MEDIA COVERAGE AND INVESTOR ATTENTION

by

Charles Clyde Gaa

B.A., Queen's University, 1996
M.A., The University of Alberta, 1999

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate Studies

(Business Administration)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

November 2008

Abstract

In this thesis, I investigate the role of investor attention in financial markets by examining the media's coverage of corporate earnings news. The first paper studies the potential impact of information in the financial press by identifying systematic differences between aggregate corporate earnings news coverage in the Financial Times, Wall Street Journal, and the New York Times, and measures of expected coverage based on contemporaneous earnings information flows as reported in I/B/E/S. I find that publication-specific estimates of "excess" aggregate positive or negative coverage exhibit strong serial correlation, consistent with media bias. Furthermore, unexplained negative (positive) weekly coverage predicts positive (negative) returns for small-stock indices and the equal-weighted NYSE, suggesting that the effects of predictability in financial news coverage are economically significant and may be related to informational inefficiency with respect to smaller firms.

The second paper examines media coverage decisions to identify the determinants of investor attention with respect to events and firms. Using ex ante predicted probability of media coverage (PMC) with respect to earnings news as a measure of attention in this context, I study the returns experienced by low-attention stocks from 1984 and 2005. As in prior studies, I find high risk-adjusted returns for "neglected" stocks, which appears to be highly consistent with, e.g., Merton's (1987) investor recognition hypothesis, or an information risk setting (Easley et al. (2002)). However, in examining the event-specific determinants of media coverage, I find evidence of a significant "negativity bias" in attention: holding other factors constant, bad news is more likely to attract coverage than is good news regarding an otherwise-identical firm. Given recent evidence in the literature regarding stock-price underreaction to low-attention events, this suggests asymmetric investor attention as a potential explanation for an apparent neglected firm premium in the cross-section of stock returns. Consistent with this hypothesis, I find that the excess returns to low-PMC portfolios are attributable to drift in the stock prices of low-attention "good news" firms, while low-attention "bad news" firms appear to be efficiently priced.

Table of Contents

List of Tables

List of Figures

Acknowledgements

I am grateful to so many people for their support as I pursued my doctoral studies. First and foremost, I would like to thank Kai Li, my committee chair, and the other members of my committee, Werner Antweiler, Adlai Fisher, and Marcin Kacperczyk, for their ongoing support and valuable feedback. This research simply would not have been possible without their generous advice and guidance. I am also very grateful to Murray Carlson, Jason Chen, Glen Donaldson, Alan Kraus, Maurice Levi, Hernan Ortiz-Molina, as well as many other UBC faculty members, fellow students, and seminar participants for their helpful comments and suggestions.

For Sherry

CHAPTER I

Introduction

Individual investors are simply unable to acquire and process all of the potentially price-relevant information that is available to them at any given point in time. In light of this fundamental limitation, financial market participants rely upon information intermediaries (such as the financial news media, among others) to help them in identifying which information events warrant closer attention, and, further, to assist in their interpretation. In this thesis, I undertake an empirical investigation of financial news media coverage in order to examine the role of investor attention in explaining stock returns. The first essay identifies persistent and predictable patterns in aggregate news coverage behavior over time, and then proceeds to study the potential impact of abnormally positive or negative media content on stock index returns. The second essay utilizes estimates of predicted financial news media coverage to create a novel measure of investor attention with respect to firms and events; in this context, asymmetric attention allocation (i.e., negativity bias) and underreaction emerge as a potential explanation for the observed neglected firm premium in the cross-section of returns.

Under even a fairly liberal definition of market efficiency, all that is required for prices to adjust to their "fundamental" values is the participation of one or more sufficiently-large market participants who are well-informed and free of cognitive bias. However, if limits to arbitrage prevent prices from quickly adjusting to eliminate mispricing, then there is scope for a group of homogenously ill-informed investors to affect market prices, at least in the short term (see, e.g., Delong et al. (1990)). But how might a group of investors come to hold mistaken beliefs in such a correlated fashion? In recent years, researchers have begun to focus on investors' constrained cognitive abilities and the resulting implications for learning and beliefs.

Peng and Xiong (2006) show that limited attention can lead investors to make predictable pricing errors through the application of simplified rules for information processing. In their model, when a representative investor is unable to process all of the information related to firms in a given sector, more attention is allocated to sector-level data; as a result, securities are effectively categorized into asset classes and portfolio allocation decisions are made on this (potentially error-prone) basis. Huang and Liu (2007) examine the role of costly information acquisition; they show that market participants may optimally choose to be less well-informed, potentially leading to predictability in asset returns as agents set prices in response to noisy signals.

More broadly, papers in the behavioral finance literature typically point to cognitive constraints as an implicit, underlying cause for biased decision-making and beliefs. For example, Barberis et al. (1998) show that pervasive information processing errors such as overconfidence and the representativeness heuristic (which we can see as natural consequences of constrained cognition, since agents presumably would not make such mistakes if they knew better) can lead to both over- and under-reaction in asset prices. Similarly, Daniel et al. (1998) point to biased self-attribution and overconfidence about the precision of private signals as potential explanations for such return anomalies.

While there continues to be significant debate in the literature regarding the existence and interpretation of various potential examples of "inefficiency" in asset prices, one of the most enduring and widely-cited examples of apparent stock-price underreaction is that of post-earnings-announcement-drift (PEAD). Ball and Brown (1968) were among the first to document that stock prices appear to "drift" following an earnings surprise, a phenomenon that has been borne out time and again by researchers over the following decades. Rather than representing a risk premium of some kind, PEAD is most commonly interpreted as a delayed price response to the information contained in earnings announcements (e.g., Bernard and Thomas (1990)). In particular, recent research has found PEAD to be most pronounced for announcements and firms where we would expect attention to be low and/or distraction to be high (Hou et al. (2006), Hirshleifer et al. (2006), and DellaVigna and Pollet (2008)).

There are numerous empirical papers linking measures of investor attention and information quality to the market's reaction to news. For example, Barber and Odean (2008) show that individual investors are more likely to purchase stocks that attract attention through noteworthy news coverage, price changes, or trading activity.[1] This thesis contributes to the

[1] See also, e.g., Brennan et al. (1993) and Hong et al. (2000) for evidence with respect to analyst coverage and slow information diffusion.

emerging literature on the role of attention in the market's reaction to news by focusing on both the determinants and potential impacts of media coverage with respect to earnings announcements.

If investors' realized attention allocations can affect their beliefs (and thereby potentially affect market prices, at least in the short term), then we might expect the information presented in high-profile sources such as the financial news media to be particularly significant in this respect. Indeed, there are numerous studies suggesting that media coverage can have significant effects on market activity and prices.[2] Looking at one particularly striking example, Huberman and Regev (2001) describe the market's reaction to a high-profile *re*-release of public information, suggesting that a broader dissemination of pre-existing information can be decisive in and of itself. They find that the biotech firm Entremed's stock price was significantly and permanently affected by the publication of a high-profile news article that contained no information that was not already in the public sphere. Looking at a much wider set of potential information events, Busse and Green (2002) show that the market reacts in real-time to firm/CEO media coverage on CNBC.

Given the strong suggestion of a link between media coverage and the market's reactions to news events, we are left with an intriguing question: How are coverage decisions actually made? Given that there isn't sufficient space and time for all of the potentially interesting stories that occur on any given day to be covered in the media, which stories tend to be chosen for publication? If coverage decisions are discretionary, then this leaves the door open for potential media bias. In particular, if certain classes of information events are consistently observed to be relatively overemphasized or underemphasized in the financial press (and, by implication, in the minds of at least some investors), then we might expect this to result in correlated pricing errors, as mentioned earlier.

Mullainathan and Shleifer (2005) present a model wherein the media caters to the expected news preferences of their readers, resulting in demand-driven bias. They show that such bias need not be eliminated in the presence of competition in the market for news, and that outlet-specific bias may in fact be exacerbated with competition as media outlets segment the market in their attempts to appeal to disparate groups of consumers. One commonly-identified

---

[2] E.g., Busse and Green, (2002), Chan (2003), Dyck and Zingales (2003), Bhattacharya et al. (2006), Tetlock (2007), Tetlock et al. (2008), Antweiler and Frank (2006), Fang and Peress (2007), and Barber and Odean (2008).

expression of coverage bias is referred to as "sensationalism": news reporters tend to seek out (and, by implication, news consumers are most often exposed to) stories that are perceived to exhibit particularly salient narrative qualities. For example, it is often observed that stories about grisly crimes and celebrity scandals seem to dominate the news headlines, despite their objectively low probabilities of occurrence and/or essential triviality from the perspective of the average reader.

While polling indicates that the majority of Americans believe that the media is biased at least some of the time (ASNE (1999)), we might suspect that such views are highly subjective. However, there is some empirical evidence that patterns of media coverage are subject to systematic bias/predictability. Groseclose and Milyo (2005) construct a measure of political bias by examining the presumed political affiliations of think-tanks cited by reporters in relation to the expected citation behavior of a hypothetical, middle-of-the-road politician; they find that media outlets differ significantly according to their "left-right" scale. Looking at the potential for supply-driven bias, Reuter and Zitzewitz (2006) find that mutual funds' advertising expenditures are linked to their performance in trade-publication industry rankings. Furthermore, in attempting to examine regularities in coverage with respect to positive and negative information flows, researchers have presented empirical evidence to suggest that the media is asymmetric in its coverage of positive and negative macroeconomic news, consistent with a negativity bias (Harrington (1989) and Soroka (2006)).

The first essay in this thesis explores the potential for bias in aggregate coverage relative to the information contained in an underlying set of contemporaneous news announcements. Applying computational linguistic tools to identify the topic and tone of news articles, I attempt to explain weekly variation in the media's coverage of earnings news with reference to the associated flow of earnings announcements made by firms. I find that weekly measures of unexplained positive and negative earnings news coverage are strongly serially correlated, suggesting that news media outlets are significantly more likely to focus on good news during certain periods, and on bad news during others. These fluctuations in coverage behavior appear to correspond broadly with stylized facts regarding changes in market sentiment and economic conditions over the sample period.[3] Furthermore, I find that estimates of unexplained news

[3] E.g., Figures 2.7, 2.8, and 2.9 show that estimates of unexplained media coverage are relatively negative in the early 1990s (potentially corresponding to the 1990-1991 U.S. recession and its aftermath), become more positive in the latter half of the 1990s (as the U.S. economy and stock markets were booming), and

coverage contain information about future returns for portfolios of smaller stocks: current media coverage that is more positive than expected predicts negative stock returns in subsequent weeks, and *vice versa.*

In addition to the hypothesized role of investor attention in mediating the responsiveness and speed of the market's reaction to news events, generally low levels of attention/recognition with respect to a firm (i.e., neglect) have long been linked empirically to high risk-adjusted returns.[4] Traditionally, researchers have explained this "neglected firm effect" by hypothesizing that such firms are subject to costly information frictions (Merton (1987)), or are simply riskier is some sense that is not completely captured by traditional asset pricing models. For example, firms with relatively poor information environments may be subject to heightened asymmetric information, resulting in higher expected returns for these stocks (Easley et al. (2002)).

In the second essay, I identify the event- and firm-specific determinants of positive and negative media coverage regarding almost 180,000 quarterly corporate earnings announcements from 1984 to 2005 in a multinomial logit setting. I utilize estimates of coverage to construct a novel measure of investor attention in the cross-section: probability of media coverage (PMC). Subsequently, I investigate potential linkages between biased attention allocation, stock return predictability, and asymmetry in the market's reactions to positive and negative corporate information flows. Focusing on the event-specific determinants of media coverage, I find that, holding other factors constant, there is a "negativity bias" in news coverage: bad news is more likely to attract attention than is good news.

This second paper contributes to the literature by bringing together two apparently disconnected sets of findings regarding investor attention: 1) low-attention events tend to elicit delayed price responses, and 2) low-attention stocks tend to earn high risk-adjusted returns. Identifying "neglected" firms as those least likely to attract media attention, my findings suggest that higher-than-expected returns for these stocks are attributable to drift in the prices of neglected firms that have recently experienced positive earnings news, while neglected firms with recent negative news appear to be efficiently priced. In other words, asymmetric inattention emerges as a potential explanation for the neglected firm premium in this context.

---

then seem to turn sharply negative once again in 2001 (contemporaneous with the bursting of the "tech bubble" and a general period of slowing in the U.S. economy).

[4] E.g., see Banz (1981), Arbel and Strebel (1982), Foerster and Karolyi (1999), Hou and Moskowitz (2005), Fang and Peress (2007).

References

American Society of Newspaper Editors, 1999, Examining our credibility: Perspectives of the public and the press, Reston, VA., www.asne.org.

Antweiler, Werner, and Murray Z. Frank, 2006, Do U.S. Stock Markets Typically Overreact to Corporate News Stories, Working paper, University of British Columbia and University of Minnesota.

Arbel, Avner, and Paul Strebel, 1982, The Neglected and Small Firm Effects, *Financial Review* 17, 201-218.

Ball, Ray, and Philip Brown, 1968, An Empirical Evaluation of Accounting Income Numbers, *Journal of Accounting Research* 6 (2), 159-178.

Banz, R.W., 1981, The relationship between return and market value of common stocks, *Journal of Financial Economics* 9, 3–18.

Barber, Brad M., and Terrance Odean, 2008, All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors, *Review of Financial Studies* 21(2), 785-818.

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307-343.

Bernard, Victor L., and Jacob K. Thomas, 1990, Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium, *Journal of Accounting Research* 27, 1-36.

Brennan, Michael J., Narasimhan Jegadeesh, and Bhaskaran Swaminathan, 1993, Investment Analysis and the Adjustment of Stock Prices to Common Information, *Review of Financial Studies* 6(4), 799-824.

Bhattacharya , Utpal, Neal Galpin, Rina Ray, and Xiaoyun Yu, 2006, The Role of the Media in the Internet IPO Bubble, *Journal of Financial and Quantitative Analysis*, forthcoming.

Busse, Jeffrey A., and T. Clifton Green, 2002, Market efficiency in real time, *Journal of Financial Economics* 65, 415-437.

Chan, Wesley S., 2003, Stock price reaction to news and no-news: drift and reversal after headlines, *Journal of Financial Economics* 70, 223-260.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor Psychology and Security Market under- and Overreactions, *Journal of Finance* 53(6), 1839-1885.

DellaVigna, Stefano, and Joshua Pollet, 2008, Investor Inattention and Friday Earnings Announcements, *Journal of Finance,* forthcoming.

DeLong, J. Bradford, Andrei Shleifer, Lawrence Summers, and Robert Waldmann, 1990, Positive feedback investment strategies and destabilizing rational speculation, *Journal of Finance* 45, 379-395.

Dyck, Alexander, and Luigi Zingales, 2003, The media and asset prices, Working paper, Harvard Business School and University of Chicago.

Easley, David, Soeren Hvidkjaer, and Maureen O'Hara, 2002, Is information risk a determinant of asset returns?, *Journal of Finance* 57(5), 2185-2221.

Fang, Lily, and Joel Peress, 2007, Media Coverage and the Cross-Section of Returns, Working paper, INSEAD.

Foerster, Stephen R., and G. Andrew Karolyi, 1999, The Effects of Market Segmentation and Investor Recognition on Asset Prices: Evidence from Foreign Stocks Listing in the United States, *Journal of Finance* 54 (3), 981-1013.

Groseclose, Tim, and Jeffrey Milyo, 2005, A Measure of Media Bias, *Quarterly Journal of Economics* 120(4), 1191-1237.

Harrington, David E., 1989, Economic news on television: The determinants of coverage, *Public Opinion Quarterly* 53, 17-40.

Hirshleifer, David, Sonya S. Lim, and Siew H. Teoh, 2006, Driven to distraction: Extraneous events and underreaction to earnings news, Working paper, University of California at Irvine, and Depaul University.

Hong, Harrison, Terence Lim, and Jeremy C. Stein, 2000, Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies, *Journal of Finance* 55 (1), 265-295.

Hou, Kewei, and Tobias J . Moskowitz, 2005, Market Frictions, Price Delay, and the Cross-Section of Expected Returns, *Review of Financial Studies* 18 (3), 981-1020.

Hou, Kewei, Lin Peng, and Wei Xiong, 2006, A tale of two anomalies: The implications of investor attention for price and earnings momentum, Working paper, Ohio State University, Baruch College, and Princeton University.

Huang, Lixin, and Hong Liu, 2007, Rational Inattention and Portfolio Selection, *Journal of Finance* 62 (4), 1999-2040.

Huberman, Gur, and Tomer Regev, 2001, Contagious speculation and a cure for cancer: A nonevent that made stock prices soar, *Journal of Finance* 56, 387-396.

Merton, Robert C., 1987, A Simple Model of Capital Market Equilibrium with Incomplete Information, *Journal of Finance* 42 (3), 483-510.

Mullainathan, Sendhil, and Andrei Shleifer, 2005, The market for news, *American Economic Review* 95 (4), 1031-1053.

Peng, Lin, and Wei Xiong, 2006, Investor attention, overconfidence, and categorical learning, *Journal of Financial Economics* 80(3), 563-602.

Reuter, Jonathan, and Eric Zitzewitz, 2006, Do ads influence editors? Advertising and bias in the financial media, *Quarterly Journal of Economics* 121(1), 197-227.

Soroka, Stuart N., 2006, Good news and bad news: Asymmetric responses to economic information, *Journal of Politics* 68 (2), 372-385.

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

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.

CHAPTER II

The Impact of Predictability in Financial News Media Coverage[5]

2.1 Introduction

Do financial reporters' selective coverage decisions result in a biased aggregate picture of corporate news events? If so, does it matter? Or are investors effectively able to "see through" any apparent distortions in media coverage? While it seems natural to think of media bias as a tool applied in furtherance of a political agenda (Groseclose and Milyo (2005) and DellaVigna and Kaplan (2007)), or an advertising objective (Reuter and Zitzewitz (2006)), the potential for systematic biases in corporate news coverage has important implications for our understanding of public information in markets. Mainstream financial news media outlets such as the Wall Street Journal and the Financial Times represent a significant source of information for millions of investors – if media coverage is seen as informative, and if the information presented is systematically skewed in some sense, then we might well expect investors' beliefs, and therefore price formation, to be significantly affected.

In this paper, I apply computational linguistic tools to develop measures of tone in media coverage regarding corporate earnings releases in the Wall Street Journal (WSJ), the Financial Times (FT), and the New York Times (NYT). In this context, media bias is identified as persistent or predictable deviations in observed coverage relative to an empirical model of expected coverage behavior. I find evidence that, while contemporaneous earnings information releases are able to explain a significant proportion of the variation in associated media output measures, there exist predictable deviations in the media information environment surrounding

[5]A version of this chapter will be submitted for publication. Gaa, C., The Impact of Predictability in Financial News Media Coverage.

this very specific type of corporate news event. The results are consistent with an interpretation of temporally-varying, demand-driven media bias with respect to positive and negative news stories, as described, for example, by Mullainathan and Shleifer (2005). I find that the estimates of unexplained coverage Granger-cause stock market index returns, in particular with respect to the securities of smaller NYSE-listed firms.

As an empirical investigation of demand-driven media bias and its potential impact upon financial markets, this paper contributes both to the growing economics literature regarding media bias, and also to the behavioral strand of research in financial economics that focuses on investors' cognitive constraints as an explanation for apparent informational inefficiency in security prices. In contrast to earlier studies exploring the media's coverage of news events, I explicitly examine the *difference* between "fundamental" information flows and contemporaneous media accounts, thereby allowing me to identify potential distortions (in a relative sense) in one information channel as compared to another.

The choice of quarterly earnings announcements as the underlying information event is crucial to the analysis, not only because such announcements are typically the single most reported-on event with respect to any given firm (Tetlock et al. (2008)), but also because these are events whose timings and underlying probabilities of occurrence are essentially exogenous with respect to the behavior of the information intermediary. NYSE firms, for example, are required by their listing agreements to publicly announce their financial results once per quarter. Regardless of whether the quarterly earnings information is seen as positive or negative for the firm's prospects, an announcement must be made when the pre-appointed day arrives. Therefore, the fundamental flow of accounting information that underlies the production of this type of news article must simply be taken as given from the perspective of the financial reporter. However, reporters tasked with writing such stories cannot report on all of the earnings announcements that are made on a given day; there is simply not enough space in the newspaper to cover them all. More to the point, the majority of these information releases will not be seen as "interesting enough" to warrant discussion in a major news outlet such as the Financial Times, the New York Times, or the Wall Street Journal.

This leads us to an important question: When reporters decide which earnings release stories are worth writing about on a given day, which ones do they choose? As a group, do they tend to select a representative sample (or even a consistently unrepresentative one) from the larger population of corporate information events? Or do they tend to over-weight or under-weight (e.g., relative to some model of expected coverage) certain kinds of stories at different points in time? If we wish to address questions related to some notion of "abnormal coverage", however, we must first arrive at some way of determining "normal coverage". The phenomenon of corporate earnings announcements is particularly well-suited in this respect as well, since it is possible here to form an independent estimate of the "quality" of earnings events based upon contemporaneous data recorded in I/B/E/S. In particular, I study and compare the respective proportions of "positive" and "negative" news in the two information channels over time. While the categorical distinction of interest here (i.e., "positive earnings news" vs. "negative earnings news") is essentially qualitative in nature, it is, nonetheless, one that is independently meaningful (as well as empirically discernable) with respect to each of the information channels we are interested in looking at. In particular, recent studies (e.g., Antweiler and Frank (2004), Tetlock (2007), and Tetlock et al. (2008)) have demonstrated the utility of computational linguistic techniques for estimating the semantic content of large-scale text databases.

The setting in this paper also provides an excellent testing ground with respect to the informational efficiency of market prices. While I find that patterns of "abnormal" positive and negative coverage (i.e., relative to an empirical model of expected coverage) are strongly persistent, consistent with time-varying media bias, it is not clear that the observation of any apparent mismatch between these two information sources should have any particular impact upon market prices. Note that all of the information used to create my measures of media bias is also available to market participants – if market prices are efficient, there should be no additional information embodied in my measures of estimated coverage bias, which are essentially linear combinations of public information measures. On the other hand, if investors are subject to cognitive constraints (e.g., Barberis et al. (1998)), and if limits to arbitrage are binding, then we might expect market prices to react more strongly, at least in the short run, to information made available via channels that entail lower cognitive costs.

Addressing this question of informational efficiency, I utilize the weekly estimates of coverage bias in a vector-autoregressive (VAR) setting. I investigate whether observed differentials between the two information channels contain information with respect to future stock market returns, and in particular for groups of firms that we would expect to vary in terms of the quality of their information environments. Here, I find that estimated "excess" negative (positive) media coverage Granger-causes significant positive (negative) weekly returns for small stock indices.

The plan for the rest of the paper is as follows. Section two discusses some of the related literature. Section three discusses hypotheses to be tested empirically. The fourth section describes the data and methodology. Section five presents the results of regressing the media coverage variables on the I/B/E/S information measures and discusses the residuals in the context of an interpretation of media bias. Section six presents the causality results. Section seven discusses some issues of interpretation and section eight concludes.

2.2 Related Literature

While researchers have long been interested in studying the economic implications of the news, it is only relatively recently that theorists (see, e.g., Mullainathan and Shleifer (2005), Gentzkow and Shapiro (2005), and Baron (2004)) have turned their attention toward a provocative concept that is familiar to nearly everyone who watches or reads the news: that of media bias. While not everyone will agree on its nature or direction, most of us perceive bias in the media at least some of the time (ASNE (1999)).

Although there has been significant recent activity with respect to media bias in the economics and political science literatures, this concept is clearly more difficult to address in the context of finance (Reuter and Zitzewitz's (2006) study of advertising and mutual fund rankings is a notable exception). When we think of media bias, the image that comes to mind is not necessarily that of the relatively staid and buttoned-down major financial news media outlets. However, theoretical models of media bias, and, in particular, the demand-driven setting described by Mullainathan and Shleifer (2005) and Gentzkow and Shapiro (2005), are equally applicable in this setting. In these models, news providers, either consciously or unconsciously, compete to supply a narrative stream (e.g., by selecting which information events are deemed newsworthy on any given day) that target audiences will find most compelling and/or credible in light of their current outlooks and beliefs. This selective tendency is referred to in news circles as the "narrative imperative", or, somewhat more pejoratively, "sensationalism". Particularly in those cases where the "true state of the world" may be highly subjective (e.g., politics), or relatively complicated in nature (e.g., the state of evolving economic conditions), competition among providers will not necessarily result in the speedy elimination of any such consumer demand-driven biases in average content.

But can media reporting behavior, in and of itself, have a significant impact on investors' beliefs? While notions of market efficiency would imply that it should not, a growing body of empirical evidence has begun to suggest otherwise. In one very striking case, Huberman and Regev (1999) document a large (and at least partially permanent) change in the biotech firm Entremed's stock price in reaction to a high-profile news article which essentially reproduced information that was already in the public sphere. While relatively clear-cut cases such as this are difficult to identify in practice, larger-scale studies have confirmed the intuition that media information flows are important. For example, Busse and Green (2002), looking at the market's real-time reaction to CNBC broadcasts, report that stock prices reliably react to firm/CEO media exposure within seconds, largely unrelated to the actual information content of the coverage itself. Similarly, Barber and Odean (2008) present evidence that media attention affects the actions of investors in systematic and predictable ways, this time by examining individual investors' portfolio trading decisions with respect to high media-attention and low media-attention stocks. While findings such as these are difficult to reconcile with classical assumptions, they are consistent with the settings described by researchers such as Barberis et al. (1998).

Recently, the application of computational linguistic techniques has allowed researchers to delve deeper into issues of media content and its effect upon markets. For example, Tetlock (2007) examines the use of negative words in a long-running daily column in the WSJ. He finds that his measures of negative language usage, which he identifies with indicators of sentiment, are able to predict market returns and trading volume at daily horizons. While this study does not seek to control for fundamental information as a determinant of specific media content, it illustrates an explicit link between media "tone" and market dynamics. Similarly, Antweiler and Frank (2004) also apply automated textual analysis and find that positive and negative sentiment in internet postings about stocks has some predictive content with respect to market movements and activity. At the same time, however, one is left to wonder: How might these findings correspond to an investigation of media bias? Should we think of linguistic media content measures as unbiased indicators of the "true" tenor of the flow of contemporaneous market news events? Rather, what if media accounts sometimes offer a distorted picture (in some observable, quantifiable sense) relative to the "true state of the world"? How would we know if such distortions were present? What would be the impact?

Empirical studies of media bias typically rely upon a comparison between observed media coverage and some baseline of "expected" media coverage in the absence of bias. In the context of political reporting, Groseclose and Milyo (2005) use an innovative approach to identify potential (relative) biases in mainstream coverage: they identify the think tanks that reporters working at different publications tend to cite in their political stories, comparing these to the think tanks mentioned in the speeches of U.S. legislators of various political stripes. Political media bias is then defined relative to the expected citation behavior of a hypothetical, middle-of-the-road, moderate politician. With a focus on the potential for supply-driven bias favoring advertisers, Reuter and Zitzewitz (2006) link mutual funds' advertising expenditures to a notion of their "abnormal" performance in industry rankings published in trade publications relative to the rankings that we might expect to see based upon objective measures of fund performance. Bhattacharya et al. (2006) attempt to identify "exaggerated" media sentiment by comparing media coverage regarding matched samples of internet and non-internet IPOs during the late 1990s. In this study, I undertake an empirical strategy that is analogous to the ones found in these papers; specifically, I identify potential media bias as the observation of predictable and persistent deviations relative to an estimated model of expected media coverage.

With respect to studies regarding earnings announcements in particular, there is ample evidence to suggest that these news events are a very important source of information regarding firms, and also that the very manner in which they are discussed in the press can be informative. Dyck and Zingales (2003) investigate potential "spin" in earnings news coverage and show that stock prices react most strongly to the information elements in the press release that reporters emphasize in their coverage. More recently, Tetlock et al., (2008) show that the tone of media coverage leading up to a firm's earnings announcement date can tell us something not only about the content of the upcoming release, but also regarding the probable nature of the market's ultimate response to it.

2.3 Hypothesis Development

In this study, I focus on media outlets' relative propensities to report on "positive" versus "negative" earnings news events over time. For practical purposes, this positive/negative categorical distinction has the advantage of being relatively easy to identify in both linguistic and numeric terms. Positive and negative news attributes are also extremely salient and relatively unambiguous in terms of their expected impacts on security prices, allowing us to make clear-cut predictions regarding the likely market impact of any such potential distortions in the information environments surrounding firms.

In this context, how might we expect media bias (e.g., as defined by Mullainathan and Shleifer (2005)) to manifest itself, if present? Given that empirical proxies for positive/negative characteristics in the respective information channels may not be completely comparable, evidence of *static* differences in levels is unlikely to be persuasive. Rather, we would expect any meaningful expression of media bias to vary somewhat over time – in other words, it would only make sense for us to interpret current media coverage as being "excessively" focused on positive events in a *relative* sense, e.g., compared to the coverage typically observed during some other period. At the same time, by their very nature, observable instances of media bias must be at least somewhat predictable within a given period – in other words, an observation of apparently random errors (relative to expected coverage) over time would be similarly unpersuasive as evidence of systematic media bias.

In short, if news coverage is subject to *persistent media bias (PMB)*, then we should expect this to reveal itself by way of predictable (although not completely time-invariant) deviations relative to the fundamental information set over time, leading us to the first testable hypothesis:

> **[PMB] Hypothesis I:** Measures of unexplained/residual media coverage are serially correlated.

If the information transmitted to investors by way of a particular information channel is "distorted" in some sense (e.g., if PMB is present), it is still by no means clear that this should have any particular impact at all upon securities' prices. Under the standard assumptions underlying market efficiency, so long as at least one sufficiently large market participant can "see through" the distortion, then prices should adjust quickly. However, even if we assume that limits to arbitrage are binding, why would even an unsophisticated investor pay attention to a distorted signal when relatively undistorted measures of the underlying information (i.e., earnings press releases, analyst forecasts, etc.) are also available? Simply, investors' information sets will become distorted under PMB if cognitive constraints prevent them from identifying the media distortion, or, equivalently, from searching out and analyzing the underlying fundamental information set for themselves. Given the vast quantity of earnings announcement information that is typically released on a daily basis, it may be entirely rational for an individual investor to simply accept the financial news media's selective daily "summary" at face value (essentially accepting the possibility that the information may be subject to a certain degree of distortion), rather than incur the cost of gathering the underlying information him- or her-self. In other words, faced with limited cognitive capacities, investors might rationally prefer to generalize information from an expedient (yet potentially inaccurate) signal in forming their beliefs about the larger "state of the world".

Under limits to arbitrage and constrained attention, what is the expected impact of PMB on securities' prices in the current setting? If the current media coverage mix is "excessively" negative (e.g., relative to some model of expected coverage based upon fundamentals), and if investors are unable to accurately identify and account for the biased information, then the implication is that stock prices may be currently under-valued. Assuming that information communicated in earnings announcements but not via the news media will eventually be revealed to markets, evidence of a current negative bias would lead us to predict positive returns in future periods, and *vice versa.*

In short, if *limited investor attention (LIA)* prevents at least a subset of investors from identifying potentially distorted information in news media coverage (and if, in addition, there are limits to arbitrage), then we should expect PMB-based distortions in coverage to contain information about future stock prices, leading us to the second testable hypothesis:

> **[LIA] Hypothesis II:** Unexplained/residual negative (positive) media coverage predicts positive (negative) stock market returns.

2.4 Data and Methodology

The empirical strategy of this study is to investigate potentially systematic differences between the respective contents of two distinct sources of information regarding one particular type of information event: quarterly corporate earnings results. I examine contemporaneous reports in the financial news media and in the I/B/E/S database, attempting to answer such questions as: Is there evidence of predictable or time-varying deviations between the information flows described by these two channels? If so, what might this tell us about the potentially systematic nature of distortions in the larger information environment with respect to these firms?

Earnings announcement and stock market data

The I/B/E/S sample data set covers all quarterly earnings announcement observations (including the preceding analyst forecasts) from January 1984 to December 2005 where a PERMNO is available and can be successfully matched with a corresponding daily observation of closing stock price and shares outstanding in CRSP. The combined data set comprises 265,323 observations of quarterly EPS announcements, analysts' expectations, and firm-day stock market variables regarding North American firms. Following, DellaVigna and Pollet (2008), I restrict the sample to those firms with stock prices greater than or equal to $1, and to observations where the absolute value of the announced earnings result is smaller than the contemporaneous stock price. Table 2.1 shows that the greatest number of earnings releases are recorded on Tuesdays, Wednesdays, and Thursdays, with only about half as many entries on Fridays and Mondays. Similar to DellaVigna and Pollet (2008), it appears that there is a significantly greater incidence of negative surprises on Friday, consistent with the idea that firms may try to time the release of bad news in such a way that an intervening weekend will serve to interrupt the potentially negative news cycle.

Daily NYSE stock index levels are obtained from CRSP. Weekly returns are the log differences of Friday-close to Friday-close observations. NYSE trading volumes are from the NYSE Group historical data. I focus on NYSE-listed securities since they represent a relatively large and diverse population of firms that, nonetheless, operate under what is advertised to be the most stringent financial reporting requirements in the North American market – essentially ensuring that they are reliable in releasing complete financial results. Perhaps even more importantly, NYSE-listed securities enjoy an institutionally-mandated minimum level of liquidity due to the specialist system. In this context, the stocks of even relatively small firms are unlikely to be highly illiquid or subject to extreme informational asymmetries due to spotty corporate reporting – given that firm size deciles are utilized, these minimum attributes may be important with respect to interpretation of some of the results in the sections to follow.

Media data

The media data set includes 68,102 distinct *Wall Street Journal* (WSJ) articles (January 1985 to December 2005), 17,289 *New York Times* (NYT) articles (January 1990 to December 2005), and 83,727 *Financial Times* (FT) articles (November 1990 to December 2005) from Factiva in plain-text format.[6] Each article is stored as a separate file and identified by a unique filename. The original articles include Factiva's indexing codes and formatted information fields; however, coded information in the raw text is stripped from the articles before text categorization software is applied.

In order to focus on the relatively small proportion of media articles that are of interest for the purposes of this study (i.e., articles reporting specifically on the content of a recent corporate quarterly earnings announcement), as a first step, I consider only those articles identified as pertaining to earnings by Factiva's *Intelligent Indexing*. However, this is a relatively broad filter, including a very large number of "false positives". For example, this identifier flags articles that report on firms' profit warnings and their releases of preliminary sales figures, as well as more general articles discussing proposed regulatory changes with respect to reporting standards, etc. Therefore, after filtering out certain easily-identified false positive article types from this initial Factiva set (details available upon request), the subset of "true" earnings release articles is identified using a computational linguistics software package known as the Rainbow Toolkit (McCallum, (1996)). Then, in a further round of automated categorization, the so-identified earnings articles are further classified into "negative" and "non-negative" categories. This methodology is described briefly in the following sub-section; for more details, including unconditional probabilities of class membership and estimates of out-of-sample classification accuracy, please see Appendix I.

Applying natural language processing to news story classification

In the context of this paper, 500 articles were chosen at random from each of the WSJ, FT and NYT corpora, resulting in three distinct training sets. Once each document in a training set has been read by me and assigned to its appropriate categories ("earnings" or "not earnings", and then, subsequently, in the case of identified "earnings" articles, "negative" or "neutral" or "positive"), natural language software is then used to index each of the training sets. Essentially, the objective is to estimate models of class membership based upon observed linguistic similarities among articles in each category, and with respect to each publication.

The text classification technique utilized here is known as a "bag of words" approach. In short, the software counts the number of times specific words (or groups of words) appear in

[6] The respective start dates for the media data sets correspond to the earliest points at which Factiva's indexing codes begin to identify significant numbers of articles in that publication (i.e., as required for estimation with respect to the methodology described in this section).

each document (omitting a standard list of common words such as: "of", "and", "the", etc.), and then relates these data to the category into which documents have been assigned. Under the Naïve Bayesian assumption that word observations within documents are independent, odds ratios with respect to inclusion in each of the pre-specified document categories are calculated for each word.

The result is an empirical model of document category membership which can be applied to new documents of unknown type by simply adding up the logs of the odds ratios corresponding to each word contained therein. For each document, the category which maximizes this sum is inferred to be its "true" category. Although the Naïve Bayesian assumption is rather strong, and other techniques (such as TFIDF, and k-nearest neighbor, etc.) are also available, this approach is seen as relatively robust to potential misspecification and has been found to work surprisingly well compared to other, more computationally-intensive, strategies (Manning and Schuetzte (1999)).

To summarize briefly, the automated document classification strategy consists of two sequential stages. In the first stage, articles are sorted into one of two categories: "earnings announcement articles" and "not announcement articles". In the second stage, the "earnings announcement articles" are further differentiated into one of two sub-categories: "negative earnings announcement articles", and "not negative earnings announcement articles".[7]

Table 2.1 provides some statistics regarding potential day-of-the-week patterns. All but a very small number of earnings articles in the WSJ are published from Monday to Friday. The successively larger totals as we move through the week, peaking on Friday, support the idea that articles may come a day, or even two days, after they show up in I/B/E/S. Notice that the proportion of negative news articles peaks on Monday, while we noted earlier that the proportion of negative surprise announcements in I/B/E/S is highest on Friday. (Accordingly, the regressions in the following section will attempt to account for these potential lags in information.) Alternately, note that the NYT does publish a significant number of earnings articles on Saturdays, and it is on this day that it exhibits the highest proportion of identified negative stories. The FT, due to its primary location in London, is effectively one day behind

[7] In the second stage of text categorization, a set of binary classification models are also estimated with respect to the two other potential categorical distinctions in the training sets (i.e., "neutral" vs. "not neutral", and "positive" vs. "not positive"). For this version of the paper, in the interests of brevity, I have concentrated upon the "negative"-"not negative" dimension of category separation – in other words, "neutral" and "positive" earnings articles in the publication training sets are treated as "not negative" in order to implement a binary classification scheme.

New York in terms of the "Wall Street" news cycle. Here again, we find a high proportion of negative articles on Saturday, and also on Monday. Before even looking at the simulated model accuracy results in Appendix I, this apparent confirmation of expected intra-weekly patterns (i.e., consistent with the findings of DellaVigna and Pollet (2008) and the observed intra-week distribution of I/B/E/S negative surprises in Table 2.1) may give us some confidence that the categorization algorithms are performing broadly as they should.

Definition of media variables

In short, I am interested in forming weekly estimates of the overall picture of corporate performance as described by financial news media articles discussing recently released quarterly corporate earnings announcements. Individually, these articles tend to be rather prosaic, factual, and to-the-point, but together they can be seen as representing a near-continuous flow of corporate information throughout the day and week.

Absolute levels of positive and negative media coverage about earnings announcements are not of primary interest here, but rather the relative proportions of positive or negative coverage as a share of relevant media output, thereby obviating concerns regarding the potential impact of secular changes in overall media attention over time. Moreover, the media measures described here (just as with the I/B/E/S information measures to be defined in the following section) are designed to be homogenous of degree zero in overall activity levels.

The normalized media measure with respect to the Wall Street Journal is defined as follows:

$$WSJ_t^{-ve} = \frac{\sum_{i=1}^{N_{1,t}} I_{i,t}^{-ve}}{N_{1,t}},$$

where $N_{1,t}$ is the number of earnings announcement articles identified by text classification model WSJ1 (see Appendix I) in period t; $I_{i,t}^{-ve}$ is an indicator variable that takes on the value of one if article i in period t has been identified as an "earnings announcement" article and then subsequently identified as "negative", by models WSJ1 and WSJ2 respectively, and zero otherwise.

Similarly, for the New York Times,

$$NYT_t^{-ve} = \frac{\sum_{i=1}^{N_{2,t}} I_{i,t}^{-ve}}{N_{2,t}},$$

where $N_{2,t}$ here is the number of earnings-related articles identified by Factiva in period t[8]; $I_{i,t}^{-ve}$ is an indicator variable that takes on the value of one if article i in period t is identified by model NYT1 as an "earnings announcement" article and then subsequently identified as "negative" by model NYT2, and zero otherwise.

Finally, with respect to the Financial Times,

$$FT_t^{-ve} = \frac{\sum_{i=1}^{N_{3,t}} I_{i,t}^{-ve}}{N_{3,t}},$$

where $N_{3,t}$ is the number of earnings announcement articles identified by model FT1 in period t; $I_{i,t}^{-ve}$ is an indicator variable that takes on the value of one if article i in period t has been identified as an "earnings announcement" article and then subsequently identified as "negative", by models FT1 and FT2 respectively, and zero otherwise.[9]

Definition of I/B/E/S news variables

Since I am interested in explaining the relative proportion of negative news in earnings coverage over a given period, it makes sense to look at the contemporaneous flow of underlying information events that could be described as negative relative to some objective, relatively unambiguous accounting-based criteria. Therefore, my first I/B/E/S-based measure is the proportion of earnings announcements that feature earnings per share (EPS) strictly less than zero (over a given observational period). While we can think of circumstances where reports of negative earnings are not necessarily disastrous (such as with early-stage growth firms, for example, where such results may be well in line with performance expectations), and indeed we might have some reason, a priori, to expect that a negative number should not really matter for

[8] Given the relatively small number of NYT articles, and the corresponding incidence of zero article periods in the NYT sample, the number of articles identified by model NYT1 is problematic for use as a binomial denominator at the weekly frequency.

[9] An earlier version of this paper also examined articles' word count-weighted versions of the WSJ, FT, and NYT measures; results are very similar to those presented here.

announcement-day returns (i.e., to the extent that the result is in line with market expectations), it would be extremely hard to interpret announcements of losses as "good news" in an absolute sense.

The second proxy for negative news in I/B/E/S is the proportion of announcements where EPS is lower than the median analyst forecast over the preceding month, reflecting the intuition that only the "surprise" portion of news shocks should matter.

But how should these proportions be defined? In particular, should each announcement in I/B/E/S be given equal weight, considering that we are hoping to explain variation in observed media attention? There are at least two simple, alternative scenarios that might help to clarify the issue of "expected" relative coverage with respect to equal- vs. value-weighting.

First, what if reporters appear to randomly choose firm-announcement subjects from among those firms that happen to release their quarterly earnings announcement on any given day? For example, imagine that reporters are attracted to some firm qualities that are completely orthogonal to short-term realizations of simple accounting profitability, but that, once the choice of which firm-announcement to write about has been taken, the "objective" quality of the announcement's associated accounting data will bear some relationship with the tone of the resulting story? In this case, we might expect more "negative" news stories to be published during a period where we see, for example, a higher-than-average incidence of companies reporting quarterly profits below market expectations.

Second, suppose that reporters pay more attention to the largest companies; imagine that sheer size makes firms intrinsically more interesting. For example, what if writers at each publication simply choose to write about the announcements associated with the largest firms, (perhaps assuming that this will attract the attention of the greatest number of readers due to these companies' larger employee and customer populations, etc.), until they have used up all of their available space in their respective newspapers on that day? In this case, we would expect the proportions of positive and negative media stories to be related to the market value-weighted incidence of firms reporting negative accounting data.

Therefore, in order to investigate these two potential effects regarding reporters' publication choices, we will consider two types of proxies (i.e., market value-weighted, and observation-weighted) with respect to each of the two concepts of "negative" I/B/E/S information described earlier.

The "negative EPS" proxies are defined as follows:

$$EARN_t^{COUNT,EPS<0} = \frac{\sum_{j=1}^{M_t} I_{j,t}^{EPS<0}}{M_t},$$

and

$$EARN_t^{MKTVAL,EPS<0} = \frac{\sum_{j=1}^{M_t} \left(MarketValue_{j,t} \cdot I_{j,t}^{EPS<0}\right)}{\sum_{j=1}^{M_t} MarketValue_{j,t}},$$

where M_t is the number of earnings announcements reported in I/B/E/S during period t; $I_{i,t}^{EPS<0}$ is an indicator variable that takes on the value of one if EPS was negative for I/B/E/S earnings announcement j in period t, and zero otherwise; $MarketValue_{j,t}$ is the equity market value (share price multiplied by shares outstanding on the day of the announcement) of the firm that made EPS announcement j during period t.

Similarly, the "negative EPS surprise" proxies are defined:

$$EARN_t^{COUNT,EPS-e(EPS)<0} = \frac{\sum_{j=1}^{M_t} I_{j,t}^{EPS-e(EPS)<0}}{M_t},$$

and

$$EARN_t^{MKTVAL,EPS-e(EPS)<0} = \frac{\sum_{j=1}^{M_t} \left(MarketValue_{j,t} \cdot I_{j,t}^{EPS-e(EPS)<0}\right)}{\sum_{j=1}^{M_t} MarketValue_{j,t}},$$

where M_t is the number of earnings announcements reported in I/B/E/S in period t; $I_{i,t}^{EPS-e(EPS)<0}$ is an indicator variable that takes on the value of one if EPS announcement j in period t was less than the median analyst forecast, and zero otherwise; $MarketValue_{j,t}$ is the equity market value (share price multiplied by shares outstanding on the day of the announcement) of the firm that made EPS announcement j in period t.

2.5 Explaining the Media's Coverage of Earnings News

In this section, I present the results of regressing the proportional media-based measures on the I/B/E/S-based fundamental earnings news proxies. In short, my purpose here is to explain as much of the observed variation in the normalized media coverage measures as possible with reference to the associated contemporaneous and recently-released accounting data.

Before going further, recall that the central objective is to examine potential relationships between our two sets of normalized weekly information flow variables: media-based (WSJ, NYT, and FT) and accounting-based (the I/B/E/S measures), respectively. As such, and given that these measures are designed to be homogenous of degree zero in activity (i.e., in total publication and announcement activity levels), most explanatory variables that might immediately spring to mind as potential additions in the regressions to follow turn out to be frankly inappropriate upon further reflection. For example, several macroeconomic data series might be attractive in the sense that they can be seen as containing additional fundamental information regarding the general state of corporate health in the economy. Unfortunately, these series are typically measured at monthly or quarterly frequencies; even so, they often relate to information regarding conditions several months in the past, so it is not immediately apparent how they could be appropriately interpreted in a weekly information flow context, particularly one as tightly focused as this.

Security prices and market activity variables, on the other hand, might be sufficiently high-frequency and contemporaneous for our purposes, but it is even less clear how we might interpret price-level, liquidity, or activity measures in terms of the objective to explain relative rates of occurrence regarding two very specific sets of corporate/media information events. Furthermore, recall that one of the ultimate goals of this study is to investigate what unexplained media information content might be able to tell us about future returns (i.e., about the current and future locations of relative informational inefficiency in the market). If market variables are incorporated at this point, it could serve to introduce considerable ambiguity regarding this question later on.

While individual results with respect to each of the three publications are described below, there are some common elements which may be more efficiently discussed up front. Since the dependent variables in these regressions are proportions between zero and one, each model is estimated as a logit with the total number of earnings articles serving as the binomial denominator. In this way, the estimation utilizes the additional information represented by the number of article observations comprising each weekly proportion observation. Also, lags of the I/B/E/S proxies are explored as potential independent variables for inclusion; as mentioned earlier, the reasoning here is that there may exist a lag of one to two days between the release of an earnings announcement and the publication of an article discussing that announcement. If an announcement occurs on a Friday or Thursday, particularly on a Friday afternoon after markets have closed, then we might not expect to see an associated news story published until the following week. Therefore, we should account for the potential impact of last week's I/B/E/S news on this week's coverage measure.

In addition, anecdotal evidence suggests that reporters may be more likely to report on a given story if it is perceived as being part of a "trend". If a string of similar results arrive in successive weeks (for example, if there is an observed trend of large companies failing to meet market expectations), reporters may be more likely to report on an announcement that appears to represent a continuation of the recently-observed pattern. While an attempt to explicitly explore this kind of dynamic in reporting patterns is beyond the scope of the current exercise, including lagged versions of the I/B/E/S measures should allow us to test for this kind of potential impact of past I/B/E/S news realizations on current coverage decisions.

One important caveat to the analysis here is that the posited link between earnings news events and contemporaneous media coverage actually aggregates two potentially distinct publication decisions: the decision of whether or not to write a story regarding a given announcement, and the subsequent (or possibly joint) decision of what kind of story to write. In short, given the aggregated nature of the measures, an apparent "mismatch" between the positive/negative tenors of weekly fundamental news and contemporaneous news coverage may be driven either by reporters' choices of which firm-announcements to report on, or by their choices of whether to interpret particular announcement news items as positive or negative given the fundamental information, or both. For the purposes of this study, I implicitly interpret systematic, time-varying deviations in either decision-making activity as evidence of persistent media bias.

The rest of this section describes results with respect to each of the three empirical publication models.

<u>Earnings news coverage in the WSJ</u>

Table 2.2 shows the results from logit regressions of the WSJ negative coverage variable on various combinations (including lags) of the I/B/E/S news proxies. In columns 3 and 4, we observe that all of the contemporaneous I/B/E/S information flow proxies, as well as some lags, are statistically significant. Standard errors are shown with Newey-West corrections for

heteroskedasticity and autocorrelation up to 5 lags. Column six, which represents the "best" base-line specification according to model selection criteria and switching regression analysis, incorporates the square of the contemporaneous negative EPS announcement variable to account for potential non-linearity (no other similar transformations of the I/B/E/S variables were found to be significant). The specification from column six of Table 2.2 is as follows:

$$WSJ_t^{-ve} = \alpha_0 + \alpha_1 EARN_t^{COUNT,EPS<0} + \alpha_2\left(EARN_t^{COUNT,EPS<0}\right)^2 + \alpha_3 EARN_t^{COUNT,EPS-e(EPS)<0} + \alpha_4 EARN_{t-1}^{COUNT,EPS-e(EPS)<0}$$
$$+ \alpha_5 EARN_t^{MKTVAL,EPS<0} + \alpha_6 EARN_t^{MKTVAL,EPS-e(EPS)<0}$$
$$+ \alpha_7 EARN_{t-1}^{MKTVAL,EPS-e(EPS)<0} + \alpha_8 EARN_{t-2}^{MKTVAL,EPS-e(EPS)<0} + w_t$$

In Table 2.2, it is interesting to note how much explanatory power the I/B/E/S information flow measures seem to have with respect to variation in the WSJ's relative coverage mix. The specification in column six, which includes the non-linear term, reveals an adjusted-R^2 of over 36%.

Earnings news coverage in the NYT

Table 2.3 presents the results from regressing the NYT news variable on the I/B/E/S information variables. Compared to the WSJ regressions above, variation in the I/B/E/S-based measures here appears to be able to explain relatively little. The specification in column six, including only the contemporaneous announcement-weighted proportion of negative surprises and the square of the value-weighted proportion of negative EPS announcements, is as follows:

$$NYT_t^{-ve} = \phi_0 + \phi_1\left(EARN_t^{MKTVAL,EPS<0}\right)^2 + \phi_2 EARN_t^{COUNT,EPS-e(EPS)<0} + n_t$$

Given the relative sparseness of the NYT dataset, evident from Table 2.1 and Figure 2.3, which includes some 13 weeks of zero-article observations, it is perhaps not surprising that the results here are relatively poor compared to those with respect to the WSJ. An adjusted R^2 of 2.67% confirms that there is relatively little explanatory power to be found here, despite the fact that two of the contemporaneous I/B/E/S-based measures are statistically significant.

Earnings news coverage in the FT

Looking at the FT media logit regression results in Table 2.4, we do not see anything like the broad significance across different measures that we observed with the WSJ. In fact, the regression in column 5 reveals significance only with respect to contemporaneous and lagged values of the value-weighted negative surprise proxy, resulting in the specification seen below.

$$FT_t^{-ve} = \gamma_0 + \gamma_1 EARN_t^{MKTVAL,EPS-e(EPS)<0} + \gamma_2 EARN_{t-1}^{MKTVAL,EPS-e(EPS)<0} + \gamma_3 EARN_{t-2}^{MKTVAL,EPS-e(EPS)<0} + \gamma_4 EARN_{t-3}^{MKTVAL,EPS-e(EPS)<0} + f_t$$

However, unlike with the NYT regressions, which similarly revealed significance only with respect to a couple of measures, this regression yields an adjusted R^2 of over 15%. Given the FT's relative focus on international markets and firms compared to the WSJ, it is not surprising that the North American I/B/E/S-based measures do less well here in explaining the FT's earnings news output.

Interpreting the media information residuals

Figures 2.7, 2.8 and 2.9 show the 13-week moving averages of the estimated residuals resulting from the three media information regressions described above (i.e., estimates of w_t, n_t, and f_t). Looking at Figures 2.7 and 2.8, in particular, it is clear that the estimated residuals are strongly persistent, a finding that is confirmed in (untabulated) tests of autoregressive terms. In addition, there appears to be some significant cyclicality. For example, Figure 2.7 implies that WSJ earnings coverage was generally more negative than average (given fundamentals) in 1986 and 1987, more positive than average in 1988 to 1990, and so on. In short, the evidence here would not appear to allow us to reject the PMB hypothesis.

The presence of autocorrelation in the media residuals bears some discussion. While this could potentially be seen as a statistical artifact, or as simply evidence of model misspecification, it is important to note that an "omitted variable" interpretation is fully consistent with an interpretation of systematic, demand-driven media bias. In this context, the potential "omitted variable" that we are attempting to isolate in the residual, i.e., representing discrepancies in media coverage relative to the flow of fundamental information, may actually be persistent itself. Moreover, this apparent persistence in the informational residuals may represent a large portion of their potential value as predictors. Therefore, while controls designed to address autocorrelation were incorporated in two types of alternative specifications (and, in fact, the causation results in the following section have been checked for robustness in this respect – results are similar), it is not clear that autocorrelation in the residuals is necessarily a "problem" that needs to be "fixed" before we can move on to the Granger-causality tests.

Consequently, for the remainder of the paper, I focus on the information residuals resulting from the right-hand-most logit regressions in Tables 2.2, 2.3, and 2.4.

In addition to serial correlation, tests of the contemporaneous relationships between, and cross-causality among, the three media residuals provide some interesting results (with respect to the common sample period beginning in November if 1990). Simply, the NYT, FT, and WSJ residuals are all significantly contemporaneously correlated with each other, which makes sense if we think of each publication as operating within the context of a larger information environment, subject to many of the same forces (e.g., common memes) that affect conditions in the information environment more generally. In the context of a simple trivariate VAR, I find (in untabulated results) that the WSJ residual Granger-causes the FT residual ($X^2(10)=38.167$, p-value < 1%), but that there exist no other significant causal relationships among the variables. This could imply that the WSJ's editorial policy (with respect to corporate earnings coverage) leads the FT's in some sense. This might be the case, for example, if the WSJ tended to spot new trends earlier than the FT (albeit potentially only with respect to this rather narrow slice of the financial news universe, of course). Alternately, this predictability might stem from differences in the fundamental news shocks affecting the firm populations of interest with respect to each of the news providers (i.e., "mostly U.S." firms for the WSJ vs. "more global" firms for the FT).

The results in this section may be somewhat puzzling from an information efficiency perspective. In particular, the patterns that we observe seem to be fully consistent with the definition of media bias that we have set out for ourselves: predictable and persistent deviations in coverage relative to a "normal" model of expected behavior. Consider, once more, Figure 2.7. Why would reporters apparently systematically over-report negative earnings news events during some periods, only to predictably under-report similar events at other times? While it might be tempting to interpret such observed variations in average coverage as evidence that reporters are simply making persistent "errors" in transmitting news relative to the fundamental information set, it seems rather more likely that our relatively simplistic categorization of events is doing a poor job of proxying for the narrative story elements that reporters perceive as newsworthy at various points in time. Clearly, evidence of systematic deviations from the econometrician's "normal" model of expected coverage needs not imply any consequent distortions in the information environment; if the real-world information environment is much more complicated than the model, and if, in addition, real-world participants understand and can account for this un-modeled complexity, then any apparent "deviations" may simply reflect misspecification. On the other hand, if there is evidence that the observed/imputed "errors" have predictive content with respect to future events, then this may go some way toward convincing us that we are observing a source (or, alternatively, a symptom) of meaningful distortion in the information environment.

Moreover, while we have seen evidence that residual negative media coverage is apparently serially correlated and cyclical in nature, implying that typical coverage patterns are at least somewhat predictable in the short run, it remains to be seen whether or not such behavior actually has an impact. For example, if markets are informationally efficient, then any systematic media "spin" or "slant" with respect to these events should not make a difference for stock return predictability, so long as at least one sufficiently large market participant has access to the underlying information set. In the following section, we explore these questions of causation and economic significance.

2.6 Unexplained Media Coverage and Stock Market Predictability

Having constructed normalized media earnings information residuals by filtering the news content measures on I/B/E/S announcement data, I can now proceed to investigate their potential relationship with respect to stock market returns and activity.

In this section, I construct a series of bivariate vector-autoregressive (VAR) systems and conduct Granger causality tests (Granger (1969)). If excessive negative media coverage (i.e., relative to fundamentals) has distortionary effects on investors' information sets, we should expect to see current positive realizations of the media bias proxies (w_t, f_t, and n_t) predict positive stock returns in the future – in other words, if current coverage is "too negative", then, so long as the current positive news that isn't being reflected in coverage eventually "comes out" at some point, and has the expected impact on security prices, then we should expect to see current negative coverage residuals predict abnormal positive returns in future periods. At the same time, we are also interested in whether past returns are able to predict future realizations of the information residuals; causality is tested in both directions.

To summarize the results in this section, I find that the media information residuals Granger-cause weekly NYSE index returns (and in particular with respect to the small size decile returns); at the same time, I find that returns do not Granger-cause the information residuals, indicating that causality flows only in one direction.

The VAR systems are defined as follows:

$$i_t = \eta_0 + \eta_1 i_{t-1} + \ldots + \eta_{k+1} R_{t-1} + \ldots + \eta_{2k+1} X_t + \eta_{2k+2} X_{t-1} + \ldots + u_t$$

$$R_t = \delta_0 + \delta_1 i_{t-1} + \ldots + \delta_{k+1} R_{t-1} + \ldots + \delta_{2k+1} X_t + \delta_{2k+2} X_{t-1} + \ldots + v_t \,,$$

where k is the number of lags; i_t is one of the information residual series corresponding to the WSJ, NYT, and FT regressions from the preceding section (i.e., w_t, n_t, or f_t); R_t is one of the stock index return series; and X_t represents a vector of exogenous control variables as suggested by the literature on return predictability, including: monthly dummy variables to account for potential seasonality, detrended log aggregate NYSE trading volumes, a proxy for recent volatility with respect to the return series, and a dummy variable to account for the week of the October 1987 stock market crash.[10]

Before going on to discuss the causality results, I will first make some general remarks regarding estimation techniques. Dickey-Fuller and Philips-Perron tests allow us to reject the hypothesis that unit roots are present. However, LM tests indicate that ARCH effects are present in each of the equations to be described below. While Liew and Chong (2005) report that ARCH errors do not affect the appropriateness of the lag-length selection criteria I utilize, Cheung and Fujii (2001) use Monte Carlo techniques to show that the power of standard causality tests are significantly and adversely affected when ARCH effects are not accounted for. For each of the regressions in this section, the appropriate ARCH error specification of u_t or v_t is determined by minimization of the Akaike information criterion (AIC) and the Bayesian information criterion (BIC).[11] An examination of the estimated error terms from the return and information residual equations reveals no patterns of statistically significant contemporaneous correlation. Moreover, since news publication decisions and equity price determination are relatively distinct processes in terms of the exogenous "shocks" that we might expect to affect

[10] Similar to the approaches utilized in Campbell et al. (1993) and Tetlock (2007), I detrend log trading volumes by subtracting the average of log trading volumes over the preceding 13 weeks (i.e., roughly one quarter). The (similarly detrended) proxy for recent volatility used here is the weekly standard deviation of returns less the average standard deviation of returns over the preceding 13 weeks.

[11] The presented results are nonetheless generally robust to varying error specification assumptions within the ARCH family.

their respective paths, there are, a priori, no obvious relationships to impose among the respective sets of ARCH processes. Therefore, as in Tetlock (2007), I assume independence between the disturbance terms with respect to each return-residual equation pair.

Does residual media coverage predict stock returns?

Regarding the predictive regressions with stock index returns as the dependent variable, lag-length specification tests yield somewhat ambiguous results – in some cases, BIC indicates that one lag is appropriate; in others, two lags are suggested. Results for both the one- and two-lag length specifications are shown in Tables 2.5 and 2.6, respectively.

Table 2.5 shows results for the one-lag specification of the predictive regressions of stock market returns on lags of returns and the WSJ, NYT, or FT residuals (plus the control variables). After controlling for lagged returns and the exogenous control variables, the coefficient on the lagged WSJ information residual is positive and statistically significant (i.e., rejecting Granger non-causality) with respect to the 1st, 2nd, 3rd, 4th, and 6th size decile return equations, as well as equal-weighted NYSE returns. In the NYT regressions, coefficients on the lags of the NYT residual have the expected positive sign but are statistically insignificant in all cases for this specification. Significance results for the regressions with lagged FT residuals are very similar to those regarding the WSJ in the cases of the smaller size deciles, but, at the largest size decile (column 10) the sign of the coefficient seems to reverse, indicating that abnormally negative earnings coverage actually predicts negative large-stock index returns in the following week. This last result, while not confirmed in the two-lag specification below, is nonetheless somewhat puzzling.

Table 2.6 presents the return predictability results corresponding to the VAR specification with a lag length of two. Again, we find that the FT and WSJ information residuals Granger-cause NYSE index returns for the smallest size deciles. Interestingly, we see some significant positive coefficients (t-2) in the middle-decile return equations with respect to the FT and NYT residuals.

There is at least one alternative explanation for these results that should be considered: If contemporaneous realizations of the information residuals are positively related to current returns, then, given the observed positive autocorrelation in the media residuals, we might be seeing apparent predictability solely as a result of the residual's ability to predict itself! In fact, this does not seem to be the case. For instance, when contemporaneous values of the WSJ residuals are added to the regressions with respect to the WSJ regression in Table 2.6 (note that this specification would no longer have the interpretation of a predictive regression), the resulting (untabulated) coefficients are negative with respect to all twelve indices, significantly so (at 5%) in all cases save for deciles 1, 6, and 7.

Do stock returns predict residual media coverage?

Specification tests here are less ambiguous, indicating an appropriate lag-length of two weeks. Tables 2.7, 2.8, and 2.9 present the results for the information residual equations. In each case, and with respect to each index return series, we are unable to reject the null hypothesis of non-causality. This implies that the causality relationships identified earlier exist in one direction only: from the media residuals to returns, and not vice-versa. Note as well the evidence of highly significant autocorrelation in all of the information residual measures, here even after the control variables have been incorporated to the specification.

Discussion

The results presented above indicate that media coverage residuals can help to predict (at least some segments of) NYSE index returns at lags of up to two weeks. In short, this does not seem to support rejection of the LIA hypothesis. But what else might be going on here?

At this stage, there may be several potential explanations. It could be argued that reporters' coverage decisions are subject to the broader influence of market-wide changes in sentiment, and that it is therefore simply a high degree of correlation with this factor that creates the illusion of apparent predictability with respect to returns. Under this interpretation, which is consistent with the explanation put forward by Tetlock (2007), abnormal coverage (i.e., implied instances of "bias" with respect to the expected coverage model) would simply be an indicator of current, "sentiment-related" misvaluation (as described in, e.g., Baker and Wurgler (2006), or Neal and Wheatley (1998)).

Alternately, if PMB and LIA are significant factors, distorted media coverage could be seen as contributing directly to specific locations of information inefficiency with respect to poor-information environment stocks in the market, suggesting a more directly causal link with respect to the observed phenomenon of stock predictability. For example, if current media coverage represents the state of corporate earnings as being "more negative" than an impartial observer might judge it to be based upon accounting fundamentals (or even, e.g., relative to an unconditional prior), this would be associated with a positive realization of the estimated media information residual in the current period. The implication would be that good earnings news (in a relative sense) is "out there", but it simply isn't being reflected in terms of concurrent media coverage. If smaller stocks, whose information environments are generally of poorer quality, are more subject to temporary misvaluation associated with the distorted information set, then we should expect to see greater predictability with respect to these prices. The key distinction between this and the previous explanation might be expressed as follows: Could systematic distortions in the information environment be playing a causal role in helping to generate observed patterns of apparent misvaluation which we have come to associate with relatively amorphous concepts of generalized investor psychology, such as market sentiment?

While I am unable to make definitive statements in this respect, there are a couple of observations that can be made. First, based upon the results presented above, it would be difficult to argue that the underlying predictability mechanism in question is simply attributable to the media's reflection of the contemporaneous flow of fundamental accounting news over time – in fact, the more of this fundamental information that we are able to filter out (e.g., as with the WSJ), the better the resulting residuals seem to do. Second, my consistent finding that the information residuals perform better at predicting the returns of smaller stocks is suggestive of a relationship to the observation of differing qualities in the information environments experienced by small and large firms.

2.7 Interpreting the Results

Economic significance

While we have seen evidence that the media residuals Granger-cause index returns with respect to small stocks on the NYSE, it is unclear at first glance whether these predicted returns might be economically significant (e.g., would a potential trading strategy based upon this observed statistical relationship be profitable after trading costs?). Following Tetlock (2007), one simple way to look at this is to calculate the predicted change in returns from a one-standard deviation innovation to the media information residuals, utilizing the coefficients reported in Tables 2.5 and 2.6.

The sample standard deviations for the WSJ, FT, and NYT residuals are .129556, .0934678, and .1236826, respectively. Back-of-the-envelope calculations for Table 2.5, therefore, yield "predicted" weekly bottom-decile returns of approximately 12 bps and 22 bps from the WSJ and FT residuals respectively.[12] These hypothetical returns would indeed seem to have potential economic significance: 12 bps in weekly returns corresponds to almost 6.4% on an annualized basis. Of course, we should be very cautious in projecting these results: while the numbers are certainly somewhat suggestive, this is by no means an appropriate (e.g., out-of-sample, ex-trading costs, risk-adjusted, etc.) test of simulated trading strategy profitability.

What about market sentiment?

As mentioned earlier, the finding that unexplained media coverage seems to Granger-cause stock returns is potentially consistent with investor sentiment as well as the PMB/LIA explanation. In particular, Tetlock (2007) finds that his measure of media sentiment has predictive power in this respect as well. While I have taken a distinct approach in attempting to identify potential distortions in media coverage relative to fundamental information flows, a market sentiment-based explanation is also potentially plausible here as well (e.g., if reporters' potential coverage biases are merely further expressions of wider market sentiment). However, while other empirical studies have found that the causal relationship between sentiment and returns is bi-directional (e.g., Tetlock (2007) or Brown and Cliff (2004)), I do not find that past returns are able to predict the media information residuals. The implication here is that the media-based information distortions may not be so simply explained as a reflection of market sentiment as it is commonly described in the literature.

2.8 Conclusion

While classical finance theory suggests that security prices should quickly incorporate all available public information, evidence has begun to accumulate that the *manner* in which, and the *degree* to which, information is transmitted to the investing public can have a significant impact upon price formation. This paper investigates the potential for distortionary biases in media coverage related to corporate earnings announcements by controlling for the contemporaneous weekly flow of fundamental earnings information. I show that publication-specific media information residuals are serially correlated, suggesting that revealed publication

[12] Note that most of the implied difference in predictive power between the WSJ and FT is due to the mismatched data sample. When the return-WSJ regressions are estimated with respect to the period from 1990 on, the resulting coefficients are much closer to the FT values.

decision policies are persistent and at least somewhat predictable over time – a phenomenon that is consistent with theoretical definitions of media bias in the economics literature.

A priori, however, it is not clear that predictable patterns in estimated coverage decision residuals (i.e., presumed evidence of "media bias") should have any particular impact upon market behavior. After all, given the empirical nature of the study, an apparent pattern in "abnormal" coverage might simply be a statistical artifact, devoid of any true economic significance. Similarly, even if identified distortions in the information environment are "real" in some sense, it is plausible that at least some market participants should be able to see through, and correct for, any observed bias in media coverage – under classical assumptions, all it takes is one sufficiently large, sophisticated investor to push the market price to its "true value".

On the contrary, however, I find that my estimated information residuals Granger-cause small stock index returns (after controlling for previously identified return anomalies), suggesting that media publication decisions do have an effect on market prices, in and of themselves; this finding is particularly important when we recognize that reporters' selective coverage decisions appear to contain a predictable component. Taken together, the findings presented here may serve to underscore the potential importance of information-based factors for our understanding of efficiency in financial markets.

Figure 2.1: Earnings news articles in the WSJ

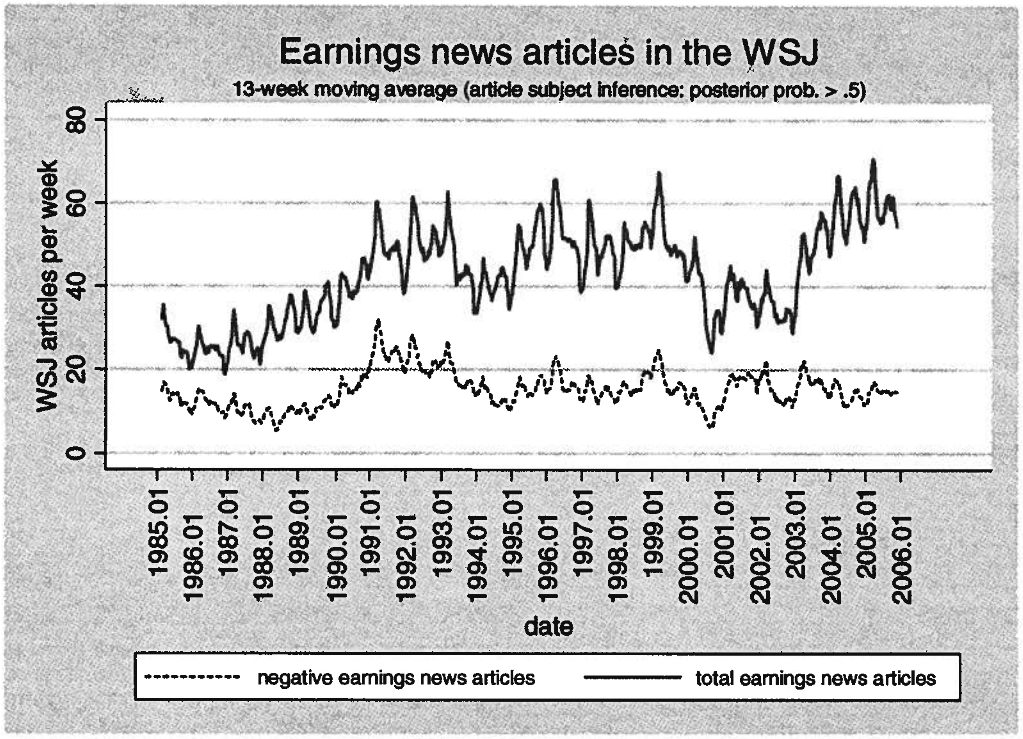

Figure 2.2: Earnings news articles in the FT

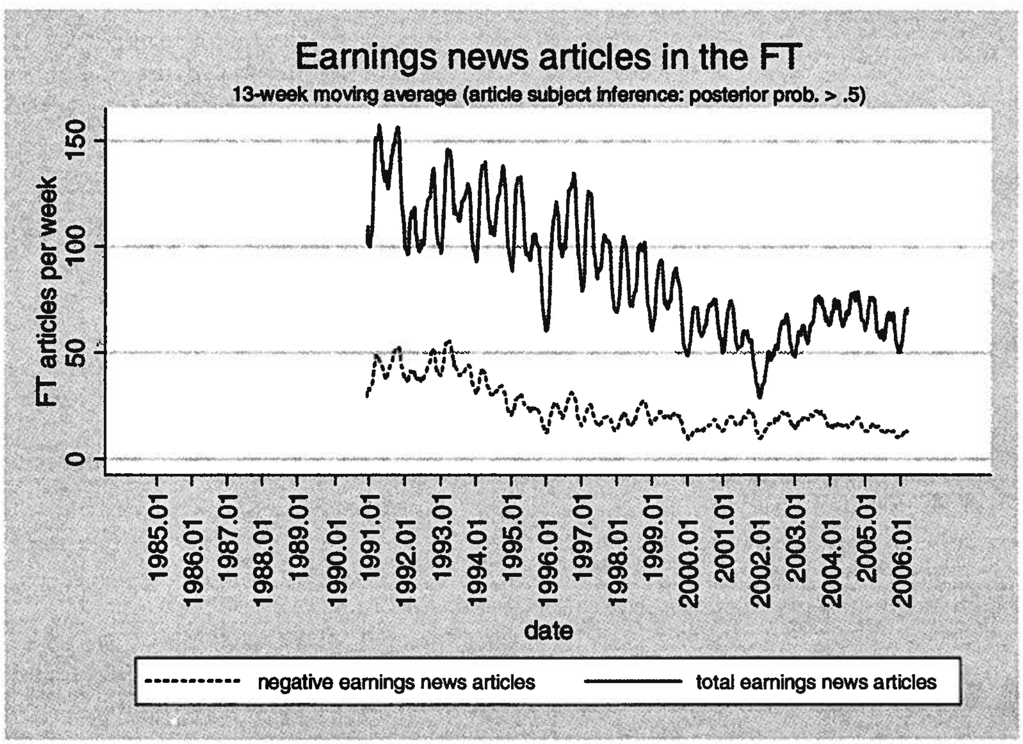

Figure 2.3: Earnings news articles in the NYT

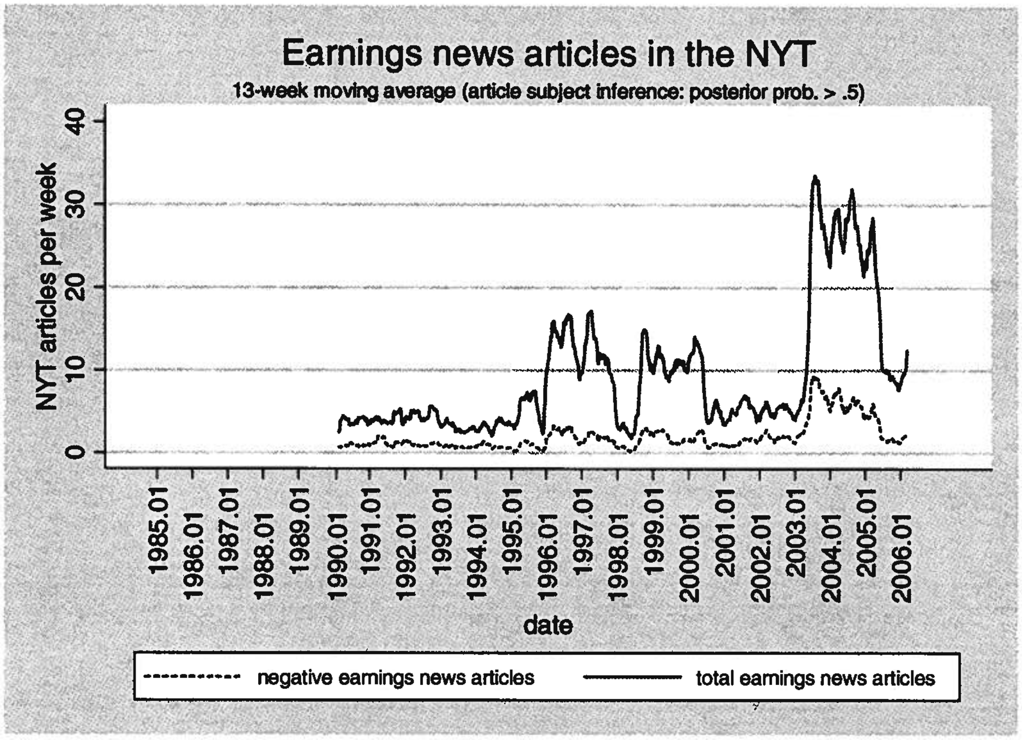

Figure 2.4: Earnings announcements in I/B/E/S

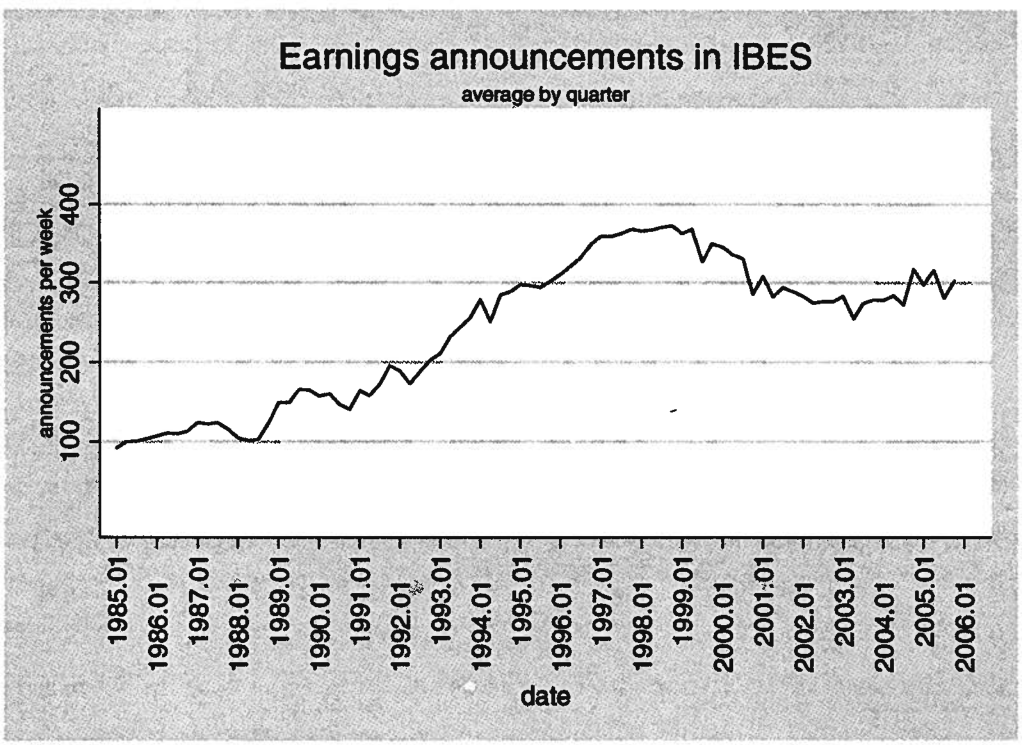

Figure 2.5: I/B/E/S earnings news – negative EPS

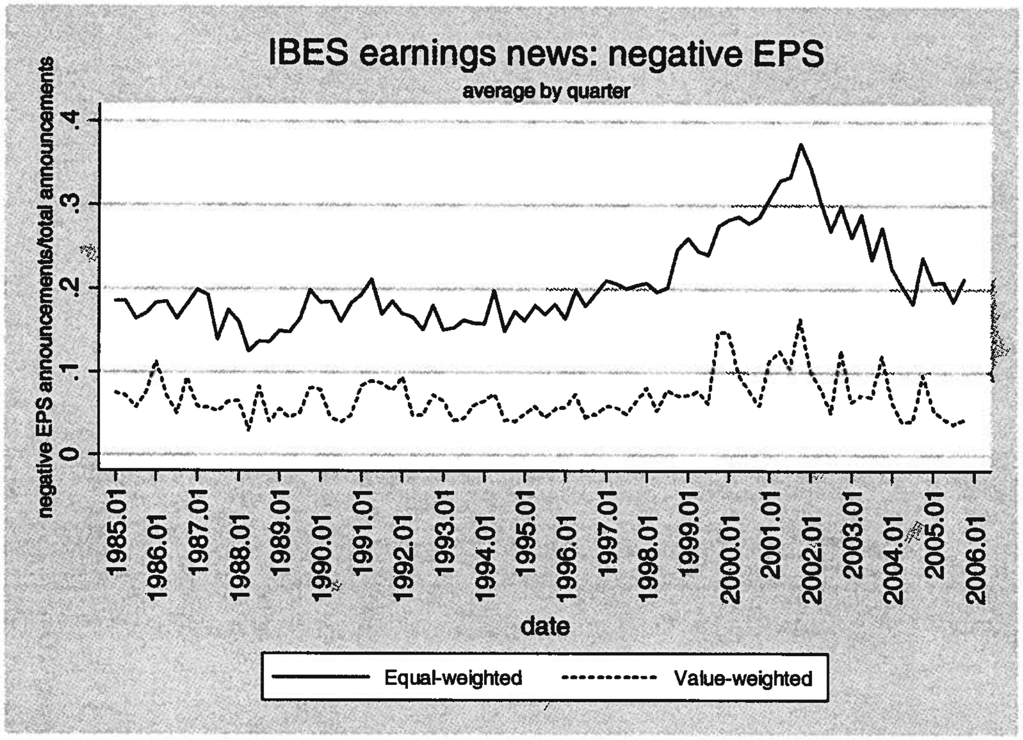

Figure 2.6: I/B/E/S earnings news – negative surprises

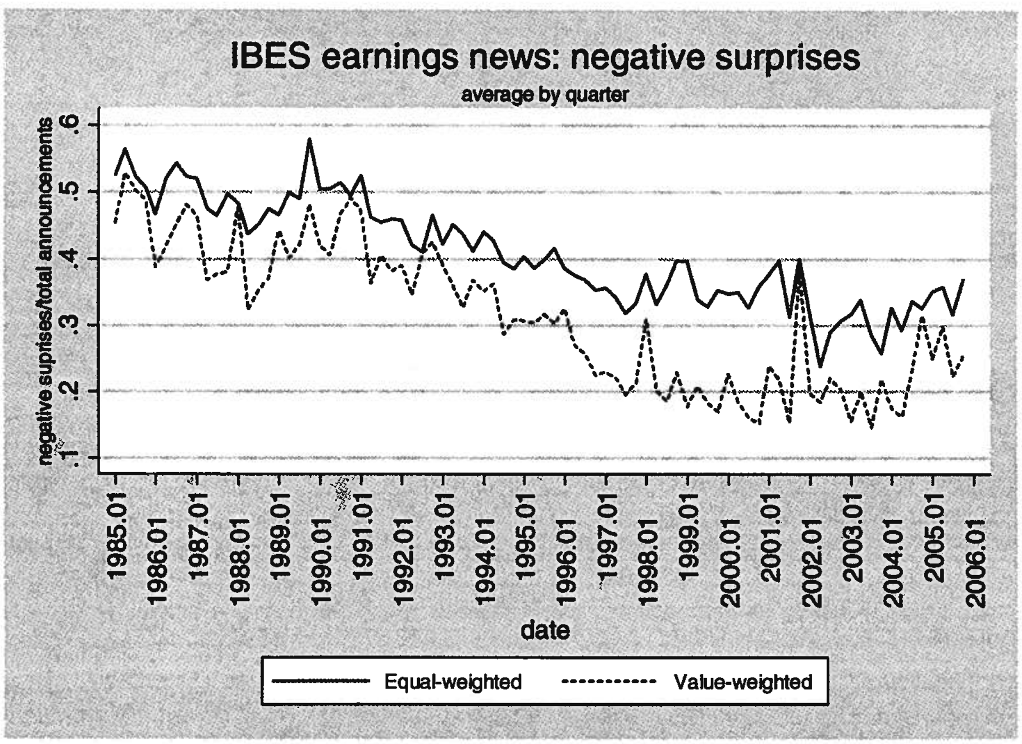

Figure 2.7: WSJ residual negative earnings news coverage

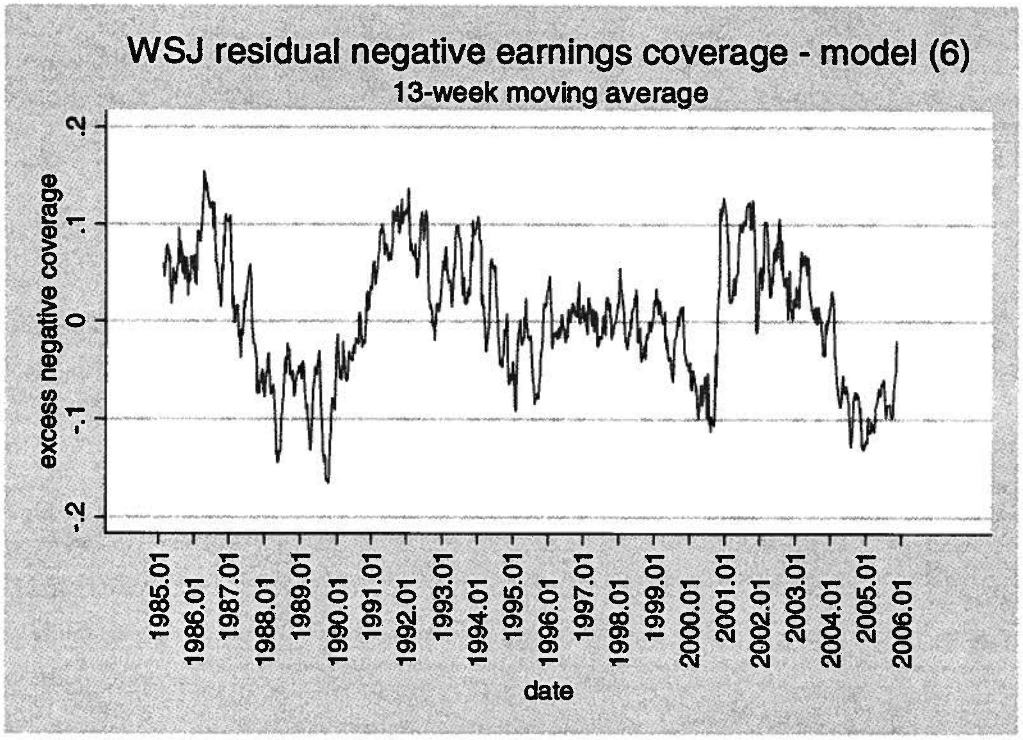

Figure 2.8: FT residual negative earnings news coverage

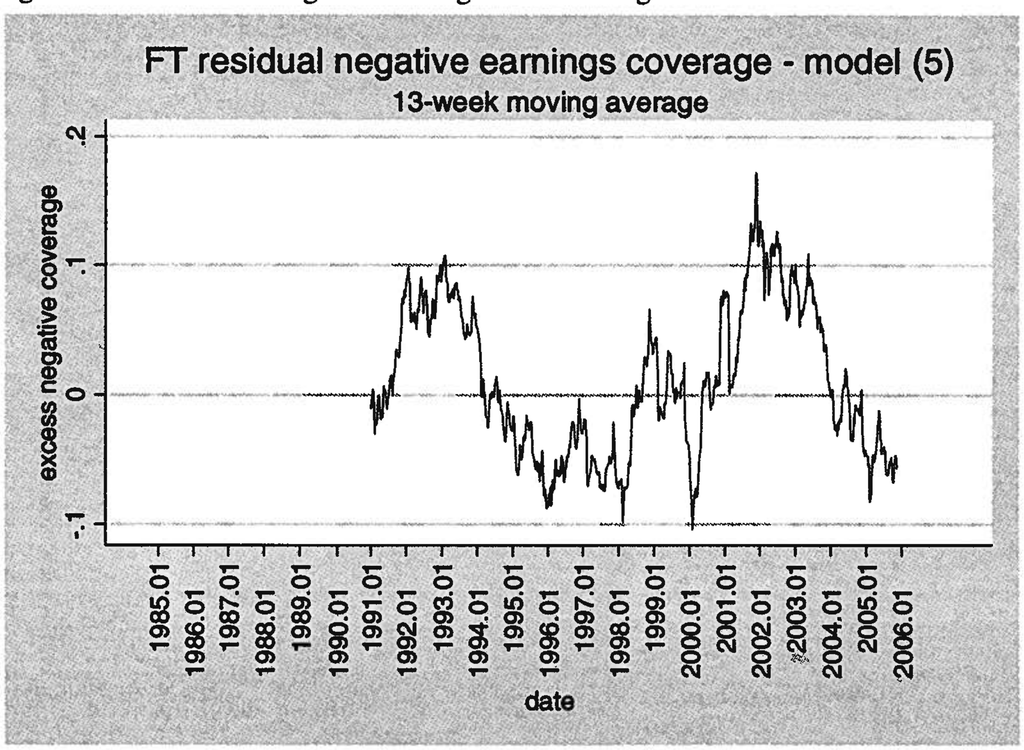

Figure 2.9: NYT residual negative earnings news coverage

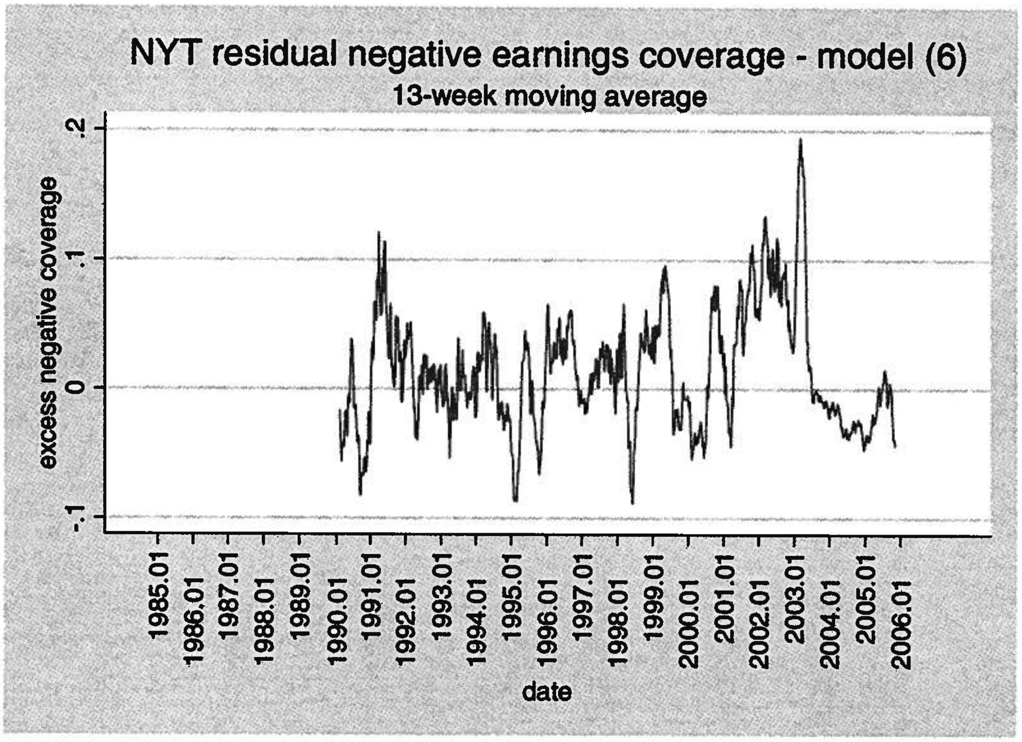

Table 2.1: Intra-week seasonality in article publications and announcements

		Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	All
WSJ	Earnings articles (total by day-of-week)	-	5,335	7,784	11,553	10,836	13,479	126	49,113
	Negative earnings articles (% of total)	-	46.86%	35.17%	33.23%	33.55%	33.61%	29.37%	35.18%
NYT	Earnings articles (total by day-of-week)	29	39	1,354	1,819	1,842	1,904	842	7,829
	Negative earnings articles (% of total)	31.03%	23.08%	19.79%	18.20%	20.85%	23.11%	29.45%	21.57%
FT	Earnings articles (#) (total by day-of-week)	24	2,359	11,402	16,301	16,683	18,442	5,624	70,835
	Negative earnings articles (% of total)	41.67%	32.09%	26.16%	25.77%	26.91%	27.38%	35.05%	27.47%
IBES	Earnings announcements (#) (total by day-of-week)	-	39,403	64,460	64,649	65,492	31,319	-	265,323
	Negative announcements (% of total) (EPS<0)	-	19.57%	18.26%	18.74%	20.27%	21.29%	-	19.43%
	Negative surprises (% of total) (EPS–e(EPS) <0)	-	37.72%	34.65%	35.23%	34.75%	45.53%	-	36.55%

Table 2.2: Regressions of WSJ proportional negative earnings news on I/B/E/S information variables

Note: columns show results from weekly logit regressions of WSJ_t^{-ve} on contemporaneous and lagged values of the I/B/E/S proportional news measures; 1090 weekly obs. from January 1985 to December 2005; Newey-West error correction up to 5 lags; ()* and ()** denote significance at 5% and 1% respectively; the binomial denominator is the total weekly number of identified WSJ earnings articles (prob > .5); $EARN_t^{COUNT,EPS<0}$ is the equal-weighted proportion of I/B/E/S announcements in week t with negative earnings; $EARN_t^{COUNT,EPS-e(EPS)<0}$ is the equal-weighted proportion of I/B/E/S announcements in week t with earnings less than the median analyst forecast during the previous month; $EARN_t^{MKTVAL,EPS<0}$ is the equity market value-weighted proportion of I/B/E/S announcements in week t with negative earnings; $EARN_t^{MKTVAL,EPS-e(EPS)<0}$ is the equity market value-weighted proportion of I/B/E/S announcements in week t with earnings less than the median analyst forecast during the previous month

	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*
$EARN_t^{COUNT,EPS<0}$	**1.640389** **(7.22)****		**1.370656** **(5.01)****	**1.702798** **(6.42)****	**1.656305** **(6.12)****	**3.817826** **(5.31)****
$EARN_{t-1}^{COUNT,EPS<0}$				-.1593196 (-0.59)		
$EARN_{t-2}^{COUNT,EPS<0}$				-.045733 (-0.19)		
$(EARN_t^{COUNT,EPS<0})^2$						**-4.679074** **(-3.31)****
$EARN_t^{COUNT,EPS-e(EPS)<0}$	**2.352448** **(9.62)****		**1.766361** **(5.78)****	**.7189133** **(2.46)***	**.7596573** **(2.46)***	**.7492063** **(2.48)***
$EARN_{t-1}^{COUNT,EPS-e(EPS)<0}$				**.9407907** **(3.64)****	**.7988183** **(3.38)****	**.7906165** **(3.37)****
$EARN_{t-2}^{COUNT,EPS-e(EPS)<0}$				-.1382251 (-0.56)		
$EARN_t^{MKTVAL,EPS<0}$		**3.082235** **(8.49)****	**1.234702** **(3.51)****	**1.222538** **(3.61)****	**1.237688** **(3.61)****	**1.376174** **(4.17)****
$EARN_{t-1}^{MKTVAL,EPS<0}$				.2728632 (0.90)		
$EARN_{t-2}^{MKTVAL,EPS<0}$				-.1068298 (-0.38)		
$EARN_t^{MKTVAL,EPS-e(EPS)<0}$		**1.103567** **(8.07)****	**.4857274** **(3.25)****	**.4378038** **(3.15)****	**.4457052** **(3.24)****	**.4125812** **(3.05)****
$EARN_{t-1}^{MKTVAL,EPS-e(EPS)<0}$				.0854624 (0.63)	.1286819 (0.98)	.1190692 (0.90)
$EARN_{t-2}^{MKTVAL,EPS-e(EPS)<0}$				**.4190474** **(3.00)****	**.3715032** **(3.04)****	**.4052816** **(3.37)****
Constant	-1.839402 (-20.53)**	-1.120327 (-21.12)**	-1.780177 (-20.40)**	-1.849769 (-16.47)**	-1.885386 (-19.08)**	-2.08693 (-18.45)**
*Adj-R^2**	0.3097	0.2183	0.3319	0.3543	0.3526	0.3652
AIC	5.874885	6.135071	5.811574	5.744985	5.740221	5.704378

Table 2.3: Regressions of NYT proportional negative earnings news on I/B/E/S information variables

Note: columns show results from weekly logit regressions of NYT_t^{-ve} on contemporaneous and lagged values of the I/B/E/S proportional news measures (i.e., $EARN_t^{COUNT,EPS<0}$, etc.); 818 weekly obs. from January 1990 to December 2005 (13 missing values in NYT); Newey-West error correction up to 5 lags; ()* and ()** denote significance at 5% and 1% respectively; the binomial denominator is the total weekly number of NYT earnings articles (Factiva C151); $EARN_t^{COUNT,EPS<0}$ is the equal-weighted proportion of I/B/E/S announcements in week t with negative earnings; $EARN_t^{COUNT,EPS-e(EPS)<0}$ is the equal-weighted proportion of I/B/E/S announcements in week t with earnings less than the median analyst forecast during the previous month; $EARN_t^{MKTVAL,EPS<0}$ is the equity market value-weighted proportion of I/B/E/S announcements in week t with negative earnings; $EARN_t^{MKTVAL,EPS-e(EPS)<0}$ is the equity market value-weighted proportion of I/B/E/S announcements in week t with earnings less than the median analyst forecast during the previous month

	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*
$EARN_t^{COUNT,EPS<0}$	.4285062 (1.02)		.6767222 (1.12)	1.210472 (1.74)		
$EARN_{t-1}^{COUNT,EPS<0}$				.0195123 (0.03)		
$EARN_{t-2}^{COUNT,EPS<0}$				**-1.521433** **(-2.47)***		
$EARN_t^{COUNT,EPS-e(EPS)<0}$	**1.811969** **(3.38)****		**1.556902** **(2.70)****	**1.482814** **(2.21)***	**1.968835** **(4.17)****	**1.950716** **(4.18)****
$EARN_{t-1}^{COUNT,EPS-e(EPS)<0}$				.4165696 (0.90)		
$EARN_{t-2}^{COUNT,EPS-e(EPS)<0}$				-.113984 (-0.20)		
$EARN_t^{MKTVAL,EPS<0}$		.1394221 (0.31)	-.3577259 (-0.64)	-.6413761 (-1.18)		
$EARN_{t-1}^{MKTVAL,EPS<0}$				-.4453578 (-0.71)		
$EARN_{t-2}^{MKTVAL,EPS<0}$				.3272031 (0.49)		
$(EARN_t^{MKTVAL,EPS<0})^2$						**-.6854775** **(-2.61)****
$EARN_t^{MKTVAL,EPS-e(EPS)<0}$		**.7513807** **(2.15)***	.2673279 (0.80)	.2370111 (0.75)		
$EARN_{t-1}^{MKTVAL,EPS-e(EPS)<0}$				-.3399442 (-1.14)		
$EARN_{t-2}^{MKTVAL,EPS-e(EPS)<0}$				.0287525 (0.09)		
Constant	-2.945249 (-16.83)**	-2.416328 (-25.64)**	-2.953303 (-16.81)**	-2.727607 (-12.44)**	-2.903743 (-16.38)**	-2.889718 (-16.62)**
N	818	818	818	816	818	818
*Adj-R^2**	0.0605	0.0269	0.0761	0.1125	0.0257	0.0267
AIC	3.244742	3.285521	3.247326	3.231725	3.245776	3.245232

Table 2.4: Regressions of FT proportional negative earnings news on I/B/E/S information variables

Note: columns show results from weekly logit regressions of FT_t^{-ve} on contemporaneous and lagged values of the I/B/E/S proportional news measures (i.e., $EARN_t^{COUNT,EPS<0}$, etc.); 789 weekly obs. from November 1990 to December 2005; Newey-West error correction up to 5 lags; ()* and ()** denote significance at 5% and 1% respectively; the binomial denominator is the total number of identified FT earnings articles (prob. > .5); $EARN_t^{COUNT,EPS<0}$ is the equal-weighted proportion of I/B/E/S announcements in week t with negative earnings; $EARN_t^{COUNT,EPS-e(EPS)<0}$ is the equal-weighted proportion of I/B/E/S announcements in week t with earnings less than the median analyst forecast during the previous month; $EARN_t^{MKTVAL,EPS<0}$ is the equity market value-weighted proportion of I/B/E/S announcements in week t with negative earnings; $EARN_t^{MKTVAL,EPS-e(EPS)<0}$ is the equity market value-weighted proportion of I/B/E/S announcements in week t with earnings less than the median analyst forecast during the previous month

	(1)	*(2)*	*(3)*	*(4)*	*(5)*
$EARN_t^{COUNT,EPS<0}$	**-.5711793** **(-2.07)***		-.4776241 (-1.61)	.0329217 (0.13)	
$EARN_{t-1}^{COUNT,EPS<0}$				-.2988414 (-1.20)	
$EARN_{t-2}^{COUNT,EPS<0}$				-.1967758 (-0.85)	
$EARN_t^{COUNT,EPS-e(EPS)<0}$	**1.267085** **(4.77)****		**.5931962** **(1.98)***	.0801171 (0.29)	
$EARN_{t-1}^{COUNT,EPS-e(EPS)<0}$				.0517574 (0.26)	
$EARN_{t-2}^{COUNT,EPS-e(EPS)<0}$				.2023058 (0.88)	
$EARN_t^{MKTVAL,EPS<0}$		.0813034 (0.30)	.3696128 (1.31)	.2610379 (0.93)	
$EARN_{t-1}^{MKTVAL,EPS<0}$				.0001718 (0.00)	
$EARN_{t-2}^{MKTVAL,EPS<0}$				-.0817912 (-0.33)	
$EARN_t^{MKTVAL,EPS-e(EPS)<0}$		**.8752747** **(6.89)****	**.639781** **(4.54)****	**.5713415** **(4.41)****	**.6485769** **(5.98)****
$EARN_{t-1}^{MKTVAL,EPS-e(EPS)<0}$				.1613865 1.31)	.1438762 (1.76)
$EARN_{t-2}^{MKTVAL,EPS-e(EPS)<0}$				**.38655** **(3.19)****	**.351953** **(3.99)****
$EARN_{t-3}^{MKTVAL,EPS-e(EPS)<0}$					**.3954751** **(4.12)****
Constant	-1.33809 (-15.31)**	-1.227063 (-28.97)	-1.308209 (-15.33)	-1.33928 (-11.37)	-1.413863 (-24.18)**
N	789	789	789	787	786
*Adj-R^2**	0.0737	0.1005	0.1138	0.1465	0.1556
AIC	7.380938	7.298827	7.264395	7.191203	7.141614

Table 2.5: Results from predictive regressions of stock market index returns on lagged WSJ, NYT, or FT news residuals (one lag)

Note: results from weekly regressions of log NYSE index returns on 1) one of the lagged news residuals, 2) lagged log index returns, 3) the exogenous control variables (lagged detrended log NYSE aggregate trading volume, lagged volatility proxies, and 19 October 1987 and monthly dummy variables), and 4) a constant, estimated with SAARCH(1,1,1) errors (regression results for the lagged returns, exogenous variables, constant, and the conditional volatility terms not shown – available upon request); z-stats in parentheses; ()* and ()** denote significance at 5% and 1% respectively; weekly (Friday-close to Friday-close) observations from January 1985 to December 2005 for the WSJ (1089 weeks), from January 1990 to December 2005 for the NYT (817 weeks, with 8 gaps), and from November 1990 to December 2005 for the FT (785 weeks); *w* is the (WSJ) residual from regression (6) of Table 2.2; *n* is the (NYT) residual from regression (6) of Table 2.3; *f* is the (FT) residual from regression (5) of Table 2.4; R_t is the log index return for week *t*

	(1) NYSE 1st Decile R_t	*(2)* NYSE 2nd Decile R_t	*(3)* NYSE 3rd Decile R_t	*(4)* NYSE 4th Decile R_t	*(5)* NYSE 5th Decile R_t	*(6)* NYSE 6th Decile R_t	*(7)* NYSE 7th Decile R_t	*(8)* NYSE 8th Decile R_t	*(9)* NYSE 9th Decile R_t	*(10)* NYSE 10th Decile R_t	*(11)* NYSE Equal-weighted R_t	*(12)* NYSE Value-weighted $R(t)$
w_{t-1}	**0.0092**	**0.0089**	**0.0064**	**0.0083**	0.0059	**0.0076**	0.0073	0.0049	0.0049	0.0018	**0.0074**	0.0026
	(3.38)**	**(3.36)****	**(2.17)***	**(2.44)***	(1.63)	**(2.01)***	(1.83)	(1.21)	(1.14)	(0.44)	**(2.16)***	(0.63)
n_{t-1}	0.0083	0.0063	0.004	0.0052	0.007	0.0094	0.0065	0.004	0.0026	0.0009	0.0042	0.0011
	(1.84)	(1.66)	(0.91)	(1.18)	(1.37)	(1.85)	(1.24)	(0.76)	(0.48)	(0.17)	(0.91)	(0.22)
f_{t-1}	**0.0238**	**0.0137**	**0.0117**	0.0102	0.0037	0.0036	0.0032	-0.0001	-0.0019	**-0.015**	0.0072	-0.0116
	(5.43)**	**(3.05)****	**(2.15)***	(1.88)	(0.54)	(0.53)	(0.46)	(-0.02)	(-0.25)	**(2.11)***	(1.21)	(-1.68)

Table 2.6: Results from predictive regressions of stock market index returns on lagged WSJ, NYT, or FT news residuals (two lags)

Note: weekly time-series regressions of log NYSE index returns on 1) lagged news residuals, 2) lagged log index returns, 3) exogenous control variables (lagged detrended log NYSE aggregate trading volume, lagged volatility proxies, and 19 October 1987 and monthly dummy variables – results not shown), and 4) a constant, with SAARCH(1,1,1) errors (regression results for the exogenous variables, the constant, and the conditional volatility terms not shown – available upon request); z-stats in parentheses; * and ** denote significance at 5% and 1% respectively; weekly (Friday-close to Friday-close) observations from January 1985 to December 2005 for the WSJ (1088 weeks), from January 1990 to December 2005 for the NYT (816 weeks, with 8 gaps), and from November 1990 to December 2005 for the FT (784 weeks); X^2 statistics refer to (joint) tests of Granger non-causality; w is the residual from regression (6) of Table 2.2; n is the residual from regression (6) of Table 2.3; f is the residual from regression (5) of Table 2.4; R_t is the log index return for week t

	(1) NYSE 1st Decile R_t	*(2)* NYSE 2nd Decile R_t	*(3)* NYSE 3rd Decile R_t	*(4)* NYSE 4th Decile R_t	*(5)* NYSE 5th Decile R_t	*(6)* NYSE 6th Decile R_t	*(7)* NYSE 7th Decile R_t	*(8)* NYSE 8th Decile R_t	*(9)* NYSE 9th Decile R_t	*(10)* NYSE 10th Decile R_t	*(11)* NYSE Equal-weighted R_t	*(12)* NYSE Value-weighted $R(t)$
w_{t-1}	**0.0059 (2.17)***	**0.0073 (2.75)****	0.0045 (1.48)	0.0063 (1.85)	0.0039 (1.08)	0.0062 (1.62)	0.0059 (1.46)	0.0039 (0.93)	0.0038 (0.87)	0.0016 (0.38)	0.0057 (1.62)	0.0019 (0.46)
w_{t-2}	**0.0114 (4.26)****	**0.0059 (1.98)***	0.0054 (1.7)	0.0062 (1.83)	0.0049 (1.33)	0.004 (1.03)	0.0039 (0.93)	0.0029 (0.65)	0.002 (0.43)	-0.0008 (-0.17)	0.0048 (1.29)	0.0005 (0.11)
$X^2(2)$	**22.43****	**12.45****	**6.01***	**7.30***	3.26	4.39	3.47	1.60	1.09	0.15	4.86	0.25
prob > X^2	**0.0000**	**0.0020**	**0.0495**	**0.0260**	0.1956	0.1112	0.1766	0.4487	0.5787	0.9270	0.0882	0.8812
n_{t-1}	0.0065 (1.39)	0.0054 (1.35)	0.0033 (0.75)	0.0034 (0.78)	0.0058 (1.14)	0.0081 (1.61)	0.0056 (1.07)	0.0039 (0.73)	0.0024 (0.43)	0.0006 (0.11)	0.0036 (0.77)	0.0008 (0.16)
n_{t-2}	0.0065 (1.45)	0.0034 0.85)	0.0087 (1.96)	**0.0107 (2.50)***	**0.0115 (2.42)***	**0.01 (2.01)***	**0.0107 (2.07)***	0.0065 (1.22)	0.0064 (1.23)	0.0075 (1.35)	0.0078 (1.75)	0.008 (1.53)
$X^2(2)$	5.24	2.84	4.38	**7.44***	**6.93***	**6.87***	5.89	2.17	1.77	1.89	3.83	2.41
prob > X^2	0.0727	0.2415	0.1121	**0.0242**	**0.0313**	**0.0322**	0.0527	0.3375	0.4134	0.3896	0.1475	0.3002
f_{t-1}	**0.018 (3.48)****	**0.0137 (2.93)****	0.0084 (1.45)	0.0036 (0.64)	-0.0022 (-0.29)	-0.0034 (-0.43)	-0.0035 (-0.47)	-0.0033 (-0.44)	-0.0043 (-0.52)	-0.0128 (-1.64)	0.0029 (0.45)	-0.0105 (-1.4)
f_{t-2}	0.0097 (1.59)	0.0028 (0.54)	0.0056 (0.93)	**0.0132 (2.39)***	0.0119 (1.63)	**0.0154 (2.06)***	**0.0142 (1.98)***	0.0066 (0.94)	0.0034 (0.46)	-0.0049 (-0.63)	0.0085 (1.36)	-0.0023 (-0.31)
$X^2(2)$	**26.48****	**11.40****	4.49	**7.92***	2.80	4.66	3.99	0.90	0.34	5.07	2.91	3.01
prob > X^2	**0.0000**	**0.0033**	0.1058	**0.0191**	0.2466	0.0975	0.1358	0.6372	0.8423	0.0791	0.2340	0.2215

Table 2.7: Results from predictive regressions of WSJ news residuals on lagged NYSE index returns

Note: weekly regressions of WSJ news residuals (model 6) on 1) lagged WSJ residuals, 2) lagged log NYSE index returns, 3) exogenous control variables (lagged detrended log NYSE aggregate trading volume, lagged volatility proxies, and 19 October 1987 and monthly dummy variables), and 4) a constant, with ARCH(2) errors (regression results for the exogenous variables, the constant, and the conditional volatility terms not shown – available upon request); z-stats in parentheses; ()* and ()** denote significance at 5% and 1% respectively; weekly (Friday-close to Friday-close) observations from January 1985 to December 2005; w is the residual from regression (6) of Table 2.2; R_t is the log index return for week t

	(1) w_t	*(2)* w_t	*(3)* w_t	*(4)* w_t
w_{t-1}	**0.1848**	**0.1821**	**0.1841**	**0.1822**
	(7.01)**	**(6.81)****	**(6.92)****	**(6.80)****
w_{t-2}	**0.2854**	**0.2784**	**0.2827**	**0.2792**
	(8.65)**	**(8.51)****	**(8.53)****	**(8.52)****
NYSE 1st Size Decile R_{t-1}	-0.236			
	(-1)			
NYSE 1st Size Decile R_{t-2}	-0.3173			
	(-1.38)			
$X^2(2)$	3.29			
prob > X^2	0.1926			
NYSE 10th Size Decile R_{t-1}		-0.164		
		(-0.84)		
NYSE 10th Size Decile R_{t-2}		-0.1489		
		(-0.89)		
$X^2(2)$		1.55		
prob > X^2		0.4615		
NYSE Equal-weighted R_{t-1}			-0.162	
			(-0.68)	
NYSE Equal-weighted R_{t-2}			-0.278	
			(-1.22)	
$X^2(2)$			2.00	
prob > X^2			0.3672	
NYSE Value-weighted R_{t-1}				-0.1741
				(-0.85)
NYSE Value-weighted R_{t-2}				-0.1961
				(-1.11)
$X^2(2)$				2.01
prob > X^2				0.3668
N	1088	1088	1088	1088

Table 2.8: Results from predictive regressions of NYT news residuals on lagged NYSE index returns

Note: weekly regressions of NYT news residuals (model 6) on 1) lagged WSJ residuals, 2) lagged log NYSE index returns, 3) exogenous control variables (lagged detrended log NYSE aggregate trading volume, lagged volatility proxies, and 19 October 1987 and monthly dummy variables), and 4) a constant, with GARCH(1,1) errors (regression results for the exogenous variables, the constant, and the conditional volatility terms not shown – available upon request); z-stats in parentheses; ()* and ()** denote significance at 5% and 1% respectively; weekly (Friday-close to Friday-close) observations from January 1990 to December 2005; n is the residual from regression (6) of Table 2.3; R_t is the log index return for week t

	(1) n_t	*(2)* n_t	*(3)* n_t	*(4)* n_t
n_{t-1}	**0.146**	**0.147**	**0.1467**	**0.1469**
	(3.61)**	**(3.62)****	**(3.63)****	**(3.62)****
n_{t-2}	**0.0767**	**0.078**	**0.0784**	**0.0782**
	(2.02)*	**(2.04)***	**(2.08)***	**(2.05)***
NYSE 1st Size Decile R_{t-1}	0.1543			
	(0.69)			
NYSE 1st Size Decile R_{t-2}	0.0679			
	(0.32)			
$X^2(2)$	0.85			
prob > X^2	0.6536			
NYSE 10th Size Decile R_{t-1}		-0.0319		
		(-0.19)		
NYSE 10th Size Decile R_{t-2}		-0.0067		
		(-0.04)		
$X^2(2)$		0.04		
prob > X^2		0.9810		
NYSE Equal-weighted R_{t-1}			0.0947	
			(0.45)	
NYSE Equal-weighted R_{t-2}			-0.014	
			(-0.07)	
$X^2(2)$			0.20	
prob > X^2			0.9050	
NYSE Value-weighted R_{t-1}				-0.0121
				(-0.07)
NYSE Value-weighted R_{t-2}				-0.0137
				(-0.08)
$X^2(2)$				0.08
prob > X^2				0.9590
N	799	799	799	799

Table 2.9: Results from predictive regressions of FT news residuals on lagged NYSE index returns

Note: weekly regressions of log index returns on 1) lagged FT residuals, 2) lagged log index returns, 3) exogenous control variables (lagged detrended log NYSE aggregate trading volume, lagged volatility proxies, and 19 October 1987 and monthly dummy variables), and 4) a constant, with SAARCH(1,1,1) errors (regression results for the exogenous variables, the constant, and the conditional volatility terms not shown – available upon request); z-stats in parentheses; ()* and ()** denote significance at 5% and 1% respectively; weekly (Friday-close to Friday-close) observations from November 1990 to December 2005; f is the residual from regression (5) of Table 2.4; R_t is the log index return for week t

	(1)	*(2)*	*(3)*	*(4)*
	f_t	f_t	f_t	f_t
f_{t-1}	**0.384**	**0.4001**	**0.378**	**0.3992**
	(11.71)**	**(10.00)****	**(11.14)****	**(9.93)****
f_{t-2}	**0.2508**	**0.2032**	**0.2569**	**0.2045**
	(5.91)**	**(5.57)****	**(6.07)****	**(5.53)****
NYSE 1st Size Decile R_{t-1}	-0.2679			
	(-1.78)			
NYSE 1st Size Decile R_{t-2}	0.2465			
	(1.54)			
$X^2(2)$	4.72			
prob > X^2	0.0943			
NYSE 10th Size Decile R_{t-1}		-0.0166		
		(-0.12)		
NYSE 10th Size Decile R_{t-2}		-0.2452		
		(-1.87)		
$X^2(2)$		3.58		
prob > X^2		0.1670		
NYSE Equal-weighted R_{t-1}			-0.1401	
			(-1)	
NYSE Equal-weighted R_{t-2}			-0.0568	
			(-0.37)	
$X^2(2)$			1.18	
prob > X^2			0.5535	
NYSE Value-weighted R_{t-1}				-0.005
				(-0.03)
NYSE Value-weighted R_{t-2}				-0.2539
				(-1.84)
$X^2(2)$				3.52
prob > X^2				0.1721
N	784	784	784	784

References

American Society of Newspaper Editors, 1999, Examining our credibility: Perspectives of the public and the press, Reston, VA., www.asne.org.

Baker, M., and F. Wurgler, 2006, Investor sentiment and the cross-section of stock returns, *Journal of Finance* 61, 1645-1680.

Antweiler, W., and M. Frank, 2004, Is all that talk just noise? The information content of internet stock message boards, *Journal of Finance* 59, 1259-1294.

Barber, B. M., and T. Odean, 2008, All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors, *Review of Financial Studies* 21(2), 785-818.

Barberis, N., A. Shleifer, and R. Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307-343.

Baron, D. P., 2004, Persistent media bias, Working paper, Stanford University.

Bhattacharya, U., N. Galpin, R. Ray, and X. Yu, 2006, "The Role of the Media in the Internet IPO Bubble," *Journal of Financial and Quantitative Analysis*, forthcoming.

Brown, Gregory W., and Michael T. Cliff, 2004, Investor sentiment and the near-term stock market, *Jqurnal of Empirical Finance* 11, 1-27.

Busse, Jeffrey A., and T. Clifton Green, 2002, Market efficiency in real time, *Journal of Financial Economics* 65, 415-437.

Cheung, Y., and E. Fujii, 2001, A note on the power of money-output causality tests, *Oxford Bulletin of Economics and Statistics* 63 (2), 247-261.

Campbell, J., S. Grossman, and J. Wang, 1993, Trading Volume and Serial Correlation in Stock Returns, *Quarterly Journal of Economics* 108 (4), 905-939

DellaVigna, Stefano, and Ethan Kaplan, 2007, The Fox News Effect: Media Bias and Voting, *Quarterly Journal of Economics* (122), 1187-1234.

DellaVigna, S., and J. Pollet, 2008, Investor Inattention and Friday Earnings Announcements, *Journal of Finance,* forthcoming.

Dyck, A., and L. Zingales, 2003, The media and asset prices, Working paper, Harvard Business School and University of Chicago.

Gentzkow, M., and J. Shapiro, 2006, Media Bias and Reputation, *Journal of Political Economy*, 114(2).

Granger, C., 1969, Investigating causal relationships by econometric models and cross-spectral methods, *Econometrica* 37, 424-438.

Groseclose, Tim, and Jeffrey Milyo, 2005, A Measure of Media Bias, *Quarterly Journal of Economics* 120(4), 1191-1237.

Huberman, G., and T. Regev, 2001, Contagious speculation and a cure for cancer: A nonevent that made stock prices soar, *Journal of Finance* 56, 387-396.

Liew, V., and Chong, T., 2005, Autoregressive lag length selection criteria in the presence of ARCH errors, *Economics Bulletin* 3(19), 1-5.

Manning, C. D., and H. Schuetze, 1999, Foundations of Statistical Natural Language Processing, MIT Press, Cambridge, MA.

McCallum, Andrew Kachites, 1996, *Bow: A toolkit for statistical language modeling, text retrieval, classification and clustering.* http://www.cs.cmu.edu/ mccallum/bow, 1996.

Mullainathan, S., and A. Shleifer, 2005, The market for news, *American Economic Review* 95 (4), 1031-1053.

Neal, R., and S. Wheatley, 1998, Do Measures of Investor Sentiment Predict Returns?, *Journal of Financial and Quantitative Analysis* 33(4), 523-547.

Reuter, J., and E. Zitzewitz, 2006, Do ads influence editors? Advertising and bias in the financial media, *American Economic Review* 121(1), 197-227.

Tetlock, P., 2007, Giving content to investor sentiment: The role of the media in the stock market, *Journal of Finance* 62 (3), 1139-1168.

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.

CHAPTER III

Good News is No News: Asymmetric Inattention and the Neglected Firm Effect[13]

3.1 Introduction

In attempting to explain persistent findings of high risk-adjusted returns for "neglected", low-recognition stocks,[14] researchers have typically focused on the identification and estimation of additional sources of risks and/or frictions that may not be fully priced in the traditional asset pricing models. In particular, one influential line of research has pointed to problems in the information environments faced by these firms (e.g., Merton (1987), Easley et al. (2002)). In this context, the observation of excess returns for neglected/low-information stocks can be explained as compensation for increased costs of information acquisition, exacerbated parameter uncertainty, or heightened asymmetric information, etc. Recently, researchers have also begun to focus on a related set of information characteristics (e.g., analyst coverage, institutional ownership, trading volumes, media coverage, etc.) as proxies for investor attention in attempting to explain patterns of apparent underreaction (and overreaction) in stock returns (Hou et al. (2006), Barber and Odean (2008), Chan (2003), Brennan et al. (1993), Hong et al. (2000)). This paper creates a link between these two strands of the literature by identifying underreaction to positive news related to asymmetric investor attention (i.e., "negativity bias") as a potential explanation for an apparent neglected firm premium in stock returns.

Constrained attention effectively prevents investors from acquiring and processing all of the potentially-relevant information that might be available at any given point in time. But how

[13] A version of this chapter will be submitted for publication. Gaa, C., Good News is No News: Asymmetric Inattention and the Neglected Firm Effect.
[14] See, e.g., evidence regarding firm size and analyst coverage (e.g., Arbel and Strebel, 1982), cross-listing (Foerster and Karolyi, 1999), "delayed" firms (Hou and Moskowitz, 2005), and financial news media coverage (Fang and Peress, 2007).

do agents decide which information items are worthy of attention, and which others will, of necessity, be ignored? While theoretical models predict that limited investor attention can lead to predictability in asset prices (see Peng and Xiong (2006) or Huang and Liu (2007)), attention allocations are extremely difficult to observe in practice. As a result, our understanding of this potentially important area of economic decision-making has continued to represent something of a "black box" for researchers.

This paper examines financial news reporters' coverage decisions in order to identify the event-specific (as well as the firm-specific and market-wide) factors that predict whether a particular earnings news event will be communicated to investors via this highly important information channel. Utilizing ex ante predicted probability of media coverage (PMC) as a new measure of investors' attention allocations in this context, I find that portfolio strategies with long positions in low-PMC stocks generate excess returns of approximately 70 bps per month after controlling for the standard risk factors identified in the literature. Insofar as this finding is consistent with the existing evidence regarding premia for "neglected" and "delayed" firms, it should not be surprising. Upon closer examination of the event-specific determinants of coverage, however, an alternative explanation emerges. In particular, I find evidence of a significant "negativity bias" in media attention: Bad news is more likely to result in coverage than is good news regarding an otherwise-identical firm. Given recent empirical evidence that market prices systematically underreact to low-attention events, and to news from low-attention firms in general, asymmetric underreaction to positive news emerges as a potential alternative to the standard information friction- and risk-based explanations for the neglected firm effect. Consistent with the asymmetric underreaction hypothesis, I find that the observed excess returns to low-PMC portfolios are attributable to high returns for low-attention "good news" firms, while low-attention "bad news" firms appear to be efficiently priced.

A growing body of research indicates that the mainstream financial news media, in particular, represents an important source of information in financial markets.[15] One explanation simply refers to the relative costs of information acquisition and processing: information in the news media is cheap and is typically presented in way that it is quickly and easily understood by non-specialists. Another, potentially more interesting, interpretation for the significance of

[15] See, e.g., Huberman and Regev (2001), Busse and Green, (2002), Chan (2003), Dyck and Zingales (2003), Barber and Odean (2008), Bhattacharya et al. (2006), Tetlock (2007), Tetlock et al. (2008), Antweiler and Frank (2006), and Fang and Peress (2007).

media information points to the fact that news media outlets actively compete with each other in attempting to anticipate reader interest. Given investors' cognitive constraints, then, it seems natural that the financial news media would play a crucial role as an information intermediary, providing relatively cheap and easy access to a sub-set of news items that the typical consumer will, on average, find most useful and interesting. I argue that financial news media coverage decisions provide a crucial window for our understanding of the determinants of investor attention. Put simply, events and firms whose characteristics predict a greater likelihood of receiving financial news media coverage are expected to attract higher levels of attention from investors.

To the present day, researchers continue to find evidence of high unexplained returns related to various proxies for the quality of a firm's information environment: the neglected firm effect. In this context, the concept of neglect is potentially quite broad. Arbel and Strebel (1982) were among the first to identify such a premium by looking at the relationship between stock returns and the number of securities analysts following a firm. More recently, Hou and Moskowitz (2005) show that firms whose stock prices exhibit significant "delay" with respect to their adjustment to common information shocks subsequently experience high returns that cannot be explained by the standard set of risk factors identified in the literature. Looking at the potential role of media coverage, in particular, Fang and Peress (2007) find that firms with no recent media coverage significantly outperform those who have experienced relatively high levels of coverage. The authors of both of these papers suggest that the identified return premia are consistent with the effects of frictions in the information environments faced by these firms. More broadly, depending upon one's definition of neglect, investor recognition (Merton (1987), Shapiro (2002), Basak and Cuoco (1998)), information risk (Easley et al. (2002)), and illiquidity (Amihud and Mendelson (1986), Pastor and Stambaugh (2003)) are all potentially consistent with the observation of a neglected firm premium in the cross-section of stock returns.

While low attention (or neglect) has traditionally been viewed primarily as a potential source of additional friction and/or risk for which investors must be compensated, researchers have also begun to examine investor attention as a potential explanation for observed variation in the speed with which different firms' stock prices react to news events. For example, Huang and Liu (2007) show that inattention to public news is rational when information acquisition is costly, potentially leading to over- or under-investment in portfolio selection when signals are noisy. If limits to arbitrage are sufficient to prevent prices from quickly adjusting to eliminate mispricing, then there is scope such phenomena to affect market prices, at least in the short term (Delong et al. (1990)). Consistent with delayed price reactions to news from neglected firms, Brennan et al. (1993) and Hong, Lim, and Stein (2000) (HLS) provide evidence that the stock prices of firms with low analyst coverage appear to respond more sluggishly to information shocks.

In this context, earnings announcements provide a particularly good testing ground for a study investigating the relationship between investor attention and potential underreaction in the market's response to an underlying news event. Going back as far as Ball and Brown (1968), post-earnings-announcement drift (PEAD) – the tendency for stock returns to exhibit continuation following an earnings surprise – has represented a significant puzzle for researchers. Bernard and Thomas (1990) ask whether PEAD represents a risk premium or a delayed response (i.e., underreaction) to earnings news; they find that their results are more consistent with the latter interpretation. If PEAD can be explained as an underreaction to earnings news, it is natural to ask whether, as one might expect, the phenomenon is stronger when investor attention is expected to be low. Indeed, there have been several studies linking the extent and/or speed with which prices react to earnings news to a number of potential proxies for attention, such as: firm size (Bamber (1987), Christensen et al. (2004)), trading volumes and overall market conditions (Hou et al. (2006)), analyst coverage (Christensen et al. (2004)), the presence of competing news events (DellaVigna and Pollet (2008), Hirshleifer et al. (2006)), and the ease with which earnings information may be processed (Engelberg (2007)).

In particular, there is recent research regarding the impact of media coverage on the market's reaction to earnings news. For example, Dyck and Zingales (2003) show that prices appear to react more strongly to the earnings numbers emphasized by reporters in their stories. Moreover, a large proportion of overall firm-specific media coverage seems to be linked to earnings news flows (Tetlock et al. (2008)), and they are often seen as the single most important piece of information regarding a firm's performance and future prospects. Tetlock (2007) shows that coverage surrounding earnings announcements contains significant incremental information regarding subsequent stock returns, while Tetlock et al. (2008) demonstrate the link to firms' fundamentals. Furthermore, Engelberg (2007) finds that the "soft", qualitative earnings information in news stories contains significant information regarding future returns, consistent with delayed reaction due to increased processing costs.

Given the aforementioned research, what is the contribution of this paper? While prior studies focus on the information content of earnings news coverage and its potential impact on stock returns, I examine financial reporters' story selection decision-making as an indicator of investor attention: Which earnings news events are most likely to be covered? In estimating these coverage decisions, I seek to identify the firm, event, and market-level characteristics that are associated with high levels of investor attention, and, conversely, "neglect". I demonstrate that negative news is more likely to be covered than positive news, ceteris paribus, implying that attention (and, therefore, the market's reaction) is potentially endogenous with respect to the information content of the underlying event itself. Finally, I study the potential impact of this negativity bias in the context of PEAD, presenting evidence that an apparent neglected firm premium in this context is actually more consistent with asymmetric underreaction to positive news from low-attention firms.

This paper's finding that attention varies with the information content of the underlying event is broadly consistent with evidence in the literature regarding asymmetric responses to positive and negative news (e.g., McQueen et al. (1996), Kothari et al. (2005), Veronesi (1999), Conrad et al. (2002), and Skinner and Sloan (2002)). In particular, researchers have documented an apparent negativity bias with respect to the media's coverage of macroeconomics news. Harrington (1989), for example, finds that the U.S. media generally pays greater attention to bad economic news, particularly in non-election years. More recently, Soroka (2006) shows that U.K news media and public opinion are more responsive to negative macroeconomic news releases than positive ones. For the first time, I present similar findings with respect to the financial news media's coverage of corporate earnings announcements.

But why would investors choose to allocate more attention to negative events? While this question is somewhat beyond the scope of the current paper, some discussion is warranted. Negativity bias, as a general phenomenon, is seen as consistent with prospect theory and loss aversion (Kahneman and Tversky (1979)). In particular, loss-averse agents are willing to expend more resources in avoiding a loss than in pursuing an equivalent-sized gain. If attention is a scarce resource, then it makes sense for loss-averse investors to expend relatively more attention in monitoring negative news, since negative events may be more likely to contain information regarding potential losses that may affect them. However, there are also other, more "rational" explanations that point to the underlying nature of the problem faced by investors in interpreting financial and accounting data in this context. In particular, if investors perceive that managers are reluctant to communicate bad news (due to career concerns, or in order to maximize the near-term value of their equity holdings in the firm, for example), it follows that markets would react more strongly to an unambiguously negative signal, inferring that it is more likely to be truthful than a positive one, on average. Also, if managers tend to engage in "big bath" accounting, then we might expect relatively infrequently-observed negative news announcements to elicit asymmetrically large responses from investors. For example, Kothari et al. (2005) argue that managers tend to withhold bad news; consistent with this, they find that price reactions to bad news disclosures are significantly larger than those precipitated by positive ones. Similarly, Skinner and Sloan (2002) argue that asymmetrically large reactions to negative earnings announcements for growth stocks are linked to periodic release of disappointing announcements in the face of baseline expectations that tend to be overoptimistic with respect to these firms.

The rest of the paper is organized as follows. Section 2 motivates and develops the methodological strategy employed in this study. Section 3 describes the data set. Section 4 investigates the firm- and event-specific determinants of media coverage. Section 5 describes the construction of my proposed measures of investor attention based on the predictable component of media coverage decisions. Section 6 examines the potential implications of biased inattention for the cross-section of stock returns and the market's reaction to news. Section 7 describes a series of robustness tests on the main predictability findings. Section 8 concludes.

3.2 Methodology and Hypothesis Development

In this section, I develop the methodological strategy of the paper. First, I apply tools from computational linguistics to identify a large data set of news articles that specifically discuss the results of corporate earnings announcements, and then subsequently to estimate the tone of coverage, if any. Second, linking observed instances of positive and negative media coverage to the underlying corporate information releases that precipitate them, I explain media coverage decisions in reference to event- and firm-specific story elements. Third, I employ probability of media coverage (PMC) as a proxy for investors' realized attention allocations, or, alternatively, as a measure of their underlying information consumption preferences. Fourth, investigating the potential impact of asymmetric attention on the cross-section of stock returns, I address the problem of potential omitted variables by utilizing ex ante estimated media coverage probabilities, rather than ex post realized observations of media coverage, to explain expected returns.

Determining the subject and tone of media coverage

A primary goal of the paper is to identify the underlying story elements that predict positive and negative media coverage with respect to specific corporate earnings announcements. In creating the media data set, I utilize tools from the field of computational linguistics in order to determine the subject and tone of potentially-relevant news media articles – specifically, I follow Antweiler and Frank (2004) in applying a Naïve-Bayesian "bag-of-words" approach using the Rainbow Toolkit (McCallum (1996)). A brief description follows (for a more detailed explanation, see Appendix A). In the first stage, a "training set" of 500 articles is randomly selected from the larger set and classified into semantic categories by hand. Next, the Rainbow text classification program builds an empirical model of category membership based on observed word frequencies for articles from the training set. ("Bag-of-words" refers to the fact that word order is not considered, only frequency of occurrence.) Each word that appears in a training set article is assigned an odds ratio based on its ability to predict membership in each semantic category. Finally, in order to classify text documents outside of the original training set, the program calculates sums of the odds ratios corresponding to the words found in each document, assigning documents to the category classification that maximizes the sum.

Which factors predict media coverage?

In this paper, I investigate the firm-specific and announcement-specific factors that predict media coverage, as well as the tone of coverage, if any. An essential requirement, therefore, is a large sample of information events with the potential to attract coverage. There must be variation in terms of coverage, and I need to observe the characteristics of events that do not receive coverage as well as those that do (i.e., related media coverage cannot be the only source of information regarding the event in question). Ideally, event characteristics should also be quantifiable and comparable across firms and over time. Since we are interested here in potential asymmetries in the attention paid to events of different kinds, the event in question should be reasonably classifiable as relatively "good" or "bad" news based on some metric that is, in particular, independent of the market's price reaction and the observed tone of coverage, if any.

In short, corporate earnings announcements provide a nearly-ideal setting for a study of this kind. Publicly-listed firms are typically required to produce quantitative accounting information on a quarterly basis, no matter whether that information might be viewed as good or bad, interesting or boring – this essentially eliminates the potential problem of endogenous information production and dissemination by firms. Earnings announcements also represent a very large set of information events, each one of which has the potential to attract media coverage ex ante (although, ex post, we know that most of them do not). Crucially, we can also interpret a particular firm's earnings release as being relatively positive or negative relative to other firms' announcements by, e.g., comparing its stated performance to analysts' pre-announcement forecasts, or simply by classifying the release as a "profit" or a "loss" in absolute terms. Therefore, in addition to non-event-specific predictors of firm coverage (e.g., as identified by Fang and Peress (2007) and Engelberg (2007)), I consider a set of factors related to the content of the information event itself. Furthermore, it is important to ensure that all of the proposed predictors of coverage represent information that was potentially available before the coverage decision was made; in particular, using contemporaneously-observed event-window CARs (rather than the earnings surprise relative to expectations) to predict coverage raises potential worries with respect to reverse-causality.

Given this data set of quarterly earnings announcements, which includes information on market conditions, the information content of the earnings releases, as well as characteristics of the firms making the announcements, I attempt to link each of the events to an identified earnings announcement-related news article in the Wall Street Journal (WSJ) regarding that firm. Approximately 13% of earnings announcements in the sample can be linked to an observation of related, contemporaneous media coverage.[16] I first examine the impact of event characteristics on the absolute probability of coverage by performing logit regressions with a dummy variable representing WSJ coverage on the left-hand-side and potential determinants of coverage on the right-hand-side. Subsequently, I investigate potential asymmetry regarding the

[16] The relatively selective nature of coverage here (elsewhere, e.g., Engelberg (2007) finds that approximately one-half of the announcements in his sample are covered by Dow Jones News Service) may reflect the fact that the WSJ stories used in this study are all "stand-alone" articles, written by reporters themselves, with the subjects presumably chosen with a closer eye to potential reader interest, while newswire stories are often summaries (or even reproductions) of company-produced press releases. Given the potential selection issues related to firm-originated news, this may also help to explain the relatively strong evidence of asymmetric coverage (i.e., negativity bias) presented in this paper.

impacts of story elements on the tone of coverage by differentiating observations of positive and negative coverage in a multinomial logit setting.

Media coverage as a proxy for investor attention and neglect

Having identified a set of factors that are significant in explaining variation in the probability and tenor of media coverage, I use predicted values from the multinomial logit estimation as a proxy for relative attention and neglect in the overall information environment. I then proceed to investigate the neglected firm effect in this context. While somewhat novel in its construction, I argue that this interpretation of neglect parallels and complements prior definitions based on characteristics such as trading activity, size, analyst coverage, etc.

Under what kinds of conditions might estimated media coverage represent an appropriate proxy for investor attention in a broader sense? First, if the Wall Street Journal (WSJ) is, in and of itself, a significant source of information for at least a sub-set of investors with constrained attention, then we would expect its coverage decisions to have a direct influence on investors' knowledge and beliefs regarding the events in question. What if, instead, the WSJ's individual effect is negligible, but the financial news media as a whole represents a significant source of information for investors? In this case, if media coverage decisions exhibit commonality across reporters and news organizations, then, given its status as an industry leader, we might expect the WSJ to serve as a relatively good proxy for media coverage patterns more widely. Under both of these interpretations, we can think of the news media as informationally important due to its role of selecting and presenting particular information items for its readers, affecting investors' marginal costs of information acquisition and processing through their publication decisions.

Finally, imagine an even more restrictive case: realized coverage decisions in the financial news media do not directly impact investors' information sets at all. In this case, media coverage will, nonetheless, be useful as an indicator of investor attention to the extent that news outlets are able to model the information consumption preferences of their customers. In other words, if reporters are successful in applying a "theory of mind" regarding their consumers' underlying information preferences in order to predict which news items their readers will find most interesting (and, in a competitive news market, we might expect that they would be, on average), then observed coverage decisions will still serve as a guide for a researcher hoping to identify the characteristics of firms and events that will tend to attract attention more broadly.

This paper departs from previous studies in its identification of neglected firms based not only on firm characteristics, but also upon the very nature of the underlying information events that do (or do not) attract attention with respect to each firm. In this context, if attention allocation is asymmetric with respect to event characteristics, and if attention affects the speed with which stock prices respond to new information, this creates the potential for a directional drift in neglected firm values that might otherwise be indistinguishable from a more standard, symmetric risk-based story. In order to test this hypothesis, I disambiguate and compare the returns of neglected "good news" firms versus those of neglected "bad news" firms – if, as is suggested in the prior literature, the neglected firm premium is due to a symmetrically-distributed risk factor (e.g., information risk), then we should expect to observe excess return premia for both types of firms.

3.3 Data

This section describes the data set, consisting of information regarding the content of earnings announcements themselves, the characteristics of the firms making the announcements, as well as measures of related news coverage in the Wall Street Journal for the sample period from October 1984 to December 2005.

Earnings announcements and analyst data

The base event sample includes all earnings announcements that appear in I/B/E/S from 1984 to 2005 where 1) there are at lease two analyst forecasts in the previous month, and 2) the absolute surprise relative to the median analyst forecast is strictly less than the announcement day stock price. The resulting sample comprises 263,627 quarterly earnings announcement observations, including the announcement date, the announced normalized EPS number (i.e., expressed as a percentage of the announcement-day stock price), the number of distinct analyst EPS forecasts in the preceding month, the standard deviation of analysts' normalized EPS forecasts, the median analyst's normalized EPS forecast over the preceding month, and the quarter end-date to which the announced earnings number pertains.

Media coverage

I focus on media coverage in the Wall Street Journal (WSJ), which is considered to be one of the world's dominant financial news outlets, reaching approximately 2.6 million paying print and online subscribers with an average household net worth of US$ 2.5 million.[17] The media sample consists of 68,102 potentially-relevant Wall Street Journal articles from 1984 to 2005 (*Factiva Intelligent Indexing* code c151: "Earnings") in text format. In a two-stage classification strategy, articles are first categorized as either "earnings-related" or "not earnings-related", and then subsequently as "negative" or "not negative" (for simplicity, referred to as "positive" hereafter) using the "bag-of-words" computational linguistics tool *Rainbow* (McCallum (1996)).

Articles identified by the first-stage classification model as pertaining to a quarterly earnings announcement with posterior probability > 0.5 (49,113 articles) are then subjected to the second-stage classification. Those articles with a calculated posterior probability > 0.5 of being a member of the "negative" category are recorded as "negative" (17,280 articles), and all others are recorded as "not negative" (31,833 articles).[18] (For a more detailed description, please refer to Appendix I.)

With respect to each earnings announcement in the I/B/E/S event sample, I search among the media article observations for an identified earnings-related article within one week of the announcement date. If at least one identified earnings article pertaining to that firm is observed within the one-week window, I record that event as having received positive or negative coverage, depending upon how that article is classified. As we might expect, most firm-events are not observed to receive coverage in the WSJ: approximately 9% of announcements are associated with an identified positive earnings story, while approximately 4% can be linked with a related negative story.

---

[17] http://www.dj.com/Products_Services/PrintPublishing/WSJ.htm

[18] While the "not negative" classification is referred to more simply as "positive" for the remainder of the paper, given the binary categorization scheme, the "not negative" category includes all earnings articles that would be classed as either "positive" or "neutral" under a ternary classification scheme. Given that double-negatives are relatively uncommon in standard English (e.g., one is unlikely to hear "Firm X *failed* to *disappoint* analysts' expectations" instead of "Firm X *exceeded* analysts' expectations"), this categorization method is expected to result in a somewhat sharper distinction between "positive" and "negative" overall semantic meanings. However, the inclusion of "neutral" articles under the *de facto* "positive" category implies that the "positive" vs. "negative" categorical distinction referred to herein should be interpreted in a relative rather than an absolute sense.

Firm characteristics

Using data from COMPUSTAT, I calculate market to book ratios for each firm relative to each earnings announcement date by dividing the previous December 31 market value of equity by the book value of equity for the fiscal year ending in the prior year. Firms with M/B less than zero are discarded from the sample. I also consider the percentage of institutional ownership of equity observed as of the previous December 31. Stock return and trading volume data for the 60 trading days prior to each announcement window (i.e., days -61 to -2) are obtained from CRSP, representing roughly the most recent quarter of daily data for each firm – average daily dollar-value of trading volume, average daily stock returns, and the standard deviation of daily returns are calculated for each of these pre-event windows. Stocks with a closing price less than $1 are discarded. Market size is shares outstanding multiplied by the stock price observed on day -2 relative to the announcement. Firms' industries are identified according to their "Fama-French 49" classifications (definitions available on Kenneth French's website). Data on institutional ownership are from Thomson Financial's CDA/Spectrum 13-F database.

After discarding firm-event observations due to missing values in CRSP, I/B/E/S, and/or COMPUSTAT, the final sample comprises 178,898 firm-events spanning a sample period from October 1984 to December 2005.

3.4 The Determinants of Media Coverage

In the first stage of the analysis, I investigate the market conditions and firm- and announcement-specific factors that predict absolute media coverage (i.e., an observation of either positive or negative coverage). I first consider factors based on firm characteristics. For example, we might predict larger firms to attract more attention – readers may be more interested in hearing about a firm with which they are already familiar, and a large firm will tend to have a greater number of people with a direct interest in the firms' prospects (e.g., employees, customers, suppliers, investors, etc.). Certain industries may also be favored, on average. Recent stock returns, analyst coverage, and trading activity are also considered.

A second set of potential predictors describe the information content of the earnings announcement itself. For example, if media attention is drawn to negative events (as we might expect to see if negative events are seen as relatively more sensational/interesting), then the announcement of an earnings number below analysts' expectations should result in a higher probability of coverage. Conversely, if media coverage is typically drawn to positive or "feel good" news events, then the observation of a loss should result in a lower probability of coverage, ceteris paribus. However, the appropriate specification to describe such a relationship is unclear. Is the size of the positive or negative surprise crucial, or does the media view such news in more categorical terms? I also investigate potential interactions – for example, might a loss combined with a negative surprise attract negative media coverage more reliably than the "sum of the parts"?

Finally, underlying economic and market-wide conditions may contribute to variation in coverage. I include recent returns on the S&P 500, and recent volatility in the S&P 500 to capture the potential effects of overall market conditions. For example, Veronesi (1999) presents a general equilibrium model wherein investors "overreact" to bad news in good times, and "underreact" to good news in bad times.

I apply a logit model to explain absolute coverage as follows:

$$\begin{aligned}\text{Prob}[\text{ABS(COVERAGE)}_{i,t} = 1] = \text{F}(&\text{SURPRISE}_{i,t}, \text{I\_NEGSURPRISE}_{i,t}, \text{I\_LOSS}_{i,t},\\ &\text{I\_NEGSURPISE}_{i,t}\cdot\text{I\_LOSS}_{i,t}, \text{ANALYSTS}_{i,t}, \text{STDEV(FORECASTS)}_{i,t}, \text{B/M}_{i,t},\\ &\text{RETURNS}_{-1,i,t}, \text{STDEV(RETURNS}_{-1})_{i,t}, \text{S\&P}_{-1,i,t}, \text{STDEV(S\&P}_{-1})_{i,t}, \text{MKTVALUE}_{i,t},\\ &\text{USFIRM}_{i,t}, \$\text{VOLUME}_{i,t}, \text{INDUSTRY}_{i}, \text{YEAR}_{i,t}, \text{MONTH}_{i,t}, \text{DAY}_{i,t}),\end{aligned}$$

where $\text{ABS(COVERAGE)}_{i,t}$ equals one if there is an identified WSJ article associated with the announcement, and zero otherwise; $\text{SURPRISE}_{i,t}$ is the announced EPS minus the median analyst forecast from the prior thirty days, divided by the stock price; $\text{I\_NEGSURPRISE}_{i,t}$ is a dummy variable equal to one if $\text{SURPRISE}_{i,t}$ is negative, and zero otherwise; $\text{I\_LOSS}_{i,t}$ is a dummy variable equal to one if announced EPS is negative, and zero otherwise; $\text{ANALYSTS}_{i,t}$ is the natural log of the number of distinct analyst forecasts observed in I/B/E/S in the 30 days preceding the announcement; $\text{STDEV(FORECASTS)}_{i,t}$ is the standard deviation of the normalized analyst EPS forecasts recorded in I/B/E/S in the 30 days preceding the announcement; $\text{MKTVALUE}_{i,t}$ is the natural log of the firm's market value of equity prior to the event window; $\text{RETURNS}_{i,t}$ is the firm's cumulative stock return over the 60 trading days prior to the event window; $\text{I\_USFIRM}_{i,t}$ is a dummy variable equal to one if the firm is identified as a U.S. firm in CRSP; $\$\text{VOLUME}_{i,t}$ is the natural log of the average value of daily stock trading over the 60 days preceding the announcement window; INDUSTRY_{i} represents a set of dummy variables for the Fama-French 49 industry classification; $\text{YEAR}_{i,t}$ represents the set of year dummy variables; and $\text{MONTH}_{i,t}$ and $\text{DAY}_{i,t}$ represent month-of-the-year and day-of-the-week dummy variables to control for potential seasonal effects (see, e.g., DellaVigna and Pollet (2008)).

Table 3.2 presents the results of a logit regression of ABS(COVERAGE)$_{i,t}$ on market conditions, firm and event characteristics, and the time and industry dummies. With respect to the event-specific determinants of unsigned media coverage, I find that losses and negative earnings surprises (relative to median analyst expectations) are more likely to attract media coverage. As mentioned earlier, this may be related to, for example, the media's often-cited propensity to focus on "sensational" and/or unexpected stories. Table 3.1 shows that accounting losses are relatively infrequently observed, so there may tend to be a certain degree of "surprise" attached to each such announcement. In particular, a firm typically attempts to present a positive (or, at the very least, an ambiguous) picture of its performance to the market. Thus, an unambiguously negative piece of news might naturally be framed as "surprising" relative to the communication lines that firms generally attempt to put forward. Note, however, that this increased probability of coverage regarding bad news is relative: negative news articles are still less frequently observed than non-negative ones in absolute terms (i.e., 4% versus 9%). Not surprisingly, I also find that large firms and those with high analyst coverage and high trading volumes are most likely to receive attention in the media.

Figure 3.1 illustrates the relationships between firm and event characteristics and probability of coverage as described in Table 3.2. Holding event characteristics constant, any news regarding large, important firms is more likely to receive coverage. On the other hand, holding firm characteristics constant, a negative event is more likely to receive media coverage than a positive one. The implication is that negative news events regarding large firms are most likely to receive media coverage, while positive news events involving small firms are least likely to receive coverage. In short, the results in Table 3.2 support the asymmetric attention hypothesis – this is a crucial finding for the analysis that follows.

Having examined the determinants of coverage in an absolute sense, I make use of the positive/negative categorical distinction. For example, the observation of a loss and/or a negative surprise might be expected to predict a negative article.

With respect to positive or negative coverage, I test the following relationship as a multinomial logit:

$$\text{Prob}[\text{COVERAGE}_{i,t} = j] = G_j(\text{SURPRISE}_{i,t}, \text{I\_NEGSURPRISE}_{i,t}, \text{I\_LOSS}_{i,t},$$
$$\text{I\_NEGSURPISE}\cdot\text{I\_LOSS}_{i,t}, \text{ANALYSTS}_{i,t}, \text{STDEV(FORECASTS)}_{i,t}, \text{B/M}_{i,t},$$

$RETURNS_{-1,i,t}$, $STDEV(RETURNS_{-1})_{i,t}$, $S\&P_{-1,i,t}$, $STDEV(S\&P_{-1})_{i,t}$, $MKTVALUE_{i,t}$, $USFIRM_{i,t}$, $\$VOLUME_{i,t}$, $INDUSTRY_i$, $YEAR_{i,t}$, $MONTH_{i,t}$, $DAY_{i,t}$),

j= -1, 0, 1

where $COVERAGE_{i,t}$ is equal to 1 if there is an identified positive earnings story related to the announcement, equal to -1 if there is an identified negative earnings story related to the announcement, and 0 otherwise.

Table 3.3 presents the results from the multinomial logit regressions. As expected, positive event characteristics (such as positive surprises relative to analysts' expectations) predict positive media coverage, and *vice versa*. Firm characteristics that would be expected to attract attention unambiguously, such as firm size, analyst coverage, and average daily dollar value of stock trading, likewise behave as expected. Somewhat more interesting is the observation that some firm characteristics appear to have asymmetric effects with respect to predicting positive or negative media coverage. For example, high analyst forecast dispersion predicts a greater likelihood of negative coverage but a lesser likelihood of positive coverage, perhaps reflecting the impact of increased uncertainty. Similarly, high realized stock volatility in the period preceding the announcement predicts negative coverage, and implies a smaller probability of positive coverage. Somewhat counter-intuitively, high recent stock returns appear to predict smaller probabilities of both positive and negative coverage. Pre-event market conditions also seem to predict coverage: recent S&P volatility seems to predict positive coverage while high recent S&P returns predict negative coverage.

It should be noted that the foregoing results on the determinants of coverage highlight a potential problem with attempting to estimate the impact of media coverage on contemporaneously-observed returns in an event study setting: many of the event characteristics that predict positive and negative coverage may also affect announcement returns. In other words, there is a very real possibility that an apparent relationship between, e.g., negative media coverage and negative stock returns could be due not to the impact of coverage, but, rather, to the fact that reporters tend to write negative stories about firms that report losses, fail to meet analysts' expectations, etc.

Discussion of results

One explanation regarding reader interest simply recognizes that news organizations are motivated to sell news, while news consumers are more likely to be interested in reading about events related to firms with which they are already familiar or have some personal interest. Firm size can be seen as a proxy for the number of customers, employees, suppliers, etc., that a firm possesses. Similarly, the average dollar value of recent stock trading should be related to the number of active investors in a stock and the intensity of their interest. Therefore, it is no surprise that these proxies are positively related to probability of coverage. Fang and Peress (2007) examine a comprehensive set of firm-related media stories which they attempt to match to the entire universe of stocks. Since they observe both when firms do and do not receive media coverage, they are able to identify firm characteristics that predict media coverage. In particular, they find that size, B/M, analyst coverage, and individual ownership are all associated with a higher probability of observing some news story with respect to that firm. At the same time, since their measure of coverage is not centered on identification of comparable events across firms, it is difficult to identify the event-specific (as opposed to firm-specific) characteristics that tend to attract attention. Engelberg (2007) presents similar evidence with respect to size and analyst coverage. Analogously, Kaniel et al. (2007) find that total net assets under management are positively related to the probability of media coverage for mutual funds.

With respect to the event-specific determinants of coverage, while a finding of negativity bias in media attention is consistent with Soroka (2006) and Harrington (1989)[19], it may be somewhat puzzling in light of results from Kaniel et al. (2007) and Engelberg (2007). For example, Kaniel et al. (2007) find that positive recent performance predicts higher levels of media coverage for mutual funds. What might account for this apparent difference in the event-specific determinants of coverage? While it is difficult to speculate, one potential explanation has to do with endogeneity in the production of the information used by reporters. If mutual funds are more likely to produce information (e.g., issue press releases) when performance has been good, and if independent sources of information regarding fund performance are relatively inaccessible, then this may naturally result in higher levels of coverage for these funds. Similarly, Engelberg (2007) presents mixed evidence that the probability of observing media coverage is positively related to contemporaneous event-window stock returns; in addition to the potential issue of simultaneity, the media coverage data in this setting (i.e., DJNS) may be more likely to include firm-originated reports, so it is possible that self-selection of positive

[19] The finding of negativity bias is also consistent, albeit somewhat indirectly, with broader evidence regarding asymmetrically small reactions to positive news and/or large reactions to negative news in, e.g., McQueen et al. (1996), Veronesi (1999), Kothari et al. (2005), Skinner and Sloan (2002), and Conrad et al. (2002).

news stories may be a factor here as well. However, the explanation for these seemingly anomalous results remains unclear – this is an area for further study.

3.5 Probability of Media Coverage (PMC)

In this section, I describe the construction of the PMC measures and discuss their distributional characteristics and statistical relationships with respect to other variables of interest.

Construction of PMC

Positive, negative, and absolute probability of media coverage with respect to each firm-event is calculated as the predicted probability of positive, negative, or any (i.e., positive or negative) coverage, respectively, from the multinomial logit regression specified in Table 3.3, column 5 (omitting the year dummies).

$$\begin{aligned} PMC_{i,t} &= \text{Prob}\{\text{Positive or negative media coverage} \mid X_{i,t}\} \\ &= 1 - \text{Prob}\{\text{COVERAGE}_{i,t} = 0 \mid X_{i,t}\} \\ &= PMC^{+}_{i,t} + PMC^{-}_{i,t}, \end{aligned}$$

$$\begin{aligned} PMC^{+}_{i,t} &= \text{Prob}\{\text{Positive media coverage} \mid X_{i,t}\} \\ &= \text{Prob}\{\text{COVERAGE}_{i,t} = 1 \mid X_{i,t}\} \end{aligned}$$

$$\begin{aligned} PMC^{-}_{i,t} &= \text{Prob}\{\text{Negative media coverage} \mid X_{i,t}\} \\ &= \text{Prob}\{\text{COVERAGE}_{i,t} = -1 \mid X_{i,t}\} \end{aligned}$$

where $X_{i,t}$ includes the identified firm and event characteristics that were found earlier to predict media coverage.

Figure 3.2 illustrates that predicted attention probabilities for most firm events lie relatively close to zero – this is as expected, given that the unconditional probability of WSJ coverage in the sample is 13.1%. Simply, there are typically many more earnings announcements made during a given day or week than could be the subject of stand-alone articles in the WSJ, even if all of them were thought to be potentially "interesting enough" to warrant such coverage. As the results in Tables 3.2 and 3.3 demonstrate, however, the observed high degree of selectivity in coverage decisions here does not appear to present a problem for the statistical identification of the determinants of coverage – the sample is more than large enough to accommodate a high percentage of zeros on the LHS. At the same time, the highly non-normal distribution of PMC might cause us to hesitate before using it as a raw input in further analysis. With this in mind, the relative rankings provided by PMC, rather than the raw values themselves, is the primary focus in the analysis that follows.

PMC and firm and event characteristics

Panel A of Table 3.4 compares average firm and event characteristics by absolute probability of media coverage decile. Given the results in Table 3.2, it is not surprising that low-attention firms are typically smaller firms with less analyst activity, smaller trading volumes, higher recent stock returns, and positive earnings surprises.

Panels B and C look at characteristics across positive and negative PMC deciles, respectively. Firms in the smallest PMC^+ decile (i.e., those least likely to receive positive news coverage) are relatively small, low-trading volume, low analyst coverage, low B/M, low institutional ownership firms that have experienced relatively negative earnings news and have low and stable recent stock returns. Firm in the smallest PMC^- decile (i.e., those least likely to receive negative news coverage), on the other hand, tend to be largely similar in profile, except with relatively positive earnings news and higher, more volatile recent returns.

Surprisingly, while we might have expected relatively low-attention events to elicit relatively weak event-window price reactions, despite the typically positive tenor of such news, this does not seem to be the case. Comparing the final row of column 1 in panel A to the final row of column 10 in panel B, we observe that the average announcement return for firms in the lowest PMC^- decile is actually higher than the average announcement window return for firms in the highest PMC^+ decile. While the difference in average event-window returns is relatively modest, and we do not account here for other firm and event characteristics that may affect, for example, the expected volatility of short-term announcement-window returns, it is nonetheless clear that the typically good news in low-PMC and low-PMC^- events is indeed "noticed" by market participants despite the very low probability of media coverage.

PMC and industry classification

Figure 3.3 presents average predicted probability of media coverage by Fama-French 49 industry classification for the top-15 and bottom-15 industries (attention probabilities calculated from the regression in the last column in Table 3.3), illustrating that firms in some industries seem to be unconditionally more likely to receive coverage than others. A casual inspection suggests that abnormally high propensities for media coverage seem to be associated with "high-profile" industries that might typically possess large customer and/or employee populations, or higher advertising expenditures, etc. (e.g., Tobacco Products, Beer and Liquor, Printing and Publishing, Recreation, Apparel, Automobiles and Trucks, and Retail) – this makes sense if we consider that reporters may prefer to report on firms with which the average reader is more likely to already be familiar; this observation seems to generally confirm the intuition discussed earlier with respect to expected reader interest and firm size, etc. Other, seemingly lower-profile, less-glamorous industries (such as Electrical Equipment, Shipping Containers, Medical Equipment, and Measuring and Control Equipment, etc.) seem to be over-represented among those industries with low average probability of media coverage. It is interesting to note also, that, while PMC^+ seems to increase mostly monotonically with absolute PMC, PMC^- does not seem to do so. This may be due to the fact that PMC^- seems to be somewhat more sensitive to event-specific (e.g., good news vs. bad news) rather than certain firm-specific (e.g., large vs. small) characteristics, which are likely to be more homogenously distributed within industry classes.

3.6 Attention and Neglect in the Cross-Section of Stock Returns

In this section, I explore the relationship between investor attention and expected returns. First, I do this by examining the time-series of returns on portfolios based on the PMC measures, controlling for commonly-identified risk factors. Second, I perform pooled cross-sectional regressions of individual returns on PMC and firm characteristics. Third, I examine the time series of monthly cross-sectional regressions on PMC decile membership in Fama-MacBeth tests.

Portfolio formation

The monthly portfolios are formed by sorting all of the stocks by their most recent PMC, PMC^+, and PMC^- observations within the prior three months. For example, assume that firm A makes an earnings announcement on February 15th, with firm and event characteristics resulting in media coverage probabilities of PMC = 0.78, $PMC^+ = 0.22$, and $PMC^- = 0.56$. Deciles are formed at the beginning of each month with respect to each set of scores. Firm A will then be included in the PMC decile portfolios formed on March 1, April 1, and May 1 based on the Feb. 15th observation. If A releases its next earnings announcement on April 23rd, then its May 1 decile assignments will instead be based on those more recent predicted probability values, etc.

However, if A does not report its next set of results until June 4th, then it is dropped from inclusion in the PMC portfolios formed at the beginning of June, before returning again in the July, August, and September portfolio assignments. Portfolio returns are equal-weighted and are based on the closing prices observed on the last trading day of the month.

Table 3.4 presents average firm and event characteristics by PMC deciles. Panel A shows the results for firms sorted by absolute PMC (i.e., probability of positive or negative coverage). Looking at the first column, we see that firm-events with the lowest predicted probability of receiving coverage are typically small, low-trading volume, low-analyst activity, low-beta, low-institutional ownership firms with better-than-expected earnings, positive and relatively volatile recent stock returns, and positive announcement-window CARs[20]. On the other hand, firm-events with the highest probability of coverage are typically the converse: larger, high-trading volume, etc., firms with lower-than-expected earnings results, and smaller (although still positive) recent stock returns and announcement-window CARs.

Panels B and C differentiate media attention into positive and negative coverage. Looking at the first columns of the two panels, we see that smaller firms with low analyst activity, etc., are least likely to attract both positive and negative coverage. However, the event-specific and recent stock performance factors show some striking differences, as we might expect. Firm-events least likely to attract positive media coverage are those with respect to firms that failed to meet analyst expectations on average, with relatively low recent stock returns and negative announcement returns. On the other hand, firm-events least likely to attract negative media coverage typically beat analysts' expectations and experienced positive announcement returns. Again, these results are as expected, given the results in Table 3.3.

Media coverage and expected returns: portfolio tests

I examine excess returns on the monthly PMC decile portfolios. Accounting for sources of return previously identified in the literature may be particularly important because we know that the PMC measures load heavily on several firm and event characteristics that are analogous to risk-factors such as size, value, momentum, etc. If, for example, significant excess returns on an PMC portfolio were observed to disappear when the Fama-French (1993) factors are added to

[20] Where $CAR_{i,t} = [R_{i,-1} - R^e_{-1}] + [R_{i,0} - R^e_0] + [R_{i,+1} - R^e_{+1}] + [R_{i,+2} - R^e_{+2}]$; $R_{i,t}$ is stock i's return on day t (relative to the announcement date) and R^e_t is the expected return calculated using coefficients Carhart (1997) 4-factor model estimated on the 60 trading days prior to the event window.

the return model, this might indicate that any apparent outperformance was actually due to the portfolio being heavily loaded with, e.g., small-sized and/or high book-to-market stocks. I begin with the CAPM model:

$$R_t = \alpha + \beta MKT_t + e_t \qquad t = 1, \ldots, T ,$$

where R_t is the monthly return on an equal-weighted PMC portfolio less the risk-free rate and MKT_t is the market return minus the risk-free rate. Panel A of Table 3.5 presents estimates of α for each of the PMC decile portfolios; in the final column, I show the results for a zero-investment portfolio that is long stocks in the smallest decile (i.e., firms least likely to attract absolute, positive, or negative media coverage) and short stocks in the largest decile (i.e., firms most likely to attract absolute, positive, or negative media coverage). Looking at the first row of panel A, we see that the portfolio that is long low-absolute attention stocks and short high-absolute attention stocks yields returns of almost 85 bps per month, results which are both statistically and economically significant. Differentiating between positive and negative coverage, the estimated alpha for the long-short PMC^+ portfolio is not significant, while the long-short PMC^- portfolio yields excess returns of approximately 154 bps per month, or over 20% per year.

Proceeding to a less parsimonious specification, panel B presents estimates from a Fama-French (1993) 3-factor model:

$$R_t = \alpha + \beta_1 MKT_t + \beta_2 SMB_t + \beta_3 HML_t + e_t \qquad t = 1, \ldots, T ,$$

where R_t and MKT_t are as defined above, and SMB_t and HML_t are the size and value risk factors available on CRSP. Here, the estimated alphas on the long-short PMC and PMC^- portfolios fall slightly, to approximately 70 bps and 123 bps per month (8.7% and 15.8% per year) respectively.

Since we have noted that the zero-investment PMC and PMC^- portfolios are effectively long stocks with relatively good recent performance and short stocks with relatively poor recent performance (both in terms of the information in the earnings announcement itself and in terms of stock returns prior to the announcement), it may be crucially important to account for price momentum, which has previously been shown to explain significant return predictability in the cross-section. Panel C presents estimates from the following regression:

$$R_t = \alpha + \beta_1 MKT_t + \beta_2 SMB_t + \beta_3 HML_t + \beta_4 WML_t + e_t \qquad t = 1, \ldots, T ,$$

where R_t, MKT_t, SMB_t, and HML_t are as described earlier and WML_t is the "winners minus losers" momentum factor from CRSP. Interestingly, while the alpha estimate on the long-short PMC^- portfolio drops to 75 bps per month (9.4% per year), excess returns on the long-short PMC portfolio remained unchanged from the 3-factor model above.

Finally, researchers have identified liquidity as an important potential risk factor explaining the cross-section of stock returns (e.g., Pastor and Stambaugh (2003)). Since low-PMC firms tend to be smaller firms with lower trading volumes, it is important to check whether the apparent excess returns might actually be due to a liquidity premium. Panel D presents estimated alphas from the following model.

$$R_t = \alpha + \beta_1 MKT_t + \beta_2 SMB_t + \beta_3 HML_t + \beta_4 WML_t + \beta_5 PS_t + e_t \quad t = 1,...,T,$$

where R_t, MKT_t, SMB_t, HML_t, and MOM_t are as defined above and PS_t is either the Pastor-Stambaugh (2003) liquidity factor (in levels) available from CRSP.[21] The addition of the liquidity factor lowers the estimated alphas by a couple of basis points, but the results are basically unchanged.

The findings reported above are consistent with Fang and Peress (2007), where they find that firms with no observed coverage subsequently outperform those with high coverage. But where might the outperformance be coming from? The results in the second and third rows of panels A, B, C, and D, where we are now able to make a distinction between positive and negative news events, give us a clear answer. While the long-short PMC^+ portfolios generate returns that are not significantly different from zero, the long-short PMC^- portfolio generates returns from a high of 154 bps per month in the CAPM model to a low of 74 bps per month in the five-factor model. In the previous section, I presented evidence for asymmetric attention with respect to positive and negative news. Here, we find support for the hypothesis that the systematic inattention to neglected "good news" firms in the information environment has strong implications for pricing: after controlling for all of the risk factors commonly cited in the literature, the long-short PMC and PMC^- portfolios yield returns that are potentially both economically and statistically significant. On the other hand, the portfolio strategy focusing on neglected "bad news" firms does not yield the significant negative abnormal returns that we would expect to see if the low-attention effect was symmetric. In short, the results here are more

[21] In untabulated results, I test alternate specifications with the Pastor-Stambaugh "innovations" version as well as Sadka's (2006) liquidity factors – the results are similar.

consistent with systematic underreaction to particular public news events, as in, e.g., Hong, Lim, and Stein (2000).

Factor loadings

Table 3.6 shows 4-factor loadings for the PMC portfolios.[22] In the final column of panel A, note that the long-short portfolio formed on absolute PMC (i.e., long firms with low expected probability of positive or negative coverage, and short firms with high expected probability of coverage) loads positively on HML and SMB. Given that relatively large, relatively high book-to-market firms are most likely to attract media coverage (Table 3.2 and Table 3.4), this is not surprising. Similarly, we saw earlier that high beta firms attract attention, so the long-short portfolio loads negatively on the market return. Interestingly, all of the PMC decile portfolios load negatively on WML, albeit with relatively small coefficients.

In panel B, we begin to see some clear differences among portfolios formed on absolute, positive, and negative attention probabilities. In contrast to the results above, the long-short PMC^+ portfolio loads negatively on HML and negatively on WML. In particular, looking at the first column of panel B, firms with the lowest probability of receiving positive coverage (i.e., the "low-profile, bad news" firms) have much smaller coefficients on HML and WML than we saw for the firms with the lowest probability of absolute coverage (the first column of panel A). Note the negative coefficient WML; this is as we might expect: by sorting on positive attention, the long-short PMC^+ portfolio is long stocks with relatively bad news in the previous month and short stocks with relatively good news.

Panel C shows the results for portfolios formed on PMC^- (expected probability of receiving negative attention). In the final column, we see that the long-short portfolio now loads positively on WML, as the portfolios with the highest probability of negative coverage load increasingly negatively on this factor. Otherwise, the factor loadings here are identical in sign to those with respect to the absolute PMC long-short portfolio in panel A.

Double-sorts by firm and event characteristics

Since the PMC measures load on a number of variables that we might otherwise think of as potentially affecting expected returns, I examine the returns of long-short PMC portfolios formed within firm- and event-characteristic quintiles. In particular, while we know that the

[22] The Pastor-Stambaugh liquidity factor is omitted from the specifications in Table 6 due to statistical insignificance – results available upon request.

portfolio strategies appear to yield positive excess returns in the aggregate, it is important to find out where the trading profits are coming from. I concentrate here on returns from portfolios based on quintiles of PMC$^-$. Panel A of Table 3.7 presents the results sorted first on market value. Interestingly, while the long-short PMC$^-$ portfolio yields significant positive excess returns only in the bottom three quintiles of market value, returns are actually highest within the middle quintile.

Panel B looks at portfolios formed within analyst coverage quintiles. Here, while abnormal returns appear to be concentrated in the lower half in terms of analyst activity, profits are highest among trades within the lowest quintile. Panel C shows the results for portfolios sorted first on book-to-market of equity. Excess returns are highest for the high B/M firms, but they are not nearly as concentrated as in the two earlier cases – even within the lowest B/M quintile, we observe significant positive returns (albeit only significant at 10%). In panel D, we see that profits are highest among low-beta firms, although, again, excess returns remain positive and (modestly) significant up to the largest quintile. Finally, panel E presents the results from sorting on recent returns: here, the long-short portfolios on PMC$^-$ are most profitable among firms with positive recent returns – this is not entirely surprising since high recent stock returns is presumably one of the most important "good news" elements that is being ignored for firms with low PMC$^-$.

<u>Earnings momentum and media coverage</u>

The profitable portfolio trading strategies described above essentially prescribe going long a set of stocks that, on average, have experienced "good earnings news" in the previous 3 months. This is a consequence of the fact that low attention probabilities are explained, in part, by the observation of positive profits and positive earnings surprises (see Tables 3.2 and 3.3). This inevitably leads to the question: Is what we are seeing simply post-earnings announcement drift (PEAD)? After all, earnings momentum strategies advise going long stocks with large positive earnings surprises and shorting stocks with large negative surprises. However, while there is obviously an element of PEAD at work here, there are two lines of evidence against a "standard", symmetric PEAD story.

First, a traditional PEAD explanation would imply that the part of the portfolio that is short negative earnings surprise stocks should also be profitable, as well as the one that is long positive surprise stocks. As we see in Table 3.4, the stocks in the largest PMC and PMC$^-$ deciles have experienced negative earnings surprises on average. However, referring to Table 3.5, we observe that the excess returns on these largest decile portfolios are never significantly negative. Finally, Hou et al. (2006) observe that earnings momentum profits are highest for low-attention stocks – however, note here that returns on the long-short PMC$^+$ portfolio (which is long the relatively low-attention, bad-news stocks) are not significantly different from zero, implying that the underreaction is asymmetric with respect to good and bad news.

Second, symmetric PEAD would predict that the PMC-based strategies should work equally well across earnings surprise quintiles – that is to say, not all that well, since forming portfolios within earnings surprise quintiles will reduce the differences among the stocks. For example, if earnings momentum were behind the results, by forming PMC portfolios within the largest earnings surprise quintile we would effectively be going long a portfolio of firms with "best of the best" surprises and going short a portfolio with "worst of the best" surprises. Obviously, we would not expect this to be as profitable as a strategy that is long the "best of the best" surprises and short the "worst of the worst". Looking at the results in Table 3.7, panel B, however, it is clear that the PMC- and PMC$^-$-based strategies are actually *more* profitable within the largest earnings surprise quintile than they were with respect to the sample as a whole (Table 3.5, panel C). Alphas from the 3rd and 4th earnings-surprise quintiles are smaller, but also positive and significant, before falling to insignificance in the two smallest quintiles.

Cross-sectional analysis

In order to test for the potential impact of PMC in the cross-section, I perform the following regression:

$$R_{i,t} = \phi_0 + \phi_1 \log(PMC)_{i,t-1} + \theta Controls_{i,t} + \varepsilon_t \qquad t = 1,\dots,T\ ;\ i = 1,\dots,N\ ,$$

where $R_{i,t}$ is the month *t* return for stock *i*, $\log(PMC)_{i,t-1}$ is the lagged natural logarithm of a PMC measure, and $Controls_i$ includes a set of characteristics that may help to explain returns (including $B/M_{i,t}$, $MKTVALUE_{i,t}$, $STDEV(RETURNS_{-1})_{i,t}$, etc., as defined earlier). Table 3.8 shows the results with standard errors robust to arbitrary heteroskedasticity and clustering by firm.[23] In the first three columns, monthly returns are regressed on lagged log(PMC), log(PMC$^+$), and log(PMC$^-$) in turn, along with a number of "standard" controls. Consistent with the earlier portfolio results, I find that low PMC and low PMC$^-$ predict higher subsequent returns. However, in the cross-section, low PMC$^+$ now appears to predict lower returns. While

[23] Results are similar with clustering by year or industry, or with firm fixed effects.

the size of the estimated coefficient on log(PMC$^+$) suggests that the effect is not quite as large as with log(PMC) and log(PMC$^-$), this finding of a significant positive coefficient seems to contradict the earlier portfolio results. While this result goes against an information risk or illiquidity-based explanation (i.e., a low expectation of positive attention should unambiguously predict higher returns if "information risk" is priced and media coverage is a proxy for such risk), there are at least two potential explanations for this finding. The first is that low-attention, bad news firms (i.e., firms with the lowest probability of attracting positive attention) subsequently experience negative returns as ignored information is slowly impounded into prices. If true, this would imply that low-attention drift is somewhat more symmetric than we would have thought based on the portfolio results.

A second potential explanation for the positive coefficient on log(PMC$^+$) is that we may be seeing residual post-earnings announcement drift: firms with the highest probability of positive coverage (who, in general, will have experienced positive earnings surprises) experience positive returns in the following month, and the converse with respect to firms with the lowest probability of positive coverage.[24] In columns 4 to 6 of Table 3.8, I include the earnings surprise (among other controls) as an explanatory variable. In this specification, the coefficient on log(PMC$^+$) falls (although we still cannot reject significance at 10%), while the coefficient on SURPRISE$_{i,t}$ is significant and positive. On the other hand, the coefficients on log(PMC) and log(PMC$^-$) in columns 4 and 6 are observed to grow larger in absolute magnitude.

Fama-MacBeth regressions

I perform Fama-MacBeth tests by regressing stock returns on PMC variables within each month and then examining the resulting time series of monthly estimated coefficients. Panel A of Table 3.9 presents the results from regressing returns on dummy variables indicating membership in the first or tenth deciles of PMC for the previous month.

$$R_{i,t} = \phi_0 + \phi_1 I\{PMCdecile\}_{i,t-1} + \theta Controls_{i,t} + \varepsilon_t \qquad t = 1,\dots,T\ ;\ i = 1,\dots,N\ ,$$

where $I\{PMCdecile\}_{i,t}$ is equal to one if firm i was in the first or tenth decile of PMC, PMC$^+$, or PMC$^-$ for the previous month, and zero otherwise; and $Controls_{i,t}$ contains the set of control variables in the first specification from Table 3.8.

[24] Furthermore, if PEAD is having an effect on the results, the estimated coefficients on PMC and PMC$^-$ may be too small (since PMA and PMA$^-$ are negatively related to earnings surprise).

Note that the average coefficients on the first decile PMC and PMC^- dummies are significant and positive, while the average coefficient on the tenth decile PMC^- dummy is significant and negative. The positive coefficients on the first decile dummies are as we might expect, given the preceding results for both the time-series tests and the pooled cross-section (i.e., low-attention, good-news firms experience high subsequent returns). The result for the tenth decile PMC^- dummy is somewhat more surprising, suggesting that bad-news high-attention firms experience downward drift.

Instead of using the binary measures, Panel B of Table 3.9 presents the results of regressing returns on the natural log of the lagged PMC variables.

$$R_{i,t} = \phi_0 + \phi_1 \log(PMC)_{i,t-1} + \theta Controls_{i,t} + \varepsilon_t \qquad t = 1,\ldots,T\ ;\ i = 1,\ldots,N\ ,$$

where the variables are as defined earlier.

Sub-period results

Are the findings sensitive to the sample period? The Fama-MacBeth results would seem to indicate that the estimates are fairly consistent over time, but we may be interested in finding out how simulated portfolio returns might have changed. Is this a strategy that might still work today? In untabulated results, I divide the sample into four sub-periods (1984-1989, 1990-1994, 1995-1999, 2000-2005), and then re-run the four-factor portfolio tests with respect to each one. The estimated monthly alphas for the long-short PMC^- portfolios are[25]: 188 bps*** for 1984-1989; 64 bps** for 1990-1994; 39 bps (no statistical significance) for 1995-1999; and 125 bps*** for 2000-2005. In the late-90s sub-period, the culprit appears to be 1999 – when this year is omitted, the estimated portfolio alpha for the remaining 1995-1998 period is 105 bps*** per month.

3.7 Neglected Firm Effect or Delayed Price Response to Positive News?

The foregoing results indicate that PMC has significant incremental predictive power with respect to stock returns. Estimated excess annualized PMC^- portfolio returns range from a high of 20% in a CAPM model to just over 9% in a five-factor model that includes both momentum and liquidity factors. These findings are consistent with (albeit significantly larger in magnitude than) previous research finding that low-media coverage predicts high returns. Fang and Peress

[25] * significant at 10%; ** significant at 5%;*** significant at 1%

(2007) identify a return premium of over 3 percent per annum on stocks that are observed not to receive media coverage versus those that receive high attention; they show that this effect cannot easily be explained by the standard set of risk factors. In a similar vein, Gadarowski (2002) finds that high news coverage predicts lower subsequent stock returns.

More broadly, looking beyond the potential impact of media coverage in particular, researchers have long sought to explain the puzzles of apparently high risk-adjusted returns for groups of firms which could be characterized as possessing relatively severe frictions or risks in their information environments, e.g., with respect to small firms (e.g., Banz (1979), and Reinganum, (1981)), firms that are "neglected" in terms of analyst coverage (e.g., Arbel and Strebel, (1982)), and "price delayed" firms (Hou and Moskowitz (2005)). Explanations have typically focused on liquidity (Amihud and Mendelson (1986)), or on potential frictions and risks in the information environments faced by these firms (e.g., Merton (1987), or Easley et al. (2002)).

Easley et al. (2002) find that probability of informed trading (PIN), which they interpret as a proxy for information risk, has predictive power with respect to the cross-section of returns. In terms of media coverage, one potential explanation for my low-attention premium is that these firms are subject to higher levels of information risk.

Sadka (2006) argues that a substantial proportion of both momentum and PEAD returns can be explained by liquidity risk. Again, I do not find here that any of the liquidity factors examined are significant in explaining PMC portfolio returns. In addition, while it is true that low-attention firms have many of the characteristics that we would typically associate with low liquidity, the liquidity explanation is hard to reconcile with the observation that only the low-attention good news firms seem to trade at a discount.

At a very basic level, this paper presents evidence of a return premium for firms that are likely to be neglected in terms of media attention. How, then, do my results differ from the aforementioned papers, in particular Fang and Peress (2007)? In this study, I construct an event-specific measure of neglect, documenting significant asymmetry in financial news media coverage decisions with respect to the content of the underlying information shock. Making use of this observation, I identify and investigate a new explanation for the finding of higher performance for low-attention (neglected) stocks in this context: asymmetric underreaction with respect to positive news for these firms.

My results support the asymmetric underreaction hypothesis, while they do not tend to support the information risk and investor recognition explanations. Specifically, when I look at propensities to attract positive and negative coverage separately, I do not find evidence that the relatively low-attention, bad news firms (i.e., firms with the lowest probability of positive coverage) experience either significantly positive excess returns (as we would expect to see if these stocks were being discounted due to information risk) or significantly negative returns (as we would expect if the low-attention effect was symmetric with respect to positive and negative news events).

While the observation of apparent underreaction to positive news is consistent with most of the prior research mentioned, it would seem to be rather at odds with the findings of HLS, who argue that "bad news travels slowly". The authors find that momentum strategies are more profitable for relatively small firms and those with low analyst coverage; furthermore, they show that the effect is more pronounced for recent losers than recent winners, which they interpret as evidence that, given managers' reluctance to communicate bad news, the presence of analysts is more important for "drawing out" negative information. One important point is that the authors' explanation is potentially consistent with a negativity bias in media coverage: if bad news is more likely to be credible/truthful, then it may be natural for information intermediaries to feature it more prominently once such an item has already been "drawn out" or made public (as in the case of earnings announcements). Also, note that the firms and the information flows being considered are potentially quite different. In this paper, I investigate potential underreaction to a particular set of underlying public events that are characterized by the fact that they are unlikely to have had a large, immediate impact upon prices. In contrast, HLS focus on drift following unspecified information flows which are actually identified (albeit implicitly) by the observation of a significant contemporaneous price movement. Finally, HLS identify their strongest evidence of asymmetric price momentum with respect to a subset of firms that are generally much smaller and may have significantly worse information environments than those included in my sample. In particular, while all of the firms considered in this study have at least two distinct analyst forecasts in the month prior to their announcements, HLS include all firms above the 20th percentile NYSE/AMEX size breakpoints, a significant proportion of whom (e.g., 41.7% in 1988 for the 20th-40th percentile by size) are observed to have no analyst coverage.

3.8 Robustness Tests

The foregoing results showing high excess returns from trading strategies based on predicted media coverage, (and, in particular, the evidence of asymmetry with respect to positive and negative information), are certainly suggestive of a link between investors' attention allocations and the market's underreaction/overreaction to news events, but we are left with several unanswered questions. In particular, what if the simulated low-PMC trading strategies are simply being rewarded for holding stocks that are, e.g., small and/or have high market betas, etc.? It is clear that PMC's ability to predict returns, by its very construction, must derive from the identified determinants of coverage in some sense. Since PMC is the estimated probability resulting from a regression of observed coverage, PMC is essentially a non-linear combination of the regressors (recall that this is one of the key advantages of this measure, in that it allows us to side-step some potential problems related to omitted variable bias and measurement error). We might hope that the four-factor setting would correct for such effects in portfolio returns with respect to these standard risk factors, but there are no guarantees that this is sufficient. In short, it is possible that PMC may simply identifying stocks using an optimal weighting of previously-identified risk factors. In particular, market size and past returns were found to be highly significant predictors of coverage. Furthermore, other variables, such as analyst attention, may also be highly correlated with size, etc.

In this section, I address this potential concern by examining the predictive content of restricted versions of expected coverage. Specifically, I examine PMC variants projected on different sets of predictors in turn, excluding size, B/M, momentum, and beta in all specifications. In order to eliminate any lingering effects from covariation with the remaining regressors, I then orthogonalize each restricted measure by taking the residual from a regression of PMC on size, B/M, momentum, and beta. Decile portfolios are then formed with respect to each measure. Table 3.10 presents the four-factor portfolio alphas based on these restricted and orthogonalized versions of the base PMC measures.

Panel A shows the results for portfolios formed on PMC1, a specification that only includes information on the earnings release itself. The results here are consistent with a generalized underreaction to positive earnings news: both the highest decile of $PMC1^{+}$ and the lowest decile of $PMC1^{-}$ (i.e., the "good news" stocks) have high predicted returns. Since positive coverage is unconditionally more likely to be observed than negative coverage, and we are controlling for no other factors, absolute PMC is positively related to earnings surprise and negatively related to the earnings loss dummy, etc., in this specification. Looking at positive and negative coverage in this simple way, without reference to any of the firm-specific or market-wide determinants that are significant in predicting coverage, results in a trading strategy that is equivalent to earnings momentum (albeit one with relatively high risk-adjusted returns).

As we proceed to include additional cross-sectional predictors of coverage in Panels B, C, D, E, and F, an interesting pattern emerges. Evidence of underreaction to positive earnings news generally persists, but it wanes with respect to the high PMC^+ stocks (high coverage good news) while staying robust for the low PMC^- firms (low coverage good news). In other words, as we move away from a simpler specification of coverage, the evidence of underreaction to high-coverage positive events disappears while the underreaction with respect to low-coverage positive events persists.

The results strongly suggest that a risk-based factor such as size or liquidity, which we would expect to impact both low PMC^+ and low PMC^- firms in similar ways, cannot explain this asymmetry between low-coverage and high-coverage events. In particular, the lowest decile of PMC^+ stocks never have alphas significantly different from zero, as we would expect to see if all low-coverage firms were being rewarded for loading on some symmetric risk factor (such as information risk. If, instead, one were to argue that high expected returns for low-coverage bad news stocks were being offset by earnings momentum, resulting in no net predictability for the low PMC^+ stocks, then we would surely expect to see evidence of such earnings-related drift in panel A, where $PMC1^+$ and $PMC1^-$ only load on positive and negative earnings news attributes. On the contrary, in each panel, we only find evidence of underreaction with respect to positive events, an effect which becomes increasingly concentrated on low-coverage events as we include additional firm-specific predictors of coverage behavior.

At the same time, as we move toward the more complete specifications of predicted coverage, the predictability of absolute PMC begins to look more like a weaker version of PMC^-, suggesting that the underreaction to low-coverage positive events exhibited by PMC^- is behind the observed predictability of absolute PMC.

Panels G and H show estimated month t+n excess returns for portfolios formed on PMC5 and PMC6, respectively. In particular, Panel H illustrates that the asymmetric predictability results are somewhat persistent over subsequent months. For example, estimated four-factor alphas for the long-short $PMC6^-$ portfolios fall relatively sharply after month t+1, but remain positive and statistically significant at 5 percent out to month t+5, addressing the potential concern that we might simply be observing return predictability based on residual short-term earnings drift stemming from announcements made towards the ends of the 3-month formation periods.

3.9 Conclusion

Utilizing estimates of financial news coverage as a proxy for investor attention, this study identifies asymmetric delays in the market's response to positive news events as a potential explanation for the neglected firm effect. I identify the cross-sectional event- and firm-specific characteristics that predict positive and negative media coverage regarding almost 180,000 quarterly earnings announcements from 1984 to 2005, showing that positive corporate news tends to go relatively unnoticed compared to negative news regarding otherwise-similar firms. If news media outlets maximize readership by attempting to publish stories that readers will find most interesting, the results suggest that the impact of cognitive constraints (i.e., limited attention) on investors' information preferences is asymmetric with respect to positive and negative news events. If limits to arbitrage are binding, this relative lack of attention with respect to some events will contribute to predictability in asset returns. Finally, if good news from neglected firms is more likely to be ignored, low media coverage will predict positive returns for these stocks, on average. In short, my findings support this hypothesis.

This study contributes to the existing literature in the following ways: First, examining the cross-sectional news story elements that predict coverage in the media, I find evidence of significant asymmetry in reporting: negative earnings information is more likely to result in media coverage than is positive information, holding other factors constant; second, I propose and apply a new measure of investor attention based on the predictable component of media coverage decisions: probability of media coverage (PMC); third, I utilize PMC to identify relatively "neglected" stocks, confirming the results of other studies which find that these firms enjoy a return premium that cannot be explained by the standard set of risk factors. In contrast to these earlier papers, however, I find that this apparent neglected firm effect is attributable to systematic underreaction to positive news for these firms (i.e., consistent with the effects of negativity bias in attention), while the stocks of relatively neglected firms with negative news appear to be efficiently priced.

Regarding this evidence of asymmetry in attention, why might the media (and, by implication, investors) be more likely to focus on negative events rather than positive ones? If the market for news is reasonably competitive, then this observed bias in coverage must, in some sense, reflect the underlying information preferences/requirements of financial news readers. While some potential explanations for negativity bias are identified in the literature, the answer to this question is beyond the scope of the current paper; this remains a potentially important line of inquiry for future research.

While the study focuses on an examination of media coverage with respect to one particular type of news event (albeit one that lends itself particularly well to the question of interest), the results underscore the potential importance of investor attention allocations for our understanding of the market's reaction to news.

Figure 3.1: Firm and event characteristics that attract (absolute) media coverage – Illustration based on the results from Table 3.2

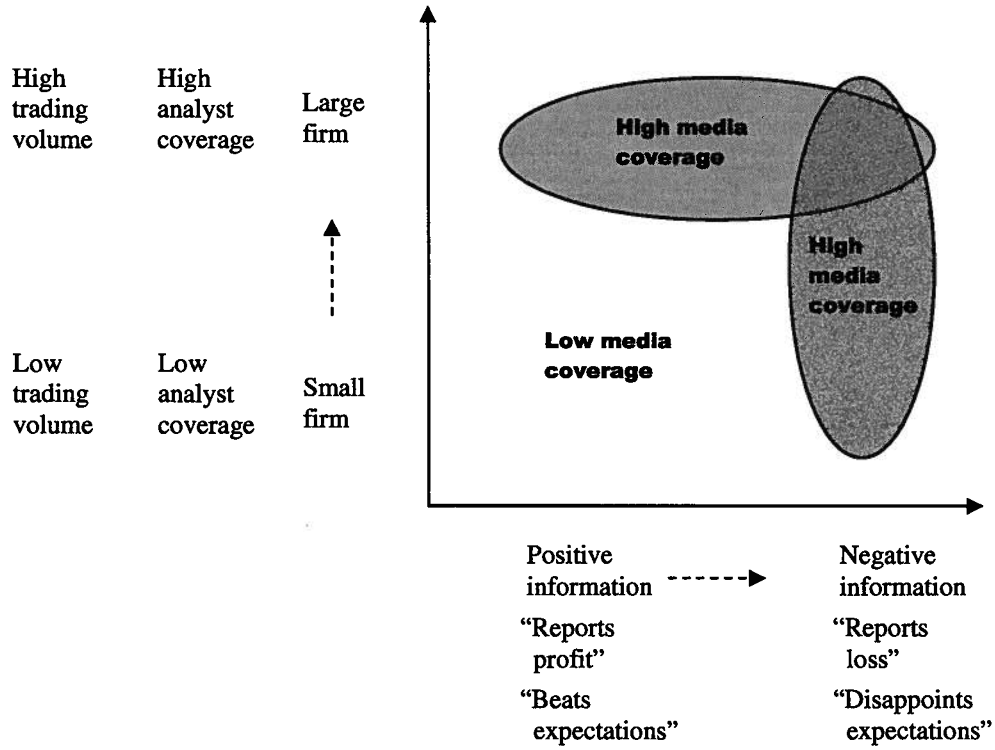

Figure 3.2: Empirical distributions of probability of media coverage (PMC)
(200 bins); Probability of coverage (PMC) is the sum of the probabilities of positive coverage (PMC^{+}) and negative coverage (PMC^{-})

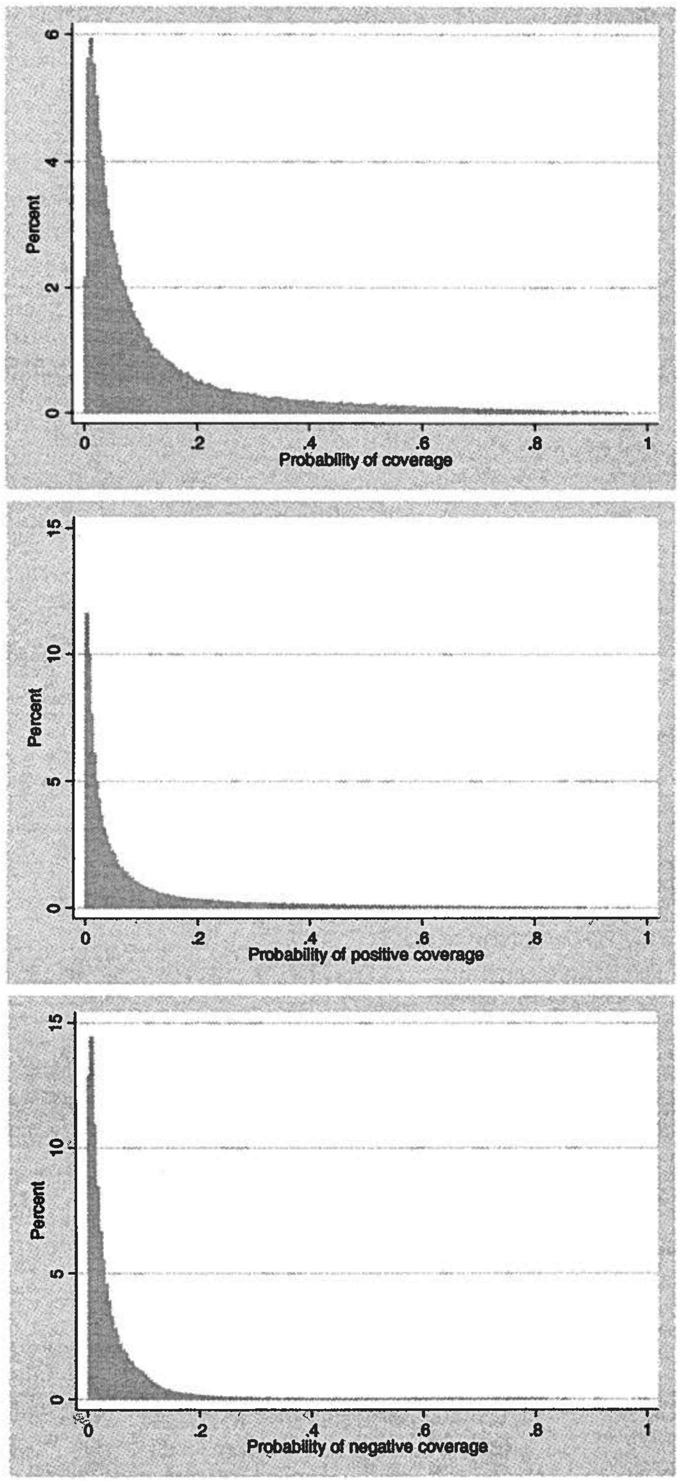

Figure 3.3: Probability of media coverage (PMC) by Fama-French 49 industry classification

This figure shows the bottom-15 and top-15 Fama-French 49 industries ranked by average PMC.

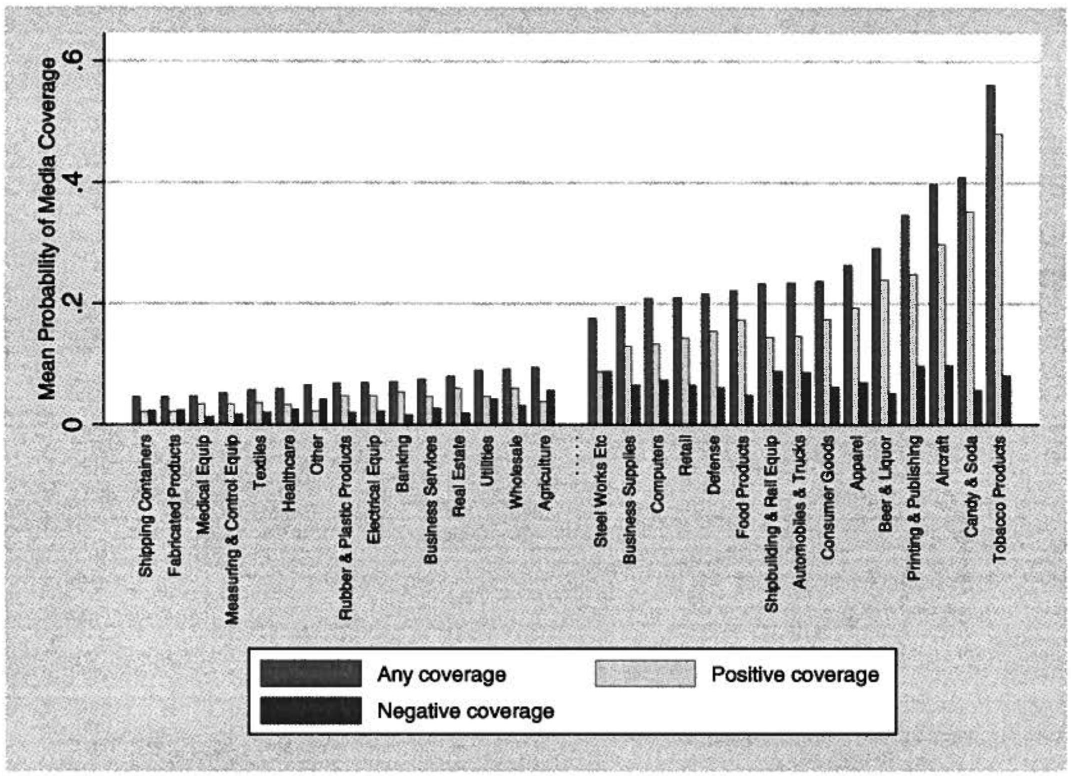

Table 3.1: Descriptive statistics of earnings announcement information variables and firm characteristics by observed media coverage

This table reports descriptive statistics for earnings announcements in the sample. *Earnings surprise* is the announced EPS less the median analyst forecast from the 30 days prior to the earnings announcement, divided by the stock price. *Analyst attention* is the number of distinct analyst EPS forecasts observed in the 30 days prior to the announcement. *stdev(forecasts)* is the standard deviation of analysts' EPS forecasts (normalized by the stock price) observed in the 30 days prior to the announcement. *Market value* is the number of shares outstanding multiplied by the closing stock price two days before the announcement. *$Trading Volume* is the firm's average dollar value of trading volume in the 60 trading days prior to the announcement. *B/M* is the book value of equity from the fiscal year ending in the previous year divided by the market value of equity from December 31 (divided by 1000). *US firm* is a dummy variable that is equal to 1 if the firm is identified as a US firm in CRSP. *Beta* is the estimated slope coefficient from a CAPM regression of daily returns from the 60 days prior to the announcement window. *Institutional ownership* is the percentage of equity held by institutions at the end of the previous calendar year. *Recent returns* is the average daily stock return during the 60 trading days prior to the announcement window. *stdev(Recent returns)* is the standard deviation of daily stock returns during the 60 trading days prior to the announcement window. *CAR* is the abnormal announcement-window stock return based on a four-factor model of daily returns estimated over the previous 60 trading days. Panel A reports statistics for all earnings announcements in the sample. Panel B reports statistics only for those announcements with an associated positive Wall Street Journal news article. Panel C reports statistics for those announcements without any associated Wall Street Journal news article. Panel D reports statistics only for those announcements with an associated negative Wall Street Journal news article.

Panel A: All announcements

variable	*N*	*mean*	*sd*	*min*	*p1*	*p5*	*p10*	*p25*	*p50*	*p75*	*p90*	*p95*	*p99*	*max*
Earnings surprise	190808	-0.00252	0.038373	-0.99429	-0.09143	-0.01474	-0.00571	-0.00078	0	0.001096	0.00354	0.006923	0.030303	0.977966
Earnings (%)	190461	0.005336	0.120014	-0.99921	-0.19033	-0.03688	-0.01031	0.002971	0.0088	0.015962	0.024583	0.032537	0.071048	25.04173
Analyst attention	190808	6.419044	5.112228	2	2	2	2	3	5	8	13	17	25	46
stdev(forecasts)	190808	0.210451	0.772615	0	0	0	0	0.02439	0.0625	0.166667	0.421053	0.833333	2.5	95.60001
Market value	190808	3124359	1.34E+07	1148.624	23174.99	50998.24	80076.49	186213.8	548087.6	1765107	5571076	1.15E+07	4.79E+07	5.79E+08
$Trading Volume	190807	1.60E+07	7.28E+07	63.11498	33728.63	109670.9	204817.3	617034.9	2339028	9021924	2.99E+07	6.15E+07	2.33E+08	4.42E+09
B/M	181249	0.016978	1.195242	3.40E-08	2.08E-05	7.38E-05	0.00012	0.000236	0.000416	0.000663	0.000977	0.001299	0.002806	186.9693
US firm	190808	0.97564	0.154163	0	0	1	1	1	1	1	1	1	1	1
Beta	190808	0.806831	0.726316	-5.0086	-0.73189	-0.1568	0.047852	0.351605	0.721025	1.16065	1.696	2.1108	3.0715	7.9255
Institutional ownership (%)	183749	46.63316	23.00282	0.000172	2.747857	9.564054	15.299	28.31198	46.83545	64.41245	77.43977	84.15045	93.88927	99.99456
Recent returns [-61 to -2]	190808	0.071631	0.407534	-3.67218	-1.09541	-0.5746	-0.36752	-0.11962	0.075913	0.269385	0.498527	0.688497	1.195283	5.973645
stdev(Recent returns)	190808	0.028847	0.017571	0	0.007895	0.010625	0.012553	0.016827	0.024312	0.035775	0.050508	0.062259	0.091454	0.35137
CAR[-1 to 3]	190808	0.091849	9.55053	-147.174	-27.8574	-14.4875	-9.67611	-3.98083	0.070352	4.3158	9.90672	14.45842	26.7681	218.69

Panel B: Announcements with positive coverage in the WSJ

variable	*N*	*mean*	*sd*	*min*	*p1*	*p5*	*p10*	*p25*	*p50*	*p75*	*p90*	*p95*	*p99*	*max*
Earnings surprise	16901	0.000234	0.010242	-0.46	-0.01226	-0.0034	-0.00134	0	0.000193	0.000879	0.002505	0.00451	0.013333	0.228571
Earnings (%)	16899	0.013149	0.117317	-0.36429	-0.00491	0.000925	0.001882	0.00456	0.009412	0.016053	0.024776	0.032363	0.061359	15
Analyst attention	16901	11.75481	6.63597	2	2	3	4	7	11	16	21	24	30	44
stdev(forecasts)	16901	0.108481	0.324927	0	0	0	0	0.020408	0.045455	0.1	0.2	0.333333	1.05	13.66667
Market value	16901	1.34E+07	3.00E+07	13793.5	93785.72	277312	527963.1	1494986	4446590	1.23E+07	3.02E+07	5.46E+07	1.50E+08	5.26E+08
$Trading Volume	16901	5.93E+07	1.41E+08	6425.677	231694.7	965674.6	1925451	6012078	1.90E+07	5.54E+07	1.40E+08	2.29E+08	6.31E+08	3.03E+09
B/M	16496	0.047821	2.458221	2.20E-07	2.72E-05	7.02E-05	0.000109	0.000199	0.000347	0.00055	0.000792	0.000987	0.001666	186.9693
US firm	16901	0.975564	0.154404	0	0	1	1	1	1	1	1	1	1	1
Beta	16901	0.935877	0.584175	-2.5804	-0.21561	0.1321	0.29873	0.56961	0.87782	1.2231	1.6167	1.9475	2.83	5.0124
Institutional ownership (%)	16545	56.3496	19.00853	0.049414	6.463413	21.03038	29.7944	44.86758	58.0683	70.10284	79.30531	84.69053	93.54565	99.98873
Recent returns [-61 to -2]	16901	0.08432	0.311557	-1.90529	-0.78797	-0.40647	-0.25956	-0.07294	0.084333	0.243917	0.418465	0.560178	0.94692	3.293368
stdev(Recent returns)	16901	0.02208	0.011839	0.001478	0.007914	0.010008	0.011457	0.014327	0.018954	0.026534	0.036365	0.044091	0.064262	0.32312
CAR[-1 to 3]	16901	0.659712	7.620515	-85.7621	-21.5445	-10.5094	-6.8927	-2.88749	0.4863	4.16483	8.86934	12.43585	22.3734	93.3546

Panel C: Announcements with no coverage in the WSJ

variable	*N*	*mean*	*sd*	*min*	*p1*	*p5*	*p10*	*p25*	*p50*	*p75*	*p90*	*p95*	*p99*	*max*
Earnings surprise	165732	-0.0024	0.038135	-0.99429	-0.08833	-0.015	-0.00585	-0.00082	0	0.001131	0.003636	0.007111	0.031189	0.977966
Earnings (%)	165423	0.0054	0.121556	-0.99921	-0.19	-0.03833	-0.01143	0.002977	0.008889	0.016058	0.024675	0.032668	0.072533	25.04173
Analyst attention	165732	5.707262	4.429997	2	2	2	2	3	4	7	11	15	22	46
stdev(forecasts)	165732	0.210688	0.781761	0	0	0	0	0.02439	0.0625	0.166667	0.421053	0.833333	2.5	95.60001
Market value	165732	1922393	9292109	1148.624	21861.39	47784	73454.42	163409.8	446050.3	1268752	3417051	6335140	2.35E+07	5.79E+08
$Trading Volume	165731	1.05E+07	5.78E+07	63.11498	30981.78	98841.34	182152.2	522420	1833569	6394109	1.91E+07	3.68E+07	1.41E+08	4.42E+09
B/M	156819	0.012222	0.861688	3.40E-08	2.06E-05	7.37E-05	0.000121	0.000238	0.000418	0.000664	0.000975	0.00129	0.002763	186.9693
US firm	165732	0.975756	0.153806	0	0	1	1	1	1	1	1	1	1	1
Beta	165732	0.78631	0.73744	-5.0086	-0.77203	-0.18597	0.024088	0.32439	0.69311	1.1436	1.7013	2.1208	3.0785	7.9255
Institutional ownership (%)	159246	45.30571	23.20005	0.000172	2.546038	8.973255	14.41449	26.68684	44.69246	63.15168	77.07803	84.12677	94.01344	99.99456
Recent returns [-61 to -2]	165732	0.073242	0.412661	-3.67218	-1.10158	-0.58314	-0.37369	-0.12196	0.076681	0.274004	0.508065	0.701447	1.216335	5.973645
stdev(Recent returns)	165732	0.029473	0.0178	0	0.00786	0.010682	0.012705	0.017229	0.025031	0.036667	0.051418	0.063262	0.092352	0.35137
CAR[-1 to 3]	165732	0.102713	9.678056	-147.174	-27.9506	-14.6799	-9.84426	-4.05451	0.05957	4.382648	10.09807	14.75812	27.25886	218.69

Panel D: Announcements with negative coverage in the WSJ

variable	*N*	*mean*	*sd*	*min*	*p1*	*p5*	*p10*	*p25*	*p50*	*p75*	*p90*	*p95*	*p99*	*max*
Earnings surprise	8175	-0.01079	0.067714	-0.98824	-0.31563	-0.05807	-0.02177	-0.0041	0	0.000898	0.00417	0.009534	0.055118	0.82
Earnings (%)	8139	-0.01218	0.087904	-0.95915	-0.44353	-0.11333	-0.04704	-0.00604	0.004378	0.012851	0.021767	0.029818	0.065494	2.289157
Analyst attention	8175	9.817859	6.358488	2	2	2	3	5	9	13	19	22	29	41
stdev(forecasts)	8175	0.41645	1.122849	0	0	0	0.021277	0.055556	0.142857	0.363636	1	1.5	4.5	33
Market value	8175	6327817	1.80E+07	2697.2	34312	94783.49	174495.8	557669.8	1808531	5281947	1.41E+07	2.46E+07	7.56E+07	4.35E+08
$Trading Volume	8175	3.67E+07	1.01E+08	5005	77214.07	327755.3	710128.9	2626058	9466614	3.05E+07	8.13E+07	1.56E+08	4.66E+08	2.32E+09
B/M	7934	0.046857	2.322482	2.08E-07	2.24E-05	9.52E-05	0.00016	0.000314	0.000553	0.000877	0.001422	0.002006	0.006169	141.275
US firm	8175	0.973456	0.160757	0	0	1	1	1	1	1	1	1	1	1
Beta	8175	0.95605	0.719325	-3.6665	-0.54913	-0.01279	0.20408	0.51837	0.87327	1.2906	1.8095	2.3072	3.2245	6.0908
Institutional ownership (%)	7958	52.99563	20.58458	0.149269	5.001407	15.36306	23.01242	39.11233	55.2866	68.79938	77.8599	83.23251	92.13445	99.49509
Recent returns [-61 to -2]	8175	0.012745	0.468732	-3.54245	-1.46057	-0.73001	-0.49402	-0.19432	0.035713	0.239437	0.47199	0.672202	1.242535	3.77118
stdev(Recent returns)	8175	0.030143	0.019721	0.002262	0.008574	0.011433	0.013281	0.017254	0.024349	0.036978	0.054154	0.067915	0.10369	0.27084
CAR[-1 to 3]	8175	-1.30241	10.36446	-115.566	-33.2927	-17.8895	-11.709	-5.2624	-0.7014	3.42588	8.63563	12.982	25.1824	146.618

Table 3.2: Which factors predict media coverage?

This table reports coefficient estimates from logit regressions of $ABS(COVERAGE)_{i,t}$ on market conditions and firm and event characteristics; $ABS(COVERAGE)_{i,t}$ is equal to 1 if there is an identified positive or negative WSJ earnings article within one week of the announcement date, and 0 otherwise; *Earnings surprise* is the announced EPS less the median analyst forecast from the 30 days prior to the earnings announcement divided by the stock price; *I{Negative surprise}* is a dummy variable equal to one if *Earnings surprise* is negative, and zero otherwise; *I{Loss}* is a dummy variable equal to one if announced EPS is negative, and zero otherwise; *log(Market value)* is the natural logarithm of the number of shares outstanding multiplied by the closing stock price two days before the announcement; *B/M* is the book value of equity from the fiscal year ending in the previous year divided by the market value of equity from December 31 (divided by 1000); *Institutional ownership* is the percentage of equity held by institutions at the end of the previous calendar year; *log($Trading Volume)* is the natural logarithm of the firm's average dollar value of trading volume in the 60 trading days prior to the announcement; *log(Analyst attention)* is the natural log of the number of distinct analyst forecasts observed in I/B/E/S in the 30 days preceding the announcement; *stdev(forecasts)* is the standard deviation of the normalized analyst EPS forecasts recorded in I/B/E/S in the 30 days preceding the announcement; *Recent returns* is the average daily stock return during the 60 trading days prior to the announcement window; *S&P500 returns* is the average daily return on the S&P500 during the 60 trading days prior to the announcement window; *stdev(S&P500 returns)* is the standard deviation of the average daily return on the S&P500 during the 60 trading days prior to the announcement window; month-of-the-year and day-of-the-week dummy variables are included in all specifications; significant standard errors in brackets, robust to arbitrary heteroskedasticity and clustering by firm; * significant at 10%; ** significant at 5%; *** significant at 1%.

	Media coverage dummy			
	(1)	(2)	(3)	(4)
Earnings surprise	-1.918***	-1.869***	-1.954***	-1.915***
	[0.256]	[0.267]	[0.271]	[0.282]
I{Negative surprise}	0.118***	0.062**	0.122***	0.074***
	[0.026]	[0.026]	[0.026]	[0.025]
I{Loss}	0.241***	0.358***	0.392***	0.509***
	[0.066]	[0.068]	[0.067]	[0.069]
I{Negative surprise}*I{Loss}	0.242***	0.202***	0.197***	0.158**
	[0.064]	[0.065]	[0.065]	[0.067]
log(Market value)	0.524***	0.470***	0.543***	0.489***
	[0.031]	[0.031]	[0.034]	[0.035]
B/M (/1000)	0.013***	0.014***	0.01	0.012*
	[0.003]	[0.003]	[0.008]	[0.007]
Institutional ownership (%)	0.002	0.004***	-0.001	0.001
	[0.001]	[0.001]	[0.001]	[0.001]
log(Value of trading volume [-61 to -2])	0.090***	0.203***	0.121***	0.234***
	[0.026]	[0.027]	[0.027]	[0.029]
log(Analyst attention)	0.513***	0.457***	0.497***	0.445***
	[0.044]	[0.044]	[0.044]	[0.044]
stdev(Analyst forecasts)	0.077***	0.067***	0.073***	0.065***
	[0.016]	[0.015]	[0.016]	[0.016]
Recent returns [-61 to -2]	-0.005***	-0.005***	-0.005***	-0.005***
	[0.000]	[0.000]	[0.000]	[0.000]
S&P500 returns [-61 to -2]	0.016***	0.006***	0.016***	0.005***
	[0.001]	[0.001]	[0.001]	[0.001]
stdev(S&P500 returns [-61 to -2])	-0.04	0.053*	-0.029	0.054*
	[0.028]	[0.028]	[0.029]	[0.029]
Year dummies	No	Yes	No	Yes
Industry dummies	No	No	Yes	Yes
Observations	178898	178898	178898	178898
Pseudo R-squared	0.21	0.22	0.24	0.25

Table 3.3: Determinants of positive and negative media coverage

This table reports coefficient estimates from multinomial logit regressions of $COVERAGE_{i,t}$ on market conditions and firm and event characteristics. $COVERAGE_{i,t}$ is equal to 1 if there is an identified positive WSJ earnings article within one week of the announcement date, -1 if there is an identified negative WSJ earnings article within one week of the announcement date, and 0 otherwise; *Earnings surprise* is the announced EPS less the median analyst forecast from the 30 days prior to the earnings announcement divided by the stock price; *I{Negative surprise}* is a dummy variable equal to one if *Earnings surprise* is negative, and zero otherwise; *I{Loss}* is a dummy variable equal to one if announced EPS is negative, and zero otherwise; *log(Market value)* is the natural logarithm of the number of shares outstanding multiplied by the closing stock price two days before the announcement; *B/M* is the book value of equity from the fiscal year ending in the previous year divided by the market value of equity from December 31 (divided by 1000); *Institutional ownership* is the percentage of equity held by institutions at the end of the previous calendar year; *log($Trading Volume)* is the natural logarithm of the firm's average dollar value of trading volume in the 60 trading days prior to the announcement; *log(Analyst attention)* is the natural log of the number of distinct analyst forecasts observed in I/B/E/S in the 30 days preceding the announcement; *stdev(forecasts)* is the standard deviation of the normalized analyst EPS forecasts recorded in I/B/E/S in the 30 days preceding the announcement; *Recent returns* is the average daily stock return during the 60 trading days prior to the announcement window; *S&P500 returns* is the average daily return on the S&P500 during the 60 trading days prior to the announcement window; *stdev(S&P500 returns)* is the standard deviation of average daily return on the S&P500 during the 60 trading days prior to the announcement window; results for year, month-of-the-year, day-of-the-week and industry dummies not shown; standard errors in brackets, robust to arbitrary heteroskedasticity and clustering by firm; * significant at 10%; ** significant at 5%;*** significant at 1%

	Negative coverage (1)	*Positive coverage* (1)	*Negative coverage* (2)	*Positive coverage* (2)	*Negative coverage* (3)	*Positive coverage* (3)	*Negative coverage* (4)	*Positive coverage* (4)	*Negative coverage* (5)	*Positive coverage* (5)
Earnings surprise	-1.958***	0.682	-1.906***	0.806	-1.874***	0.911	-1.838***	1.102	-1.859***	0.973
	[0.243]	[0.958]	[0.246]	[0.982]	[0.254]	[0.953]	[0.256]	[0.968]	[0.255]	[0.932]
I{Negative surprise}	0.701***	-0.065**	0.630***	-0.123***	0.672***	-0.047	0.612***	-0.095***	0.616***	-0.101***
	[0.035]	[0.030]	[0.035]	[0.030]	[0.036]	[0.030]	[0.036]	[0.029]	[0.036]	[0.029]
I{Loss}	1.574***	-1.674***	1.645***	-1.497***	1.726***	-1.534***	1.785***	-1.359***	1.772***	-1.382***
	[0.070]	[0.134]	[0.072]	[0.134]	[0.070]	[0.134]	[0.072]	[0.134]	[0.069]	[0.133]
I{Negative surprise}*I{Loss}	-0.356***	0.486***	-0.367***	0.421***	-0.403***	0.453***	-0.410***	0.392***	-0.394***	0.425***
	[0.067]	[0.147]	[0.068]	[0.148]	[0.068]	[0.148]	[0.069]	[0.149]	[0.067]	[0.145]
log(Market value of equity)	0.407***	0.581***	0.339***	0.460***	0.412***	0.603***	0.347***	0.483***	0.335***	0.473***
	[0.037]	[0.043]	[0.038]	[0.044]	[0.038]	[0.043]	[0.039]	[0.044]	[0.037]	[0.042]
B/M (/1000)	0.017*	0.015***	0.019**	0.016***	0.015	0.012	0.017	0.013		
	[0.009]	[0.004]	[0.008]	[0.005]	[0.013]	[0.009]	[0.011]	[0.009]		
Institutional ownership (%)	0.001	0.001	0.003**	0.004***	-0.001	-0.002	0.001	0		

	Negative coverage	Positive coverage	Negative coverage	Positive coverage	Negative coverage	Positive coverage	Negative coverage	Positive coverage	Negative coverage	Positive coverage
	(1)		(2)		(3)		(4)		(5)	
	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]		
log(Value of trading volume [-61 to -2])	0.108***	0.119***	0.227***	0.301***	0.159***	0.146***	0.276***	0.329***	0.283***	0.338***
	[0.031]	[0.035]	[0.032]	[0.038]	[0.031]	[0.035]	[0.033]	[0.038]	[0.030]	[0.035]
US firm dummy	0.134	-0.034	-0.094	-0.353**	0.166	-0.043	-0.066	-0.373**	-0.066	-0.351**
	[0.203]	[0.164]	[0.198]	[0.159]	[0.192]	[0.168]	[0.193]	[0.174]	[0.180]	[0.162]
log(Analyst attention)	0.568***	0.465***	0.495***	0.404***	0.542***	0.449***	0.469***	0.392***	0.484***	0.383***
	[0.049]	[0.049]	[0.050]	[0.050]	[0.049]	[0.050]	[0.049]	[0.050]	[0.047]	[0.048]
stdev(Analyst forecasts)	0.109***	-0.009	0.099***	-0.023	0.102***	-0.001	0.094***	-0.004	0.092***	-0.004
	[0.019]	[0.028]	[0.019]	[0.029]	[0.019]	[0.025]	[0.019]	[0.025]	[0.018]	[0.024]
Stock returns [-61 to -2]	-0.007***	-0.005***	-0.007***	-0.004***	-0.007***	-0.004***	-0.007***	-0.004***	-0.007***	-0.004***
	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]
stdev(Stock returns [-61 to -2])	5.008***	-5.410***	4.064**	-10.656***	7.666***	-6.798***	7.034***	-12.038***	5.766***	-12.615***
	[1.393]	[1.848]	[1.588]	[2.217]	[1.426]	[1.999]	[1.597]	[2.333]	[1.557]	[2.248]
S&P500 returns [-61 to -2]	0.013***	0.017***	0.009***	0.003**	0.012***	0.016***	0.008***	0.003	0.009***	0.002
	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]
stdev(S&P500 returns [-61 to -2])	-0.018	-0.026	-0.025	0.173***	-0.044	-0.005	-0.069	0.192***	-0.059	0.194***
	[0.044]	[0.038]	[0.049]	[0.039]	[0.044]	[0.038]	[0.050]	[0.041]	[0.048]	[0.040]
Year dummies	No		Yes		No		Yes		Yes	
Industry dummies	No		No		Yes		Yes		Yes	
Observations	178898		178898		178898		178898		190807	
Pseudo R-squared	0.21		0.23		0.24		0.26		0.26	

Table 3.4: Firm and event characteristics by coverage probability decile

Panel A: Averages by absolute coverage (PMC) decile

	decile 1	*decile 2*	*decile 3*	*decile 4*	*decile 5*	*decile 6*	*decile 7*	*decile 8*	*decile 9*	*decile 10*
PMC	0.008413	0.017302	0.027334	0.039771	0.05629	0.079044	0.112622	0.16724	0.2703	0.527344
Earnings surprise	0.003717	-0.00032	-0.00212	-0.00237	-0.00293	-0.00388	-0.00405	-0.00469	-0.00501	-0.00338
Earnings (%)	-0.00524	-0.00251	0.001936	0.000777	0.000943	0.000186	0.003364	-0.00165	-0.00049	0.003925
Analyst attention	2.492233	3.027241	3.529878	4.125421	4.789786	5.632891	6.752067	8.322849	10.75336	14.55817
Market value	106886.6	206547.7	313184.5	454681.7	642810.9	901031.8	1379579	2234809	4301189	2.04E+07
$Trading volume	416132.9	1004774	1677842	2585093	3762556	5384633	8162267	1.35E+07	2.65E+07	9.54E+07
B/M (/1000)	0.003478	0.007449	0.005894	0.011443	0.028207	0.018266	0.021493	0.020908	0.022844	0.027833
Beta	0.515972	0.657666	0.729581	0.765989	0.802008	0.837004	0.86359	0.923508	0.95262	1.012684
Institutional ownership (%)	27.86756	35.44424	39.93365	43.33074	46.72175	49.53634	51.67275	54.58726	57.68943	58.00283
Recent returns	0.116278	0.097904	0.088688	0.079582	0.071903	0.0595	0.05621	0.04866	0.049333	0.050344
stdev(Recent returns)	0.034606	0.033447	0.032075	0.030665	0.029466	0.028437	0.027283	0.026306	0.024594	0.021796
CAR [-1 to 3]	0.676465	0.196768	-0.05085	-0.09165	-0.07879	0.024408	0.127856	-0.05123	0.063268	0.10237

Panel B: Averages by positive coverage (PMC^+) decile

	decile 1	*decile 2*	*decile 3*	*decile 4*	*decile 5*	*decile 6*	*decile 7*	*decile 8*	*decile 9*	*decile 10*
PMC^+	0.002361	0.006238	0.01131	0.018388	0.028197	0.042436	0.064677	0.102894	0.182529	0.410472
Earnings surprise	-0.01904	-0.00533	-0.0024	-0.00091	-9.3E-05	4.50E-06	0.000282	0.000421	0.000415	0.000278
Earnings (%)	-0.08633	-0.01245	0.004044	0.009397	0.012284	0.014246	0.015079	0.015418	0.012773	0.01064
Analyst attention	2.790695	3.271821	3.653985	4.230664	4.771599	5.487378	6.491396	8.021265	10.45043	14.44459
Market value	104382.9	197716	300878.6	436736.3	643220.1	919497.4	1386050	2199159	4059168	2.04E+07
$Trading volume	717795.8	1394021	2077001	3033687	4484215	6045202	8428405	1.32E+07	2.45E+07	9.28E+07
B/M (/1000)	0.007896	0.006333	0.007145	0.016943	0.024341	0.013944	0.018165	0.023894	0.011168	0.037184
Beta	0.679528	0.717723	0.737084	0.762249	0.780005	0.799122	0.815879	0.866857	0.911758	0.981299
Institutional ownership (%)	28.86881	34.01651	38.31315	42.45004	46.14008	49.60592	52.20011	54.83252	58.33189	58.2843
Recent returns	0.031128	0.060192	0.07302	0.078432	0.080355	0.0785	0.080836	0.078425	0.07621	0.077095
stdev(Recent returns)	0.04575	0.035586	0.032185	0.030013	0.027894	0.026498	0.024949	0.024113	0.022576	0.020481
CAR [-1 to 3]	-0.3868	-0.10396	-0.03111	0.032021	-0.01119	0.240443	0.221948	0.281683	0.341351	0.282786

Panel C: Averages by negative coverage (PMC^-) decile

	decile 1	*decile 2*	*decile 3*	*decile 4*	*decile 5*	*decile 6*	*decile 7*	*decile 8*	*decile 9*	*decile 10*
PMC^-	0.00289	0.006128	0.009698	0.014065	0.019719	0.027435	0.038257	0.054596	0.081806	0.178857
Earnings surprise	0.004752	0.001861	0.000766	-0.00016	-0.00144	-0.0022	-0.00317	-0.0041	-0.00615	-0.01584
Earnings (%)	0.01818	0.010028	0.002737	0.00098	0.003628	0.002024	-0.00062	-0.00348	-0.00674	-0.02656
Analyst attention	2.603615	3.23384	3.850433	4.462886	5.272751	6.247223	7.477658	8.733793	10.70223	11.82548
Market value	162139.9	308005.6	465044.8	671939.6	1009273	1710435	3234825	5160103	9273032	9484960
$Trading volume	644585.3	1465221	2390599	3585826	5261388	8320019	1.46E+07	2.32E+07	4.73E+07	5.44E+07
B/M (/1000)	0.00639	0.008788	0.010718	0.01534	0.015739	0.013679	0.029765	0.034414	0.016829	0.017235
Beta	0.513952	0.655139	0.708466	0.753413	0.789071	0.821657	0.875899	0.909814	0.971595	1.081843
Institutional ownership (%)	30.32216	38.06048	41.99662	45.00486	47.48159	49.20917	51.41433	53.14553	54.55903	54.90458
Recent returns	0.143468	0.121738	0.109492	0.090536	0.080553	0.065855	0.053646	0.047855	0.030247	-0.03035
stdev(Recent returns)	0.030146	0.030229	0.029934	0.029529	0.028772	0.028419	0.027824	0.027333	0.0268	0.029479
CAR [-1 to 3]	0.96151	0.581536	0.248212	0.125852	-0.0518	-0.02108	-0.02652	-0.05419	-0.18949	-0.69694

Table 3.5: Media coverage probability trading profits – Time-series tests

This table reports estimates of excess returns from OLS regressions of monthly PMC-portfolio returns (minus the risk-free rate) on a constant plus risk factors identified in the literature. Portfolios are formed based upon the most recent predicted probability of positive, negative, or no coverage in the preceding three months from the multinomial regression described in column 5 of Table 3.3 (omitting the year dummy variables). Portfolios are re-balanced monthly and returns are equal-weighted; * significant at 10%; ** significant at 5%; *** significant at 1%; Newey-West standard errors.

Panel A: Monthly portfolio alphas – CAPM model (MKT-Rf)

	Media coverage probability decile										
	1 (low)	2	3	4	5	6	7	8	9	10 (high)	1 minus 10
Pr{Positive or Negative coverage | X}	0.00622**	0.00291	0.0013	0.00048	-0.0009	-0.0018	-0.00128	-0.00171	-0.00116	-0.00223***	0.00845***
	[0.00295]	[0.00240]	[0.00205]	[0.00196]	[0.00172]	[0.00171]	[0.00152]	[0.00141]	[0.00129]	[0.00083]	[0.00300]
Pr{Positive coverage | X}	-0.00392	-0.00023	0.00051	0.00019	0.00075	0.0014	0.00073	0.0004	0.00102	0.00076	-0.00468
	[0.00347]	[0.00244]	[0.00211]	[0.00192]	[0.00167]	[0.00168]	[0.00171]	[0.00150]	[0.00156]	[0.00076]	[0.00367]
Pr{Negative coverage | X}	0.00769***	0.00567**	0.00259	0.00124	0.00012	-0.00032	-0.00237	-0.00181	-0.00338**	-0.00770***	0.01539***
	[0.00287]	[0.00230]	[0.00194]	[0.00190]	[0.00183]	[0.00161]	[0.00151]	[0.00133]	[0.00135]	[0.00214]	[0.00378]

Panel B: Monthly portfolio alphas – Fama-French 3-factor model (MKT-Rf, SMB, HML)

	Media coverage probability decile										
	1 (low)	2	3	4	5	6	7	8	9	10 (high)	1 minus 10
Pr{Positive or Negative coverage | X}	0.00443***	0.00193	0.0002	0.00001	-0.00148	-0.00265**	-0.00161	-0.00223*	-0.00183	-0.00254***	0.00697***
	[0.00165]	[0.00130]	[0.00105]	[0.00114]	[0.00103]	[0.00125]	[0.00131]	[0.00128]	[0.00126]	[0.00087]	[0.00164]
Pr{Positive coverage | X}	-0.00299	-0.00042	0.00047	-0.00076	-0.00025	0.00001	-0.00072	-0.00097	-0.00055	0.00024	-0.00323
	[0.00226]	[0.00182]	[0.00146]	[0.00126]	[0.00100]	[0.00100]	[0.00112]	[0.00099]	[0.00108]	[0.00070]	[0.00230]
Pr{Negative coverage | X}	0.00555***	0.00410***	0.00135	0	-0.00092	-0.001	-0.00277**	-0.00223**	-0.00320**	-0.00673***	0.01227***
	[0.00162]	[0.00121]	[0.00101]	[0.00105]	[0.00113]	[0.00091]	[0.00116]	[0.00113]	[0.00142]	[0.00228]	[0.00293]

Panel C: Monthly portfolio alphas – Carhart 4-factor model (MKT-Rf, SMB, HML, WML)

	Media coverage probability decile										
	1 (low)	2	3	4	5	6	7	8	9	10 (high)	1 minus 10
Pr{Positive or Negative coverage | X}	0.00662***	0.00424***	0.00248**	0.00216*	0.00064	-0.00015	0.00081	0.0003	0.00023	-0.00035	0.00697***
	[0.00163]	[0.00115]	[0.00114]	[0.00113]	[0.00090]	[0.00090]	[0.00114]	[0.00123]	[0.00112]	[0.00082]	[0.00171]
Pr{Positive coverage | X}	0.00379	0.00399**	0.00319**	0.00151	0.00106	0.00117	0.0005	0.00016	0.00042	0.00113*	0.00266
	[0.00272]	[0.00197]	[0.00125]	[0.00095]	[0.00083]	[0.00087]	[0.00106]	[0.00097]	[0.00103]	[0.00067]	[0.00277]
Pr{Negative coverage | X}	0.00641***	0.00508***	0.00286***	0.00182	0.00095	0.00106	-0.00047	0.00033	-0.00002	-0.00111	0.00752***
	[0.00150]	[0.00115]	[0.00091]	[0.00111]	[0.00096]	[0.00105]	[0.00098]	[0.00093]	[0.00138]	[0.00213]	[0.00282]

Panel D: Monthly portfolio alphas – Pastor-Stambaugh 5-factor model (MKT-Rf, SMB, HML, WML, PS)

	Media coverage probability decile										
	1 (low)	2	3	4	5	6	7	8	9	10 (high)	1 minus 10
Pr{Positive or Negative coverage | X}	0.00625***	0.00394***	0.00196	0.00174	0.00066	-0.00023	0.00007	-0.00027	-0.00064	-0.00036	0.00661***
	[0.00167]	[0.00126]	[0.00127]	[0.00121]	[0.00089]	[0.00101]	[0.00109]	[0.00112]	[0.00109]	[0.00081]	[0.00185]
Pr{Positive coverage | X}	0.00389	0.00369*	0.00280**	0.00147	0.00047	0.0011	-0.00029	-0.00078	-0.00035	0.00105*	0.00285
	[0.00281]	[0.00191]	[0.00124]	[0.00102]	[0.00091]	[0.00099]	[0.00120]	[0.00096]	[0.00103]	[0.00063]	[0.00291]
Pr{Negative coverage | X}	0.00585***	0.00486***	0.00237**	0.0017	0.00038	0.00069	-0.00117	0.00007	-0.00015	-0.00152	0.00738**
	[0.00152]	[0.00128]	[0.00106]	[0.00125]	[0.00106]	[0.00101]	[0.00106]	[0.00084]	[0.00143]	[0.00237]	[0.00323]

Table 3.6: Media coverage probability portfolios – Factor loadings (MKT-Rf, SMB, HML, WML)

This table reports estimates of excess returns from OLS regressions of monthly PMC-portfolio returns (minus the risk-free rate) on a constant plus MKT-Rf, SMB, HML, and WML. Portfolios are formed based upon the most recent predicted probability of positive, negative, or no coverage in the preceding three months from the multinomial regression described in column 5 of Table 3.3 (omitting the year dummy variables); Portfolios are re-balanced monthly and returns are equal-weighted; * significant at 10%; ** significant at 5%; *** significant at 1%; Newey-West standard errors.

Panel A: Portfolios formed on $Pr\{\text{Positive or Negative coverage} \mid X\}_{t-1}$

	Absolute (i.e. Positive or Negative) Coverage Probability Deciles										
	1 (low)	2	3	4	5	6	7	8	9	10 (high)	1 minus 10
MKT-Rf	1.00392***	1.08844***	1.07511***	1.09574***	1.13996***	1.15116***	1.12191***	1.17927***	1.18616***	1.13669***	-0.13278***
	[0.04081]	[0.03274]	[0.02602]	[0.02285]	[0.02694]	[0.02706]	[0.02971]	[0.02450]	[0.03170]	[0.02054]	[0.04356]
SMB	0.89104***	0.90576***	0.89238***	0.80243***	0.71071***	0.62436***	0.56984***	0.42530***	0.31365***	0.05382**	0.83722***
	[0.07104]	[0.06555]	[0.06355]	[0.04530]	[0.05301]	[0.06345]	[0.04597]	[0.04790]	[0.04445]	[0.02729]	[0.06943]
HML	0.32118***	0.18504***	0.20392***	0.09423**	0.10812**	0.14231***	0.0516	0.07408*	0.09910**	0.02286	0.29833***
	[0.08418]	[0.05785]	[0.05721]	[0.04234]	[0.04448]	[0.04643]	[0.04404]	[0.03937]	[0.04105]	[0.02876]	[0.08594]
Momentum	-0.21324***	-0.22622***	-0.22272***	-0.21025***	-0.20681***	-0.24391***	-0.23633***	-0.24636***	-0.20027***	-0.21400***	0.00076
	[0.06179]	[0.03717]	[0.04544]	[0.03082]	[0.03027]	[0.02674]	[0.03809]	[0.04794]	[0.03645]	[0.02520]	[0.07111]
alpha	0.00662***	0.00424***	0.00248**	0.00216*	0.00064	-0.00015	0.00081	0.0003	0.00023	-0.00035	0.00697***
	[0.00163]	[0.00115]	[0.00114]	[0.00113]	[0.00090]	[0.00090]	[0.00114]	[0.00123]	[0.00112]	[0.00082]	[0.00171]
Observations	256	256	256	256	256	256	256	256	256	256	256
R-squared	0.84	0.91	0.93	0.94	0.95	0.94	0.94	0.94	0.94	0.96	0.53

Panel B: Portfolios formed on $Pr\{\text{Positive coverage} \mid X\}_{t-1}$

	Positive Coverage Probability Deciles										
	1 (low)	2	3	4	5	6	7	8	9	10 (high)	1 minus 10
MKT-Rf	1.14264***	1.13650***	1.09520***	1.09535***	1.13327***	1.13535***	1.10124***	1.12067***	1.15004***	1.07008***	0.07255
	[0.05870]	[0.03744]	[0.04250]	[0.02636]	[0.02846]	[0.02696]	[0.02725]	[0.02122]	[0.02827]	[0.01850]	[0.06533]
SMB	1.29866***	1.00005***	0.91523***	0.79823***	0.67610***	0.55080***	0.45549***	0.34661***	0.20285***	-0.04702	1.34568***
	[0.07918]	[0.07033]	[0.05502]	[0.06183]	[0.05412]	[0.06470]	[0.06799]	[0.05289]	[0.05108]	[0.02936]	[0.08546]
HML	-0.18167*	0.02468	0.01973	0.17288***	0.18954***	0.24921***	0.25215***	0.23494***	0.26161***	0.07230**	-0.25396**
	[0.10872]	[0.07008]	[0.06129]	[0.04544]	[0.05750]	[0.06756]	[0.06279]	[0.04923]	[0.06196]	[0.03341]	[0.12669]
Momentum	-0.66268***	-0.43049***	-0.26536***	-0.22111***	-0.12803***	-0.11280***	-0.11937***	-0.11079***	-0.09455***	-0.08761***	-0.57507***
	[0.11655]	[0.09361]	[0.04832]	[0.03538]	[0.03180]	[0.03823]	[0.03688]	[0.03224]	[0.02708]	[0.02398]	[0.12477]
alpha	0.00379	0.00399**	0.00319**	0.00151	0.00106	0.00117	0.0005	0.00016	0.00042	0.00113*	0.00266
	[0.00272]	[0.00197]	[0.00125]	[0.00095]	[0.00083]	[0.00087]	[0.00106]	[0.00097]	[0.00103]	[0.00067]	[0.00277]
Observations	256	256	256	256	256	256	256	256	256	256	256
R-squared	0.84	0.89	0.93	0.94	0.95	0.93	0.92	0.93	0.92	0.96	0.68

Panel C: Portfolios formed on $Pr\{\text{Negative coverage} \mid X\}_{t-1}$

	Negative Coverage Probability Deciles										
	1 (low)	2	3	4	5	6	7	8	9	10 (high)	1 minus 10
MKT-Rf	0.99986***	1.06314***	1.09610***	1.10765***	1.08789***	1.10545***	1.11802***	1.12094***	1.16994***	1.31132***	-0.31146***
	[0.04295]	[0.02730]	[0.02535]	[0.02249]	[0.02836]	[0.02044]	[0.02979]	[0.02690]	[0.02905]	[0.05351]	[0.08613]
SMB	0.77203***	0.78715***	0.73100***	0.66560***	0.66586***	0.62559***	0.52462***	0.43760***	0.41571***	0.55980***	0.21223*
	[0.08121]	[0.06973]	[0.05154]	[0.07171]	[0.05220]	[0.04745]	[0.04880]	[0.04238]	[0.05278]	[0.06377]	[0.11111]
HML	0.39421***	0.29671***	0.22880***	0.22022***	0.18732***	0.12032***	0.06376	0.05711	-0.05322	-0.21614*	0.61035***
	[0.09377]	[0.08047]	[0.05910]	[0.05813]	[0.05311]	[0.03914]	[0.04687]	[0.03535]	[0.05498]	[0.11451]	[0.19054]
Momentum	-0.08458	-0.09590*	-0.14696***	-0.17769***	-0.18265***	-0.20137***	-0.22477***	-0.24974***	-0.31133***	-0.54903***	0.46445***
	[0.06440]	[0.04941]	[0.03767]	[0.03645]	[0.02677]	[0.03763]	[0.02843]	[0.03376]	[0.05907]	[0.10147]	[0.14250]
alpha	0.00641***	0.00508***	0.00286***	0.00182	0.00095	0.00106	-0.00047	0.00033	-0.00002	-0.00111	0.00752***
	[0.00150]	[0.00115]	[0.00091]	[0.00111]	[0.00096]	[0.00105]	[0.00098]	[0.00093]	[0.00138]	[0.00213]	[0.00282]
Observations	256	256	256	256	256	256	256	256	256	256	256
R-squared	0.84	0.9	0.93	0.94	0.94	0.94	0.94	0.95	0.93	0.89	0.43

Table 3.7: Media coverage probability portfolios – Double-sorts

This table reports estimates of excess returns from OLS regressions of monthly PMC-portfolio returns (minus the risk-free rate) on a constant plus MKT-Rf, SMB, HML, and WML. For each month, stocks are sorted first into quintiles based on firm or event characteristics, and then, within each of these quintiles, stocks are further sorted into quintiles based on lagged PMC. * significant at 10%; ** significant at 5%;*** significant at 1%; Newey-West standard errors.

Panel A: Portfolios formed within Market Value quintiles (Four-factor alphas)

		Probability of Negative Media Coverage (PMC) Quintiles					
		1(low)	2	3	4	5(high)	1 minus 5
	(low)1	0.00580***	0.00669***	0.00570***	0.00412	-0.00088	0.00668
		[0.00215]	[0.00213]	[0.00210]	[0.00329]	[0.00426]	[0.00406]
	2	0.00702***	0.00346**	0.00138	-0.00167	-0.00039	0.00741**
Market value		[0.00139]	[0.00154]	[0.00149]	[0.00137]	[0.00291]	[0.00343]
Quintiles	3	0.00593***	0.00286*	0.00085	0	-0.0029	0.00867***
		[0.00125]	[0.00146]	[0.00114]	[0.00115]	[0.00195]	[0.00250]
	4	0.00264*	0.00177	-0.00063	-0.00021	-0.00117	0.00381
		[0.00137]	[0.00129]	[0.00127]	[0.00122]	[0.00196]	[0.00249]
	(high)5	0.00149	0.00094	0.00057	0.00001	-0.00101	0.00233
		[0.00096]	[0.00123]	[0.00088]	[0.00083]	[0.00126]	[0.00190]

Panel B: Portfolios formed within Analyst attention quintiles (Four-factor alphas)

		Probability of Negative Media Coverage (PMC) Quintiles					
		1(low)	2	3	4	5(high)	1 minus 5
	(low)1	0.00710***	0.00599***	0.00444***	-0.00137	-0.00560**	0.01270***
		[0.00180]	[0.00117]	[0.00105]	[0.00176]	[0.00247]	[0.00305]
	2	0.00382**	0.00181	0.00287**	0.00048	-0.00268	0.00650**
Analyst attention		[0.00171]	[0.00125]	[0.00136]	[0.00178]	[0.00229]	[0.00289]
Quintiles	3	0.00508***	0.00354**	0.00182	0.00056	-0.00071	0.00579*
		[0.00161]	[0.00142]	[0.00121]	[0.00151]	[0.00248]	[0.00327]
	4	0.00343**	0.0002	0.00217*	-0.00055	-0.00039	0.00382
		[0.00154]	[0.00145]	[0.00117]	[0.00114]	[0.00237]	[0.00305]
	(high)5	0.00326*	0.00032	0.00244*	0.00226**	0.00105	0.00221
		[0.00176]	[0.00147]	[0.00125]	[0.00106]	[0.00218]	[0.00294]

Panel C: Portfolios formed within B/M quintiles (Four-factor alphas)

		Probability of Negative Media Coverage (PMC) Quintiles					
		1(low)	2	3	4	5(high)	1 minus 5
	(low)1	0.00522***	0.00353	0.00015	-0.00002	0.00037	0.00485
		[0.00175]	[0.00252]	[0.00230]	[0.00246]	[0.00319]	[0.00352]
	2	0.00607***	0.00011	-0.00075	0.00114	-0.00044	0.00651***
B/M		[0.00191]	[0.00127]	[0.00130]	[0.00163]	[0.00145]	[0.00199]
Quintiles	3	0.00410**	0.00118	0.00101	-0.00190*	-0.00042	0.00527***
		[0.00170]	[0.00098]	[0.00119]	[0.00115]	[0.00119]	[0.00190]
	4	0.00674***	0.00229**	0.00059	0.00123	-0.00153	0.00827***
		[0.00166]	[0.00112]	[0.00124]	[0.00137]	[0.00149]	[0.00198]
	(high)5	0.00739***	0.00372**	0.00189	-0.00025	-0.00041	0.00806***
		[0.00149]	[0.00144]	[0.00162]	[0.00188]	[0.00235]	[0.00257]

Panel D: Portfolios formed within Beta (market) quintiles (Four-factor alphas)

		Probability of Negative Media Coverage (PMC) Quintiles					
		1(low)	2	3	4	5(high)	1 minus 5
	(low)1	0.00604***	0.00243**	0.00052	-0.00384***	-0.00867***	0.01471***
		[0.00195]	[0.00102]	[0.00121]	[0.00124]	[0.00200]	[0.00225]
	2	0.00613***	0.00263**	0.00082	-0.00082	-0.00341***	0.00955***
Beta		[0.00162]	[0.00103]	[0.00132]	[0.00115]	[0.00122]	[0.00162]
Quintiles	3	0.00594***	0.00144	0.00118	-0.00059	-0.00149	0.00723***
		[0.00159]	[0.00144]	[0.00101]	[0.00118]	[0.00150]	[0.00171]
	4	0.00579***	0.00166	0.00299**	0.0016	-0.00033	0.00612**
		[0.00188]	[0.00156]	[0.00150]	[0.00158]	[0.00231]	[0.00236]
	(high)5	0.00756***	0.00404	0.00251	0.00458*	0.00305	0.00480*
		[0.00202]	[0.00247]	[0.00275]	[0.00271]	[0.00309]	[0.00274]

Panel E: Portfolios formed within Recent Returns quintiles (Four-factor alphas)

		Probability of Negative Media Coverage (PMC) Quintiles					
		1(low)	2	3	4	5(high)	1 minus 5
	(low)1	0.00631***	0.00463*	0.00401*	0.00433	0.00184	0.00447
		[0.00186]	[0.00254]	[0.00240]	[0.00267]	[0.00335]	[0.00287]
	2	0.00473***	0.00513***	0.00205	0.00301*	0.00016	0.00456***
Recent Returns		[0.00155]	[0.00154]	[0.00143]	[0.00156]	[0.00110]	[0.00167]
Quintiles	3	0.00524***	0.00221	0.00143	-0.00008	-0.00131	0.00668***
		[0.00144]	[0.00141]	[0.00099]	[0.00140]	[0.00102]	[0.00184]
	4	0.00456***	0.00011	-0.00210*	-0.00298***	-0.00531***	0.00987***
		[0.00143]	[0.00110]	[0.00113]	[0.00083]	[0.00106]	[0.00201]
	(high)5	0.00804***	0.00424***	-0.00192	-0.00244	-0.00363*	0.01146***
		[0.00184]	[0.00138]	[0.00208]	[0.00197]	[0.00195]	[0.00285]

Panel F1: Portfolios formed within Earnings Surprise quintiles (Four-factor alphas)

		Probability of Negative Media Coverage (PMC) Quintiles					
		1(low)	2	3	4	5(high)	1 minus 5
	(low)1	-0.00428**	-0.00461***	-0.00208	-0.00143	-0.00355	-0.00074
		[0.00168]	[0.00169]	[0.00218]	[0.00202]	[0.00321]	[0.00344]
	2	-0.00097	-0.00190*	-0.00225**	-0.00232**	-0.00073	-0.00024
Earnings surprise		[0.00129]	[0.00100]	[0.00103]	[0.00113]	[0.00130]	[0.00199]
Quintiles	3	0.00874***	0.00509***	0.00368***	0.00171	0.00346**	0.00528*
		[0.00181]	[0.00133]	[0.00137]	[0.00122]	[0.00162]	[0.00275]
	4	0.00965***	0.00506***	0.00270*	0.0014	0.00096	0.00854***
		[0.00186]	[0.00126]	[0.00144]	[0.00157]	[0.00175]	[0.00236]
	(high)5	0.00987***	0.01116***	0.00767***	0.00377*	0.00186	0.00874***
		[0.00195]	[0.00175]	[0.00219]	[0.00213]	[0.00277]	[0.00287]

Panel F2: Portfolios formed within Earnings Surprise quintiles (Four-factor alphas)

		Probability of Media Coverage (PMC) Quintiles					
		1(low)	2	3	4	5(high)	1 minus 5
	(low)1	-0.00414**	-0.00365*	-0.00285	-0.00302	-0.00232	-0.00182
		[0.00188]	[0.00201]	[0.00180]	[0.00203]	[0.00286]	[0.00294]
	2	-0.00075	-0.00202**	-0.00225*	-0.00248**	-0.00067	-0.00008
Earnings Surprise		[0.00115]	[0.00095]	[0.00117]	[0.00108]	[0.00090]	[0.00139]
Quintiles	3	0.01002***	0.00413***	0.00279**	0.00205**	0.00395***	0.00606***
		[0.00173]	[0.00133]	[0.00136]	[0.00103]	[0.00101]	[0.00194]
	4	0.00832***	0.00642***	0.00255*	0.00143	0.00123	0.00716***
		[0.00165]	[0.00150]	[0.00147]	[0.00157]	[0.00121]	[0.00164]
	(high)5	0.01323***	0.00986***	0.00733***	0.00369**	0.00039	0.01264***
		[0.00210]	[0.00237]	[0.00214]	[0.00182]	[0.00158]	[0.00234]

Table 3.8: Cross-sectional regressions

Pooled OLS regressions of monthly stock returns minus the risk-free rate on the natural log of the previous month's value of PMC, lagged firm-event characteristics, and year and industry dummy variables; standard errors in brackets, robust to arbitrary heteroskedasticity and clustering by firm; * significant at 10%; ** significant at 5%;*** significant at 1%.

	$R_{i,t+1}-r_f$	$R_{i,t+1}-r_f$	$R_{i,t+1}-r_f$	$R_{i,t+1}-r_f$	$R_{i,t+1}-r_f$	$R_{i,t+1}-r_f$
log(PMC)	-0.00784***			-0.01534***		
	[0.00098]			[0.00136]		
log(PMC⁺)		0.00306***			0.00156*	
		[0.00074]			[0.00091]	
log(PMC⁻)			-0.00518***			-0.00690***
			[0.00056]			[0.00066]
log(Market value)	0.00417***	-0.00358***	0.00158***	0.00550***	-0.00178**	0.00071
	[0.00072]	[0.00060]	[0.00040]	[0.00084]	[0.00073]	[0.00064]
B/M (/1000)	0.00035	0.00034	0.00035	0.00038	0.00036	0.00036
	[0.00032]	[0.00031]	[0.00032]	[0.00033]	[0.00032]	[0.00033]
Beta [-61 to -2]	-0.00324***	-0.00378***	-0.00324***	-0.00374***	-0.00380***	-0.00360***
	[0.00074]	[0.00074]	[0.00074]	[0.00076]	[0.00076]	[0.00076]
Recent returns [-61 to -2]	-0.00019***	-0.00013***	-0.00018***	-0.00020***	-0.00013***	-0.00017***
	[0.00002]	[0.00002]	[0.00002]	[0.00002]	[0.00002]	[0.00002]
stdev(Recent returns [-61 to -2])	0.07166	0.0413	0.10907**	0.11255**	0.07215	0.14988***
	[0.04714]	[0.04788]	[0.04889]	[0.05114]	[0.05434]	[0.05285]
Institutional ownership (%)				0.00013***	0.00012***	0.00012***
				[0.00002]	[0.00002]	[0.00002]
log($Trading volume [-61 to -2])				-0.00002	-0.00196***	-0.00089
				[0.00057]	[0.00057]	[0.00055]
log(Analyst attention}				0.00966***	0.00265***	0.00654***
				[0.00091]	[0.00084]	[0.00081]
Earnings surprise				0.00028	0.04554**	0.01704
				[0.01936]	[0.01882]	[0.01908]
stdev(forecasts)				0.00082	-0.00063	0.00051
				[0.00055]	[0.00053]	[0.00054]
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Industry dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	180480	180480	180480	178144	178144	178144
R-squared	0.01	0.01	0.01	0.02	0.01	0.02

Table 3.9: Cross-sectional (Fama-MacBeth) regressions

Results are based on time series of estimated coefficients obtained from monthly regressions of returns on PMC values from the preceding month and lagged control variables; Fama-MacBeth standard errors in brackets; * significant at 10%; ** significant at 5%;*** significant at 1%.

Panel A

	Absolute Coverage Decile 1 dummy	*Absolute Coverage Decile 10 dummy*	*Positive Coverage Decile 1 dummy*	*Positive Coverage Decile 10 dummy*	*Negative Coverage Decile 1 dummy*	*Negative Coverage Decile 10 dummy*
Avg. coefficient	0.00638***	0.00002	-0.00031	0.00013	0.00632***	-0.00506**
F-M std. error	[0.00192]	[0.00199]	[0.00248]	[0.00194]	[0.00150]	[0.00203]
Number of monthly coefficients	257	257	257	257	257	257

Panel B

	log of Absolute Coverage Probability	*log of Positive Coverage Probability*	*log of Negative Coverage Probability*
Avg. coefficient	-0.01201*	-0.00091	-0.04254***
F-M std. error	[0.00673]	[0.00755]	[0.01242]
Number of monthly coefficients	257	257	257

Table 3.10: How does probability of media coverage predict stock returns?
Monthly four-factor alphas for decile portfolios formed on restricted versions of the PMC measures; each measure is orthogonalized with respect to cross-sectional estimates of size, B/M, beta, and momentum; Newey-West standard errors in brackets; * significant at 10%; ** significant at 5%;*** significant at 1%.

Panel A: PMC1 portfolio returns – expected media coverage conditional on the earnings surprise, the negative surprise dummy, the loss dummy, and the interaction between loss and negative surprise (orthogonalized with respect to size, B/M, beta, and momentum)

	four-factor alpha		
	PMC Decile 1 (low)	PMC Decile 10 (high)	PMC Decile 1 – Decile 10
PMC1 (probability of any coverage)	0.00178	0.00800***	-0.00622***
	[0.00226]	[0.00209]	[0.00224]
$PMC1^+$ (probability of positive coverage)	-0.00217	0.00914***	-0.01131***
	[0.00182]	[0.00165]	[0.00222]
$PMC1^-$ (probability of negative coverage)	0.00870***	-0.00398*	0.01268***
	[0.00151]	[0.00215]	[0.00216]

Panel B: PMC2 portfolio returns – expected media coverage conditional on the seasonality, industry, and year dummies (orthogonalized with respect to size, B/M, beta, and momentum)

	four-factor alpha		
	PMC Decile 1 (low)	PMC Decile 10 (high)	PMC Decile 1 – Decile 10
PMC2 (probability of any coverage)	0.00189	0.00042	0.00148
	[0.00133]	[0.00164]	[0.00179]
$PMC2^+$ (probability of positive coverage)	0.00105	0.00094	0.00011
	[0.00113]	[0.00169]	[0.00187]
$PMC2^-$ (probability of negative coverage)	0.00337**	-0.00019	0.00355
	[0.00166]	[0.00157]	[0.00223]

Panel C: PMC3 portfolio returns – expected media coverage conditional on the control dummies (PMC2) and the earnings information variables (PMC1) (orthogonalized with respect to size, B/M, beta, and momentum)

	four-factor alpha		
	PMC Decile 1 (low)	PMC Decile 10 (high)	PMC Decile 1 – Decile 10
PMC3 (probability of any coverage)	0.00273*	0.00101	0.00172
	[0.00162]	[0.00139]	[0.00192]
$PMC3^+$ (probability of positive coverage)	-0.00177	0.00482***	-0.00659**
	[0.00214]	[0.00149]	[0.00266]
$PMC3^-$ (probability of negative coverage)	0.00620***	-0.0036	0.00980***
	[0.00162]	[0.00234]	[0.00282]

Panel D: PMC4 portfolio returns – expected media coverage conditional on: the earnings information variables (PMC1), plus log(analyst attention) and stdev(analyst forecasts) (orthogonalized with respect to size, B/M, beta, and momentum)

	four-factor alpha		
	PMC Decile 1 (low)	PMC Decile 10 (high)	PMC Decile 1 – Decile 10
PMC4 (probability of any coverage)	0.00234**	0.00208	0.00026
	[0.00094]	[0.00233]	[0.00268]
$PMC4^{+}$ (probability of positive coverage)	-0.00134	0.00245	-0.00379
	[0.00181]	[0.00164]	[0.00251]
$PMC4^{-}$ (probability of negative coverage)	0.00559***	0.00055	0.00504*
	[0.00123]	[0.00246]	[0.00294]

Panel E: PMC5 portfolio returns – expected media coverage conditional on: the control dummies and earnings information variables (PMC3), plus log(analyst attention) and stdev(analyst forecasts) (orthogonalized with respect to size, B/M, beta, and momentum)

	four-factor alpha		
	PMC Decile 1 (low)	PMC Decile 10 (high)	PMC Decile 1 – Decile 10
PMC5 (probability of any coverage)	0.00302***	0.00107	0.00195
	[0.00104]	[0.00199]	[0.00225]
$PMC5^{+}$ (probability of positive coverage)	-0.00085	0.00198	-0.00283
	[0.00157]	[0.00157]	[0.00227]
$PMC5^{-}$ (probability of negative coverage)	0.00513***	-0.00065	0.00578**
	[0.00128]	[0.00226]	[0.00274]

Panel F: PMC6 portfolio returns – expected media coverage conditional on: the control dummies, earnings information variables, and analyst variables (PMC5), plus value of trading volume, institutional ownership, and the U.S. firm dummy (orthogonalized with respect to size, B/M, beta, and momentum)

	four-factor alpha		
	PMC Decile 1 (low)	PMC Decile 10 (high)	PMC Decile 1 – Decile 10
PMC6 (probability of any coverage)	0.00515***	-0.00158	0.00673***
	[0.00131]	[0.00180]	[0.00195]
$PMC6^{+}$ (probability of positive coverage)	0.00185	0.00063	0.00121
	[0.00168]	[0.00151]	[0.00243]
$PMC6^{-}$ (probability of negative coverage)	0.00694***	-0.0033	0.01024***
	[0.00153]	[0.00223]	[0.00278]

Panel G: Long-range PMC5 portfolio returns
Month t+n four-factor alphas for long-short portfolios formed on the estimated media coverage probability observed in month t; expected media coverage conditional on: the control dummies and earnings information variables (PMC3), plus log(analyst attention) and stdev(analyst forecasts) (subsequently orthogonalized with respect to size, B/M, beta, and momentum); * significant at 10%; ** significant at 5%;*** significant at 1%.

Month	*four-factor alpha* Any coverage (PMC5) portfolio Decile1-Decile10	Positive coverage (PMC5+) portfolio Decile1-Decile10	Negative coverage (PMC5-) portfolio Decile1-Decile10
t+1	0.002	-0.0028	0.00578**
	[0.00224]	[0.00226]	[0.00274]
t+2	0.00049	-0.00282	0.00364
	[0.00223]	[0.00243]	[0.00278]
t+3	0.00133	-0.00265	0.00359
	[0.00238]	[0.00243]	[0.00286]
t+4	0.00052	-0.00099	0.0009
	[0.00247]	[0.00242]	[0.00275]
t+5	0.00028	-0.00181	0.00159
	[0.00234]	[0.00241]	[0.00254]
t+6	0.00156	-0.00089	0.0004
	[0.00240]	[0.00237]	[0.00283]

Panel H: Long-range PMC6 portfolio returns
Month t+n four-factor alphas for long-short portfolios formed on the estimated media coverage probability observed in month t; expected media coverage conditional on: the control dummies, earnings information variables, and analyst variables (PMC5), plus value of trading volume, institutional ownership, and the U.S. firm dummy (subsequently orthogonalized with respect to size, B/M, beta, and momentum); * significant at 10%; ** significant at 5%;*** significant at 1%.

Month	*four-factor alpha* Any coverage (PMC5) portfolio Decile1-Decile10	Positive coverage (PMC5+) portfolio Decile1-Decile10	Negative coverage (PMC5-) portfolio Decile1-Decile10
t+1	0.00676***	0.00124	0.01024***
	[0.00195]	[0.00243]	[0.00278]
t+2	0.00487**	0.00007	0.00624**
	[0.00200]	[0.00257]	[0.00294]
t+3	0.00722***	0.00288	0.00747**
	[0.00217]	[0.00259]	[0.00302]
t+4	0.00473**	0.00384	0.00582**
	[0.00188]	[0.00273]	[0.00281]
t+5	0.00429**	0.00392	0.00644**
	[0.00204]	[0.00278]	[0.00266]
t+6	0.00374*	0.00349	0.00442
	[0.00206]	[0.00253]	[0.00284]

References

Amihud, Yakov, and Haim Mendelson, 1986, Asset Pricing and the Bid-Ask Spread, *Journal of Financial Economics* 17(2), 223-249.

Antweiler, Werner, and Murray Z. Frank, 2004, Is all that talk just noise? The information content of internet stock message boards, *Journal of Finance* 59, 1259-1294.

Antweiler, Werner, and Murray Z. Frank, 2006, Do U.S. Stock Markets Typically Overreact to Corporate News Stories, Working paper, University of British Columbia and University of Minnesota.

Arbel, Avner, and Paul Strebel, 1982, The Neglected and Small Firm Effects, *Financial Review* 17, 201-218.

Ball, Ray, and Philip Brown, 1968, An Empirical Evaluation of Accounting Income Numbers, *Journal of Accounting Research* 6 (2), 159-178.

Bamber, Linda Smith, 1987, Unexpected Earnings, Firm Size, and Trading Volume Around Quarterly Earnings Announcements, *Accounting Review* 62 (3), 510-532.

Banz, R.W., 1981, The relationship between return and market value of common stocks, *Journal of Financial Economics* 9, 3–18.

Barber, Brad M., and Terrance Odean, 2008, All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors, *Review of Financial Studies* 21(2), 785-818.

Basak, Suleyman, and Domenico Cuoco, 1998, An Equilibrium Model with Restricted Stock Market Participation, *Review of Financial Studies* 11 (2), 309-341.

Bernard, Victor L., and Jacob K. Thomas, 1990, Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium, *Journal of Accounting Research* 27, 1-36.

Bhattacharya , Utpal, Neal Galpin, Rina Ray, and Xiaoyun Yu, 2006, The Role of the Media in the Internet IPO Bubble, *Journal of Financial and Quantitative Analysis*, forthcoming.

Brennan, Michael J., Narasimhan Jegadeesh, and Bhaskaran Swaminathan, 1993, Investment Analysis and the Adjustment of Stock Prices to Common Information, *Review of Financial Studies* 6(4), 799-824.

Busse, Jeffrey A., and T. Clifton Green, 2002, Market efficiency in real time, *Journal of Financial Economics* 65, 415-437.

Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57-82.

Chan, Wesley S., 2003, Stock price reaction to news and no-news: drift and reversal after headlines, *Journal of Financial Economics* 70, 223-260.

Christensen, Theodore E., Toni Q. Smith, and Pamela S. Stuerke, 2004, Public Predisclosure Information, Firm Size, Analyst Following, and Market Reactions to Earnings Announcements, *Journal of Business Finance & Accounting* 31 (7-8), 951-984.

Conrad, Jennifer, Bradford Cornell, and Wayne R. Landsman, 2002, When Is Bad News Really Bad News?, *Journal of Finance* 57 (6), 2507-2532.

Conrad, Jennifer, Bradford Cornell, Wayne R. Landsman, and Brian Rountree, 2004, How do Analyst Recommendation Respond to Major News?, Working paper, University of North Carolina – Chapel Hill, University of California – Los Angeles, and Rice University.

DellaVigna, Stefano, and Joshua Pollet, 2008, Investor Inattention and Friday Earnings Announcements, *Journal of Finance,* forthcoming.

DeLong, J. Bradford, Andrei Shleifer, Lawrence Summers, and Robert Waldmann, 1990, Positive feedback investment strategies and destabilizing rational speculation, *Journal of Finance* 45, 379-395.

Dyck, Alexander, and Luigi Zingales, 2003, The media and asset prices, Working paper, Harvard Business School and University of Chicago.

Easley, David, Soeren Hvidkjaer, and Maureen O'Hara, 2002, Is information risk a determinant of asset returns?, *Journal of Finance* 57(5), 2185-2221.

Engelberg, Joseph, 2007, Costly Information Processing: Evidence from Earnings Announcements, Working paper, Kellogg School of Management, Northwestern University.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3-56.

Fang, Lily, and Joel Peress, 2007, Media Coverage and the Cross-Section of Returns, Working paper, INSEAD.

Foerster, Stephen R., and G. Andrew Karolyi, 1999, The Effects of Market Segmentation and Investor Recognition on Asset Prices: Evidence from Foreign Stocks Listing in the United States, *Journal of Finance* 54 (3), 981-1013.

Gadarowski, Christopher, 2002, Financial press coverage and expected stock returns, Working paper.

Harrington, David E., 1989, Economic news on television: The determinants of coverage, *Public Opinion Quarterly* 53, 17-40.

Hirshleifer, David, Sonya S. Lim, and Siew H. Teoh, 2006, Driven to distraction: Extraneous events and underreaction to earnings news, Working paper, University of California at Irvine, and Depaul University.

Hong, Harrison, Terence Lim, and Jeremy C. Stein, 2000, Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies, *Journal of Finance* 55 (1), 265-295.

Hou, Kewei, and Tobias J . Moskowitz, 2005, Market Frictions, Price Delay, and the Cross-Section of Expected Returns, *Review of Financial Studies* 18 (3), 981-1020.

Hou, Kewei, Lin Peng, and Wei Xiong, 2006, A tale of two anomalies: The implications of investor attention for price and earnings momentum, Working paper, Ohio State University, Baruch College, and Princeton University.

Huang, Lixin, and Hong Liu, 2007, Rational Inattention and Portfolio Selection, *Journal of Finance* 62 (4), 1999-2040.

Huberman, Gur, and Tomer Regev, 2001, Contagious speculation and a cure for cancer: A nonevent that made stock prices soar, *Journal of Finance* 56, 387-396.

Kahneman, Daniel, and Amos Tversky, 1979, Prospect Theory: An Analysis of Decision under Risk, *Econometrica* 47, 263-291.

Kaniel, Ron, Laura Starks, and Vasudha Vasudevan, 2007, Headlines and Bottom Lines: Attention and Learning from Media Coverage of Mutual Funds, Working paper, Duke University, and University of Texas at Austin.

Kothari, S. P., Susan Shu, Peter Wysocki, 2005, Do Managers Withhold Bad News?, Working paper, MIT Sloan and Boston College.

Manning, Christopher D., and Hinrich Schuetze, 1999, Foundations of Statistical Natural Language Processing, MIT Press, Cambridge, MA.

McCallum, Andrew Kachites, 1996, *Bow: A toolkit for statistical language modeling, text retrieval, classification and clustering*. http://www.cs.cmu.edu/ mccallum/bow, 1996.

McQueen, Grant, Michael Pinegar, and Steven Thorley, 1996, Delayed Reaction to Good News and the Cross-Autocorrelation of Portfolio Returns, *Journal of Finance* 51 (3), 889-919.

Merton, Robert C., 1987, A Simple Model of Capital Market Equilibrium with Incomplete Information, *Journal of Finance* 42 (3), 483-510.

Pastor, Lubos, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111(3), 642-685.

Peng, Lin, and Wei Xiong, 2006, Investor attention, overconfidence, and categorical learning, *Journal of Financial Economics* 80(3), 563-602.

Reinganum, M. R., 1981, Misspecification of Capital Asset Pricing: Empirical Anomalies Based on Earning Yields and Market Values, *Journal of Financial Economics* 9, 19-46.

Sadka, Ronnie, 2006, Momentum and post-earnings announcement drift anomalies: The role of liquidity risk, *Journal of Financial Economics* 80(2), 309-349.

Shapiro, Alexander, 2002, The Investor Recognition Hypothesis in a Dynamic General Equilibrium: Theory and Evidence, *Review of Financial Studies* 15 (1), 97-141.

Skinner, Douglas J., and Richard G. Sloan, 2002, Earnings Surprises, Growth Expectations, and Stock Returns or Don't Let an Earnings Torpedo Sink Your Portfolio, *Review of Accounting Studies* 7, 289-312.

Soroka, Stuart N., 2006, Good news and bad news: Asymmetric responses to economic information, *Journal of Politics* 68 (2), 372-385.

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

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.

Veronesi, Pietro, 1999, Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model, *Review of Financial Studies* 12 (5), 975-1007.

CHAPTER IV

Conclusion

This thesis undertakes an empirical study of media coverage to examine the role of investor attention in financial markets. First, I investigate the existence and potential impact of bias in media coverage of positive and negative events relative to fundamental information flows. Second, I study coverage decisions to identify the event- and firm-specific determinants of attention and neglect; I suggest a new potential explanation for the long-standing empirical puzzle of the neglected firm premium based on asymmetric inattention.

In the first essay, I present evidence of time-varying predictability in corporate news coverage in the Financial Times, the Wall Street Journal, and the New York Times. Furthermore, I find that estimates of unexplained positive or negative news coverage have significant predictive content with respect to future stock index returns, particularly for smaller firms. In short, the results indicate that 1) holding other factors constant, the news media is significantly more likely to focus on either positive or negative corporate news stories at different points in time, and 2) this aggregate coverage behavior appears to contain information about returns for stocks where we would expect information frictions to be most severe. However, while these findings are certainly suggestive of the importance of media coverage (and, by implication, investor attention), it should be noted that a causal relationship between news coverage and prices cannot be firmly established in this context. In particular, media coverage decisions may simply be correlated with the broader attention allocation decisions undertaken by investors. Indeed, I present evidence that the coverage decisions of news outlets are correlated each other – it could simply be the case that the identified common variation in coverage with respect to positive or negative news is simply a reflection of deeper, underlying patterns in market attention more broadly. For example, it is possible that media coverage is merely (or partially) an indicator of the events that investors and market participants would be paying more attention to in any case, even in the absence of the news media altogether. Finally, although the VAR setting attempts to control for a host of other variables that might potentially be expected to predict weekly stock returns, there always remains the possibility that the news differentials are simply correlated with some unobserved factor.

Furthermore, even if we are convinced that media coverage (Granger-)causes future stock returns, the results in the first essay are potentially ambiguous with respect to the precise mechanism at work. On the one hand, we could view return predictability as stemming from systematic underreactions to the subset of news events that are not covered; for example, if "excessive" negative coverage is observed in the current week, this could imply that the impact of the uncovered positive news is simply delayed to future weeks. Alternatively, the results are also potentially consistent with temporary stock-price overreactions to current media coverage that corrects itself in subsequent weeks. This latter interpretation would be consistent with, e.g., Antweiler and Frank (2006).

The finding of predictable patterns in unexplained coverage has potentially interesting implications for the literature on media bias. In particular, I find that there appear to be long stretches of time over which media outlets tend to focus their coverage on one type of corporate news story over another. While other studies have focused either on potential bias on the "supply-side" (e.g., in favor of advertisers, as in Reuter and Zitzewitz (2006), or on differences in contemporaneous coverage behavior among competing news outlets (e.g., with respect to left/right political bias, as in Groseclose and Milyo (2005), this thesis identifies time-series variation that appears to be common across news outlets. For example, I find that the WSJ, FT, and NYT all exhibit higher-than-expected levels of negative coverage in the early '90s, relatively more positive coverage in the mid-90's, and then generally more negative coverage again after 2000. In short, while the type of media bias discussed here is more consistent with a relatively passive, demand-driven catering to market sentiment, (rather than, e.g., conscious, supply-side manipulation on the part of news providers), the return predictability findings suggest that the resulting impact on the information environment and investors' beliefs may nonetheless be highly significant.

In the second essay, I identify the cross-sectional characteristics of individual firms and events and associate them with subsequent returns, potentially allowing us to distinguish between underreaction and overreaction in the market's response to high attention and low attention events. Specifically, I am able to distinguish between high-attention firms and low-attention firms, and between firms with good earnings news and those with bad earnings news – the results here tend to support the underreaction hypothesis in this context, consistent with DellaVigna and Pollet (2008), Hirshleifer et al. (2006), and Hou et al. (2006). At the same time, it should be noted that, since the focus here is on using coverage to identify neglected firms and events, I do not make direct use of realized coverage observations in this study, so I cannot exclude the possibility that prices do, indeed, overreact to earnings announcements that are actually covered in the media.[26] This remains a potentially interesting avenue for future research.

Applying media coverage estimates to identify relatively well-attended and neglected firms, this thesis introduces asymmetric inattention as a potential explanation for the finding of high risk-adjusted returns for "neglected firms". Incorporating the event-specific predictors of attention and neglect, I show that the finding of an apparent neglected firm premium in this context (which initially appears to support, for example, Merton's (1987) investor recognition hypothesis, or Easley et al. (2002) information risk setting) is actually consistent with delayed price responses combined with a negativity bias in investor attention. In particular, this implies a potentially new and different interpretation regarding the results of Fang and Peress (2007), and, somewhat more indirectly, Arbel and Strebel (1982), Foerster and Karolyi (1999), and Hou and Moskowitz (2005), etc. This being said, a few caveats and limitations to the analysis should be noted. In particular, since I use constructed estimates of coverage as a proxy for investor attention, an unambiguous test of the asymmetric inattention hypothesis is not possible in this setting. However, in that I am able to differentiate between low-attention "good news" firms and low-attention "bad news" firms in the sample, and link an apparent neglected firm premium to the former group and not the latter, it is somewhat difficult to argue that a more traditional, symmetric friction or risk premium is actually at work. Specifically, one would have to believe that, e.g., low-attention firms that have recently experienced good news are perceived to be riskier than those that have recently experienced bad news – at a very basic level, volatility leverage would seem to suggest that, if anything, the opposite should be true.

At the same time, it should also be noted that the findings here do not necessarily exclude such factors as information risk as potential contributors to high expected returns for neglected firms more broadly. In particular, note that the lowest PMC^{+} decile portfolios in Panels C and D of Table 3.5 do exhibit positive (albeit statistically insignificant) alphas, indicating that even the

[26] Please see Appendix II for a preliminary analysis of the impacts of realized coverage in event-time.

low-attention "bad news" stocks may enjoy some premium that cannot be explained in a standard 4- or 5-factor model. In other words, while the asymmetric underreaction effect appears to be quite significant in and of itself, it may well represent only a partial explanation for observations of neglected firm premia more broadly in the cross-section of returns.

While an attempt to explain the observed negativity bias in coverage is somewhat beyond the scope of this thesis, some discussion is warranted. Why might investors be more interested in negative news events than positive ones? One interpretation comes to us via prospect theory. With loss-aversion, a negative outcome is weighted more strongly than the prospect of a comparable positive outcome (Kahneman and Tversky (1979)) – therefore, individuals might be especially attentive to news of a potential loss that might affect them (or, in the case of financial news, the value of their portfolio of stocks). For example, Soroka (2006) finds that negative economic news has a greater impact on public attitudes than positive news. Alternately, it may be that investors are accounting for the incentives of managers to delay or minimize the reporting of negative news about the firm whenever possible – with respect to the business press, most of the news regarding companies comes from firm-originated press releases, and, perhaps not coincidentally, most of this news is rather positive in tone. Hence, when an unambiguously negative news story does come along, it may be seen as more believable, therefore attracting greater attention and resulting in larger market responses. Whatever the mechanism, the observation of asymmetrically large responses to negative information is potentially consistent both with theory[27], as well as the empirical literature.[28] However, as discussed in the second essay, we are nonetheless faced with two recent papers (Kaniel et al. (2007 and Engelberg (2007)) that fail to identify a negativity bias in coverage decisions. While this remains an issue for future research, one potential explanation points to differences in the selectivity of coverage. In particular, if firms are able to more effectively "push" the publication of positive news stories in some settings, we would expect this to potentially counteract an off-setting demand-driven bias for negative news.

One potential problem typically associated with identifying the impact of media coverage on returns is that of omitted variable bias. Specifically, while evidence of strong correlation between observed coverage patterns and subsequent market returns is suggestive of a relationship between the two, it does not allow us to exclude the possibility that we are seeing

[27] e.g., Veronesi (1999) and Kothari et al. (2005).
[28] e.g., Conrad et. al (2002), Skinner and Sloan (2002), McQueen et al. (1996), and Kothari et al. (2005).

the effects of an unidentified factor (for example, industry classification, public relations expenditures, or some unobserved contemporaneous information shock, etc.) that is both correlated with media coverage and has some predictive content with respect to the cross-section of expected stock returns.[29]

The identification strategy employed in the second essay is to utilize firms' ex ante predicted coverage to form portfolios for use in asset pricing tests. The reasoning here is similar to that underlying the use of instrumental variables estimation in a situation where endogeneity is suspected to be a problem. In particular, since predicted coverage probabilities are explicitly calculated based upon a known set of factors observed to predict coverage, any potential impacts from omitted variables on realized coverage should not carry over into portfolio selection. Similarly, any potential effects from measurement error related to the algorithmic identification and classification of news stories will be lessened. Finally, even if we suspect that apparent return predictability might be related to one of the identified factors which happen to load heavily as an explanatory variable in the media coverage regressions, this methodology offers the added advantage of allowing us to explicitly test such hypotheses by dropping suspected explanatory variables from the coverage probability regressions. The results suggest that, in particular, the predictive ability of PMC with respect to returns is not due to simple weightings on identified risk factors such as size, B/M, beta, or momentum.

[29] Kaniel et al. (2007) examine a similar point related to endogenous media coverage decisions and their potential impact on mutual fund flows.

References

Antweiler, Werner, and Murray Z. Frank, 2006, Do U.S. Stock Markets Typically Overreact to Corporate News Stories, Working paper, University of British Columbia and University of Minnesota.

Arbel, Avner, and Paul Strebel, 1982, The Neglected and Small Firm Effects, *Financial Review* 17, 201-218.

Conrad, Jennifer, Bradford Cornell, and Wayne R. Landsman, 2002, When Is Bad News Really Bad News?, *Journal of Finance* 57 (6), 2507-2532.

DellaVigna, Stefano, and Joshua Pollet, 2008, Investor Inattention and Friday Earnings Announcements, *Journal of Finance,* forthcoming.

Easley, David, Soeren Hvidkjaer, and Maureen O'Hara, 2002, Is information risk a determinant of asset returns?, *Journal of Finance* 57(5), 2185-2221.

Engelberg, Joseph, 2007, Costly Information Processing: Evidence from Earnings Announcements, Working paper, Kellogg School of Management, Northwestern University.

Fang, Lily, and Joel Peress, 2007, Media Coverage and the Cross-Section of Returns, Working paper, INSEAD.

Foerster, Stephen R., and G. Andrew Karolyi, 1999, The Effects of Market Segmentation and Investor Recognition on Asset Prices: Evidence from Foreign Stocks Listing in the United States, *Journal of Finance* 54 (3), 981-1013.

Groseclose, Tim, and Jeffrey Milyo, 2005, A Measure of Media Bias, *Quarterly Journal of Economics* 120(4), 1191-1237.

Hirshleifer, David, Sonya S. Lim, and Siew H. Teoh, 2006, Driven to distraction: Extraneous events and underreaction to earnings news, Working paper, University of California at Irvine, and Depaul University.

Hou, Kewei, and Tobias J . Moskowitz, 2005, Market Frictions, Price Delay, and the Cross-Section of Expected Returns, *Review of Financial Studies* 18 (3), 981-1020.

Hou, Kewei, Lin Peng, and Wei Xiong, 2006, A tale of two anomalies: The implications of investor attention for price and earnings momentum, Working paper, Ohio State University, Baruch College, and Princeton University.

Kahneman, Daniel, and Amos Tversky, 1979, Prospect Theory: An Analysis of Decision under Risk, *Econometrica* 47, 263-291.

Kaniel, Ron, Laura Starks, and Vasudha Vasudevan, 2007, Headlines and Bottom Lines: Attention and Learning from Media Coverage of Mutual Funds, Working paper, Duke University, and University of Texas at Austin.

Kothari, S. P., Susan Shu, Peter Wysocki, 2005, Do Managers Withhold Bad News?, Working paper, MIT Sloan and Boston College.

McQueen, Grant, Michael Pinegar, and Steven Thorley, 1996, Delayed Reaction to Good News and the Cross-Autocorrelation of Portfolio Returns, *Journal of Finance* 51 (3), 889-919.

Merton, Robert C., 1987, A Simple Model of Capital Market Equilibrium with Incomplete Information, *Journal of Finance* 42 (3), 483-510.

Reuter, Jonathan, and Eric Zitzewitz, 2006, Do ads influence editors? Advertising and bias in the financial media, *Quarterly Journal of Economics* 121(1), 197-227.

Skinner, Douglas J., and Richard G. Sloan, 2002, Earnings Surprises, Growth Expectations, and Stock Returns or Don't Let an Earnings Torpedo Sink Your Portfolio, *Review of Accounting Studies* 7, 289-312.

Soroka, Stuart N., 2006, Good news and bad news: Asymmetric responses to economic information, *Journal of Politics* 68 (2), 372-385.

Veronesi, Pietro, 1999, Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model, *Review of Financial Studies* 12 (5), 975-1007.

Appendix I: Classification of news articles

In building the media data set, 68,102 Wall Street Journal news articles with Intelligent Indexing Code c151 ("earnings") were obtained from Factiva in text format from October 1984 to December 2005.[30] Articles were then classified using the computational linguistics program Rainbow (McCallum (1996)), applying a two-stage, Naïve Bayesian, bag-of-words methodology. In order to "train" the classification algorithm, 500 articles were selected at random for hand-classification. Of these, 129 were identified as "not earnings articles" (i.e., articles that did not discuss the content of a specific recent corporate earnings announcement), and 371 were identified as "earnings articles" (stage 1). Further categorizing the set of "earnings articles", 147 were classified as "negative", and 224 as "not negative" (stage 2). The classification software was then applied to the articles, resulting in predictive models of category membership based on word frequencies. The first binary classification model (WSJ1) distinguishes between "earnings" and "not earnings" articles using odds-ratios from the top 300 bigrams ranked by infogain. The second binary classification model (WSJ2) distinguishes between "negative" and "not negative" among identified "earnings" articles using odds-ratios from the top 25 unigrams ranked by infogain.

Model accuracy was tested by randomly excluding 100 articles from the training set, re-estimating the model, and then examining accuracy with respect to the classifications predicted for the excluded articles. In the course of 100 such trials, model WSJ1 had an average accuracy of 88.52%, while model WSJ2 implied an average accuracy of 83.42%. This estimated rate of accuracy is comparable to hand-classification methods using paid assistants.

Models WSJ1 and WSJ2 were then applied to the full data set of uncategorized articles to calculate estimated probabilities of category membership for each. Articles with predicted probability >0.5 of being an "earnings" article were recorded as such. Of these, articles with a predicted probability >0.5 from WSJ2 of being a "negative" article were identified as "negative earnings", with all others recorded as "not negative earnings". The categorization results are shown in Table AI.1.

[30] Additional filters were also applied in excluding the most common types of "false positive" articles resulting from the Factiva search – details are available upon request.

Example of a WSJ article classified as "Positive earnings announcement news":

HD Walgreen Co. Earnings Rise 17% on Sales Gain Of Prescription Drugs
WC 194 words
PD 4 January 2002
LP DEERFIELD, Ill. -- Walgreen Co., boosted by strong prescription sales, said fiscal first-quarter earnings jumped 17%.
The drugstore chain reported net income of $185.9 million, or 18 cents a share, compared with $158.4 million, or 15 cents a share, a year earlier. Analysts surveyed by Thomson Financial/First Call had forecast earnings of 17 cents a share for the period ended Nov. 30.
TD The recent quarter's results included a $5.5 million pretax gain from the final payment of an antitrust settlement regarding brand-name prescription drugs. Walgreen said prescriptions, which accounted for 60% of first-quarter sales, leapt 22% overall and 17% on a same-store basis. Total sales for the period climbed 17% to $6.6 billion from $5.6 billion. Chairman L. Daniel Jorndt said the chain planned to open 475 new stores and two new distribution centers this year. As of Nov. 30, the company operated 3,623 drugstores in 43 states and Puerto Rico. Shares of Walgreen rose $1.33 to $34.41 as of 4 p.m. in New York Stock Exchange composite trading.

Example of a WSJ article classified as "Negative earnings announcement news":

HD Business Brief -- Playtex Products Inc.: Loss of $959,000 Is Posted Amid a 7.9% Decrease in Sales
WC 185 words
PD 10 February 2004
LP Playtex Products Inc. posted a fourth-quarter loss as the consumer-products company was hurt by restructuring charges, competition in the tampon market and weather that cut into sales of sun-care products. The Westport, Conn., company reported a fourth-quarter loss of $959,000, or two cents a share, including restructuring charges of five cents. A year earlier, the company reported a profit of $7.2 million, or 12 cents a share. Playtex, which makes household items such as Playtex tampons, Wet Ones wipes and Banana Boat sunscreen, said fourth-quarter sales fell 7.9% to $146.7 million from $159.2 million a year earlier. Playtex, which faces intense competition from Procter & Gamble Co., had lowered its 2003 and 2004 outlook on Jan. 30, projecting a fourth-quarter net loss of two cents a share. The results were issued after the close of regular trading. In 4 p.m. composite trading yesterday on the New York Stock Exchange, Playtex was unchanged at $6.30.

Table AI.1: Full and training sample article classification

Panel A: WSJ article classification

	Not Earnings	Earnings	
Training Sample – Classified by me n=500	129 (25.8%)	371 (74.2%)	
		Not Negative	Negative
		224 (60.38%)	147 (39.6%)

	Not Earnings	Earnings	
Full Sample – Classified by algorithm (WSJ1 and WSJ2) n=68,102	19,989 (29.352%)	49,113 (70.648%)	
		Not Negative	Negative
		31,833 (64.816%)	17,280 (35.184%)

Panel B: NYT article classification

	Not Earnings	Earnings	
Training Sample – Classified by me (n=500)	281 (56.2%)	219 (43.8%)	
		Not Negative	Negative
		149 (68.04%)	70 (31.96%)

	Not Earnings	Earnings	
Full Sample – Classified by algorithm (NYT1 and NYT2) n=17,289	9,460 (54.717%)	7,829 (45.283%)	
		Not Negative	Negative
		6,140 (78.426%)	1,689 (21.574%)

Panel C: FT article classification

<table>
<tr><td></td><td>Not Earnings</td><td colspan="2">Earnings</td></tr>
<tr><td rowspan="3">Training Sample – Classified by me (n=500)</td><td rowspan="3">104 (20.8%)</td><td colspan="2">396 (79.2%)</td></tr>
<tr><td>Not Negative</td><td>Negative</td></tr>
<tr><td>285 (71.97%)</td><td>111 (28.03%)</td></tr>
</table>

<table>
<tr><td></td><td>Not Earnings</td><td colspan="2">Earnings</td></tr>
<tr><td rowspan="3">Full Sample – Classified by algorithm (FT1 and FT2) n=83727</td><td rowspan="3">12,892 (15.398%)</td><td colspan="2">70,835 (84.602%)</td></tr>
<tr><td>Not Negative</td><td>Negative</td></tr>
<tr><td>51,375 (72.528%)</td><td>19,460 (27.472%)</td></tr>
</table>

Text classification model details

WSJ1 (earnings vs. not earnings): naïve bayes; accuracy of 100 trials (100 articles randomly excluded from training set in each trial): 88.52 +/- 0.28%

WSJ2 (negative vs. not negative): naïve bayes; accuracy of 100 trials (100 articles randomly excluded from training set in each trial): 83.42 +/- 0.33%

NYT1 (earnings vs. not earnings): naïve bayes, top 1000 quadrigrams by infogain; accuracy of 100 trials (100 articles randomly excluded from training set in each trial): 87.87 +/- 0.26%

NYT2 (negative vs. not negative): naïve bayes; accuracy of 100 trials (100 articles randomly excluded from training set in each trial): 80.06 +/- 0.40%

FT1 (earnings vs. not earnings): naïve bayes; accuracy of 100 trials (100 articles randomly excluded from training set in each trial): 90.02 +/- 0.23%

FT2 (negative vs. not negative): naïve bayes; accuracy of 100 trials (100 articles randomly excluded from training set in each trial): 82.76 +/- 0.32%

Appendix II: The effects of realized and abnormal media coverage

While the main focus of this paper lies in the identification of expected coverage patterns as a proxy for investor attention, it is natural to ask about the potential impact of realized or abnormal coverage on returns, both during the announcement window and in subsequent periods. At a basic level, it may be particularly important to test whether the media events that I identify are actually associated with significant, abnormal price innovations at the time of announcement, controlling for other factors. Additionally, we may choose to exploit the empirical model of expected coverage developed in Section 5 to identify the unexpected portion of coverage and attempt to study return predictability on that basis.

Event study: The significance of model-identified media coverage events

Given that the identification and classification of media stories in this paper is performed by (supervised) algorithm, an event study analysis may give us some additional assurance that the Naïve-Bayesian identification of positive and negative earnings stories is indeed performing broadly as expected. While it is impossible to assign a direction of causality in this context (e.g., announcement-window returns may affect coverage decisions, vice versa, and/or both may be responding to some unidentified variable), strong evidence of correlation with unexplained returns minimally suggests that media coverage is associated with information events that are abnormally positive or negative in some significant sense.

As a first step, we must attempt to account for the non-media related factors that predict announcement window returns. Research in the accounting literature has shown that the market's response to earnings news may be both asymmetric and non-linear (e.g., Milev (2004)). For example, there is often a significant negative price reaction when a firm fails to meet analysts' expectations, regardless of the extent or size of the relative underperformance. Furthermore, we might expect markets to react differently to a negative surprise if the firm reports a profit compared to the reaction to if it reports a loss. Finally, if investors can be thought of as Bayesian updaters, then we might expect the market reaction to decline in proportional terms as the surprise gets larger. In order to account for these potential nonlinearities, I impose the following (non-parametric) functional form in estimating the earnings response coefficient (ERC)[31]:

$$[\text{I\_NEGSURPRISE}_{i,t} + \text{LOSS}_{i,t} + \text{I\_NEGSURPISE}_{i,t}\cdot\text{LOSS}_{i,t}]*[1 + \text{SURPRISE}_{i,t} + (\text{SURPRISE}_{i,t})^2]$$

Firm characteristics and market conditions may also be important for explaining announcement returns. For example, investors may discount the stock of a relatively risky firm heading into an uncertain information release, resulting in an apparently abnormal positive return as the risk related to the event is resolved in one way or another. Firm risk characteristics may be proxied by, e.g., the volatility of recent stock returns, dispersion in analyst forecast, recent stock performance, analyst coverage, etc. To the extent that potentially important firm characteristics may be correlated by industry and over time, the industry and time dummy variables will also help in this respect.

For the event study results, I apply a daily Fama-French 4-factor model to generate excess returns.[32] Specifically, with respect to each firm-announcement, I estimate an expected returns model by regressing the sixty trading days of stock returns prior to the event window on the daily market, SMB, HML, and Momentum factors from Ken French's website. I then use the resulting coefficient estimates to generate expected returns for the announcement window (day -1 to +2).

$$\begin{aligned}\text{CAR}_{i,t} = {} & \alpha_0 + \alpha_1\cdot\text{NEGATIVE}_{i,t} + \alpha_2\cdot\text{POSITIVE}_{i,t} + \alpha_3\text{I\_NEGSURPRISE}_{i,t} + \alpha_4\cdot\text{I\_LOSS}_{i,t} \\ & + \alpha_5\cdot \text{I\_NEGSURPISE}_{i,t}\cdot\text{I\_LOSS}_{i,t} + \alpha_6\cdot\text{SURPRISE}_{i,t} + \alpha_7\cdot(\text{SURPRISE}_{i,t})^2 + \\ & \alpha_8\cdot\text{I\_LOSS}_{i,t}\cdot\text{SURPRISE}_{i,t} + \alpha_9\cdot\text{I\_LOSS}_{i,t}\cdot (\text{SURPRISE}_{i,t})^2 + \\ & \alpha_{10}\cdot\text{I\_NEGSURPRISE}_{i,t}\cdot\text{SURPRISE}_{i,t} + \alpha_{11}\cdot\text{I\_NEGSURPRISE}_{i,t}\cdot(\text{SURPRISE}_{i,t})^2 + \\ & \alpha_{12}\cdot\text{I\_NEGSURPRISE}_{i,t}\cdot\text{I\_LOSS}_{i,t}\cdot\text{SURPRISE}_{i,t} + \\ & \alpha_{13}\cdot\text{I\_NEGSURPRISE}_{i,t}\cdot\text{I\_LOSS}_{i,t}\cdot(\text{SURPRISE}_{i,t})^2.+ \alpha_{14}\cdot\#\text{ANALYSTS}_{i,t} + \\ & \alpha_{15}\cdot\text{STDEV(FORECASTS)}_{i,t} + \alpha_{16}\cdot\text{B/M}_{i,t} + \alpha_{17}\cdot\text{RETURNS}_{-1,i,t} + \\ & \alpha_{18}\cdot\text{STDEV}(\text{RETURNS}_{-1})_{i,t} + \alpha_{19}\cdot\text{S\&P}_{-1,i,t} + \alpha_{20}\cdot\text{STDEV}(\text{RETURNS}_{-1})_{i,t} + \\ & \alpha_{21}\cdot\text{INSTITOWN}_{i,t} + \alpha_{22}\cdot\text{I\_USFIRM}_{i,t} + \alpha_{23}\cdot\text{\$VOLUME}_{i,t} + \\ & \textstyle\sum\gamma[\text{TIME\_DUMMIES}_{i,t}] + \sum\delta[\text{INDUSTRY\_DUMMIES}_{i,t}] + \varepsilon_{i,t}\end{aligned}$$

---

[31] An illustration of the proposed functional form for ERC (based on estimated coefficients from this restricted model) is presented in Figure AII.1. Moreover, this non-parametric functional form implicitly assumes that investors' reference points with respect to categorical earnings surprises and losses lie at 0. In this iteration, I also ignore the potential for systematic analyst forecast biases.

[32] CAPM and 3-factor excess returns were also examined – the results are not sensitive to this choice.

Table 3.10 presents the results of regressing (4-factor risk-adjusted) earnings announcement window returns on the negative and positive coverage dummies, controlling for firm characteristics and measures of the information content of the announcement itself. Looking at the results in column 4, I find that positive (negative) coverage is associated with 58 (-77) bps in unexplained returns over the 5 day window. However, while these results may be interesting in a correlative sense, it should be noted once more it is impossible to make a statement regarding causality in this setting: prices may be reacting to the tone of coverage, coverage decisions may be affected by early-window market reactions, and/or both. We can also imagine circumstances in which the determinants of coverage might be strongly correlated with event window return-relevant firm characteristics. For example, we might think that larger firms will tend to have less volatile returns and earnings, and so investors may require smaller premiums going into an uncertain information event. Similarly, firms with higher recent trading activity might be seen as more uncertain, resulting in higher event-return premia. Smaller firms might also be unconditionally more likely to experience accounting losses, and we would surely expect this to affect announcement returns.

<u>Abnormal media coverage and the market's reaction to earnings news</u>

Having identified a set of event and firm specific predictors for media coverage in Section 5, I investigate the potential impact of unexplained coverage on future stock returns. For each event, I identify abnormal absolute, negative, and positive media coverage using predictions from the multinomial logit regression described in Table 3.3, column 5. As described in Section 3.7, decile portfolios are formed at the beginning of each month based on the most recent observation of residual coverage in the previous three months. Portfolios are held for one month before rebalancing, and returns are equal-weighted.

The first two columns of Table AII.2 present four-factor alphas for the first and tenth decile portfolios based on each measure of unexpected media coverage. The third column describes a portfolio that is short stocks with low abnormal coverage and short stocks with high abnormal coverage. We observe that high abnormal negative coverage predicts significantly positive future returns. Looking at the third row, third column of Table AII.2, the implied zero-investment portfolio alpha based on such a strategy is estimated at 63 bps per month. What could explain this result? Table AII.1 shows that realized negative media coverage is associated with abnormal negative event window returns. The results in Table AII.2 are consistent with a temporary overreaction to stories that receive abnormal negative coverage, followed by subsequent reversal.

Figure AII.1: Illustration of the estimated earnings response function (non-parametric)

Fitted curve based on estimated coefficients from the following regression:

$$CAR_{i,t} = \alpha_0 + \alpha_1 \cdot I\_NEGSURPRISE_{i,t} + \alpha_2 \cdot SURPRISE_{i,t} + \alpha_3 \cdot (SURPRISE_{i,t})^2 + \alpha_4 \cdot I\_NEGSURPRISE_{i,t} \cdot SURPRISE_{i,t} + \alpha_5 \cdot I\_NEGSURPRISE_{i,t} \cdot (SURPRISE_{i,t})^2 + e_{it}$$

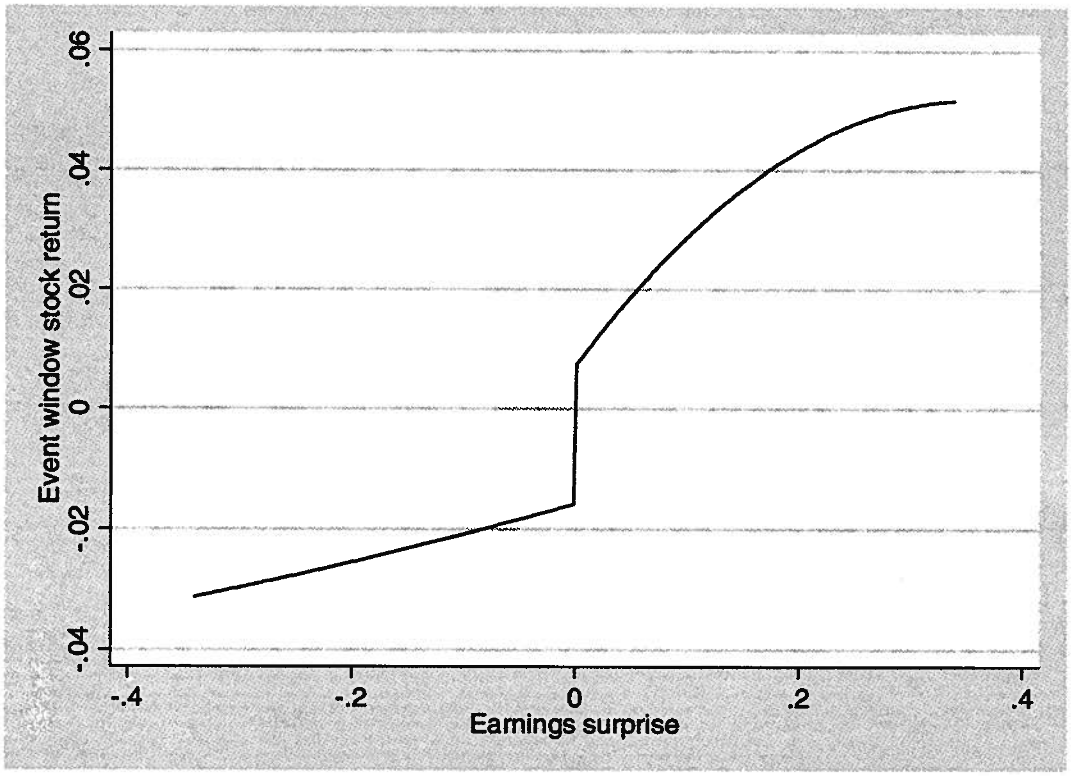

Table AII.1: Media coverage and stock returns around earnings announcements
CAR[-1 to 3] is the excess cumulative return for days -1 to +3, with expected returns calculated using a Carhart (1997) 4-factor model estimated on the 60 trading days prior to the event window; *Earnings surprise* is the announced EPS less the median analyst forecast from the 30 days prior to the earnings announcement divided by the stock price; *I{Negative surprise}* is a dummy variable equal to one if *Earnings surprise* is negative, and zero otherwise; *I{Loss}* is a dummy variable equal to one if announced EPS is negative, and zero otherwise; *log(Market value)* is the natural logarithm of the number of shares outstanding multiplied by the closing stock price two days before the announcement; *B/M* is the book value of equity from the fiscal year ending in the previous year divided by the market value of equity from December 31 (divided by 1000); *Institutional ownership* is the percentage of equity held by institutions at the end of the previous calendar year; *log($Trading Volume)* is the natural logarithm of the firm's average dollar value of trading volume in the 60 trading days prior to the announcement; *log(Analyst attention)* is the natural log of the number of distinct analyst forecasts observed in I/B/E/S in the 30 days preceding the announcement; *stdev(forecasts)* is the standard deviation of the normalized analyst EPS forecasts recorded in I/B/E/S in the 30 days preceding the announcement; *Recent returns* is the average daily stock return during the 60 trading days prior to the announcement window; *S&P500 returns* is the average daily return on the S&P500 during the 60 trading days prior to the announcement window; *stdev(S&P500 returns)* is the standard deviation of average daily return on the S&P500 during the 60 trading days prior to the announcement window; results for year, month-of-the-year, day-of-the-week and industry dummy variables not shown; significant standard errors in parentheses, robust to arbitrary heteroskedasticity and clustering by firm; * significant at 10%; ** significant at 5%; *** significant at 1%

	CAR[-1 to 3] (1)	*CAR[-1 to 3]* (2)	*CAR[-1 to 3]* (3)	*CAR[-1 to 3]* (4)	*CAR[-1 to 3] - w/ firm fixed effects -* (5)
Negative coverage dummy	-0.694***	-0.770***	-0.700***	-0.769***	-1.246***
	[0.123]	[0.123]	[0.123]	[0.123]	[0.123]
Positive coverage dummy	0.562***	0.561***	0.579***	0.581***	0.293***
	[0.071]	[0.071]	[0.073]	[0.073]	[0.093]
Earnings surprise	50.772***	50.089***	49.601***	49.012***	53.740***
	[9.051]	[9.079]	[9.048]	[9.083]	[4.248]
(Earnings surprise)^2	-61.400***	-60.404***	-60.116***	-59.244***	-69.518***
	[11.626]	[11.678]	[11.579]	[11.641]	[6.413]
I{Negative surprise}	-2.689***	-2.741***	-2.713***	-2.760***	-2.745***
	[0.061]	[0.062]	[0.061]	[0.062]	[0.058]
I{Loss}	-0.977***	-0.967***	-0.996***	-0.991***	-0.515***
	[0.144]	[0.145]	[0.148]	[0.148]	[0.126]
I{Negative surprise}*I{Loss}	0.078	0.106	0.094	0.122	-0.196
	[0.191]	[0.191]	[0.191]	[0.191]	[0.148]
I{Loss}*(Earnings surprise)	-38.564***	-38.670***	-37.829***	-38.001***	-38.554***
	[12.448]	[12.473]	[12.444]	[12.471]	[5.819]
I{Loss}*(Earnings surprise)^2	44.375***	44.254***	43.654***	43.623***	54.100***
	[16.821]	[16.880]	[16.800]	[16.862]	[9.032]
I{Negative surprise}*(Earnings surprise)^2	159.193***	159.393***	158.838***	158.964***	162.889***
	[20.829]	[20.972]	[20.814]	[20.950]	[12.717]
I{Negative surprise}*(Earnings surprise)	8.219	9.647	9.934	11.155	6.095
	[12.948]	[13.006]	[12.996]	[13.056]	[7.222]

	CAR[-1 to 3] (1)	CAR[-1 to 3] (2)	CAR[-1 to 3] (3)	CAR[-1 to 3] (4)	CAR[-1 to 3] - w/ firm fixed effects - (5)
I{Negative surprise}*I{Loss}*(Earnings surprise)^2	-137.12***	-137.96***	-137.27***	-138.04***	-137.41***
	[24.937]	[25.070]	[24.928]	[25.052]	[14.506]
I{Negative surprise}*I{Loss}*(Earnings surprise)	-10.905	-11.256	-12.031	-12.242	-7.291
	[15.803]	[15.852]	[15.842]	[15.890]	[8.306]
log(Market value)	-0.686***	-0.737***	-0.735***	-0.780***	-3.012***
	[0.042]	[0.043]	[0.043]	[0.043]	[0.070]
B/M (/1000)	0.028	0.027	0.028	0.028	0.025
	[0.023]	[0.023]	[0.021]	[0.022]	[0.020]
Institutional ownership (%)	0	0	0	0	-0.016***
	[0.001]	[0.001]	[0.001]	[0.001]	[0.002]
log($Trading volume [-61 to -2])	0.407***	0.465***	0.437***	0.492***	1.303***
	[0.036]	[0.037]	[0.036]	[0.037]	[0.045]
US firm dummy	-0.686***	-0.769***	-0.611***	-0.685***	
	[0.142]	[0.144]	[0.151]	[0.153]	
log(Analyst attention}	-0.043	-0.090*	-0.026	-0.077	-0.124*
	[0.049]	[0.050]	[0.051]	[0.051]	[0.063]
stdev(forecasts)	0.201***	0.195***	0.195***	0.190***	0.140***
	[0.039]	[0.039]	[0.039]	[0.039]	[0.030]
Recent returns [-61 to -2]	-0.091***	-0.093***	-0.091***	-0.093***	-0.090***
	[0.002]	[0.002]	[0.002]	[0.002]	[0.001]
S&P500 returns [-61 to -2]	0.098***	0.111***	0.098***	0.111***	0.113***
	[0.004]	[0.005]	[0.004]	[0.005]	[0.004]
stdev(S&P500 returns [-61 to -2])	0.274***	0.202**	0.248***	0.174*	-0.163*
	[0.068]	[0.103]	[0.070]	[0.104]	[0.090]
Year dummies	No	Yes	No	Yes	Yes
Industry dummies	No	No	Yes	Yes	No
Observations	178898	178898	178898	178898	178898
R-squared	0.07	0.07	0.07	0.07	0.08
Number of firms					8715

Table AII.2: Abnormal media coverage – portfolio returns

Monthly four-factor alphas for decile portfolios formed on lagged residual media coverage; expected coverage is the estimate from the multinomial regression shown in Table 3.3, column 5; * significant at 10%; ** significant at 5%;*** significant at 1%.

	four-factor alpha		
	Decile 1 (low)	Decile 10 (high)	Decile 1 – Decile 10
Unexpected absolute coverage	-0.00061	0.00158*	-0.00219**
	[0.00085]	[0.00092]	[0.00102]
Unexpected positive coverage	0.00089	0.00157*	-0.00069
	[0.00079]	[0.00094]	[0.00104]
Unexpected negative coverage	-0.00163	0.00470***	-0.00632***
	[0.00188]	[0.00124]	[0.00215]