability models – forecasting journalists arrivals

This table presents R^2 for the following linear probability models:

$$\begin{aligned} \text{Model 1 : } Journalist_{j,t} &= c_j + \sum_{i=1}^{38} \lambda_i \cdot Y_{i,t} + \epsilon_{j,t} \\ \text{Model 2 : } Journalist_{j,t} &= c_j + \sum_{i=1}^{38} \lambda_i \cdot Y_{i,t} + \eta \cdot Market\ Variables_t + \epsilon_{j,t} \\ \text{Model 3 : } Journalist_{j,t} &= c_j + \sum_{i=1}^{38} \lambda_i \cdot Y_{i,t} + \eta \cdot Market\ Variables_t + \\ &\quad \sum_{i=1}^{38} (\psi_i \cdot Y_{i,t} \times Journalist_{j,t-1} + \rho_i \cdot Y_{i,t} \times Journalist_{j,t-7} + \zeta_i \cdot D_t \times Y_{i,t}) + \epsilon_{j,t} \end{aligned}$$

for all $j = 1 \dots 25$. Y_i is a vector of year fixed effects for year i , while D_t is a matrix of day of the week dummies. *Market Variables* represents five lags of DJIA returns, five lags of volume, and five lags of DJIA squared residuals as previously defined. The variable $Journalist_{j,t-7}$ represents an indicator variable that equals 1 if journalist j wrote on the same day the previous week, and 0 otherwise. Each regression is run using only data for the period during which the corresponding journalist was actively writing. Column 5 presents the p -value for an F -test of $H_0 : \eta = 0$, and Column 6 reports the OLS p -value for the F -test of $H_0 : \psi_i = \rho_i = \zeta_i = 0, \forall i$.

	Model 1	Model 2	Model 3	$H_0 : \eta = 0$	$H_0 : \psi_i = \rho_i = \zeta_i = 0$
	R^2	R^2	R^2	p -value	p -value
Hillery	0.038	0.062	0.389	0.592	0.000
O'Brien	0.028	0.088	0.429	0.434	0.000
Talley	0.287	0.318	0.618	0.847	0.000
Marcial	0.004	0.029	0.192	0.985	0.000
Garcia	0.103	0.118	0.476	0.735	0.000
Smith	0.273	0.281	0.508	0.965	0.000
Wilson	0.086	0.137	0.411	0.272	0.000
Browning	0.020	0.076	0.458	0.060	0.000
Pettit	0.072	0.117	0.320	0.675	0.000
Sease	0.081	0.125	0.446	0.002	0.000
Rosenberg	0.011	0.100	0.613	0.437	0.000
Kansas	0.073	0.174	0.462	0.028	0.000
McLean	0.023	0.111	0.586	0.355	0.000
Raghavan	0.050	0.081	0.174	0.855	0.001
Ip	0.085	0.102	0.350	0.990	0.000
Gonzalez	0.000	0.043	0.100	0.907	0.779
Metz	0.040	0.044	0.192	0.759	0.000
Levingston	0.122	0.176	0.349	0.864	0.000
Pasha	0.044	0.088	0.353	0.027	0.000
Rose	0.021	0.033	0.580	0.801	0.000
Steiner	0.224	0.241	0.406	0.010	0.000
Dorfman	0.008	0.029	0.355	0.465	0.000
Bauman	0.082	0.091	0.232	0.986	0.000
McGee	0.067	0.115	0.526	0.048	0.000
Elia	0.018	0.031	0.169	0.741	0.000

Table 6: Multivariate return regressions with IV

This table presents coefficient estimates and bootstrapped t -statistics for

$$r_t = c + \sum_{i=1}^{25} \{ \beta_{i,1} \cdot IVJournalist_{i,t-1} + \beta_{i,0} \cdot IVJournalist_{i,t} + \gamma_{i,1} \cdot r_{t-2} \times IVJournalist_{i,t-1} + \gamma_{i,0} \cdot r_{t-1} \times IVJournalist_{i,t} \} + \eta \cdot Controls_t + \epsilon_t.$$

where $IVJournalist$ is defined as the fitted fitted values from the LPM model of Table 5, column 4. The $Controls$ vector is as in Table 3. The t -statistics are calculated using the bootstrap method to correct for using fitted values from the first-stage regression in our present specification. This table also presents the number of observations in the regression, the R -squared for an unreported regression with no journalist fixed effects, R_{noJFE}^2 , and the R -squared for the reported regression that includes the journalist fixed effects, R_{JFE}^2 . Also recorded are p -values and F -statistics from F -tests testing the following null hypotheses: $\beta_{i,1} = \beta_{i,0} = 0, \forall i$; $\gamma_{i,1} = \gamma_{i,0} = 0, \forall i$; and $\beta_{i,1} = \beta_{i,0} = \gamma_{i,1} = \gamma_{i,0} = 0, \forall i$. The t -statistics that are statistically significant at the 1% level are indicated by ***, at the 5% level by **, and at the 10% level by *.

	publication date		day after		publication date		day after	
	β_0	t -stat	β_1	t -stat	γ_t	t -stat	γ_1	t -stat
Hillery	-0.033	-0.4	0.058	0.6	-0.047	-0.5	0.170*	1.7
O'Brien	-0.194	-1.3	0.118	0.9	-0.067	-0.6	0.066	0.6
Talley	0.061	0.5	0.007	0.1	-0.140	-1.1	0.028	0.2
Marcial	0.092	0.7	-0.037	-0.3	0.091	0.7	0.156	1.2
Garcia	0.251	1.4	-0.155	-0.9	-0.411	-1.5	0.648*	1.8
Smith	-0.222	-1.4	0.385***	2.6	0.067	0.4	0.154	0.9
Wilson	-0.156	-1.0	-0.066	-0.4	0.034	0.2	-0.026	-0.2
Browning	0.201	1.0	-0.079	-0.4	-0.193	-0.9	0.130	0.8
Pettit	-0.159	-1.2	0.164	1.1	0.057	0.3	-0.042	-0.2
Sease	-0.184	-0.7	0.294	1.4	0.314	1.4	-0.297	-1.4
Rosenberg	-0.089	-0.6	-0.287**	-2.0	-0.109	-0.7	0.457***	3.0
Kansas	-0.122	-0.8	0.069	0.5	0.029	0.1	0.193	0.9
McLean	0.007	0.0	0.195	1.1	0.258	1.6	0.213	1.0
Raghavan	0.288	1.0	0.052	0.2	0.063	0.1	0.035	0.1
Ip	-0.391	-1.2	0.196	0.7	0.537	1.6	-0.019	-0.1
Gonzalez	-0.259	-1.1	0.175	0.8	0.754*	1.8	-0.551	-1.5
Metz	-0.243	-0.7	0.062	0.2	0.021	0.0	-0.572	-1.4
Levingston	0.030	0.2	-0.113	-0.7	-0.186	-0.6	-0.043	-0.2
Pasha	0.291	0.9	-0.484	-1.3	-0.556**	-2.0	0.229	0.7
Rose	-0.183	-0.8	0.113	0.5	-0.736**	-2.5	0.569**	2.0
Steiner	-0.077	-0.3	0.129	0.5	0.372	1.1	-0.303	-0.8
Dorfman	0.029	0.1	-0.562***	-3.3	0.335	1.4	0.347**	2.2
Bauman	-0.547	-1.5	0.156	0.5	0.733*	1.8	0.412	1.0
McGee	0.358	1.1	0.712***	3.1	-0.509	-1.1	-0.601**	-2.1
Elia	-0.549	-1.1	-0.434	-1.1	0.588	1.1	-0.579*	-2.0
Observations	9592							
R_{noJFE}^2	0.028							
R_{JFE}^2	0.054							
H_0 :	$\beta_1 = \beta_0 = 0$		$\gamma_1 = \gamma_0 = 0$		$\beta_1 = \beta_0 = \gamma_1 = \gamma_0 = 0$			
p -value/ F -stat BOOT	0.001	1.9	0.000	3.1	0.000	2.9		

Table 7: Article characteristics and journalists

This table reports coefficient estimates from the following regression:

$$Characteristic_{j,t} = c + \sum_{i=1}^{25} \beta_i \cdot Journalist_{i,t} + \eta \cdot Controls_t + \epsilon_{j,t}$$

where $Characteristic \in [Syllables, WPS, \%Complex, Fog, Pessimism]$. Syllables is the average number of syllables for all words in an article. WPS is the average number of words per sentence. $\%Complex$ is the average percentage of words in each article with three or more syllables. Fog reports the average Fog readability score. Pessimism represents the average amount of sentiment in each article as measured by the percentage of negative words minus the percentage of positive words per article. Pessimism is normalized to have a zero mean and unit variance. In the last columns, the interaction of journalists fixed effects and lagged returns are included as regressors in the above specification and the coefficients on these interactions are reported. The table also presents the number of observations, the R -squared for an unreported regression with no journalist fixed effects, R_{noJFE}^2 , and the R -squared for the model that includes the journalist fixed effects, R_{JFE}^2 . We also present the p -values from an F -test of the following null hypothesis: $\beta_i = 0, \forall i$. For each test, p -values and F -statistics are reported for F -tests calculated using the OLS variance/covariance matrix, a heteroscedasticity robust variance/covariance matrix (WHITE), and a Newey-West variance/covariance matrix using five lags (NW5). t -statistics that are statistically significant at the 1% level are indicated by ***, at the 5% level by **, and at the 10% level by *.

	Syllables		WPS		% Complex		Fog		Pessimism (Unconditional)		Pessimism (Interaction)	
	β_0	t -stat	β_0	t -stat	β_0	t -stat	β_0	t -stat	β_0	t -stat	β_0	t -stat
Hillery	0.002	0.5	-0.462**	-2.9	-0.092	-0.9	-0.223**	-3.0	-0.167**	-3.1	0.038	1.4
O'Brien	0.051***	15.5	-0.435***	-2.8	0.838***	8.2	0.164**	2.2	0.267***	5.1	-0.261***	-8.9
Talley	0.038***	11.0	-2.578***	-15.9	0.435***	4.1	-0.853***	-11.3	0.084	1.6	-0.265***	-7.0
Marcial	0.004	0.9	1.888***	9.6	0.079	0.6	0.785***	8.5	-0.293***	-4.5	0.051	1.2
Garcia	-0.006	-1.3	-0.164	-0.7	-0.411***	-2.8	-0.228**	-2.2	0.095	1.3	0.005	0.1
Smith	-0.017***	-3.9	-0.072	-0.3	-0.615***	-4.5	-0.276***	-2.9	0.297***	4.3	-0.110*	-1.8
Wilson	0.000	0.1	0.984***	4.8	-0.306**	-2.2	0.269***	2.8	-0.006	-0.1	-0.107*	-1.7
Browning	-0.082***	-19.3	2.667***	13.2	-2.203***	-16.5	0.189**	2.0	0.050	0.7	0.244***	4.2
Pettit	0.004	0.7	-1.023***	-4.2	0.161	1.0	-0.348***	-3.1	0.478***	5.9	-0.437***	-4.4
Sease	-0.025***	-5.0	0.857***	3.7	-0.353**	-2.3	0.200*	1.8	0.249***	3.2	0.097	1.6
Rosenberg	0.018***	2.9	0.694**	2.4	0.425**	2.2	0.447***	3.3	0.034	0.3	-0.033	-0.4
Kansas	0.029***	4.7	0.049	0.2	0.600***	3.1	0.254*	1.9	-0.076	-0.8	-0.230*	-1.8
McLean	0.008	1.1	-0.944***	-2.9	0.332	1.6	-0.249*	-1.7	-0.186*	-1.7	-0.248***	-2.6
Raghavan	-0.014*	-1.9	0.227	0.7	0.079	0.4	0.115	0.7	0.019	0.2	-0.284*	-1.7
Ip	0.009	1.3	0.994***	3.2	0.600***	3.0	0.643***	4.5	0.072	0.7	0.184**	2.3
Gonzalez	-0.001	-0.1	0.568	1.6	0.617***	2.6	0.471***	2.8	0.231**	2.0	-0.425**	-2.5
Metz	-0.006	-0.9	-0.091	-0.3	-0.435**	-2.1	-0.213	-1.4	-0.065	-0.6	0.032	0.3
Levingston	0.013*	1.8	-1.174***	-3.5	0.467**	2.1	-0.291*	-1.9	0.152	1.4	-0.144	-1.1
Pasha	0.003	0.4	-1.079***	-3.3	-0.166	-0.8	-0.494***	-3.2	0.261**	2.4	0.060	0.8
Rose	-0.006	-0.8	-0.598	-1.6	-0.404	-1.6	-0.395**	-2.2	-0.050	-0.4	-0.373***	-2.8
Steiner	0.015*	1.9	1.209***	3.3	0.882***	3.6	0.829***	4.8	-0.075	-0.6	-0.680***	-4.3
Dorfman	0.008	1.1	1.120***	3.0	0.379	1.6	0.599***	3.5	0.127	1.0	0.045	0.4
Bauman	0.029***	3.6	1.281***	3.4	0.932***	3.7	0.890***	5.0	-0.156	-1.2	-0.301*	-1.9
McGee	-0.015*	-1.8	1.994***	5.1	-0.340	-1.3	0.664***	3.6	-0.052	-0.4	0.086	0.9
Elia	0.004	0.5	-1.201***	-3.1	-0.158	-0.6	-0.548***	-3.0	0.556***	4.3	0.197	1.6
Observations	9552		9552		9552		9552		9552		9552	
R_{noJFE}^2	0.199		0.391		0.190		0.319		0.295		0.295	
R_{JFE}^2	0.297		0.461		0.245		0.365		0.309		0.331	
p -value/ F -stat OLS	0.000	52.7	0.000	49.6	0.000	27.7	0.000	27.2	0.000	7.5	0.000	7.3
p -value/ F -stat WHITE	0.000	46.8	0.000	45.4	0.000	28.2	0.000	24.2	0.000	8.1	0.000	8.2
p -value/ F -stat NW5	0.000	40.3	0.000	39.5	0.000	26.1	0.000	21.3	0.000	7.5	0.000	7.5

Table 8: Multivariate return regressions and robustness

This table presents F -test p -values and F -statistics for the same regression reported in Table 4 only in this instance using different return series or additional control variables. In particular, the *Open-Close* column uses DJIA open-to-close daily excess returns, the *S&P 500* column uses daily S&P 500 excess returns, the *CRSP VWTD* column uses daily CRSP value-weighted excess returns, and the *GARCH-Adj.* column uses GARCH-adjusted returns which are defined as DJIA close-to-close returns divided by the estimated daily volatility from a GARCH(1,1) model estimated on the same return series. In addition, to different return series the last two columns also use slightly different regressors. The results in the *Winsorized* column use returns, volume, and all non-dummy variables that are winsorized at the 5 percent level, and the results in the *Ten Authors* column use only journalist indicators for the ten most prolific writers in our sample. The *Calendar* column includes controls for a variety of calendar-time anomalies, including: a dummy variable that equals one if the trading-day falls within the first 15 days of the month (Ariel (1987), Lakonishok and Smidt (1988)), a dummy variable that equals one on trading days that fall before scheduled NYSE trading holidays (Ariel (1990)), weekly dummies for the first four weeks and the last four weeks of the year instead of the January dummy, and year-month fixed effects instead of year fixed effects. The *Interactions* column includes as regressors the same calendar-time controls used in the *Calendar* column and interactions of these variables with the previous day's return, r_{t-1} . Additionally, this specification includes the previous day's detrended volume interacted with the previous day's return to control for return autocorrelation on days of intense trading similar to Campbell, Grossman, and Wang (1993).

	Open-Close	S&P 500	CRSP VWTD	GARCH-Adj.	Winsorized	Ten Authors	Calendar	Interactions
Panel A: $\beta_1 = \beta_0 = 0$								
p -value OLS	0.000	0.000	0.001	0.000	0.000	0.005	0.000	0.000
p -value WHITE	0.000	0.000	0.000	0.000	0.000	0.004	0.000	0.000
p -value NW5	0.000	0.000	0.000	0.000	0.000	0.002	0.000	0.000
F -statistic OLS	1.9	1.9	1.8	2.1	2.0	2.0	1.9	1.9
F -statistic WHITE	2.1	2.0	2.0	2.3	2.3	2.1	2.2	2.2
F -statistic NW5	2.3	2.2	2.2	2.5	2.5	2.2	2.3	2.2
Panel B: $\gamma_1 = \gamma_0 = 0$								
p -value OLS	0.000	0.000	0.000	0.004	0.000	0.000	0.000	0.000
p -value WHITE	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.012
p -value NW5	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.001
F -statistic OLS	2.4	3.3	3.8	1.6	2.4	4.0	2.3	3.8
F -statistic WHITE	2.1	2.9	3.0	2.0	2.5	2.8	2.1	1.5
F -statistic NW5	3.1	4.5	4.4	2.3	3.1	2.8	2.8	1.8
Panel C: $\beta_1 = \beta_0 = \gamma_1 = \gamma_0 = 0$								
p -value OLS	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
p -value WHITE	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
p -value NW5	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
F -statistic OLS	2.2	2.6	2.8	1.9	2.2	3.0	2.1	2.9
F -statistic WHITE	2.2	2.5	2.6	2.1	2.4	2.4	2.1	1.8
F -statistic NW5	2.8	3.4	3.4	2.5	2.8	2.4	2.7	2.1

Table 9: Quantifying pessimism with word counts

Panel A of this table presents OLS coefficient estimates and their corresponding t -statistics for

$$r_t = c + \sum_{i=1}^{25} \{\beta_{i,1} \cdot Journalist_{i,t-1} + \beta_{i,0} \cdot Journalist_{i,t}\} + \zeta_1 \cdot Pessimism_{t-1} + \zeta_0 \cdot Pessimism_t + \eta \cdot Controls_t + \epsilon_t$$

Panel B of this table presents OLS coefficient estimates and their corresponding t -statistics for

$$r_t = c + \sum_{i=1}^{25} \{\beta_{i,1} \cdot Journalist_{i,t-1} + \beta_{i,0} \cdot Journalist_{i,t}\} + \sum_{j=1}^{25} \{\xi_{j,0} \cdot Journalist_{j,t} \times Pessimism_t\} + \eta \cdot Controls_t + \epsilon_t$$

In both regressions the *Controls* vector is as defined in Table 3. Also presented for both regressions is the number of observations in each regression, the R -squared for an unreported regression with no journalist fixed effects, R_{noJFE}^2 , and the R -squared for the reported regression which includes the journalist fixed effects, R_{JFE}^2 . The p -values and F -statistics from F -tests for the following null hypothesis is also recorded: $\beta_{i,1} = \beta_{i,0} = 0, \forall i$. F -test results using the OLS variance/covariance matrix, a heteroscedasticity robust variance/covariance matrix (WHITE), and a Newey-West variance/covariance matrix using five lags (NW5) are reported. t -statistics that are statistically significant at the 1% level are indicated by ***, at the 5% level by **, and at the 10% level by *.

	Panel A: Pessimism				Panel B: Interactions	
	publication date		day after		day after	
	β_0	t -stat	β_1	t -stat	ξ_0	t -stat
Hillery	0.037	0.6	-0.009	-0.1	-0.065**	-2.3
O'Brien	-0.117*	-1.9	0.202***	3.2	-0.007	-0.3
Talley	0.010	0.2	0.041	0.6	0.029	0.9
Marcial	0.075	1.0	-0.107	-1.4	-0.037	-1.0
Garcia	0.205**	2.0	-0.119	-1.1	0.036	0.8
Smith	-0.021	-0.3	0.141*	1.7	0.037	0.6
Wilson	-0.031	-0.4	-0.135	-1.5	0.136*	1.8
Browning	0.064	0.8	0.024	0.3	-0.021	-0.3
Pettit	-0.092	-0.9	0.110	1.1	-0.121*	-1.7
Sease	0.015	0.2	0.068	0.7	0.072	0.8
Rosenberg	-0.045	-0.4	-0.307***	-2.6	0.085	0.7
Kansas	0.029	0.2	0.051	0.4	-0.068	-0.7
McLean	0.089	0.7	-0.029	-0.2	-0.323***	-3.2
Raghavan	0.174	1.3	0.038	0.3	-0.011	-0.1
Ip	0.111	0.9	0.254**	2.1	-0.143	-1.1
Gonzalez	-0.074	-0.5	0.207	1.5	-0.022	-0.2
Metz	0.073	0.6	-0.084	-0.7	-0.103	-0.8
Levingston	0.040	0.3	-0.067	-0.5	-0.055	-0.4
Pasha	0.490***	3.4	-0.295**	-2.0	0.088	0.8
Rose	0.285	1.3	-0.324	-1.5	0.083	0.7
Steiner	-0.135	-0.9	0.029	0.2	0.037	0.3
Dorfman	0.052	0.4	-0.325**	-2.2	-0.283	-1.5
Bauman	-0.019	-0.1	-0.132	-0.8	-0.211	-1.5
McGee	0.383**	2.5	0.535***	3.4	0.154	0.8
Elia	-0.109	-0.7	0.056	0.4	-0.212	-1.1
	ζ_0	t -stat	ζ_1	t -stat		
Pessimism	-0.014	-1.2	0.004	0.4		
Observations	9592				9592	
R_{noJFE}^2	0.028				0.028	
R_{JFE}^2	0.038				0.041	
H_0 :	$\beta_1 = \beta_0 = 0$				$\beta_1 = \beta_0 = 0$	
p -value/ F -stat OLS	0.000	1.9			0.000	1.8
p -value/ F -stat WHITE	0.000	2.0			0.000	1.9
p -value/ F -stat NW5	0.000	2.1			0.000	2.1

Table 10: Multivariate return regression and falsification

This table presents OLS coefficient estimates and t -statistics for

$$r_t = c + \sum_{i=0}^{25} \{ \beta_{i,2} \cdot Journalist_{i,t-2} + \beta_{i,1} \cdot Journalist_{i,t-1} + \beta_{i,0} \cdot Journalist_{i,t} + \beta_{i,-1} \cdot Journalist_{i,t+1} + \beta_{i,-2} \cdot Journalist_{i,t+2} \} + \eta \cdot Controls_t + \epsilon_t$$

The *Controls* vector is as in Table 3. Coefficient estimates are recorded on the left-hand side of each column and their corresponding t -statistics are presented in parentheses on the right-hand side. Also recorded are p -values from F -tests of the following null hypotheses: $\beta_{i,k} = 0, \forall i$ and $k \in \{-2, -1, 0, 1, 2\}$. For each test, p -values are reported for F -tests calculated using the OLS variance/covariance matrix, a heteroscedasticity robust variance/covariance matrix (WHITE), and a Newey-West variance/covariance matrix using five lags (NW5). t -statistics that are statistically significant at the 1% level are indicated by ***, at the 5% level by **, and at the 10% level by *.

	2 days before		1 day before		publication date		1 day after		2 days after	
	$\beta_{i,-2}$	t -stat	$\beta_{i,-1}$	t -stat	$\beta_{i,0}$	t -stat	$\beta_{i,1}$	t -stat	$\beta_{i,2}$	t -stat
Hillery	0.113*	1.7	-0.028	-0.4	-0.005	-0.1	-0.057	-0.9	0.132**	2.0
O'Brien	0.005	0.1	0.122*	1.9	-0.130**	-2.0	0.229***	3.6	-0.026	-0.4
Talley	0.009	0.1	0.103	1.5	-0.013	-0.2	0.033	0.5	0.069	1.0
Marcial	0.101	1.3	0.019	0.2	0.033	0.4	-0.136*	-1.7	0.074	0.9
Garcia	0.039	0.4	0.020	0.2	0.226*	1.9	-0.058	-0.5	-0.142	-1.3
Smith	-0.171**	-2.0	0.114	1.3	-0.013	-0.1	0.129	1.5	-0.008	-0.1
Wilson	0.079	0.9	0.052	0.5	-0.064	-0.7	-0.118	-1.2	-0.070	-0.8
Browning	-0.043	-0.5	0.041	0.5	0.024	0.3	-0.012	-0.1	0.002	0.0
Pettit	-0.031	-0.3	0.015	0.1	-0.094	-0.9	0.105	1.0	0.031	0.3
Sease	-0.112	-1.2	0.096	1.0	0.027	0.3	0.068	0.7	0.044	0.5
Rosenberg	0.218*	1.8	0.035	0.3	-0.113	-0.9	-0.355***	-3.0	0.168	1.4
Kansas	-0.026	-0.2	0.180	1.5	0.008	0.1	0.045	0.4	-0.073	-0.6
McLean	0.209	1.6	-0.039	-0.3	0.032	0.2	-0.068	-0.5	0.058	0.4
Raghavan	-0.162	-1.2	0.146	1.1	0.177	1.3	0.046	0.3	0.023	0.2
Ip	-0.172	-1.4	-0.018	-0.1	0.088	0.7	0.248**	2.0	-0.323***	-2.6
Gonzalez	-0.006	0.0	0.069	0.5	-0.083	-0.6	0.219	1.5	-0.026	-0.2
Metz	0.189	1.4	0.220*	1.7	-0.003	0.0	-0.113	-0.9	-0.064	-0.5
Levingston	-0.110	-0.8	0.042	0.3	0.059	0.4	-0.081	-0.6	0.064	0.5
Pasha	0.088	0.6	0.097	0.6	0.452***	2.9	-0.257	-1.6	-0.008	-0.1
Rose	0.094	0.4	0.014	0.1	0.270	1.0	-0.354	-1.3	-0.001	0.0
Steiner	-0.005	0.0	0.181	1.1	-0.164	-1.0	-0.008	0.0	-0.029	-0.2
Dorfman	0.098	0.7	-0.064	-0.4	0.022	0.1	-0.352**	-2.4	0.041	0.3
Bauman	-0.035	-0.2	0.131	0.8	-0.017	-0.1	-0.145	-0.9	0.057	0.4
McGee	-0.264*	-1.7	-0.057	-0.4	0.303*	1.9	0.461***	2.9	-0.188	-1.2
Elia	-0.008	0.0	0.173	1.1	-0.142	-0.9	0.033	0.2	0.011	0.1
Observations	9592									
R_{noJFE}^2	0.028									
R_{JFE}^2	0.044									
p -value/ F -stat OLS	0.814	0.7	0.842	0.7	0.077	1.4	0.001	2.2	0.672	0.9
p -value/ F -stat WHITE	0.671	0.9	0.583	0.9	0.153	1.3	0.000	2.4	0.541	0.9
p -value/ F -stat NW5	0.659	0.9	0.519	1.0	0.115	1.3	0.000	2.6	0.580	0.9

Appendix

We consider the following model of asset returns:

$$r_t = c + \sum_{i=1}^{25} (\beta_{i,1} \cdot Journalist_{i,t-1} + \beta_{i,0} \cdot Journalist_{i,t}) + \eta \cdot Controls_t + \epsilon_t,$$

i.e., the specification from Table 3. We construct the $Controls_t$ vector which includes five lags of detrended squared DJIA residuals which proxy for volatility, day-of-the-week dummies, and a dummy variable for the month of January. This is different than the model used in the paper, since we do not include lagged volume. This is done purely for simplicity, as adding another time-series clearly increases the numerical burden of the simulations we report next.

In order to study the finite sample properties of the F -statistics behind the main tests in the paper, we proceed as follows. First, we estimate (5) under the null, i.e. setting $\beta_{i,0} = \beta_{i,1} = 0$ for all i . This gives us estimates for the control vector $\hat{\eta}$. Let the residuals of this regression be $\hat{\epsilon}_t$. We construct a simulated data set by letting $r_t = \hat{\eta} \cdot Controls_t + u_t$, where u_t is sampled (with replacement) from the empirical distribution $\hat{\epsilon}_t$.¹⁵ This generates a placebo time-series of stock returns.

To each simulated time-series r_t of 9592 observations, that cover the sample period 1970-2007 in the paper, we fit the model in (5), now with the journalists dummies included. For each simulation we save four F -statistics and their associated p -values for the null hypothesis that $\beta_{i,1} = \beta_{i,0} = 0$, i.e., the main joint test from Table 3. Specifically, for each bootstrapped sample, the F -test is performed using four different estimates for the variance-covariance of the estimated coefficients: (a) standard OLS, (b) two versions of White (1980) (HC2 and HC3, as implemented in R), and (c) Newey-West standard errors with five-lags. We do a total of 20,000 simulations.

Figure 4 plots histograms of the the p -values from such a test.¹⁶ Turning to the first panel in the graph, we see that standard OLS statistics are slightly biased, with heavy tails at both ends of the range of p -values. The bias for the lowest 2 percentile is on the other of 40%, i.e., we observe p -values of less than 2% (using the standard OLS asymptotic or Gaussian assumptions) about 2.8% of the time. We should note that while this evidence should warn the reader about potential biases

¹⁵We burn 100 runs to get over “initial conditions” generated from the existence of lags in (5).

¹⁶The histograms have 50 bins, so each one should contain 2% of the observations. The histogram is plotted as a density, so deviations from 1 can be interpreted as the percentage of bins with excess/lack of simulated values.

in this type of studies, the results in the paper are clearly not driven by finite sample biases: the p -values reported in Table 3 are on the order of 10^{-4} .

The other three panels offer a similar picture: clear biases in the estimators, with the White (1980) HC3 performing decently, although not much better than OLS, and the Newey-West procedure producing significantly bigger biases. This last fact has some interest of its own, as it highlights the dangers of using estimators that may “overfit” and, as a consequence when using asymptotic arguments, be over-optimistic on the precision of the estimated coefficients (Chesher and Jewitt (1987)).

Figure 4: Histograms of p -values for the test $\beta_0 = \beta_1 = 0$.

