Observing the Choices of a Delegated Information Monitor*

Roberto C. Gutierrez Jr.[a,†]
and
Charles Gaa[b]

March 11, 2023

Abstract

The information landscape is vast. Agents must make choices. As a delegated information monitor, *The Wall Street Journal* (WSJ) makes coverage decisions to provide value to its subscribers. The marginal benefits of reporting earnings news, as perceived by the WSJ, change over time. WSJ earnings coverage varies counter to the business cycle and increases with the equity risk premium, consistent with theories of optimal information acquisition. The choices of the WSJ reveal dynamics in investors' information structures. Increased earnings coverage signals a broad increase in the information content of small-stock prices, as well as an increase in the trading performance of mutual funds within the small-stock universe.

*We are grateful for comments from Ioannis Branikas, Jennifer Conrad, Diane Del Guercio, Michal Dzielinski, Eric Kelley, Kyle Peterson, Rob Ready, Brady Twedt, and seminar participants at the EFA Meeting (Austria), Pacific Northwest Finance Conference (University of Washington), University of Tennessee, and Tulane University. Gutierrez acknowledges financial support from the Cameron Center for Finance and Securities Analysis at the University of Oregon and the James E. Bramsen Faculty Fellowship. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.

[a]Lundquist College of Business, University of Oregon.(rcg@uoregon.edu)

[†]Corresponding author.

[b]Bank of Canada. (cgaa@bank-banque-canada.ca)

Information drives financial markets. Yet no single agent can fully monitor the myriad information flows. Attention must be allocated; choices must be made. Theories predict that the marginal benefits of acquiring information should be dynamic, changing as risk and the price of risk change. Therefore, the information structure in the marketplace is also predicted to be dynamic. Measuring changes in an unobservable information structure however is a difficult task.

Recent research emphasizes the role that information intermediaries can play in the economy. Nimark and Pitschner (2019) show that delegating the monitoring of information to a news agent is beneficial to subscribers because delegating conveys more information than subscribers can gather on their own. As a prominent delegated monitor of business news, the *Wall Street Journal* (WSJ) selects which information to report to their subscribers, who are primarily financial professionals. Decisions on whether or not to cover any given item of news are determined by the perceived marginal benefits of reporting that item relative to other items. By assuming the WSJ seeks to maximize its subscribers' utility, and ignoring any non-financial reporting biases, we can exploit the observable choices of the WSJ to reveal their perceptions of the marginal benefits to their subscribers acquiring earnings information.

In this paper, we lean on recent information theory to guide our examination of the decisions made by the WSJ to cover firms' quarterly earnings reports, and what their choices may be signaling. Earnings releases are widely viewed as important informational events. Many firms, sometimes hundreds, release their earnings information within days of each other. How many of these firms' reports should be covered? In our sample, the WSJ chooses to cover on average only about 11% of firm-earnings reports, but there is sizable time variation in the extent of their coverage. We focus in this study on that temporal variation.

We begin by estimating the likelihood of the WSJ covering each firm-earnings report based on firm and earnings characteristics. Filtering these static determinants of coverage choice reveals time variation in the WSJ perception of the marginal benefits of their subscribers acquiring earnings information.[1] When the marginal benefits are higher, more firms are covered.

Under what conditions would the marginal benefits of earnings information increase? Bansal and Shaliastovich (2011), Andrei and Hasler (2015), and Kacperczyk, Van Nieuwerburgh, and Veldkamp (2016) show that the optimal gathering of information about risky assets will vary countercyclically with the economy. In other words, information is more valuable to investors when the uncertainty of asset payoffs and the price of risk increase.[2]

We find that the WSJ expands its earnings coverage when output gap (detrended industrial production) is lower and when the consumption-wealth ratio of Lettau and Ludvigson (2001) is greater. Further support for the perceived marginal benefits of earnings information being countercyclical comes from the fact that the extent of coverage covaries positively with the equity risk premium. A one standard deviation increase in the extent of WSJ earnings coverage is associated with a roughly 2% increase in the excess stock market return over the next six months.[3]

Interestingly, all the signals that WSJ earnings coverage provide for stock returns are incremental to a broad set of macroeconomic measures of market conditions. Hence, tracking the state-dependent choices of a delegated monitor reveals stock market dynamics that are otherwise not observable.

We then turn to examining whether the extent of WSJ earnings coverage covaries with the information content of stock prices. We expect the WSJ view of the marginal benefits to

[1] Fang and Peress (2009), Solomon (2012), Ahern and Sosyura (2014), and others find newspaper coverage of firms to tilt toward larger firms, while Solomon and Soltes (2012) find coverage to tilt toward firms with earnings surprises. We also control for industry, analyst coverage, dispersion of analysts' forecasts, recent stock returns, and seasonalities.

[2] Schmalz and Zhuk (2019) and Loh and Stulz (2018) provide related evidence. They respectively find that stock prices respond more strongly to earnings surprises and to analysts' earnings revisions when the economy is weaker.

[3] We control for the aggregate monthly tone of the WSJ articles. Tetlock (2007), Tetlock et al. (2008), Solomon (2012), Gurun and Butler (2012), and others find that the tone (positive/negative) of media coverage covaries with future stock returns. The aggregate tone of the WSJ coverage is unrelated to all of the temporal dynamics in stock returns that we examine.

providing earnings information to be positively correlated with how actively their subscribers are acquiring earnings information. Recall that only a small fraction of firms actually receive earnings coverage, so this is not a firm-level signal but a broad one. Given our (perhaps unsurprising) finding that small firms are predominantly the marginal firms by which the WSJ adjusts its earnings coverage, we investigate the returns of small stocks for evidence of time variation in the informativeness of prices.

First, the cross-sectional dispersion in stock returns should increase when investors are more actively acquiring information and updating their subjective beliefs about valuations (e.g., Peress (2014), Kacperczyk et al. (2016)). We find that the cross-sectional return dispersion of small stocks is roughly 25% greater when the WSJ earnings coverage is highest compared to when it is lowest. Second, numerous studies employ post-earnings-announcement drift in returns (PEAD) as an inverse measure of how fully prices are impounding earnings news at their announcements. We find that high WSJ earnings coverage indicates times when PEAD in small stocks is low. Specifically, a long-short portfolio formed in months when the extent of coverage is lowest earns an abnormal return of nearly 6% over 60 days, while a long-short portfolio formed in the months when the extent of coverage is greatest earns roughly 2%.

In sum, this study shows that WSJ earnings coverage displays dynamics that support theories of state-dependent information choices. Greater earnings coverage is a signal of a weaker economy, a greater equity risk premium, and an increase in the informativeness of small-stock prices. We argue that WSJ coverage achieves its status as an indicator through its role as a delegated monitor of business news, whereby coverage decisions are made to maximize WSJ profitability by maximizing the utility of its subscribers. We show that dynamics in the information choices of the WSJ are linked to dynamics in stock returns.

Our findings support the nascent research on financial media as a delegated monitor of business news, rather than merely a disseminator of news. For example, Martineau and

Mondria (2022) consider the endogeneity of both media and investor choices and show conditions when media coverage triggers the acquisition and processing of private information by investors rather than crowding it out.

This study is also related to research that demonstrates complexities in the process through which information is impounded into stock prices. See, for example, the papers by Cziraki, Mondria, and Wu (2021), Liu, Peng, and Tang (2022), and Ben-Rephael, Carlin, Da, and Israelsen (2021), who note that firm-level information acquisition differs across subsets of investors. Our findings indicate that the marginal benefit to these investors from acquiring any particular bit of information should be state dependent.

Lastly, we apply the WSJ-based signal to examine time variation in mutual fund performance. Studies of mutual fund performance focus nearly exclusively on average performance or cross-sectional variation in performance. Moskowitz (2000), Glode (2011), and Kosowski (2011) provide notable exceptions. They show that fund performance is greater during economic recessions. Identifying such dynamics in fund performance is important for understanding the services that mutual funds provide and their means of production.

Given that WSJ earnings coverage displays traction for the informativeness of small-stock prices, we examine whether mutual funds' trading of small stocks covaries with informational conditions. We examine the return spread between the small stocks that mutual funds buy and the small stocks they sell, and find this spread to increase with the extent of WSJ earnings coverage. We leave the investigation of the underlying mechanisms of this result for future research. In related work though, Kacperczyk, Van Nieuwerburgh, and Veldkamp (2014, 2016) document variation in mutual fund strategies across economic conditions. Also, Nimark and Pitschner (2019) note how delegated monitoring of information can serve as a coordination mechanism among attention-constrained investors. The endogeneity of information acquisition and trading activities across investor types warrants better understanding.

The paper proceeds as follows. The next section details the data sources and how we measure the extent of WSJ earnings coverage. We then examine adjustments that the WSJ makes along the extensive margin in section 2. How the extent of coverage covaries with the macroeconomic state is investigated in section 3. Variation in the information content of stock prices is investigated in section 4. Section 5 examines the trading performance of mutual funds within the small-stock universe. Section 6 provides concluding thoughts.

1 Methodology and data

1.1 Corporate earnings events

The base dataset of earnings and *The Wall Street Journal* (WSJ) articles were collected and used by Gaa (2008). The sample of quarterly earnings reports includes 233,348 firm-earnings events from I/B/E/S, covering the period from October 1984 to December 2005. Twenty-one years of WSJ articles is a large sample by the standards of extant studies of media coverage that employ textual analysis. The starting point of the sample period is determined by data availability. The ending point of the sample precedes a change in ownership of the WSJ that potentially introduces nonstationarity into the editorial selection model of WSJ news coverage.[4]

To model the probability of each firm-earnings release being covered by the WSJ, we gather data from I/B/E/S, CRSP, and Compustat on firm-level characteristics, such as size, analyst coverage, recent stock performance, book-to-market ratio of equity (BE/ME), and industry, as well as earnings-specific characteristics, such as earnings surprise and pre-announcement forecast dispersion. Additionally, we omit earnings reports from firms having

[4]News Corp acquired the WSJ in 2007. Prior research offers evidence of a structural shift in coverage decisions. Archer and Clinton (2018) and the Pew Research Center (2011) document changes in WSJ front-page content after the acquisition. Using coverage decisions of the *The New York Times* to control for temporal changes in the information environment, these studies find that the WSJ increased its front-page coverage of politics and of lifestyle after 2007 and decreased its coverage of business news.

a negative BE/ME or having a stock price less than one dollar two days prior to the earnings announcement date. The complete set of variables and their sources are provided in Appendix A.1.

1.2 WSJ coverage of earnings

Our strategy of using the observed choices of the WSJ to signal changes in the information acquisition activities of their subscribers relies on the assumption that WSJ coverage decisions are made to maximize the utility of their subscribers. We follow Nimark and Pitschner (2019), Chahrour et al. (2021), and Martineau and Mondria (2022) who employ this same assumption to emphasize the state-dependence of coverage decisions made by information intermediaries. They show how this state dependency can influence what subscribers learn from observing coverage decisions of a delegated monitor. We contribute to this nascent literature by empirically examining state dependencies in WSJ earnings coverage and what that dependency can convey about stock market conditions.

Our measure of the extent to which the WSJ covers earnings is based on how many firm-earnings releases are selected by the WSJ to receive a news article. We begin with 68,102 "earnings" news articles from Factiva (code: c151) having at least 100 words. Requiring 100 words ensures that the earnings article represents material information acquisition. The computational linguistics program Rainbow by McCallum (1996) is then used to identify the articles which are about a specific firm's quarterly earnings release, as some of these earnings articles flagged by Factiva are about industry-level earnings trends, regulatory changes, restatements, accounting scandals, etc. The Naive-Bayesian text categorization is trained on a set of 500 articles and uses a unigram and bigram "bag of words" approach. Only articles with a posterior probability greater than 0.5 of being about a firm's earnings release remain in the final sample of 49,113 articles.

Although the linguistic tone of the WSJ articles is not of primary concern in this study, we want to distinguish the WSJ choice on the extent of their earnings coverage from the news content of their coverage. Controlling for tone is useful as we examine how the extent of earnings coverage covaries with aspects of stock returns. Prior research indicates that tone has explanatory power for stock returns (e.g., Tetlock (2007), Tetlock et al. (2008), Solomon (2012), and Gurun and Butler (2012)). To control for tone, we apply another round of text categorization to disambiguate "negative" and "nonnegative" articles, again trained on a set of 500 articles. We classify articles to be negative if the posterior probability of being negative is greater than 0.5, while all other articles are classified as nonnegative.[5]

Finally, for each of the earnings release, we search for a related article published in the WSJ within two days on either side of the I/B/E/S announcement date. If at least one article corresponding to a given earnings release is found within this 5-day window, that earnings release is considered to be "covered". We observe coverage for approximately 11% of the announcements in our sample, implying that the typical firm's quarterly earnings release is viewed by the WSJ to be insufficiently valuable to their subscribers.

1.3 Measuring the extent of earnings coverage

Our goal is to measure time variation in the marginal benefits of earnings information, as viewed by the WSJ. To begin, Table 1 provides summary statistics that illuminate what typical earnings coverage looks like. On average, 11% of earnings reports released in a given month are covered by the WSJ (97 articles out of 915 events). The monthly number of earnings reports that receive coverage has a mean of 97 with a relatively large standard deviation of 80, which is largely driven by the monthly changes in the volume of earnings reports released each month.

While normalizing the number of firms receiving earnings coverage by the number of earnings releases in that month can account for the volume changes, such a normalization implicitly assumes that each earnings report is equally likely to receive coverage in the

[5] We test Rainbow's classification accuracy by randomly excluding 100 articles from the training set, re-estimating the model, and then checking accuracy for the excluded articles. Across 100 trials, the average accuracy is 88.5% for the first-stage classification ("earnings/not-earnings") and 83.4% for the second-stage ("negative/non-negative").

WSJ. However, this is clearly not the case. Newspaper coverage is biased toward larger firms, certain industries, and firms with earnings surprises (e.g., Fang and Peress (2009), Solomon (2012), Solomon and Soltes (2012), and Ahern and Sosyura (2014)). To identify truly marginal choices in earnings coverage given these biases, we want to control for the characteristics of the firms that release earnings and the characteristics of the earnings information.[6]

We confirm that this filtering is important. The time series of the simple fraction of earnings reports receiving WSJ coverage displays little covariance with the return dynamics we examine in later sections. Hence, adjusting for static biases in WSJ coverage is necessary to increase the signal-to-noise ratio of WSJ coverage decisions.

To filter the monthly flow of earnings releases, we employ a multinomial logistic regression where the three responses are: negative coverage, nonnegative coverage, and no coverage. Classifications of each article as negative or nonnegative are done using the Rainbow program discussed in section 1.2. The set of firm and earnings characteristics we use as determinants are detailed in the Appendix section A.1. Output of the multinomial logit is shown in Appendix A.2. Using this model, we estimate the probability of coverage for each firm-earnings report. In short, firm and earnings characteristics provide a good deal of information about the probability of WSJ coverage. The (McFadden) pseudo R^2 is 0.27.

The primary determinants of WSJ coverage are firm size, industry, and the number of analysts covering the firm. These three variables alone account for a pseudo R^2 of 0.22. Each of these firm characteristics intuitively reflects greater perceived value for WSJ subscribers. Larger firms are generally more important to investors. Stocks of larger firms typically have greater weightings in investment portfolios, and larger firms typically have a greater number of shareholders. Also, the scale and scope of large-firms' operations makes them more likely to be considered "bellwethers" for other firms in their industries. Relatedly, particular industries, whether cyclical or defensive, can be better indicators of the economy's

[6]Nimark and Pitschner (2019) and Martineau and Mondria (2022) make similar appeals to filter editorial preferences from coverage outcomes.

condition. Additionally, some industries may contain firms whose performances are more highly correlated with each other, and hence, require fewer firm-level signals to understand. Lastly, the number of sell-side analysts that cover a firm reveals the outcomes of the cost-benefit analysis made by these research divisions in assessing whether to gather information on a particular firm. Their assessments are shown to be positively correlated with the WSJ assessments to cover a firm.

We remove the static determinants of WSJ earnings coverage to reveal the marginal dynamics in the extent of the coverage. The metric used is the percentage deviation of the actual number of firm-earnings events that receive coverage from the number predicted to receive coverage conditional on firm and earnings characteristics. We label this percentage residual as EXT, for the extent of earnings coverage.

$$EXT_t = \frac{\sum_{k=1}^{K_t}(C_k - \hat{C}_k)}{\sum_{k=1}^{K_t}\hat{C}_k} \tag{1}$$

where K_t is the total number of earnings releases observed in month t, C_k is an indicator variable equal to one if earnings report k is associated with a WSJ article and zero if no coverage is observed, and $\hat{C}_k$ is the sum of the predicted probabilities of negative and nonnegative coverage respectively over the earnings releases in month t.

For some perspective, in 1985, the first full year of our sample, the mean monthly number of earnings releases is 410; the mean number of releases receiving WSJ coverage is 43 per month and the predicted number is 40. In 2005, the last year of our sample, the mean monthly number of earnings reports released is 1099, while the mean number of releases receiving WSJ coverage is 104 per month and the predicted number is 155. In Table 1, we see that the mean monthly value of EXT is close to zero at 0.03, but EXT varies a good deal as indicated by its monthly standard deviation of 0.26.

Figure 1 plots EXT (the solid line). We can see the variability in EXT, and we can also see that EXT displays some persistence. The AR(1) coefficient of EXT is 0.62.[7] This is far below the near unit-root behaviors of the macroeconomic variables shown in Table 1, except one. Many studies raise concerns about spurious predictability of stock returns based on such highly persistent measures (e.g., Boudoukh et al. (2008)). We address these concerns by conducting Monte Carlo simulations to assess potential size distortions in our test statistics arising from examining a variable with an AR(1) coefficient of 0.62.

To control for the aggregate tone of the WSJ coverage of earnings, i.e. the outcome of the WSJ's information processing beyond the earnings characteristics that we observe, we form a residual measure of tone as:

$$TONE_t = \frac{\sum_{k=1}^{K_t}(C_k^{NNeg} - \hat{C}_k^{NNeg}) - \sum_{k=1}^{K_t}(C_k^{Neg} - \hat{C}_k^{Neg})}{\sum_{k=1}^{K_t}\hat{C}_k} \tag{2}$$

where C_k^{NNeg} is a dummy variable equal to one if earnings report k is associated with a nonnegative article, and zero otherwise, while C_k^{Neg} is a dummy variable equal to one if earnings report k is associated with a negative article and zero otherwise. Since each article is characterized to be either negative or nonnegative, $C_k = \left(C_k^{NNeg} + C_k^{Neg}\right)$. The predicted number of firms to receive either negative or non-negative WSJ coverage each month is labeled with a "hat". TONE declines as the percentage residual tone of coverage in month t becomes more negative.

EXT has a 0.30 correlation with TONE. Why the extent of earnings coverage is positively correlated with the tone is not pursued in this study. For the current purposes, we want only to control for the WSJ outcome of processing earnings news in our testing. In all regression tests, we include both EXT and TONE. In the single-variable exercises, such as the figures, we examine EXT orthogonalized with respect to TONE. For ease of exposition, however,

[7] The plot of EXT displays small monthly reversals around an underlying lower-frequency cycle. Recall that we include calendar month controls in the logit model. We have verified that EXT is centered near zero across the first, second, and third months of quarters, respectively. Hence, the monthly reversals in EXT are not due to a monthly nor a within-quarter seasonality. The monthly reversals in EXT are consistent with a binding capacity constraint for WSJ coverage.

we will still refer in the text to the orthogonalized version of the extent metric simply as "EXT".

2 Adjustments along the extensive margin

To better understand the dynamics in the WSJ extent of earnings coverage, we examine WSJ decisions to cover earnings across firm size (market value of equity). Given the increased propensity for the WSJ to cover large firms, we can reasonably expect smaller firms to comprise more of the marginal decisions to cover or not. Also, given that many asset-pricing anomalies are well-documented to be stronger for small firms, this dimension of the extensive margin of coverage seems fruitful to investigate.

To begin, we separate sample months into three bins based on aggregate EXT. Months with EXT above the 75th percentile are labeled "high," months below the 25th percentile are labeled "low," and remaining months are labeled "normal." Each month we recalculate the extent of earnings coverage within each quintile of stock size. In other words, we decompose aggregate EXT into its size-quintile components. We determine the number of firms within a given quintile of size for which the full-specification full-sample multinomial logit predicts coverage, and then form the extent of coverage as the percentage deviation of actual coverage from predicted coverage for that given size quintile.

Panel A of Table 2 reports the mean monthly values of the extent of WSJ earnings coverage for each size quintile across the low, normal, and high periods of coverage. The most striking result is the dramatic increase in the extent of coverage for the smaller stocks in high-EXT months. Coverage of earnings releases for the smallest stocks in these months balloons to 91%. The extent of coverage across the remaining four quintiles falls steadily down to 18% in the largest stocks.

Panel B shows the number of firm-earnings reports per month that are actually covered by the WSJ. The smallest firms receive the least amount of actual coverage, with only 10 firms on average in the high-EXT state. The coverage of one additional firm from the smallest quintile is a much larger deviation from the baseline probability of coverage. On the other end, the largest firms receive much greater coverage on average, with the actual number of firms in the highest quintile averaging 45 per month in the high-EXT state.

Analogously, in the low-EXT state, the extent of coverage recedes the least for the largest firms. In the low state, the largest firms average negative 19%, and the smallest firms average negative 37%. While the variation across the size quintiles in the low state is not as dramatic as that in the high state, nor monotonic, the extent of coverage is generally contracting as firm size declines.

In short, Table 2 shows that smaller firms are the marginal firms for WSJ earnings coverage. The smallest stocks experience a 91% increase in residual coverage in the high-EXT months and a 37% decrease in the low-EXT months. This spread between high and low states falls monotonically with firm size.

Most importantly, the correlation between the full-sample EXT and the extent of coverage within the lower four size quintiles is 0.95, while the correlation between the full-sample EXT and the extent of coverage within the largest quintile only is just 0.57. That is, the WSJ-based metric that we develop is essentially a proxy for changes in the marginal benefits of acquiring small-firm earnings information.

3 Extent of coverage and macroeconomic conditions

3.1 Countercyclical marginal benefits of information

The endogenous information models of Bansal and Shaliastovich (2011), Andrei and Hasler (2015), and Kacperczyk et al. (2016) predict that optimal information acquisition by investors moves countercyclically to the economy, as the variance of asset payoffs increases and the price of risk increases. In other words, information should be more beneficial to

WSJ subscribers during weaker economic conditions. In this section, we examine how the extensive margin of WSJ earnings coverage covaries with the economy.

Only two brief recessions in economic activity are identified by the NBER during our sample, July 1990 to March 1991 and March 2001 to November 2001. Hence, we rely more heavily on two other continuous measures of macroeconomic conditions. Output gap is measured as the deviation from a linear and quadratic trend in the log of industrial production (Cooper and Priestley (2009)). The consumption-wealth ratio (CAY) of Lettau and Ludvigson (2001) is more of a forward-looking measure of economic conditions. Lettau and Ludvigson (2001, 2002) find that CAY forecasts future stock returns and future growth in corporate investment (private nonresidential fixed investment), suggesting that CAY tracks risk premia and varies countercyclically.[8] The plots of EXT against output gap and against CAY are shown in Figures 1 and 2 respectively (along with NBER recessions as shaded bars). The comovement of EXT with each measure is impressive.

To examine the explanatory power of these variables statistically, we regress EXT on CAY and output gap. We also include other variables as regressors to provide additional stylized facts. Recall from Table 1 that EXT is autocorrelated. Our approach to dealing with the statistical concerns of spurious relations is two-fold. First, we employ the Newey-West estimator to produce the standard errors. EXT displays autocorrelation up to five lags. We use a six-month lag length for the Newey-West estimator. Second, we adjust the p-values of the t-statistics for spurious rejection rates using Monte Carlo simulations. We form 10,000 simulated samples of an independently and normally distributed random variable with an AR(1) coefficient that matches EXT. We regress each simulated variable on the actual sample of the macro variables. We then calculate the frequency of observing Newey-West t-statistics with six lags that are greater than a given t-statistic found in the

[8] CAY is the transitory deviation from a common stochastic trend in logs of aggregate consumption and wealth (that includes human capital). The measure arises from the log-linearization of the intertemporal budget constraint of investors.

actual sample (or less than a given t-statistic that is negative). We multiply the frequency by two to arrive at a simulated p-value for a two-tailed test.

In Table 3, we see that CAY and output gap robustly display marginal covariances with EXT, with signs that indicate a countercyclical pattern in the extensive margin. We can also observe that the set of macro variables used in the regression captures a large fraction of the temporal variation in WSJ earnings coverage, with an R-squared of 0.41. The remaining variables display no statistical relation with EXT except for dividend yield and the 90-day T-bill rate (RF). While short interest rates tended to move procyclically before the start of this sample period (Fama, 1990), Duffee (1996) shows that T-bill rates display increasing idiosyncratic movements starting in the 1980s. Hence, the interpretation of the sign of RF is a bit muddled. Interestingly, the sign on the dividend yield is significantly negative; however, the sign switches to positive when estimating a simple correlation between EXT and dividend yield. There is obviously some multicollinearity amongst these right-hand side measures. Which covariates are expected to be stronger and whether nonlinearities between economic conditions and information acquisition should be expected are interesting avenues for future research.

We now turn to examining the relation between the extent of earnings coverage and aggregate stock pricing. Does EXT move with the equity risk premium, which is also expected to increase both with the uncertainty in payoffs as well as the price of risk? To investigate, we regress various windows of the log of future returns of the CRSP VW stock index in excess of the T-bill rate on EXT and TONE. Table 4 finds that EXT tracks the equity premium. EXT covaries positively with excess stock returns over months (+1,+6) at the ten-percent level and over months (+1,+12) at the five-percent level (using simulated p-values).[9]

[9] When regressions include both EXT and $TONE$, the Monte Carlo simulations employ two independently drawn variables with AR(1) coefficients matching EXT and $TONE$ respectively. Readers can observe that the simulated p-values are sometimes notably greater than the p-values implied by Newey-West standard errors, consistent with HAC estimators performing poorly in finite samples when series display substantial autocorrelation.

Examining the explanatory power of EXT jointly across multiple horizons provides a more powerful and more stringent test (Boudoukh et al. (2008)). Panel B of Table 4 shows the simulated p-value testing the null hypothesis that the coefficients on EXT are jointly zero over months (+1,+6) and (+1,+12) to be 2.9%. Hence, EXT robustly comoves with the equity risk premium.

The economic magnitude of this relation is notable. The reported coefficients in Table 4 are with respect to standardized coefficients and can be easily interpreted. A one standard deviation increase in EXT is associated with an increase in excess stock returns of 1.89% over the following six months and 3.66% over the following twelve months, which are economically meaningful. Figure 3 plots means of excess stock returns across quintiles of EXT. The explanatory power of EXT for future stock returns over months (+1,+6) and (+1,+12) is visible across the full range of quintiles. Six-month returns vary from roughly 2% to 7%, and twelve-month returns vary from roughly 2% to 14%.

In sum, the WSJ extends its earnings coverage across more firms during weaker economic times. This behavior is consistent with the WSJ serving as a delegated monitor of information. The WSJ provides more earnings information to their subscribers when the additional information should be more beneficial to subscribers.

Lastly, the above findings are inconsistent with earnings coverage of the WSJ serving as a proxy for the attention of irrational retail traders. The "ostrich effect" predicts that attention is procyclical, i.e., greater when stock returns have been high. By measuring the number of logins to investment accounts, Sicherman et al. (2016) provide evidence that retail traders seek less information about their portfolios as the returns of the stock market decline, possibly because investors want to delay realizing a bad outcome.[10] Regressions like those in Table 4 with lagged excess stock returns as the dependent variable, instead of future returns,

[10]Retail trading has been linked to Google search volumes and to articles in local newspapers by Da et al. (2011) and Engelberg and Parsons (2011), respectively. In contrast, Yuan (2015) notes that WSJ subscribers are "primarily financial professionals."

indicate that the extent of WSJ earnings coverage does not covary with lagged returns over either a $[-6, -1]$ or a $[-12, -1]$ window.

3.2 Extensive margin as an incremental signal

In the prior section, we examined how EXT covaries with the state of the economy. The fact that several measures are used speaks to the difficulty in measuring economic conditions. One just has to look at the mosaic of indicators that are commonly used to diagnose the health of the economy to recognize the complexity of filtering estimates of the underlying latent state. The information channel is a promising avenue for expanding how we characterize economic conditions.

For example, instability over time in many of the interrelations among macro variables is well documented (Welch and Goyal, 2008; Stock and Watson, 2003). Recognizing that information must be acquired and that agents choose when and how to gather information can potentially explain some amount of parameter instability. Bansal and Shaliastovich (2011) make such a point explicitly. They argue that jumps in stock prices coincide with investors' endogenously choosing to gather information about the state of the economy at that particular point in time. The acquired information produces a discrete change in their expectations of fundamentals, which produces a jump in prices. The price jumps are disconnected from the presumed smooth movements in underlying fundamentals. Hence, choices on information acquisition can generate instability in covariances between prices and fundamentals.

To provide future research with additional stylized facts to consider, we examine whether EXT provides incremental linear information about the equity premium beyond that of the set of macro variables considered in Table 3. Recall that the set of macro variables correlates sizably with the temporal variation in the extent of earnings coverage, with an R-squared of 0.41. To assess the incremental covariance of EXT with the equity risk premium, we rerun the regressions of excess stock market returns on EXT and TONE adding the set of eleven continuous macroeconomic variables.

The right half of Table 4 displays the results of this larger regression for the two horizons. EXT provides linear incremental information about the equity premium over months $(+1, +6)$ at a ten-percent level of significance using the simulated p-value, but statistically provides no incremental information over the $(+1, +12)$ window. The joint test across both windows in Panel B remains significant at a 5% level. Hence, the extensive margin of WSJ earnings coverage offers information on the equity risk premium that is not linearly spanned by the set of macro variables. Moreover, the dynamics in earnings coverage seems to provide a shorter-run signal of the equity premium than the macro set provides.

4 Extensive margin and the information content of prices

Does the expansion and contraction of earnings coverage by the WSJ covary with the informativeness of stock prices? Prices will reflect the information that investors possess. When the mass of investors in the stock market is acquiring less information about firms' earnings, stock prices will embed less of the earnings information, all else equal. In this section, we examine how prices react to earnings surprises as the extent of the WSJ earnings coverage changes. Given the finding in section 2 that small firms are the marginal firms upon which the WSJ adjusts its earnings coverage, we focus on the information content of small-stock prices.

4.1 Post-earnings announcement drift

The tests herein consider the speed with which prices update to earnings news. One common assessment of how quickly stock prices impound earnings surprises is the post-earnings announcement drift in stock returns (hereafter "PEAD"; Bernard and Thomas (1989)). To the extent that PEAD is due to the market's slow reaction to earnings news, rather than to mismeasurement of expected returns, we expect the drift in returns to be smaller for earnings announced in months when the extent of earnings coverage is greater.

Each month we sort the firms announcing earnings in that month according to their standardized unexpected earnings (SUE), defined as the announced earnings per share minus the median analyst forecast within thirty days prior to the announcement divided by the pre-announcement stock price. We then form a PEAD portfolio by taking a long position in the stocks in the highest decile of SUE and a short position in the stocks in the smallest decile of SUE. Daily abnormal returns to each stock are adjusted for size and book-to-market effects using 5×5 benchmark portfolios, with the quintile breakpoints determined from NYSE stocks only.[11] We then cumulate the abnormal daily returns (CAR) for each stock over a window beginning two days after the announcement and ending sixty days after the announcement. The CAR are then equally weighted within each month's long and short SUE portfolios respectively.

We sort calendar months over the sample period into quintiles based on EXT and then report the mean abnormal profits of the long-minus-short SUE portfolios. The left panel of Figure 4 reveals that the average PEAD profits within the smallest stocks decrease from nearly 6% over the $[+2, +60]$ window for earnings surprises occurring in months with the least extensive coverage down to roughly 2% for earnings surprises occurring in months with the most extensive coverage. In short, the extent of WSJ earnings coverage gets a good deal of traction in explaining temporal variation in PEAD.[12]

To assess the statistical significance of this negative relation, we employ monthly regressions. Each month we regress the cross section of cumulative abnormal stock returns over days $[+2, +60]$ on the SUE for each firm announcing earnings in that month. Then, we regress the time series of monthly cross-sectional coefficients from the first-stage regression on EXT and TONE. Table 5 shows the second-stage results. We see that PEAD profits for

[11] We thank Kenneth French for providing the benchmark data on his website.

[12] The nonlinearity displayed by PEAD profits in quintile 3 may indicate a nonlinearity between asset pricing and information acquisition that is a byproduct of endogenizing information choices. Equilibria in these models are typically complex as they are outcomes of constrained optimizations that involve cost-benefit thresholds, corner solutions, dynamic feedback, complementarities, and other features.

small stocks decrease with EXT at a 1% level of significance (using simulated p-values).[13] In sum, when the extent of earnings coverage is greater, the price of a small stock immediately after the firm's earnings announcement tends to embed the earnings surprise more fully.

The strong comovement of the PEAD of small stocks with the extent of earnings coverage is strikingly different from the nearly flat PEAD plot for large stocks on the right side of Figure 4. Consistent with the findings of section 2, that small stocks lie on the extensive margin, the expansion and contraction of earnings coverage by the WSJ apparently provides insights into the pricing of small stocks, but not large stocks.

Our ability to examine the PEAD of large stocks is however a bit hindered because there are months in which only a few large stocks release earnings. As a consequence, we examine quintiles of SUE to produce the right side of figure 4, rather than deciles as we do for small stocks on the left side. We also require a minimum of 10 firm-earnings releases each month for the long and short portfolios. This filter reduces the number of usable large-stock sample months from 255 to 170. As figure 4 shows, large-stock PEAD does not temporally covary with the extent of WSJ earnings coverage.

We further investigate a relation between large-stock PEAD and EXT using the cross-sectional regression approach. A potential benefit with the regression approach is that we can extract more information when confronted with just a few stocks each month than is possible with the portfolio approach (since long and short stocks are pooled together each month in the regressions). Unfortunately, SUE is closer to zero for large stocks which can produce greater variability in first-stage coefficients. Although we do not tabulate the results, we can report that no relation between large-stock PEAD and EXT is detected. The failure to detect a covariance between the extent of earnings coverage and the return dynamics of large stocks is a pervasive finding throughout this study. In short, adjustments to the extent of earnings coverage are silent about the temporal dynamics of the returns of large stocks, but they do speak to the return dynamics of small stocks. Again, this is consistent with the

[13] The lag length for the Newey-West estimator is reduced to three because the AR(1) coefficient of PEAD is only 0.21.

finding in section 2 that changes over time in WSJ earnings coverage essentially reflect the marginal benefits of acquiring earnings information on small stocks, not large stocks.

Finally, as done in the equity premium analysis, we provide additional stylized facts by including the set of eleven continuous macroeconomic variables. The results are reported in the second column of Table 5. With respect to PEAD, EXT also provides linearly incremental information over the span of the macro variables. This may be related to earnings information acquisition being more of a shorter-run signal for stock returns, as noted earlier.

The finding here that high WSJ earnings coverage signals times when small-stock prices more quickly embed earnings news complements firm-level studies of attention to earnings (e.g., Peress (2008); Hirshleifer et al. (2009); DellaVigna and Pollet (2009); Drake et al. (2015); Boulland et al. (2017)). Our findings differ in two ways though. First, recall that only eleven percent of firms that release earnings in a given month are covered on average by the WSJ, and then only a mere handful of those are from the smallest quintile of firms. Hence, by construction, EXT is not *directly* measuring attention to the set of small firm's earnings news. Instead, EXT is shown to signal times when investors are broadly acquiring more information on small stocks. Secondly, and more importantly, we are measuring dynamics in the perceived marginal benefits of acquiring earnings information. This implies that stock investors' information structures are dynamic, which the PEAD finding and findings in subsequent sections support.

4.2 Return dispersion

An underpinning of information models is that an agent's beliefs are updated when information is acquired, and in this study, we are considering updates to stock valuations. Our previous finding for PEAD is consistent with investors acquiring more information about the earnings of small stocks when the WSJ is increasing the extent of its earnings coverage. With more information being impounded into prices, the cross-sectional dispersion of returns should be greater as the valuations for a greater number of firms are being updated.[14] In this section, we examine the relation between stock-return dispersion and the extent of WSJ earnings coverage.

We begin by sorting months into quintiles based on EXT. For each quintile, the left panel of Figure 5 reports the mean monthly standard deviation of the cross section of returns for firms in the smallest quintile of market value of equity. We see that return dispersion among small stocks increases notably with EXT, from about 19% per month in the months with the lowest EXT to 24% in the months with the greatest EXT. To assess statistical significance, we regress the monthly time-series of return dispersion on EXT and TONE. We see in Table 6 that the relation between EXT and return dispersion in small stocks is statistically strong, with simulated p-values below one percent.

For completeness, Figure 5 examines return dispersion among the largest stocks. The range of return dispersion for large stocks is from a little below 8% per month to a little above 9%. So once again, EXT gets no traction with large-stock return dynamics. We can report that the regression test finds no relation between EXT and the return dispersion of large stocks.

In sum, the increased return dispersion of small stocks when EXT is greater is consistent with dynamics in the information acquisition of investors. When investors are acquiring information about more small stocks, price adjustments are larger, which leads to greater cross-sectional dispersion in returns. Again, we find that WSJ earnings coverage provides a signal that is incremental to the set of macroeconomic variables.

5 Dynamics in the performance of mutual funds

Results of the prior sections combine to suggest that EXT identifies periods when investors are more actively updating small-stock prices to better reflect information about earnings.

[14] Any heterogeneity across investors' signals (and beliefs) would seemingly amply this effect.

In this section, we examine how the trading performance of a particular group of investors, mutual funds, varies over time with WSJ earnings coverage.

Several elements of the literature motivate this investigation of whether mutual fund performance covaries with the information dynamics that we observe. First, much research has sought to understand what value mutual funds may provide to their shareholders. A long-standing puzzle of active management is that these funds on average underperform passive index funds (Gruber, 1996). We will not attempt to summarize this vast literature, but will highlight one feature of fund performance that has been receiving more attention.

Moskowitz (2000) and Kosowski (2011) show that fund performance on average is greater during economic recessions. Glode (2011) notes that this dynamic in fund performance should be valued by fund shareholders whose marginal utility is greater during recessions. Kacperczyk et al. (2014, 2016) argue that it is optimal for fund managers to change their strategies — and information acquisition — over the business cycle. Our contribution is to link the perceived changes in the marginal benefits of earnings information to changes in fund performance.

That the WSJ-based metric we develop focuses on earnings information is also a motivation for examining mutual fund performance. Baker et al. (2010) show that mutual funds possess skill at forecasting firm earnings. They find that the buying and selling of stocks by mutual funds leads future earnings surprises. Jiang and Zheng (2018) find that the ability to forecast earnings surprises varies across funds. They also show that funds enjoy better returns from their earnings forecasting when the forecast dispersion of sell-side analysts is high. Jiang and Zheng interpret this to mean that fund returns increase as stock prices move closer to fundamental values and that the effect is larger when other investors disagree more on a stock's value.

The properties that the WSJ-based metric displays in the prior sections suggest that it may covary with mutual funds' trading performance within the small-stock universe. Namely, when EXT is high, the prices of small stocks are updating more actively and reflecting earnings news more quickly.

Following Edelen, Ince, and Kadlec (2016), we estimate the change in aggregate mutual fund demand for a given stock as the change in the number of funds that are holding that stock from last quarter to this quarter. This measure aggregates across the individual fund-level buy/sell signals, ignoring the sizing of the changes which would otherwise weight the aggregation toward reflecting the signals of the largest funds. Also, this measure captures only entry and exit decisions of funds, which reasonably reveal more misvaluation-based trading than adjustments to ongoing holdings do (Baker et al., 2010).

We collect data from Thomson Reuters on the quarterly stock holdings of actively managed domestic equity mutual funds.[15] Following the methodology of Edelen, Ince, and Kadlec (2016), we scale the quarterly change in the number of funds that hold a given stock by the mean number of mutual funds that are holding the stocks of small firms at the end of the prior quarter. This scaling captures the rate of change in the number of funds that hold a stock based on the typical number of funds that are holding a similar stock, which facilitates cross-sectional and temporal comparisons. Finally, the scaled measure is winsorized quarterly at the 1% and 99% values.

Because changes in fund holdings are quarterly, we recalculate EXT and TONE at quarterly intervals so that the state of earnings information acquisition will match the horizon over which we measure the trades of funds. To maintain power for our tests, we still use monthly observations for stock returns.[16] To avoid spurious rejections, we again employ simulations to adjust the p-values for the serial autocorrelations in EXT and TONE (and

[15] We identify domestic equity funds using Lipper Prospectus objective codes equal to EI, EIEI, ELCC, G, GI, LCCE, LCGE, LCVE, LSE, MC, MCCE, MCGE, MCVE, MLCE, MLGE, MLVE, MR, S, SCCE, SCGE, SCVE, SESE, SG, or Strategic Insight objective codes equal to AGG, GMC, GRI, GRO, ING, SCG, or Weisenberger objective codes equal to GCI, IEQ, IFL, LTG, MCG, SCG, G, G-I, G-I-S, G-S, G-S-I, GS, I, I-G, I-G-S, I-S, I-S-G, S, S-G-I, S-I, S-I-G. We remove funds ever having an index fund flag of D in the CRSP dataset. Only holdings reported for the ends of calendar quarters are used. Funds with TNA less than ten million dollars are removed. Shares are adjusted for stock splits occurring between the report date and the filing date.

[16] The lag length for the Newey-West estimator is increased to twelve in this section because the quarterly AR(1) coefficient of EXT increases to 0.83.

the macro variables when included). Stocks with prices below five dollars at the end of the quarter are removed.

Our assessment of how well mutual funds are trading small stocks begins by estimating a cross-sectional regression of a given month's stock returns (adjusted for size and book-to-market effects using 5×5 matched portfolios) on the most recent calendar quarter's change in the number of funds holding each stock. The time series of monthly first-stage coefficients is then regressed on EXT measured over the prior quarter, contemporaneous with the observed changes in quarterly fund holdings. In other words, the first-stage regression captures the quarter $t+1$ return spread between the small stocks that mutual funds bought the most in quarter t and the small stocks mutual funds sold the most. The second stage examines how this future buy-sell return spread changes conditional on WSJ earnings coverage.

Table 7 reports the second-stage results. The trading performance of mutual funds within the small-stock universe increases with EXT, at the five-percent level of significance (using simulated p-values). To provide an economic interpretation of this interaction, we turn to portfolios. We sort small stocks each quarter into quintiles based on changes in the number of funds holding a given stock. Stocks in the highest quintile are held long, and stocks in the lowest quintile are held short. Monthly abnormal profits of long-short portfolios over each of the next three months (one quarter) are calculated. We then compare the profits across high and low states of WSJ earnings coverage, sorting the 85 calendar quarters into two bins based on the median value of EXT over the sample period. Finally, to also consider changes over time in investing styles, the time series of long-short profits is filtered through a four-factor model of returns using Mkt-RF, SMB, HML, and UMD, allowing for loadings to vary across information states.[17]

The mean monthly abnormal profit of this trade-mimicking portfolio is 16 basis points per month when the portfolio is formed in quarters with high EXT, and -0.34 basis points per month when formed in quarters with low EXT. This is a sizable spread in alpha across the

[17] We thank Kenneth French for providing these data on his website.

information states, even given the variability in small-stock returns. The t-statistic testing that these means are equal is 1.95 (not tabulated).

Our findings here suggest a need to better understand the interactions of mutual fund activities, their return performance, and investors' information structures. Are other investors making different information choices than mutual funds? Are their information structures converging to those of mutual funds when EXT is high? Are the actions of delegated information monitors such as the WSJ serving as a coordinating mechanism across investor types (Nimark and Pitschner, 2019), and does this coordination vary over time? We leave these pursuits for future research.

Finally, we observe a few more empirical facts surrounding EXT. First, Panel A of Table 7 shows that the covariation between EXT and mutual-fund performance is incremental to the linear span of the eleven macro variables. So conditioning on the state of the macroeconomy is once again not fully capturing dynamics of the information state. Including the NBER recession indicator as a twelfth explanatory variable does not alter this conclusion.

Also, the long-short portfolio only loads on the UMD factor, and this loading varies across states, declining from 0.20 in the low EXT state to 0.10 in the high EXT state (at a ten-percent level of significance, not tabulated). Momentum trading may have an information channel, with past returns serving less as a trade signal when information acquisition activities increase.

6 Conclusion

Measuring the state of the information structures in the marketplace is an important avenue for future research. We extend studies of limited attention and information choices by observing the actions of one prominent delegated monitor of information, *The Wall Street Journal*. Given the flow of firms' earnings reports, how many of these firms will the WSJ choose to cover? We measure changes in the extensive margin of WSJ earnings coverage over time, revealing the WSJ perceptions of marginal changes in the benefits of earnings information.

We show how the choices of the WSJ can be employed to track meaningful changes in the state of the economy and the state of the information environment for stock investors. After removing static biases in WSJ coverage, we find that the decision of the WSJ to expand its earnings coverage indicates a weaker economy, a larger equity risk premium, and an increase in the informativeness of the stock prices of small firms. This WSJ-based signal is incremental to a set of commonly used macroeconomic measures. Relying on the observable choices of a delegated information monitor to extract stock-market signals is novel and arises from state-dependency in the information choices of the WSJ.

We utilize the extent of WSJ earnings coverage to better understand the services that mutual funds provide their shareholders. Glode (2011) recognizes that providing alpha in weak economic times is a valued benefit to fund shareholders. Kacperczyk et al. (2014, 2016) recognize that some funds change their strategies between economic recessions and expansions. We complement these studies by identifying dynamics between mutual fund performance and the information environment. Mutual funds trade small stocks more successfully when the extent of WSJ earnings coverage is high. Further examinations of the actions of delegated information monitors and the information structures of investors seem fruitful.

A Appendix

A.1 Variable definitions

A.1.1 Firm and earnings characteristics

Earnings announcement dates, actual earnings, and analysts' forecasts of earnings are from I/B/E/S. Stock prices, returns, and trading volume are from CRSP. Book value of equity is from the CRSP/Compustat Merged Database. Institutional ownership is obtained from Thomson Reuters.

UE quantiles are indicator variables. Each quarterly earnings release is assigned to one of 11 quantiles based on its earnings surprise, where the surprise is the announced EPS minus the median analyst earnings forecast within 30 days prior to the announcement normalized by the closing stock price two days prior. Quantiles 1 to 5 rank the negative surprise announcements into equal-sized quintiles. Quantile 6 consists of zero-surprise announcements where announced earnings equal the median of analysts' forecasts. Quantiles 7 to 11 rank the positive surprise announcements into equal-sized quintiles. Indicator variables are formed for each quantile other than the zero-surprise quantile 6, which serves as the base case.

Loss is a dummy variable equal to 1 if the announced earnings is negative.

Stdev(analysts' forecasts) is the standard deviation of analysts' EPS forecasts during the 30 calendar days prior to the announcement, with each forecast normalized by the closing stock price two days prior to the announcement. At least two covering analysts are required.

log(analysts' coverage) is the natural logarithm of one plus the number of distinct analysts' forecasts observed over the 30 calendar days prior to the announcement.

$log(ME)$ is the natural logarithm of the number of shares outstanding multiplied by the firm's closing stock price two days prior to the announcement.

$log(value\ of\ trading)$ is the natural logarithm of a stock's mean dollar value of daily trading volume over the 60 trading days prior to the announcement.

$Beta$ is the estimated coefficient from a regression of a firm's daily stock return on the S&P 500 return over the 60 days prior to the announcement.

$Recent\ returns$ is a stock's mean daily return over the 60 trading days prior to the announcement.

$Stdev(recent\ returns)$ is the standard deviation of a stock's daily returns over the 60 trading days prior to the announcement.

BE/ME is the firm's book value of equity from the fiscal year ending in the previous calendar year divided by its market value of equity from December 31 minus the value-weighted average book-to-market ratio of all announcing firms over the rolling three month period ending in the current month.

$Distraction$ is the announcement day's decile rank (in a given calendar quarter) based on the number of earnings announcements released from other firms on the same day. See Hirshleifer et al. (2009).

$Institutional\ ownership$ is the percentage of shares held by institutions at the end of the previous calendar year obtained from 13F filings.

$Seasonality\ and\ industry\ dummies$ are three indicator variables for month-of-the-year, day-of-the-week, and the 49 Fama-French industries, respectively. Firms are assigned to industries using SIC codes from Compustat.

A.1.2 Macroeconomic Variables

$PDND$ is the value-weighted dividend premium from Baker and Wurgler (2004). (Downloaded from Jeffrey Wurgler's website.)

$log(Div.\ Yield)$ is the natural logarithm of the market dividend yield (aggregate dividends for months t to $t-11$, divided by total market capitalization in month t for NYSE, AMEX, and Nasdaq firms). (Downloaded from Michael Roberts' website.)

$log(Net\ Payout\ Yield)$ is the natural logarithm of the total net payout yield (where the equity issuance yield is aggregate net equity issues for months t to $t-11$, divided by total market capitalization in month t). (See Boudoukh et al. (2007). Downloaded from Michael Roberts' website.)

$Risk\text{-}free\ rate$ is the US 90-day T-Bill rate. (Downloaded from Michael Roberts' website.)

CAY is the estimated quarterly deviation from the long-run log aggregate consumption wealth ratio. (See Lettau and Ludvigson (2001). Downloaded from Sidney Ludvigson's website.)

B/M is the book value of equity divided by its market value for the Dow Jones Industrial Average. (Downloaded from Amit Goyal's website.)

$Default\ Spread$ is the difference between the BAA and AAA corporate bond yields from FRED. (Downloaded from Amit Goyal's website.)

$Term\ Spread$ is the yield on the 10-year Treasury bond minus the yield on the 3-month Treasury bill. (Downloaded from Amit Goyal's website.)

$Equity\ Share\ of\ New\ Issues$ is the dollar amount of equity new issues divided by the dollar amount of total new issues (debt plus equity) described in Baker and Wurgler (2000). (Downloaded from Jeffrey Wurgler's website.)

Sentiment is the sentiment index from Baker and Wurgler (2006), which is based on the principal component of 6 sentiment proxies. (Downloaded from Jeffrey Wurgler's website.)

Output Gap is the estimated residual from a linear and quadratic monthly time trend in the natural logarithm of US Industrial Production over the sample period. See Cooper and Priestley (2009).

A.2 Logit Model Results

Table A.1 Predicting WSJ coverage

Below are the estimated coefficients from the multinomial logit regression of $c_{i,t} \in (-1, 0, 1)$ on various explanatory variables, where $c_{i,t}$ equals -1 when the earnings report for firm i in month t receives negative coverage, 1 when it receives nonnegative coverage, and 0 when it receives no coverage (the base case). Variable definitions are in section A.1. The t-statistics (in parentheses) are computed using standard errors clustered by firm; * indicates significance at 10%; ** indicates significance at 5%;*** indicates significance at 1%.

	WSJ Coverage	
	-1	1
UE quantile 11	0.805***	0.530***
	(11.12)	(7.52)
UE quantile 10	0.411***	0.336***
	(6.40)	(6.05)
UE quantile 9	0.113*	0.248***
	(1.74)	(4.96)
UE quantile 8	-0.113*	0.137***
	(-1.77)	(3.00)
UE quantile 7	-0.462***	0.00240
	(-6.77)	(0.06)
UE quantile 5	0.112*	-0.0185
	(1.76)	(-0.39)
UE quantile 4	0.603***	-0.000930
	(9.47)	(-0.02)
UE quantile 3	0.916***	0.186***
	(14.22)	(3.04)
UE quantile 2	1.048***	0.260***
	(14.91)	(3.56)
UE quantile 1	1.272***	0.367***
	(16.05)	(3.32)
Loss dummy	1.041***	-1.449***
	(19.01)	(-17.63)
Stdev(analysts' forecasts)	0.0624***	-0.0905**
	(4.75)	(-2.53)
log(analysts' coverage)	0.410***	0.356***
	(10.57)	(8.18)
log(ME)	0.442***	0.598***
	(12.38)	(14.95)
log(value of trading)	0.226***	0.210***
	(7.82)	(6.49)
Beta	-0.0986***	-0.0146
	(-4.16)	(-0.55)

Recent returns	-0.00226***	-0.00177***
	(-4.55)	(-3.47)
Stdev(recent returns)	0.640	-10.01***
	(0.53)	(-6.10)
BE/ME	0.570***	0.339***
	(20.79)	(8.47)
Distraction	-0.0346***	-0.0637***
	(-2.99)	(-5.13)
Seasonality and industry (FF-49) dummies	Yes	Yes
Observations	233348	
Pseudo R^2	0.269	

References

Ahern, Kenneth R., and Denis Sosyura, 2014, Who writes the news? Corporate press releases during merger negotiations, *The Journal of Finance* 69, 241–291. [page 2], [page 8]

Andrei, Daniel, and Michael Hasler, 2015, Investor attention and stock market volatility, *Review of Financial Studies* 28 (1), 33–72. [page 2], [page 12]

Archer, Allison M., and Joshua Clinton, 2018, Changing owners, changing content: Does who owns the news matter for the news?, *Political Communication* 35, 353–370. [page 5]

Baker, Malcolm, Lubomir Litov, Jessica A. Wachter, and Jeffrey Wurgler, 2010, Can mutual fund managers pick stocks? Evidence from their trades prior to earnings announcements, *Journal of Financial and Quantitative Analysis* 45, 1111–1131. [page 22], [page 23]

Baker, Malcolm, and Jeffrey Wurgler, 2000, The equity share in new issus and aggregate stock rteurns, *Journal of Finance* 55, 2219–2257. [page 29]

Baker, Malcolm, and Jeffrey Wurgler, 2004, A catering theory of dividends, *Journal of Finance* 59, 1125–1165. [page 29]

Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor sentiment and the cross-section of stock returns, *Journal of Finance* 61 (4), 1645–1680. [page 30]

Bansal, Ravi, and Ivan Shaliastovich, 2011, Learning and asset-price jumps, *Review of Financial Studies* 24 (8), 2738–2780. [page 2], [page 12], [page 16]

Ben-Rephael, Azi, Bruce I Carlin, Zhi Da, and Ryan D Israelsen, 2021, Information consumption and asset pricing, *The Journal of Finance* 76, 357–394. [page 4]

Bernard, Victor L, and Jacob K Thomas, 1989, Post-earnings-announcement drift: Delayed price response or risk premium?, *Journal of Accounting Research* 27, 1–36. [page 17]

Boudoukh, Jacob, Roni Michaely, Matthew Richardson, and Michael R. Roberts, 2007, On the importance of measuring payout yield: Implications for empirical asset pricing, *Journal of Finance* 62 (2), 877–915. [page 29]

Boudoukh, Jacob, Matthew Richardson, and Robert F. Whitelaw, 2008, The myth of long-horizon predictability, *Review of Financial Studies* 21 (4), 1577–1605. [page 10], [page 15]

Boulland, Romain, François Degeorge, and Edith Ginglinger, 2017, News dissemination and investor attention, *Review of Finance* 21, 761–791. [page 20]

Chahrour, Ryan, Kristoffer Nimark, and Stefan Pitschner, 2021, Sectoral media focus and aggregate fluctuations, *American Economic Review* 111, 3872–3922. [page 6]

Cooper, Ilan, and Richard Priestley, 2009, Time-varying risk premiums and the output gap, *Review of Financial Studies* 22 (7), 2801–2833. [page 13], [page 30] Cziraki, Peter, Jordi Mondria, and Thomas Wu, 2021, Asymmetric attention and stock returns, *Management Science* 67, 48–71. [page 4]

Da, Zhi, Joseph Engelberg, and Pengjie Gao, 2011, In search of attention, *Journal of Finance* 66 (5), 1461–1499. [page 15]

DellaVigna, Stefano, and Joshua M. Pollet, 2009, Investor inattention and Friday earnings announcements, *Journal of Finance* 64 (2), 709–749. [page 20]

Drake, Michael S., Darren T. Roulstone, and Jacob R. Thornock, 2015, The determinants and consequences of information acquisition via edgar, *Contemporary Accounting Research* 32, 1128–1161. [page 20]

Duffee, Gregory R., 1996, Idiosyncratic variation of treasury bill yields, *The Journal of Finance* 51, 527–551. [page 14]

Edelen, Roger M., Ozgur S. Ince, and Gregory B. Kadlec, 2016, Institutional investors and stock return anomalies, *Journal of Financial Economics* 119, 472 – 488. [page 23]

Engelberg, Joseph E.., and Christopher A. Parsons, 2011, The causal impact of media in financial markets, *Journal of Finance* 66 (1), 67–97. [page 15]

Fama, Eugene F., 1990, Term-structure forecasts of interest rates, inflation and real returns, *Journal of Monetary Economics* 25, 59 – 76. [page 14]

Fang, Lily, and Joel Peress, 2009, Media coverage and the cross-section of stock returns, *The Journal of Finance* 64, 2023–2052. [page 2], [page 8]

Gaa, Charles Clyde, 2008, *Media coverage and investor attention*, Ph.D. thesis, University of British Columbia. [page 5]

Glode, Vincent, 2011, Why mutual funds "underperform", *Journal of Financial Economics* 99, 546 – 559. [page 4], [page 22], [page 26]

Gruber, Martin J, 1996, Another puzzle: The growth in actively managed mutual funds, *The Journal of Finance* 51, 783–810. [page 22]

Gurun, Umit G., and Alexander W. Butler, 2012, Don't believe the hype: Local media slant, local advertising, and firm value, *The Journal of Finance* 67, 561–598. [page 2], [page 7]

Hirshleifer, David, Sonya Seongyeon Lim, and Siew Hong Teoh, 2009, Driven to distraction: Extraneous events and underreaction to earnings news, *Journal of Finance* 64 (5), 2289–2325. [page 20], [page 28]

Jiang, Hao, and Lu Zheng, 2018, Active Fundamental Performance, *The Review of Financial Studies* 31, 4688–4719. [page 22]

Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp, 2014, Time-varying fund manager skill, *The Journal of Finance* 69, 1455–1484. [page 4], [page 22], [page 26] Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp, 2016, A rational theory of mutual funds' attention allocation, *Econometrica* 571–626. [page 2], [page 3], [page 4], [page 12], [page 22], [page 26]

Kosowski, Robert, 2011, Do mutual funds perform when it matters most to investors? US mutual fund performance and risk in recessions and expansions, *Quarterly Journal of Finance* 01, 607–664. [page 4], [page 22]

Lettau, Martin, and Sydney Ludvigson, 2001, Consumption, aggregate wealth, and expected stock returns, *Journal of Finance* 56, 815–849. [page 2], [page 13], [page 29], [page 38]

Lettau, Martin, and Sydney Ludvigson, 2002, Time-varying risk premia and the cost of capital: An alternative implication of the q theory of investment, *Journal of Monetary Economics* 49, 31 – 66. [page 13]

Liu, Hongqi, Lin Peng, and Yi Tang, 2022, Retail attention, institutional attention, *Journal of Financial and Quantitative Analysis* 1–34. [page 4]

Loh, Roger K., and Rene M. Stulz, 2018, Is sell-side research more valuable in bad times?, *The Journal of Finance* 73, 959–1013. [page 2]

Martineau, Charles, and Jordi Mondria, 2022, News selection and asset pricing implications, working paper, University of Toronto. [page 3], [page 6], [page 8]

McCallum, Andrew K., 1996, Bow: A toolkit for statistical language modeling, text retrieval, classification and clustering, `http://www.cs.cmu.edu/~mccallum/bow`. [page 6]

Moskowitz, Tobias J., 2000, Discussion, *The Journal of Finance* 55, 1695–1703. [page 4], [page 22]

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 703–708.

Nimark, Kristoffer P, and Stefan Pitschner, 2019, News media and delegated information choice, *Journal of Economic Theory* 181, 160–196. [page 1], [page 4], [page 6], [page 8], [page 25]

Peress, Joel, 2008, Media coverage and investors' attention to earnings announcements, working paper. [page 20]

Peress, Joel, 2014, The media and the diffusion of information in financial markets: Evidence from newspaper strikes, *The Journal of Finance* 69, 2007–2043. [page 3]

Pew Research Center, 2011, The Wall Street Journal under Rupert Murdoch, Washington, D.C. [page 5]

Schmalz, Martin C, and Sergey Zhuk, 2019, Revealing downturns, *The Review of Financial Studies* 32, 338–373. [page 2]

Sicherman, Nachum, George Loewenstein, Duane J. Seppi, and Stephen P. Utkus, 2016, Financial attention, *Review of Financial Studies* 29, 863–897. [page 15]

Solomon, David, and Eugene Soltes, 2012, Managerial control of business press coverage, working paper. [page 2], [page 8]

Solomon, David H., 2012, Selective publicity and stock prices, *The Journal of Finance* 67, 599–638. [page 2], [page 7], [page 8]

Stock, James H., and Mark W. Watson, 2003, Forecasting output and inflation: The role of asset prices, *Journal of Economic Literature* 41, 788–829. [page 16]

Tetlock, Paul C., 2007, Giving content to investor sentiment: The role of media in the stock market, *Journal of Finance* 62 (3), 1139–1168. [page 2], [page 7]

Tetlock, Paul C., Maytal Saar-Tsechansky, and Sofus Macskassy, 2008, More than words: Quantifying language to measure firms' fundamentals, *Journal of Finance* 63 (3), 1437–1467. [page 2], [page 7]

Welch, Ivo, and Amit Goyal, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21 (4), 1455–1508. [page 16]

Yuan, Yu, 2015, Market-wide attention, trading, and stock returns, *Journal of Financial Economics* 116, 548–564. [page 15]

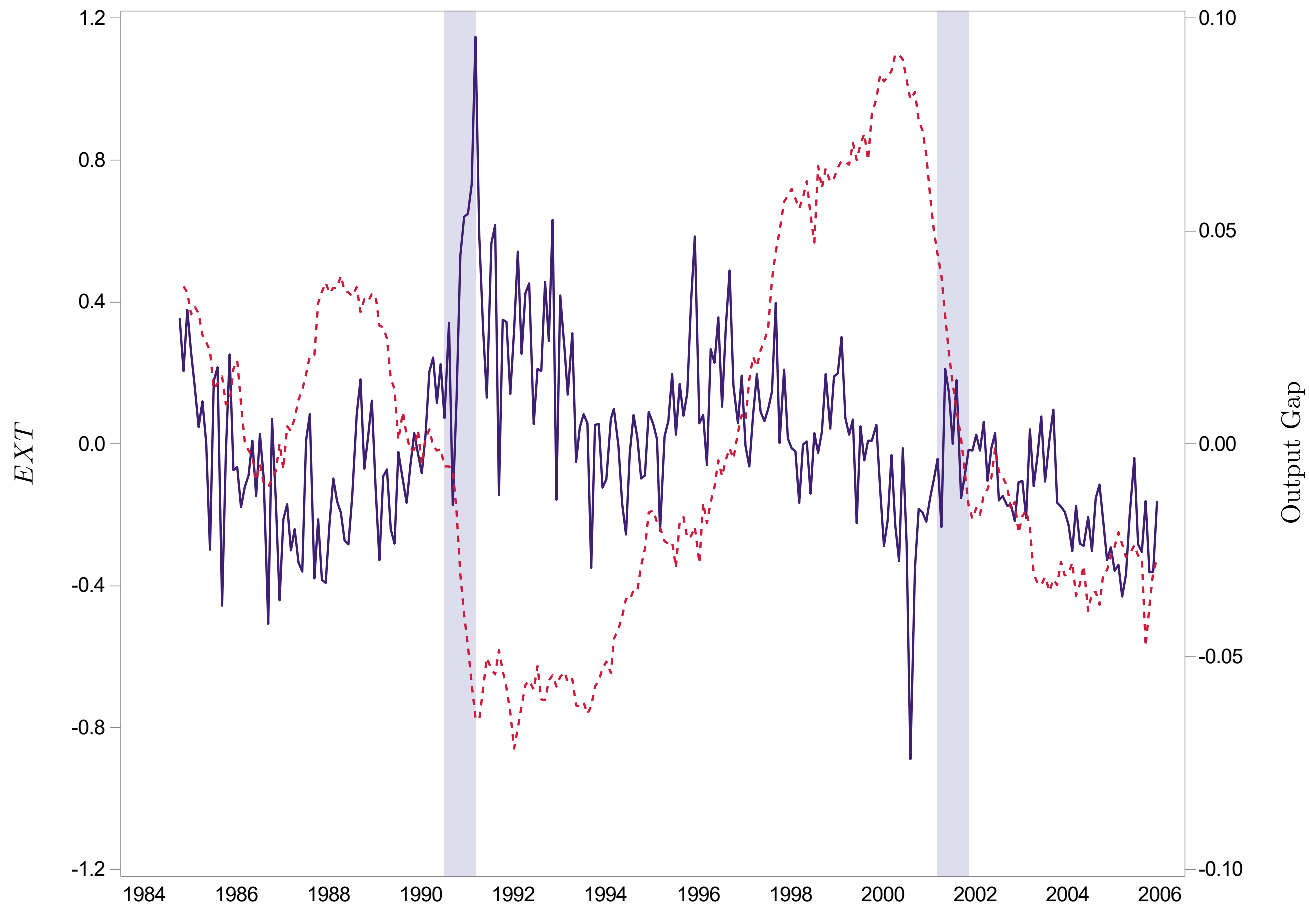

Figure 1: Extent of Coverage Covaries Negatively With Output Gap. The monthly EXT series (solid line) is plotted against the monthly output gap series (dotted line). EXT is orthogonalized with respect to tone (TONE). Output gap is the residual from a linear and quadratic time trend in the natural logarithm of industrial production. The shaded regions are NBER recession periods.

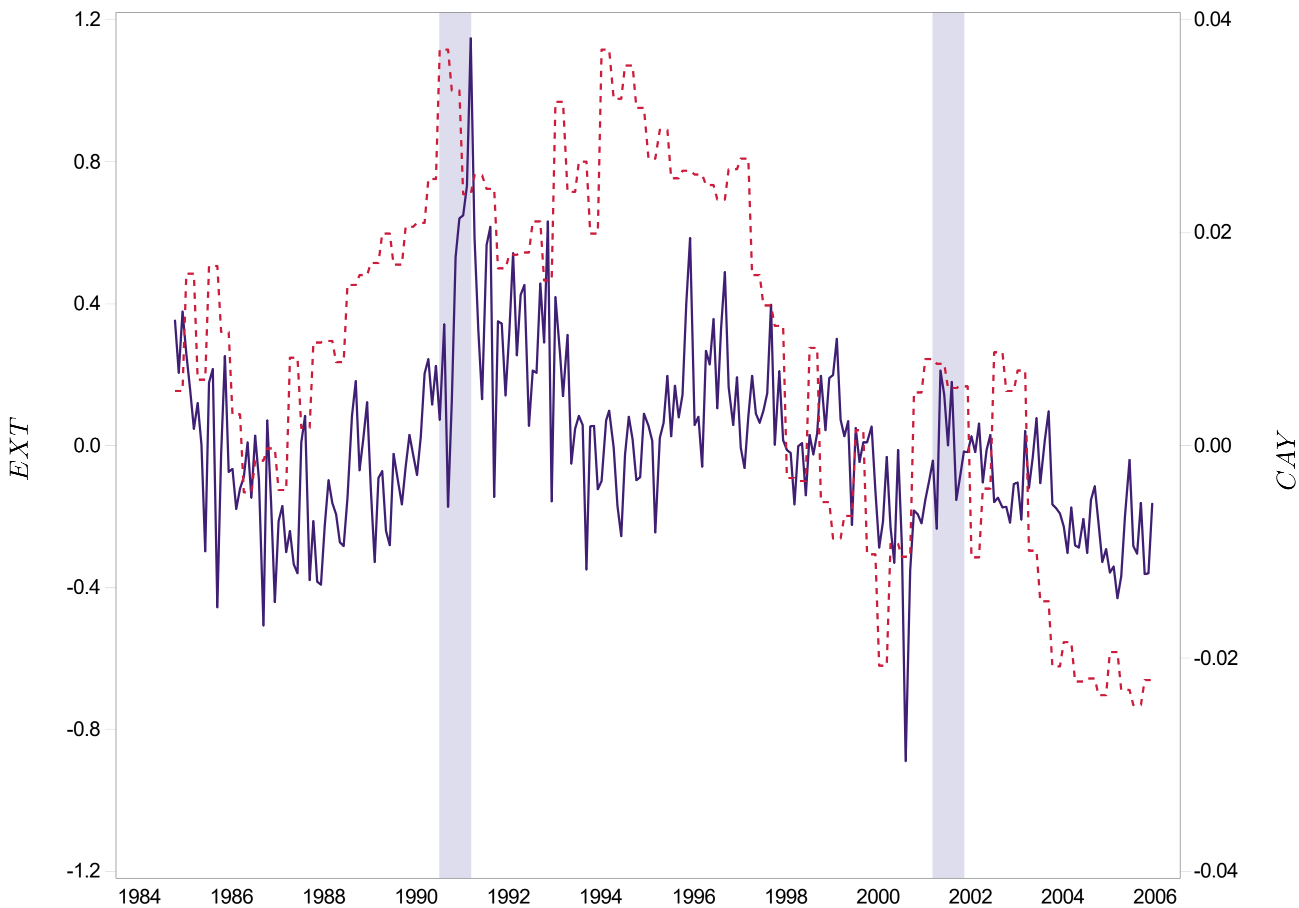

Figure 2: Extent of Coverage Covaries With CAY. The monthly EXT series (solid line) is plotted against the quarterly CAY series (dotted line). EXT is orthogonalized with respect to tone (TONE). CAY is the consumption-wealth ratio of Lettau and Ludvigson (2001). The shaded regions are NBER recession periods.

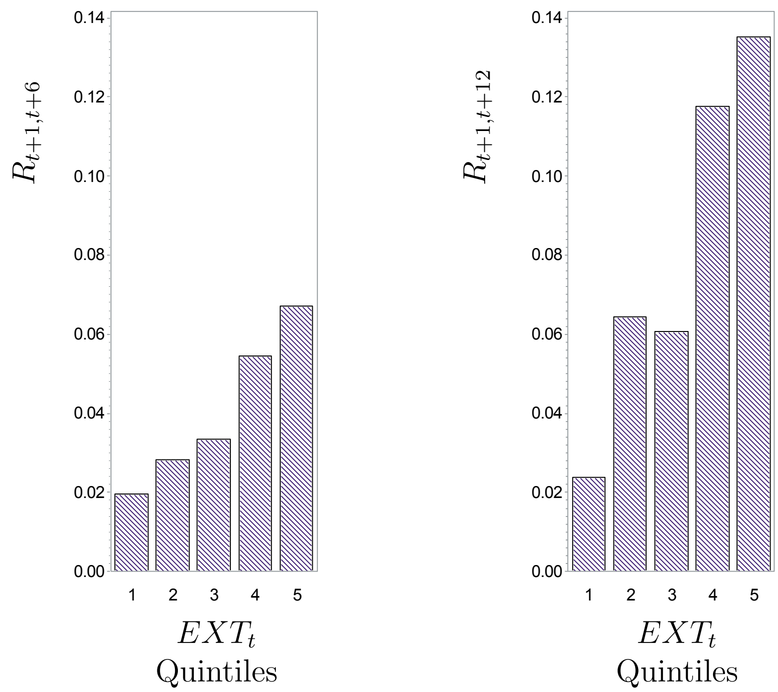

Figure 3: Extent of Coverage and Equity Risk Premium. Months are sorted into quintiles based on EXT, orthogonalized with respect to tone (TONE). Means of excess stock returns within each quintile are plotted for months (+1,+6) in the left panel and for months (+1,+12) in the right panel.

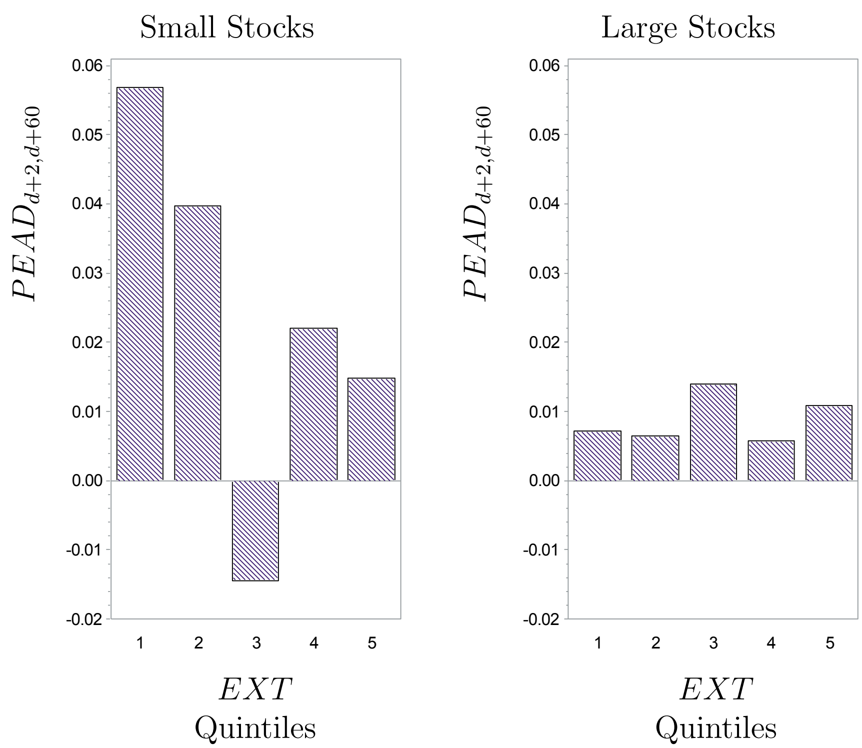

Figure 4: Extent of Coverage and Post-Earnings Announcement Drift in Returns. Months are sorted into quintiles based on EXT, orthogonalized with respect to tone (TONE). Means of PEAD profits within small stocks for days (+2,+60) are plotted in the left panel ($\leq 20^{th}$ percentile using NYSE breakpoints) and within large stocks in the right panel ($> 80^{th}$ percentile using NYSE breakpoints).

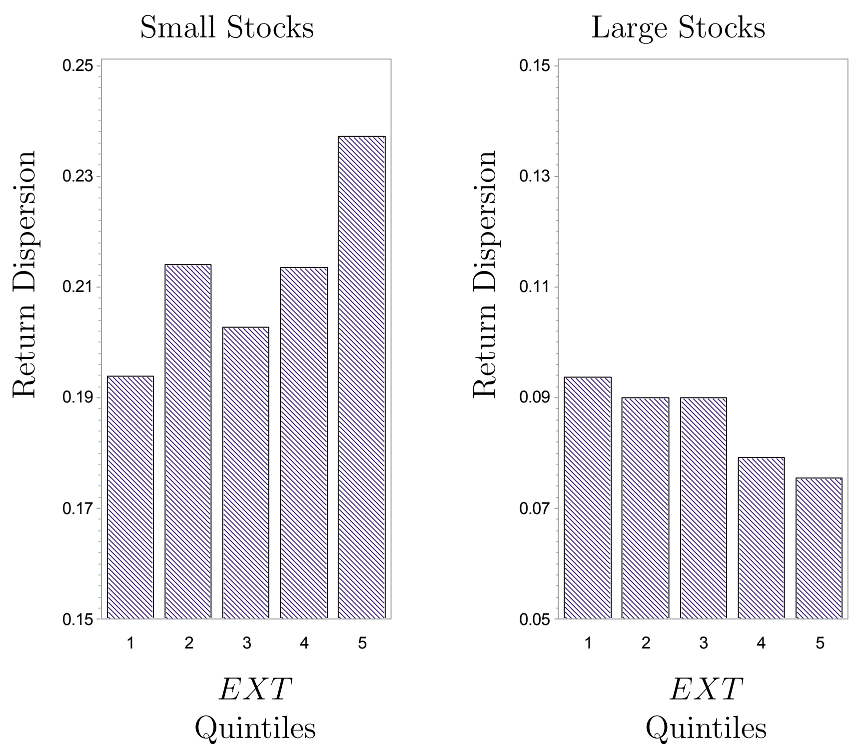

Figure 5: Extent of Coverage and Return Dispersion. Months are sorted into quintiles based on EXT, orthogonalized with respect to tone (TONE). Means of monthly cross-sectional standard deviation of returns within small stocks are plotted in the left panel ($\leq 20^{th}$ percentile of market cap using NYSE breakpoints) and within large stocks in the right panel ($> 80^{th}$ percentile of market cap using NYSE breakpoints). Different axis scales are used for small and large stocks.

Table 1
Summary Statistics

Below are summary statistics for various monthly measures of the WSJ's coverage of earnings and of macroeconomic measures from October 1984 to December 2005. The expected number of nonnegative articles and of negative articles are estimated with a multinomial logit (see section 1.3). EXT is the deviation of the actual number of firm-earnings covered in a given month from the expected number to be covered, while TONE is the deviation of the actual net tone of the coverage from the expected net tone (see equations 1 and 2). Equity premium is the monthly return on the CRSP value-weighted index minus the one-month T-Bill return in percentage terms. The macroeconomic measures in the lower portion are defined in Appendix A.1.

	Mean	Std. Dev.	AR(1)
Number of Earnings Released	915.09	706.12	−0.14
Number of Earnings Covered	96.60	79.56	−0.17
Number of Nonnegative Articles	66.39	59.99	−0.18
Number of Negative Articles	30.22	24.09	0.00
E[Number of Nonnegative Articles]	66.39	60.34	−0.16
E[Number of Negative Articles]	30.22	22.51	−0.01
Extent of Earnings Coverage (EXT)	0.03	0.26	0.62
TONE	0.00	0.19	0.43
Equity Premium	0.45	4.38	0.04
PDND	−0.12	0.11	0.93
Dividend Yield	0.02	0.01	0.99
Net Payout Yield	0.10	0.02	0.99
Risk-free rate	0.05	0.02	0.98
CAY	0.01	0.02	0.98
B/M	0.33	0.15	0.98
Default Spread	0.01	0.00	0.95
Term Spread	1.76	1.15	0.97
Equity Share of New Issues	0.12	0.06	0.61
Sentiment	0.10	0.59	0.94
Output Gap	0.00	0.04	0.99

Table 2

Extent of Earnings Coverage within Size Quintiles

Months from October 1984 to December 2005 are ranked based on EXT into the lowest 25%, highest 25%, and remaining "normal" months. Within each size quintile, we calculate the extent of coverage in a given month, from smallest stocks (quintile 1) to largest stocks (quintile 5). The full-sample logit specification is used to determine the likelihood of coverage. Panel A reports the means of the extent of coverage for the months with high EXT, normal EXT, and low EXT. The means of the actual number of earnings reports covered are shown Panel B.

EXT	(Smallest) 1	2	3	4	(Largest) 5
Panel A: Within-Quintile Extent of Coverage					
Low	−0.37	−0.48	−0.41	−0.25	−0.19
Normal	0.01	−0.07	0.04	0.06	0.02
High	0.91	0.62	0.41	0.32	0.18
Panel B: Actual Number of Firms Covered per Month					
Low	4.4	5.2	8.5	18.3	44.4
Normal	7.6	9.0	13.2	22.6	46.9
High	10.0	12.4	16.0	23.1	45.3

Table 3

Extent of Earnings Coverage and Macro Conditions

EXT is regressed on a set of macroeconomic variables, which are defined in section A.1.2 of the Appendix. The t-statistics (in parentheses) are calculated using Newey-West standard errors. The p-values shown in brackets are determined using 10,000 randomly generated samples of independent normally distributed variables with first-order serial correlation coefficients equal to that of EXT. Asterisks correspond to the simulated p-values, with * indicating significance at 10%,** indicating significance at 5%, and *** indicating significance at 1%.

PDND	0.004	Default Spread	3.709
	(1.80)		(0.37)
	[0.21]		[0.80]
log(Div. Yield)	−0.977***	Term Spread	0.064
	(−4.96)		(1.98)
	[0.00]		[0.15]
log(Net Payout Yield)	−0.052	Equity Share of New Issues	0.435
	(−0.22)		(1.57)
	[0.87]		[0.24]
RF	11.037**	Sentiment	−0.112
	(3.12)		(−1.92)
	[0.02]		[0.21]
CAY	7.523***	Output Gap	−2.959**
	(4.56)		(−2.92)
	[0.00]		[0.04]
B/M	0.927	Recession	0.164
	(2.14)		(1.92)
	[0.16]		[0.26]
Constant	−4.657		
	(−4.85)		
Observations	255		
R^2	0.41		

Table 4

Extent of Earnings Coverage and the Equity Premium

The dependent variable is the log return of the CRSP value-weighted index minus the one-month T-Bill rate across the future six and twelve months respectively. EXT_t and $TONE_t$ are the month t measures of the extent of coverage and of the tone of coverage, respectively, defined in equations (1) and (2). Both of these explanatory variables are standardized. The row labeled "Macro Controls" indicates whether or not the set of macroeconomic variables (shown in section A.1.2) are included as regressors. The t-statistics (in parentheses) are calculated using Newey-West standard errors. The p-values shown in brackets are determined using 10,000 randomly generated samples of two independent normally distributed variables with first-order serial correlation coefficients respectively equal to that of EXT and TONE. Asterisks correspond to simulated p-values, with * indicating significance at 10%,** significance at 5%, and *** significance at 1%. Panel B reports the simulated p-values of the joint hypothesis that the coefficients on EXT in Panel A are both zero across the six-month and twelve-month horizons.

A. Regressions of Excess Stock Returns

	(+1,+6)	(+1,+12)	(+1,+6)	(+1,+12)
EXT	0.0189*	0.0366**	0.0214*	0.0208
	(1.97)	(2.51)	(2.36)	(1.63)
	[0.097]	[0.047]	[0.078]	[0.234]
TONE	0.0169	0.0109	0.0092	0.0066
	(1.62)	(0.59)	(1.03)	(0.78)
	[0.168]	[0.608]	[0.412]	[0.550]
Macro Controls	No	No	Yes	Yes
Constant	0.035	0.069	0.731	−0.489
	(2.84)	(3.43)	(1.21)	(−0.80)
Observations	255	255	255	255
R^2	0.08	0.08	0.33	0.50

B. Joint-Horizon Tests of Coefficients on EXT

Horizons	Simulated p-value of joint significance
$(+1, +6) \cap (+1, +12)$	0.029
$(+1, +6) \cap (+1, +12)$ with Macro Controls	0.048

Table 5

Post-Earnings Announcement Drift of Small Stocks

The left-panel dependent variable is the time-series of monthly coefficients from a cross-sectional regression of small-stock cumulative abnormal returns from day +2 to day +60 (PEAD) on standardized unexpected earnings (SUE). Regressions of the monthly first-stage coefficients on EXT and TONE measured in the month of the announcement are reported below. Small stocks are defined as those with a percentile of market capitalization less than or equal to the 20^{th} percentile of NYSE stocks. The row labeled "Macro Controls" indicates whether or not the set of macroeconomic variables listed in section A.1.2 are included as regressors. The t-statistics (in parentheses) are computed using Newey-West standard errors. The p-values shown in brackets are determined using 10,000 randomly generated samples of two independent normally distributed variables with first-order serial correlation coefficients matching those of EXT and TONE respectively. Asterisks correspond to the simulated p-values; * indicates significance at 10%,** significance at 5%, and *** significance at 1%.

	PEAD	
EXT	-0.140^{***}	-0.184^{***}
	(-2.96)	(-3.68)
	[0.007]	[0.001]
TONE	0.050	0.075
	(1.33)	(1.52)
	[0.204]	[0.163]
Macro controls	No	Yes
Constant	0.173	-1.038
	(3.95)	(-0.36)
Observations	255	255
R^2	0.04	0.13

Table 6
Return Dispersion of Small Stocks

Monthly standard deviation of the cross section of small-stock returns is regressed on EXT, TONE, and macroeconomic controls. Small is defined to be less than or equal to the 20^{th} percentile of market capitalization using NYSE breakpoints in the prior month. The row labeled "Macro Controls" indicates whether or not the set of macroeconomic variables listed in section A.1.2 are included as regressors. The t-statistics (in parentheses) are computed using Newey-West standard errors. The p-values shown in brackets are determined using 10,000 randomly generated samples of two independent normally distributed variables with first-order serial correlation coefficients matching EXT and TONE respectively. Asterisks correspond to simulated p-values. * indicates significance at 10%,** significance at 5%, and *** significance at 1%.

	Return Dispersion	
EXT	0.016***	0.019***
	(2.76)	(4.35)
	[0.007]	[0.000]
TONE	−0.008	−0.006
	(−1.72)	(−1.71)
	[0.113]	[0.137]
Macro Controls	No	Yes
Constant	0.212	−0.021
	(37.66)	(−0.11)
Observations	255	255
R^2	0.07	0.24

Table 7

Trade Performance of Mutual Funds within Small-Stock Universe

Each month, the cross section of quarterly small-stock abnormal returns is regressed on lagged quarterly changes in the number of mutual funds that hold a given stock. The first-stage coefficients are then regressed on EXT and TONE measured over the same horizon as the changes in fund holdings. The t-statistics (in parentheses) are computed using Newey-West standard errors. The p-values shown in brackets are determined using 10,000 randomly generated samples of two independent normally distributed variables with first-order serial correlation coefficients matching EXT and TONE respectively. Asterisks correspond to simulated p-values; * indicates significance at 10%,** significance at 5%, and *** significance at 1%.

	Trading Performance	
EXT	0.002**	0.002*
	(3.35)	(2.64)
	[0.05]	[0.07]
TONE	0.000	−0.001
	(0.93)	(−1.21)
	[0.64]	[0.41]
Macro controls	No	Yes
Constant	−0.001	0.057
	(−0.84)	(1.66)
Observations	252	252
R^2	0.05	0.12