The Causal Impact of Media in Financial Markets*

Joseph Engelberg

Christopher A. Parsons‡

October 2009

Abstract: It is challenging to disentangle the causal impact of media reporting from the impact of the information being reported. We solve this problem by comparing the behaviors of investors with access to different media coverage of the same information event. First, we use zip codes to identify 19 mutually exclusive trading regions, corresponding to 19 large U.S. cities and local newspapers (e.g., the Houston Chronicle). For all earnings announcements of S&P 500 Index firms, we find that *local* media coverage strongly predicts *local* trading, after controlling for characteristics of the earnings surprise, firm, local investors, and reporting newspaper(s). Reverse causation does not explain our findings. The local coverage-local trading effect: 1) holds for firms unlikely to be of local interest (e.g., remotely located, sparsely held by local investors) and 2) disappears entirely during extreme weather events, which leaves media content unchanged, but disrupts transmission to investors. The evidence supports the idea that media -- apart from the information they transmit -- affect investor behavior.

* We thank Nick Barberis, Matthew Gentzkow, Jay Hartzell, David Hirshleifer, Paul Tetlock, Sheridan Titman, and Luigi Zingales for helpful feedback and discussions. We thank Terry Odean for providing the large discount brokerage data.

‡ Contact: Joseph Engelberg, Kenan-Flagler Business School, University of North Carolina at Chapel Hill, (Email) joseph_engelberg@kenan-flagler.unc.edu, (Tel) 919-962-6889 and Christopher Parsons, Kenan-Flagler Business School, University of North Carolina at Chapel Hill, (Email) chris_parsons@kenan-flagler.unc.edu, (Tel) 919-962-4132. Please address all correspondence to Joseph Engelberg.

I. Introduction

Traditional asset pricing models have made little distinction between the production of information and its transmission to investors. In such a framework, information dissemination is complete and immediate. However, a number of recent theories have relaxed this assumption, admitting a direct role for information intermediaries -- like financial media -- whose packaging and transmission of information may affect investor behavior directly, apart from any information effects. Yet, it is precisely this distinction that makes media effects so difficult to measure. A television, radio or newspaper account of a firm is typically accompanied by a release of information. Therefore, when the market reacts to a media report it is difficult to determine whether the market has responded to the underlying information (the "information effect") or to the medium in which it is delivered (the "media effect").

Researchers have pursued two approaches to solve this identification problem. The first is an attempt to identify a set of "non-informative" media events and argue that any observed responses to these must be media-driven. Although such examples undoubtedly exist, clinical dissection of the event, as in Huberman and Regev (2001), seems necessary to completely rule out information-based explanations. More generally, purging media events of "soft" information (e.g., Tetlock et al. (2008), Mayew and Venkatachalam (2008)) would appear considerably more difficult, if possible at all.

The second approach is purely cross-sectional. The basic idea is to take two groups of agents, and for the same information event, “treat” them differently in terms of media coverage. Recent examples of this approach are DellaVigna and Kaplan (2007), who show that voting patterns could be predicted by whether Fox News was carried by an area’s local cable provider. Likewise, Gerber et al. (2009) randomly assign households in the Washington D.C. Area to receive a subscription of either the Times or the Post, finding that political views (elicited via survey) were altered. Compared to the first approach, the key advantage is that there can be no information effects, in the sense that the total amount of information available to be reported is fixed by construction. Thus, any observed differences can only be ascribed to differences in media treatment of the underlying events.

Our study takes the second approach and is, to our knowledge, the first to do so within the study of financial markets. Using retail brokerage accounts, we first identify nineteen local, non-overlapping trading markets coinciding with major U.S. cities. Then, for each of these markets, we identify a local information source: the daily newspaper of that city. The identifying assumption is straightforward: investors near Minneapolis (for instance) are more likely to read the *Minneapolis Star Tribune* than local papers in other regions, such as the *Seattle Post Intelligencer*. This linkage and the fact that local media outlets often differ in their coverage of the same underlying event affords a powerful test of the media's effects on financial market participants.

Our main result is that for an earnings announcement by a given S&P 500 Index firm, trading in each of the nineteen local markets is strongly related to whether the local paper covers the announcement. Importantly, this result is distinct from home bias (e.g., Coval and Moskowitz (1999, 2001), Ivkovic and Weisbenner (2005)) whereby investors prefer to trade stocks of local firms. Our regressions control for the proximity between an investor base and the company's headquarters, for the size of the earnings surprise, firm characteristics, and paper (local and national) and city fixed effects. Despite these controls, we find that when local media outlets report firm news, local trading volume increases by roughly 48%. By contrast, the most extreme earnings surprises increase trading by no more than 15%. This comparison makes clear that in our sample even events with high information content generate far less activity than the media's coverage of the event.

We interpret our main result causally, with media coverage stimulating local trading activity. The main alternative is reverse causation, whereby local media coverage arises because of pre-existing local interest in certain stocks.[1] For example, we might imagine that the *Star Tribune* is more likely to cover a local stock like Minnesota-based ADC Telecom because local investors are interested in it. We address this possibility several ways. First, we specifically exclude all local stocks from our analysis, which presumably are of particular interest to local investors. This does little to the media effect; even to stocks headquartered in remote cities, investors respond to media coverage about them in their local newspapers. Second, we infer pre-existing interest of local traders from their portfolios holdings. As expected, we can predict which cities will trade a given stock from the fraction of that city's investor base that hold or trade that stock. However, the media effect survives the addition of these controls, strongly predicting local trading volume.

Our most precise tests identify media effects from the correlation between local trading and *exogenous* variation in local media coverage, fluctuation that cannot plausibly be related to other determinants of trading activity. We develop two approaches. The first exploits the fact that our sample period (1991-1996) mostly predates the Internet, so that investors depended on physical delivery of their print media sources. We collect data on the daily weather in each local market, and identify weather events severe enough to likely disrupt or delay delivery of the local newspapers: blizzards and hailstorms. Although a blizzard in Minneapolis may make it less likely that the Star Tribune will reach its readers, it is clearly unrelated to the Star Tribune's financial reporting -- perhaps a story on the earnings of Starbucks some 1500 miles away -- or preferences of Minneapolis investors. On such days, we find that the relation between local trading and local media coverage is severed. Crucially, this is not simply an extreme weather effect; extreme weather is unrelated to local retail trading. It is the interaction with media coverage (which is similar on extreme weather days) that predicts investor behavior.

[1] Note that local media sources having superior information about certain firms (e.g., Coval and Moskowitz (2001), Massa and Simonov (2006)) cannot explain cross-sectional differences in trading. With immediate and complete dissemination of information, both local and non-local investors would immediately receive this information. Differential access to information – e.g., differential media coverage - is required to sustain local information advantages.

The second proxy also uses a feature of the pre-Internet era: a story can only be reported in *tomorrow's* printed newspaper if it is written by *today's* print deadline. Such deadlines allow sufficient time for the lengthy technological steps required for printing to be completed. Presumably, this time varies by newspaper due to, e.g., printing technology, different time zones, space for other stories, etc. Regardless of the specific reason, we verify that earnings announcements are not uniformly reported the day directly after their release by firms. For example, the earnings of Dell Computer Inc. released late Wednesday afternoon might meet the *Seattle Post-Intelligencer*'s print deadline for Thursday's morning edition, but may miss it for the *Minneapolis Star Tribune* (ahead by two time zones), which carries the story on Friday. We examine events like these, asking whether such micro-variation in media coverage leads to micro-variation in local trading. It does. Local trading is most sensitive to media coverage *on the specific day it occurs*, which like the extreme weather tests, is difficult to ascribe to correlated omitted variables.

Our cross-sectional framework is inspired by recent work that examines media and voting behavior (e.g., DellaVigna and Kaplan (2007)) and more generally relates to an expanding literature exploring the media's influence on real outcomes (e.g., Stromberg (2004), Gentzkow and Shapiro (2004), Gentzkow (2006), and Gerber et al. (2009)). Given evidence that the media can alter political attitudes, it is perhaps unsurprising that financial market participants are similarly susceptible. On the other hand, this study constitutes the first systematic decomposition of media and information effects, and lays the groundwork for future work that explores in more detail the specific mechanisms underlying the media's influence.

Specifically, our evidence demonstrates *that* the media matters for investor behavior, not *why* it does. There are several plausible explanations: media may help investors allocate their limited attention; investors may "trust" some media more than others (e.g., Dyck, Volchkova and Zingales (2008)); investors may be unable to distinguish genuine news from "stale" information (e.g., Tetlock (2009)); or perhaps financial media are biased (e.g., Butler and Gurun (2009)), and can, like the political media, alter investor behavior as a result. Regardless of the reason, our study lays the empirical foundation for each of these studies by establishing causality between media and investor behavior.

The remainder of the paper is organized as follows. In section II we briefly describe the data and its sources. Next, we examine trading of retail investors in a number of local markets, with a particular eye on local media coverage in Section III. We then consider how the endogeneity of local media coverage affects our conclusions and designs several related tests in Section IV.

II. Data

Our analysis requires data of three types: 1) earnings announcements, 2) local media coverage (around those earnings events), and 3) trading of retail investors.

We first collect all earning announcements dates from S&P 500 firms between January 1991 and December 2007. Earnings announcement dates and S&P 500 membership are taken from COMPUSTAT. Because we are concerned about the appropriate timing of earnings announcement dates, we cross-check the dates in COMPUSTAT with those in I/B/E/S and only keep the dates in which they match. For each earnings announcement we calculate the earnings surprise (standard unexpected earnings, or SUE) based on the random walk model with price as the deflator (Livnat and Mendenhall (2006)).

We download our media coverage data from ProQuest's newspaper database. We have articles from the following cities (newspapers): Boston (*Globe*), Denver (*Post*), Detroit (*News*), Houston (*Chronicle*), Las Vegas (*Review Journal*), New York (*Times*), Pittsburgh (*Post Gazette*), San Antonio (*Express News*), San Diego (*Union Tribune*), San Francisco (*Chronicle*), Seattle (*Post Intelligencer*), St. Louis (*Post Dispatch*), St. Petersburg (*Times*), Minneapolis (*Star Tribune*), Atlanta (*Journal Constitution*), Sacramento (*Bee*), Washington (*Post*), and New

Orleans (*Times Picayune*). We also collect data from two newspapers with national audiences: the *USA Today* and the *Wall Street Journal*.

Articles about firms in ProQuest are indexed by company name. Using the COMPUSTAT firm name we manually match each S&P 500 firm to its indexed firm name in ProQuest (some firms are linked to multiple ProQuest company names). Using the indexed firm names and the advanced search option in ProQuest, we search and download all articles in the newspaper database between 1991 and 2007 for each firm.[2] The result is a database of newspaper articles linked by GVKEYs to S&P 500 firms. When we say that a newspaper X "covers" firm Y's earnings announcement, we mean that there is an article in newspapers X on day 0, 1 or 2 after firm Y's earnings announcement in which firm Y is indexed by ProQuest in the article.

The trading data come from the well-known large discount brokerage database used by Barber and Odean (2000), with data available between 1991 and 1996. The database consists of the holdings and trading behavior of 77,795 households of which 54,297 have valid zip code information. Among those 54,297 households, 43,198 hold at least one common stock for which we have matched COMPUSTAT/CRSP information. Among those 43,198 households, 15,951 households are located within 100 km of any of our 19 local newspapers. Portfolio holdings are available monthly between January 31, 1991, and December 31, 1996, while account-level trading data are available between January 1, 1991 and November 30, 1996. We use the trading data as the dependent variable in the majority of our tests, and use the holdings data when we consider investor predisposition to investing in certain stocks.

Table 1 provides some summary statistics from our data. Unsurprisingly, we have the largest sample of households with accounts in well-populated areas like San Francisco (4,076), Los Angeles (1,913) and New York (2,808) with generally more accounts in the western United States (Zhu (2002)). However, we still have a considerable number of brokerage accounts in other major U.S. cities like Boston (635), Washington D.C. (983) and Houston (607). From the

[2] Some articles in ProQuest are full-text while others provide only the headline and a summary or abstract.

third column of the top panel in Table 1, we see that, given that a household holds common stock, the average number of stocks held is about 2 in every city. (Of course, this does not include mutual fund ownership or household holdings in other brokerage accounts.)

Table 1 also demonstrates a strong local bias in nearly every city. To measure local bias we pool the accounts of all investors in a city and compute the number of local stocks (within a 100-km radius of the local newspaper) held by this set; we then scale this number by the number of *available* local stocks (inferred from COMPUSTAT). We compute a similar statistic for non-local firms (non-local stocks held scaled by the available number of non-local stocks). Comparing the statistics for local and non-local stocks demonstrates a strong local bias that has been documented in other papers using the same database (Ivković and Weisbenner (2005), Seasholes and Zhu (2005) and Zhu (2002)). For example, the Minneapolis portfolio holds 21.8% of local stocks compared to 3.6% of non-local stocks, and the San Francisco portfolio holds 44.4% of local stocks compared to 15.1% of non-local stocks.

Local papers also report disproportionately about local firms (Butler and Gurun (2009)). For example, Table 1 indicates that the San Antonio Express News covers 3.5% of earnings announcements reported by non-local firms, but over 40% for those by firms located near San Antonio. This relationship holds for each of our 19 papers and, like the San Antonio example, the differences in frequency are often quite large. The paper exhibiting the least local bias is the *New York Times*, who reports the earnings of New York firms twice as often as those located remotely. The *Times Picayune*, at the other extreme, is over 100 times less likely to cover earnings of a firm located outside the New Orleans area.

III. Local Trading Responses to Local Media

The appeal of local media coverage is that it allows us to examine the behavior of traders subjected to different media coverage of the same information event. Presumably, this situation describes virtually every trader each time a piece of information is released into the market.

However, almost never do we observe a given investor's information sources (television, radio, print media, internet, personal advice, etc.), and even then, we often do not observe his or her responses. This greatly limits our ability to make inferences about the media's ability to influence financial market participants.

Our empirical setting is fortunate in allowing us to identify specific groups of investors who are more (or less) likely to receive coverage by specific media outlets. Identification here requires only the assumption that an investor living in a given metro area is more likely to read the local newspaper than another regional paper (i.e., a different city's local paper). Importantly, identification does not require investors to rely solely on local papers for their financial information, although some may. More exclusive reliance on local media would increase the power of our tests, but is not necessary to achieve identification.

As seen in Table 1, the mean value of *Local Media Coverage* is only 2.8% (standard deviation of 16.6%), indicating that roughly one in thirty-five quarterly announcements is reported in local newspapers. By contrast, this is over an order of magnitude less frequent than the *Wall Street Journal*, which accompanies an earnings release with a story or report 29% of the time in our sample (1991-1996) and 33% of the time over the period 1991-2007.

A. Methodology and Specification

The relation between local media coverage and local trading is evident from simple correlations. On days in which earnings are reported in the local paper, the average absolute dollar volume of local trading in the mentioned stock is $2,200. By contrast, the average local trading on "non-news" days at the local level is only $290. While a sizeable difference, such a simple comparison limits the conclusions we can make. Local trading is likely determined by a number of other factors, some of which may be correlated with local media coverage. Thus, we estimate the following multivariate regression:

$$\begin{aligned}&Log(Local\,Trading\,Volume_{i,j,t})\,|\,Earnings_{j,t} = \alpha + \delta \bullet Local\,Media\,Coverage_{i,t} + \\&\phi \bullet Firm\,Attributes_{i,t} + \zeta \bullet Earnings\,Surprise_{j,t} + \\&\gamma \bullet Media\,Fixed\,Effects_{i} + \sigma \bullet City\,Fixed\,Effects + \varepsilon_{i,t}\end{aligned} \tag{1}$$

The dependent variable measures trading responses of retail investors in each of the nineteen local markets, *i*, to each S&P 500 firm *j*'s earnings released at time *t*. All local market are mutually exclusive, so that trading of firm *j*'s stock may occur in local market 1, but may not occur in local market 2, and so on. The goal of equation (1) is to understand what firm, media, investor, and earnings characteristics explain these cross-sectional differences in local trading behavior.

The dependent variable is constructed over a three-day trading window, so that if firm *j*'s earnings are released on Wednesday, local trading is recorded if it occurs on that Wednesday, Thursday, or Friday. In each specification of Table 2, the trading variable is the natural logarithm of one plus the absolute dollar trading volume in firm *j*'s stock, aggregated across all investors (for which we have records) in region *i* within three days (0, 1 or 2) of the earnings announcement.

The main explanatory variable of interest is *Local Media Coverage*, which takes a value of one if region *i*'s local newspaper reports firm *j*'s earnings (also within three days of its announcement), and zero otherwise.

Crucially, Equation (1) also includes paper fixed effects for each of the nineteen local newspapers. This means that *Local Media Coverage* is identified solely from the differential responses between a newspaper's local readers and its non-local readership. For example, *Media Fixed Effects* includes a dummy variable for the Houston Chronicle that takes a value of one, for each of the nineteen local markets, whenever the Chronicle covers an earnings announcement. Because there are nineteen such dummy variables, each local paper is allowed to have a differential influence on trader behavior. However, only for the 1/19th of the observations corresponding to Houston investors does *Local Media Coverage* equal one, allowing the Houston paper to have an *additional* impact on those investors most likely to be exposed to the story (local investors).

We also include controls for media mentions in either of two national media: *The Wall Street Journal* and *USA Today*. Both papers have national readership, and consequently, resist the linkage to specific investors that is possible for local papers. Nonetheless, their inclusion in our regressions control for any omitted correlation between local and national media, affording us the ability to uniquely identify media effects through local channels.

Other important control variables include *Firm Attributes*, which includes each firm's market capitalization (measured at the end of the most recent fiscal year) as well its Fama-French 30 industry classification. By clustering residuals by firm, we compute standard errors that allow for each firm i to have its own unobserved effect on the likelihood that a newspaper covers its earnings, and allow such firm-specific heterogeneity to change across time (Petersen (2009)).[3] Some of our robustness tests include firm fixed effects, firm-city fixed effects, and other similarly constructed control variables.

The *Earnings Surprise (SUE)* control variables account for the fact that some earnings events lead investors to revise their expectations more than others, and therefore, are more likely to generate trade. To capture these differences, we calculate *SUE* as the difference between actual earnings and earnings four quarters ago divided by price (see Livnat and Mendenhall (2006) for a detailed discussion of *SUE* construction). We form quintiles of this surprise variable, after pooling all earnings announcements.[4]

---

[3] Clustering by firm allows standard errors to be correctly computed in the presence of a firm fixed effect or a temporary effect (e.g. a firm effect which decays over time). From Petersen (2009): "The standard errors clustered by firm are unbiased and produce correctly sized confidence intervals whether the firm effect is permanent or temporary." In unreported results, we also find that the inclusion of firm fixed effects to the estimation of equation (1) makes no qualitative difference in our results.

[4] The results are unchanged if we use alternative ranking criteria, such as within-industry, within industry-year, and forming deciles rather than quintiles. The results are also unchanged is we use a

B. Results

Table 2 presents the results of linear regressions of equation (1). The first and second columns include only firm clustering and date fixed effects. As seen, *Local Media Coverage* increases local trading volume in the typical market by almost 75%. Column 2 adds industry controls and firm size, and although larger firms are associated with more trading (p=.000), this has only a trivial effect on both the economic and statistical significance of *Local Media Coverage*.

Controls for the magnitude of the *Earnings Surprise* reveals an intuitive finding. The omitted dummy is the middle quintile of "least surprising" earnings announcements, relative to the earnings four quarters ago. Column 3 indicates that trading is highest following the most surprising earnings events, both positive and negative. Both coefficients on SUE_1 and SUE_5 suggest increases in trading volume of about 7-8%, both effects highly significant. In contrast, those corresponding to less extreme earnings surprises, SUE_2 and SUE_4, are much less statistically and economically significant. A rational interpretation is that *SUE* captures an investor's surprise about cash flows, so that SUE_1 and SUE_5 events are associated with significant portfolio rebalancing; it is also consistent with a number of behavioral explanations. In either case, the coefficient of interest on *Local Media Coverage* remains essentially unchanged.

In unreported results, we have also run a number of regressions with an analyst-based definition of earnings surprise (Livnat and Mendenhall (2006)) and finer breakpoints, e.g., deciles instead of quintiles. The magnitude on any *SUE* variable never exceeds 0.15 in any regression specification. In other words, when we compare the effect of the information in SUE with the effect of media coverage (*Local Media Coverage*), the media effect is consistently three

definition of earnings surprise based on the median analyst forecast (rather than the random walk model).

to ten times larger than the information effect. At least for retail investors, demands appear far more sensitive to media coverage than to the underlying information.

An important consideration is that so far, our regressions have not controlled for cross-sectional differences between newspapers. Given that many of the local papers have substantial national readerships (e.g., *The New York Times*, *Washington Post*), stories broken by particularly influential papers may have the power to increase local trading relative to their less widely read counterparts. Similar reasoning applies to the two national papers in our sample. The specific concern is that without paper controls, the coefficient on *Local Media Coverage* may be picking up correlation between local coverage and non-local (or national) coverage, with the latter generating any trading effects. For example, suppose that a Detroit investor bases her trading decisions exclusively on stories reported in the *Wall Street Journal (WSJ)*, to the exclusion of the local *Detroit Free Press (DFP).* Without a control for a story's mention in the *WSJ*, the estimated coefficient on *Local Media Coverage* (*DFP* for a Detroit investor) will still be positive, because 72% of the 3,841 earnings announcements reported by the *Free Press* are simultaneously reported by the *Journal.*

Column 4 addresses this concern by including a dummy variable for each of the nineteen local and two national media outlets, allowing them to affect trading in any market. The marginal effects shown for each paper are relative to the benchmark case when neither a local nor national media outlet covers the earnings announcement.[5] As seen, most coefficients for these individual papers are not statistically different from zero, indicating that most have only a trivial effect on their non-local investor community. The few exceptions include *The New York Times* and the *San Francisco Chronicle*, newspapers widely read outside of their immediate catchment area.

Of the national newspapers, only the *USA Today* is associated with abnormal trading volume (16.5%, p=0.00); in contrast, the coefficient on the *Wall Street Journal* is nearly

[5] Note that the "dummy trap" does not apply among the set of paper fixed effects, as they do constitute a mutually exclusive set.

identically zero. While this may at first appear puzzling, the summary statistics for each paper (Table 1) make clear the reason for this disparity. The *WSJ* covers almost one-third of all earnings announcements, while the *USA Today* mentions fewer than one in twenty. Not being a financial newspaper, it is therefore unsurprising that the *USA Today* would cover only earnings announcements likely to be of particular interest (similar to the *Earnings Surprise* coefficients). Nonetheless, the main coefficient of interest remains highly significant, although fixed effects for each paper cuts its magnitude roughly in half.

In the final column of Table 2, we include fixed effects for each of the nineteen local markets, allowing average trading volumes to vary across cities. This set of controls is clearly important, increasing the R^2 by almost 2 percentage points. The results here are intuitive. We see that relative to investors in Denver (the arbitrarily omitted city), local trading volume is higher in larger markets such as Boston, Houston, New York, San Francisco, and San Diego, and is lower in smaller ones such as Las Vegas, San Antonio, and St. Louis. Moreover, with such city controls, the coefficient on *Local Media Coverage* is now identified purely off of differences-in-differences between city-newspaper pairs. That is, the coefficient on *Local Media Coverage* picks up *only* the average marginal impact of the media source being local, not the average effects due to larger cities, more influential papers, etc. Taking the final column as the most indicative of the underlying behavior, local media is associated with almost a 50% increase in trading activity.

C. Robustness

The estimates in Table 2 are presented for all transactions lumped together. In Table 3, we disaggregate them, showing the results for buys and sells separately in Panels A and B respectively. A priori, one might expect the effects to be stronger on the buying side, given the evidence that retails investors are unlikely to sell stocks short (Barber and Odean (2008)). If retail investors do not sell short, then they can only sell when they hold the stock, and thereby presents fewer occasions upon which to respond negatively to local media coverage.

As expected, the coefficients in this table approximately sum to those in the previous table. However, Panel A shows that although buying is responsible for more than half of all trading activity, with coefficients ranging from .46 (column 1) to 0.29 (column 5), the disparity over selling activity is small. In Panel B, we observe slightly lower magnitudes (.38 to .26) for the selling/shorting regressions, but *Local Media Coverage* remains highly significant in each specification. Table 3 shows clearly that the effects of the media, at least the local media, are pervasive across both additions and subtractions to investors portfolios.

A second issue of robustness concerns the way in which we have defined our dependent variable. Table 4 presents two alternatives. In the first three columns, our aim is to reduce the potential for a few large orders to generate the results we find in the first two tables. Instead of trading volume, we define as the dependent variable the natural logarithm of the *number* of local accounts that trade the stock of interest. This in turn changes the interpretation of the local media coefficient. Instead of the percentage change in absolute trading volume, the coefficient on *Local Media Coverage* represents the increases in the number of local traders who trade, regardless of the size of their order.

The first column, for example, indicates that mention in the local newspaper increases by approximately 4.67% the number of accounts trading the stock of interest within three days of its earnings being locally reported. The mean number of households in each local market is 840 (standard deviation=1028), ranging from only 134 accounts in New Orleans to 4,078 in the San Francisco Bay Area. Thus, local news mentions affect, on average, as few as two traders (4.67%*37=1.7) to well over one hundred. The second and third columns paint a similar picture to Table 3; both the percentage change in the number of "buyers" or "sellers" orders are significantly related to *Local Media Coverage*, and in approximately equal magnitudes.

The next specification reduces further any residual size effects. Here, we map the underlying trading behavior to a discrete representation. The dependent variable takes a value of one if there is *any* local trading in the stock corresponding to the earnings announcement, regardless of either the number of traders or size of their trades. For example, suppose that after IBM released its earnings on a given date, and that the following day one hundred traders in San Francisco traded some $100,000 worth of IBM stock, while in New Orleans, only two trades summing to $500 are observed. Despite the substantial size differences, the dependent variable takes a value of one in each case. As before, all estimates here include the full assortment of control variables, as shown in the final columns of Tables 2 and 3.

Columns 4-6 of Table 4 show the results of linear probability models for the discrete variable specification (although probits present similar estimates).[6] As in column 1, column 4 aggregates all transactions together, indicating that *Local Media Coverage* increases the probability of observing any local trade by 5%. The second and third columns break up this effect into buy and sell transactions where, also as before, the effect is relatively equal between the two, and highly significant.

IV. Endogeneity

The primary advantage of our cross-sectional approach is that by construction, it eliminates the typical concern that information, rather than media coverage, is driving the result. In this section, we discuss a second type of endogeneity particularly relevant for the type of tests we run. Specifically, because we are analyzing a number of local newspapers that likely cater to the interests of local investors, the possibility arises that local media may simply *reflect*, rather than *cause* the trading patterns we observe. For example, consider the hypothetical case where a local newspaper polls its readers, asking of them which stocks they would like the paper to cover. If the paper heeded these suggestions, then the observed correlation between local

[6] We use the LPM specification to accommodate the large number (1,350) of date fixed effects in our sample.

coverage and local trading would not be spurious, but the causation would not support the identification of media effects on financial markets.

Fortunately, our data are well suited to address the potential for such reverse causation. The analysis in this section is organized into three parts. First, we collect a number of additional control variables designed to measure each local market's pre-existing – i.e., unrelated to media coverage – interest in certain stocks. As we will see, many of these measures are very precise, allowing us to argue that any remaining relation between trading and media coverage can be interpreted in the desired (causal) way. Our second and third tests allow for even more precise identification. We identify two characteristics that cause interference with the *transmission* of media coverage to investors, but leaves unchanged both any underlying content (e.g., information, media spin), as well as the pre-existing demands of investors. As we will see, such exogenous variation strongly predicts trading, posing a significant challenge to alternative interpretations.

A. Local Demand

Our first set of tests is motivated by the observation that both retail and institutional investors appear to tilt their portfolios toward geographically local stocks at the expense of their remote counterparts (Coval and Moskowitz (1999, 2001), Ivković and Weisbenner (2005), Zhu(2002), Seasholes and Zhu (2005)). Regardless of why such "home bias" exists, the concern is apparent.[7] In addition to being more widely held by local investors, local firms are more likely to be covered by local newspapers. Thus, what we interpret causally as a media effect may reflect little more than the tendency of both local papers and local investors to pay attention to local stocks.

[7] Coval and Moskowitz (1999, 2001) argue for an information-based explanation of the home bias among institutional investors. Ivković and Weisbenner (2005) make a similar argument for retail traders. Zhu(2002) and Seasholes and Zhu (2005) provide a behavioral explanation of the home bias of retail traders based on familiarity.

To address this possibility, for each of the nineteen regions *i*, we augment equation (1) to include a dummy variable indicating whether firm *j*'s headquarters is within 100 km of the local newspaper.[8] This designation allows us to identify stocks likely to be of particular interest to each cohort of local investors. Confirming the home bias found in previous studies, we find that investors are more likely to both hold (0.43% vs. 0.16%) and trade (9.9% vs. 2.0%) the stocks of local firms following earnings announcements. Local coverage is tilted even more toward local firms (19.2% vs. 2.7%).

Even so, column 1 of Table 5 indicates that even when all local firms (<100 km of the paper's headquarters) are excluded, *Local Media Coverage* remains strongly related to local trading. Here, the thought experiment is to compare the trading patterns between a Houston and San Antonio investor, after Ohio-based Proctor and Gamble (P&G) releases its quarterly earnings. From the perspective of each investor, P&G is a non-local firm, with headquarters removed by over 1,000 miles from each. However, if the *Houston Chronicle* reports P&G's earnings while the *San Antonio Express News* does not, trading volume in Houston increases by 28.5% relative to its normal volume, while no similar increase is seen in San Antonio. The second column presents this analysis in a slightly different way, including a dummy variable for a local firm, and returns to the original, full sample. Adding this variable permits us to observe a "local firm effect" of nearly 50%, confirming the previous findings of (Ivković and Weisbenner (2005), Zhu(2002), Seasholes and Zhu (2005)), who document a home bias for retail traders using the same database. However, our primary interest is in the *Local Media Coverage* variable; the second column reveals that it strengthens slightly to over .3, and remains highly significant.

The third and fourth columns consider more precise proxies for the pre-existing demands of local traders. Although a firm's geographic proximity may influence an investor's willingness to hold or trade its stock, other factors may generate cross-region differences in

[8] Alternative breakpoints (e.g., 200 km, 500 km) give nearly identical results.

investor interest. For example, it would be unsurprising for investors in St. Louis (home to beer bottler Anheuser-Busch) to be interested in the fortunes of other brewers or distillers, even those located remotely. Similar arguments can be made for other cities, whose investors may cluster along certain stocks or industries. If local media outlets understand such preferences, and consider them in their coverage decisions, then the same identification remains - media coverage may simply reflect existing investor interest, rather than cause it.

To address this concern, we examine the individual portfolio holdings of each investor within each of the nineteen local markets. This exercise allows us to identify stocks likely to be traded in each market, based directly on which stocks are already held or frequently traded. For each region i, we construct two additional local variables, for stock j: 1) the fraction of investors within region i that own stock j it at the beginning of the trading month, and 2) the fraction of investors within region i that traded stock j the previous month.

The third column shows the results of the regression, once we add the first control variable. As seen, although *Fraction Locally Held* is observed as being strongly related to local trading, its inclusion actually strengthens the local media variable, which remains highly significant. The fourth column is similar, except that *Fraction Locally Traded* is the key control variable. It has a similarly strong effect (as expected), although the coefficient on *Local Media Coverage* remains strong. The final column includes all three controls for pre-existing local demand simultaneously, which shows an R^2 of over 8%, some fifty percent larger than the base regression in the final column in Table 2. The increase in explanatory power, along with the extremely high significance of each control, suggests that we have controlled for a large portion of the pre-existing local demand for each stock j and region i. If so, then the coefficient on *Local Media Coverage* (roughly 28%) represents the pure effects of media coverage.

In unreported results, we have conducted a number of robustness checks similar to the ones reported here. For example, varying the definition of a "local" firm (e.g., 200 km, 500 km) makes very little difference. Similarly, in the portfolio analysis, we have experimented with a number of alternatives, including longer time periods, e.g., averaging the fraction of local investors that hold a given stock over several months up to a year, using dummy variables instead of continuous variables in the regressions, etc. The main results even survive the inclusion of *city-firm pair fixed effects*, which control for any time-invariant, unobserved preferences any local market may have for any one stock. Results from these tests and those reported in Table 5 indicate quite clearly that although media coverage appears, in part, to cater to reflect the interests of its readers, this tendency is unable to explain the media-trading connection.

B. Exogenous Variation in Transmission of Media Content

The strategy in the previous section was to control explicitly for each market's existing demand for certain securities, so that any remaining relation between trading and media can be interpreted causally. Here, we attempt to achieve identification through different means. The goal of this analysis is to find exogenous variation in the *transmission* of media stories to investors, but crucially, to leave both the content of the media story and the demands of local traders constant. Moreover, we look for high-frequency variation, which is particularly convenient in attempting to hold constant investor demands, which may change over time.

In our search for exogenous variation in the dissemination of media content to investors, we are aided by the fact that our sample period (1991-1996) is largely pre-Internet. This implies that in contrast to the modern electronic era: 1) printed content had to be processed and transferred to an actual printed media (the paper), and 2) the printed newspaper needed to be physically delivered, usually door to door by a delivery person, or more likely, a "paper boy." Disruptions in either step would delay or prevent altogether media content from reaching investors, but because both are unrelated to either the content or investor demand for financial securities, they represent the ideal type of exogenous variation that makes identification feasible.

Our first source of variation is extreme weather. For each of our nineteen local markets, we collect weather data. Weather data are taken from the National Climatic Data Center at ftp://ftp.ncdc.noaa.gov/pub/data/gsod, which provide daily data for each weather station as well as the coordinates (longitude and latitude) of each station. For each trading area, we use weather from the station closest to the local newspaper's zip code. The weather data include high and low temperatures , inches of snow, precipitation and indicators of extreme weather.

From these data, we identify two types of weather events likely to impede or significantly delay delivery of the local, daily newspaper: hailstorms/freezing rain and blizzards. Following Loughran and Schultz (2004), we define a *Blizzard* as greater than or equal to eight inches of new snow, although other definitions give similar results. *Blizzard* and *Hail* are both dummy variables that take a value of one if the extreme weather occurs on either the earnings announcement day or that immediately proceeding. Extreme weather events, as expected, are clustered in the Northeast and Midwest (e.g., Boston, Minneapolis, Detroit). Across all markets, we identify 736 observations that include a *Blizzard*, and 1,188 that include *Hail*. There is one overlap, so their union, *Extreme Weather*, includes 1,923 observations.

The first column of Table 6 shows the results of estimating (1), with two addition variables: the *Hail* variable and *Hail**Local Media Coverage*. The usual set of full controls is deployed. With the presence of the interaction term, the coefficient on *Local Media Coverage* now corresponds to the effect of local media on local trading on days without *Hail*. As seen, this is nearly identical to the baseline case, shown the final column of Table 2, with a point estimate of 0.48 (p=0.000). As seen, the presence of *Hail* does not have a meaningful impact on local retail trading; only the weakest inference can be made regarding the positive point estimate (p=.17). However, we are interested in the sign on the interaction term, which is negative and highly significant, even more so than the original *Local Media Coverage* coefficient itself. This indicates that on extreme weather days, the relation between media coverage and trading is severed entirely.

Importantly, the interaction here has nothing to do with *Hail* lowering the probability of local financial reporting, as indicators for mention in each of our twenty-one papers are still included. Instead, the negative sign on the interaction demonstrates that local media content pertaining to the firm of interest, *on days when it is unlikely to be delivered to local investors*, has no effect on trading. This finding, we argue, isolates precisely a media effect – it holds the underlying information constant (as do our previous regressions), but also keeps constant the unobservable demands of local traders.

The second column presents the same analysis, except for our second type of extreme weather, *Blizzards*. The findings and their interpretation are nearly identical. We see that *Blizzards* alone are not observed to be associated with local retail trading. That is, on days where the local paper doesn't report any news about the firm of interest (recall that the presence of the interaction leads to this interpretation of *Blizzards*), the presence or absence of a *Blizzard* makes no difference on trading decisions.[9] However, the interaction is, as above, strongly negative, more than offsetting the positive effects of *Local Media Coverage*. The final column aggregates all instances of *Extreme Weather*, and delivers nearly identical results. On days when the delivery of local news is exogenously delayed or prevented, investors do not respond to any content. This evidence strongly suggests that the media is not passively correlated with a pre-existing relation between local trading and the information contained in media reports.

Our final test exploits a second source of exogenous variation in the transmission of media coverage to investors – "time-to-print." In contrast to the modern, electronic information age when web-based news outlets can collect and disseminate stories in hours or minutes, in our time period, print newspaper require substantial lead time in order for information on day *t* to be printed on day *t+1*. Fortunately, this requirement introduces a further source of variation

[9] Note that this finding implies at least some reliance on local newspapers for financial information. Instance of extreme weather (e.g., *Blizzards* and *Hail*) are temporary disruptions in the delivery of local news sources, but are unlikely to be correlated with, for example, weekly subscriptions to financial magazines. The fact that inclement weather only dampens trading on days when local papers report news thus suggests that local papers fill this function for at least some retail investors.

that is, in theory, perfectly uncorrelated with either firm-side or investor-side determinants of trading decisions.

Unfortunately, we of course lack data on each paper's deadline times; even if we did, we do not observe precise times corresponding to the release of earnings information. However, we can infer these (noisily) from the empirical distribution of earnings dates and story dates. We are specifically interested in whether micro-variation (e.g., one day) across local papers predicts similar micro-variation in local trading. That is, if San Diego's *Union Tribune* reports Home Depot's earnings on Tuesday, and the *Boston Globe* reports it on Wednesday, do we observe Home Depot stock with abnormal trading volume in San Diego on Tuesday, and abnormal trading volume in Boston on Wednesday? Like the tests of extreme weather in the sections above, such idiosyncratic, high-frequency variation is difficult to claim as correlated with unobserved determinants of local trading.

To conduct this analysis, we take our sample of all earnings announcements, and designate their initial release date as t. Then, we examine the distributions of media coverage in each market, both on the day immediately following ($t+1$), and the day after that ($t+2$). Here, our identifying assumption is that reasons that would cause a paper to carry an earnings story on day $t+1$ vs. $t+2$ are unrelated to any demand differences local investors may have between these dates. (It is unclear what underlying preferences could possibly generate demand differences over such a short interval, and across cities.)

Empirically, we observe substantial distribution across newspapers in the time between earnings announcements and news coverage. Averaged across all outlets, the ratio of $t+1$ stories to $t+2$ stories is 3.21 to 1, meaning that over 75% of all reported stories occur on the day immediately following their announcement day as recorded by COMPUSTAT and I/B/E/S. However, there is substantial variation, depending on the specific media source. Papers on the West Coast have the benefit of a later time zone (PST), and would be expected to "pick up" late day announcements more easily than their peers on the East Coast. For example, the $t+1$ to $t+2$

ratio for three large PST newspapers, the *Seattle Post Intelligencer*, *San Francisco Chronicle*, and *Los Angeles Times* is 4.25, 3.63, and 4.75 respectively. By comparison, the same ratios for the EST *New York Times, Atlanta Journal Constitution*, and *Pittsbrugh Post Gazette are* 3.09, 2.76, and 3.81 respectively. However, there are presumably paper-specific issues unrelated to time zones that also influence deadlines.[10] Regardless of the reason, our tests simply require that such differences be exogenous to the underlying content, as well as to local investor preferences for certain stocks.

Table 7 presents the results of running regressions very similar to equation (1), with two important changes. First, rather than pool all *t+1* and *t+2* observations as we do in Tables 2-6, we run regressions separately for each of these event-day subsamples. That is, we run equation only looking at trading/no-trading decisions on day *t+1*, and then a separate regression for those on day *t+2*. This decomposition does not change the total number of observations as each previous observation is now split into two, but does reduce the number of media and trading events within each group. The second change is that we decompose *Local Media Coverage* into current (i.e., same day) and past coverage. In the *t+1* regression, coverage can occur on day *t* (past) or day *t+1* (current). Similarly, in the *t+2* regression, coverage can occur on day *t* (past), day *t+1* (past), or day *t+2* (current). This framework allows the time-to-print to vary across newspapers, and more importantly, allows us to detect whether trading patterns in each city are predicted from this variation.

The first column of Table 7 shows the results for the *t+1* set of observations. As seen, trading on day *t+1* is much more sensitive to media coverage on day *t+1* than on day *t*. The estimated coefficients, .28 and .14 for *t+1* and *t* respectively, are different at the 1% significance level. In the second column, the determinants of trading two days removed (*t+2*) are examined. Here, the coefficients between *Past* and *Current Media Coverage* are seen to be virtually identical. However, in the final column when all observations are aggregated together, we again

[10] For example, the San Diego *Union Tribune* has the highest percentage of stories reported on day *t+2* in our sample, despite being located on the West Coast.

see evidence that trading patterns can be shifted a day either way, depending on when the story shows up in the local newspaper. Unless such micro-variation is correlated with high-frequency changes in local investor demand, this results provides complementary evidence to the weather tests shown before.

V. Conclusion

We exploit the geographic variation of local paper readership to design and implement empirical tests allowing the media's effect on investors to be identified, apart from any concurrent information effects. Analyzing the simultaneous reactions of investors in nineteen local markets to the same set of information events (earnings releases of S&P 500 Index firms), we find that the presence or absence of local media coverage is strongly related to the probability and size of local trading. Although we would expect for stories reported by local newspapers to reflect the investment interests of local readers, this alone does not reconcile the evidence. When we examine the individual portfolios of each local investor, we find that the local media coverage-local trading relation remains strong for non-local firms, as well as for those sparsely held or traded by investors in the local market.

Perhaps the strongest causal evidence comes from examining exogenous shocks to the transmission of media coverage to local investors. On days when extreme weather (hailstorms and/or blizzards) is likely to disrupt the normal delivery of daily newspapers, the link between media coverage and trading is broken. This is an important test, because weather shocks are neither correlated with any underlying content, nor with unobservable determinants of investor demand. A complementary test takes advantage of small (one or two day) differences between days in which earnings stories are carried by newspapers, presumably because of differences in time zones, printing technology, and so forth. We find that trading patterns are strongly related to the local patterns of media coverage. For example, an earnings reported by the *San Francisco Chronicle* on Wednesday stimulates trade in the Bay Area on Wednesday; the entire set of events is shifted in Atlanta if the *Journal Constitution* reports the same event Thursday. This pair of tests is difficult to reconcile with alternatives to a pure media effect on financial market participants.

The cross-sectional nature of our empirical design largely eliminates information effects; however, the corresponding improvement in identification is not free. By carving up the trading space into small regions (recall that the unit of observation is defined by firm-date-city triple), we have hampered the ability for any one group to directly influence aggregate market outcomes. To see this, consider that our largest local market is the San Francisco Bay Area, for which we observe the actions of some 4,000 individual investors. But on any given day, most of these are not active, even less so any particular stock. Considering that our universe is the set of large, liquid firms comprising the S&P 500 Index, and it is clear that we cannot link our results directly to prices, liquidity, etc.

On the other hand, it seems equally clear that the effects we identify at the local level should apply generally, i.e., to national media outlets with audiences large enough to meaningfully impact capital allocation. Specifically, beginning with Tetlock (2007) and Tetlock et al. (2008), a number of papers have shown that news stories in national newspapers are associated with substantial price responses. Here, identification usually focuses on what kind of information a story conveys – i.e., about a firm's cash flows, risk or sentiment. Less explored is the possibility that the story's very existence – a media effect – may generate a response independent of these channels. One could imagine, for example, decomposing the "news response coefficients" estimated in such studies into "media effects" and "content effects."

While we stop short of formally attempting such a decomposition, we can say something about their relative sizes. Throughout our tests we have included the earnings surprise (SUE) as a control in the regression, and we find extreme earnings surprises are related to the volume of retail trade. However, the media effect we identify is several times larger than this information effect no matter how we define our earnings surprise. Simply put, in our setting, media is much more likely to drive trade than information.

If this generalizes even partially to the national level, our results are clearly relevant for how we interpret the market's response to news stories. Moreover, our results call for careful research meant to better understand the incentives of financial media and the genesis of media coverage. If media effects are substantial, we should understand how they come about. First, given that firms manage access to their managers, how do the incentives of journalists affect the favorability of their reporting? Butler and Gurun (2009) and Reuter and Zitzewitz (2006) provide evidence that firms and mutual funds can tilt media coverage in their favor by purchasing advertizing space, indicative of a broad class of incentive problems (Dyck and Zingales (2003)). Second, as media have become increasingly electronic, the number of media sources has grown exponentially. How does this increasing competition over finite sources affect the quality of coverage? Third, given that information producers like analysts and institutional investors hold valuable private information about firms, how does the relationship between information producers and journalists affect media coverage? Some work along these lines has begun (e.g. Dyck, Volchkova and Zingales (2008)), but more is needed.

Nevertheless, the point of departure for each of these inquiries is that media coverage matter for investor behavior. The tests herein are the first to link individual investors to their respective media sources and therefore provide conclusive evidence that media effects are real and substantial.

References

Barber, B., & Odean, T., (2000). Trading is hazardous to your wealth: the common stock investment performance of individual investors. *Journal of Finance*, 55 (2), 773–806.

Barber, B., & Odean, T., (2001). Boys will be boys: gender, overconfidence, and common stock investment. *Quarterly Journal of Economics*, 116 (1), 261–292.

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

Butler, A., & Gurun, U. (2009). Don't believe the hype: local media slant, local advertising, and firm value. Working Paper. University of Texas at Dallas.

Campbell, J., Grossman, S., & Wang, J. (1993). Trading volume and serial correlation in stock returns, *Quarterly Journal of Economics*, 108, 905-939.

Chan, W. (2003). Stock price reaction to news and no-news: Drift and reversal after headlines. *Journal of Financial Economics, 70*(2), 223-260.

Coval, J., and T. Moskowitz, 1999, Home Bias at Home: Local Equity Preference in Domestic Portfolios, *Journal of Finance* 54, 2045-2073.

Coval, J., and T. Moskowitz, 2001, The Geography of Investment: Informed Trading and Asset Prices, *Journal of Political Economy*, 109, 811-841.

DellaVigna, S., & Kaplan, E. (2007). The Fox News effect: Media bias and voting. *Quarterly Journal of Economics*, 122(3), 1187–1234.

Dyck, A., & Zingales, L., (2003). The Media and Asset Prices, Working Paper. HBS and University of Chicago.

Dyck, A., Volchkova, N., & Zingales, L., (2008). "The corporate governance role of the media: Evidence from Russia," *Journal of Finance*, 63(3), 1093-1135.

Gentzkow, M., & Shapiro, J. (2004). Media, Education, and anti-Americanism in the Muslim World. *Journal of Economic Perspectives*, 18(3),117–133.

Gerber, A., Karlan, D., & Bergan, D. (2009). Does the media matter? A field experiment measuring the effect of newspapers on voting behavior and political opinions. *American Economic Journal: Applied Economics*, 1(2), 35-52.

Huberman, G., & Regev, T. (2001). Contagious speculation and a cure for cancer: A nonevent that made stock prices soar. *The Journal of Finance, 56*(1), 387-396.

Ivković, Z., & Weisbenner, S. (2005). Local does as local is: information content of the geography of individual investors' common stock investments, *Journal of Finance*, 60, 267-306.

Loughran, T. and P. Schultz, 2004, Weather, Stock Returns, and the impact of localized trading behavior, *Journal of Financial and Quantitative Analysis* 39, 343-364.

Livnat J., & Mendenhall, R. (2006). Comparing the post-earnings announcement drift for surprises calculated from analyst and time series forecasts, *Journal of Accounting Research*, 44, 177-205.

Massa, M., & Simonov, A. (2006). Hedging, Familiarity and Portfolio Choice", *Review of Financial Studies*, 19 (2), 633-685.

Mayew, W., & M. Venkatachalam. (2008). The Power of Voice: Managerial Affective States and Future Firm Performance. Working Paper, Duke University.

Merton, R. (1987). A simple model of capital market equilibrium with incomplete information, *Journal of Finance*, 42, 483-510.

Miller, G. S. (2006). The press as a watchdog for accounting fraud. *Journal of Accounting Research, 44*(5), 1001-1033.

Morris, S., & Shin, H. (2002). Social value of public information. *American Economic Review, 92(4)*, 1521-1534.

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

Odean, T., (1998). Are investors reluctant to realize their losses? *Journal of Finance*, 53 (5), 1775–1798.

Odean, T., (1999). Do investors trade too much? *American Economic Review*, 89 (5), 1279–1298.

Petersen, M. (2009). Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches, Review of Financial Studies. *Review of Financial Studies, 22*, 435–480.

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

Seasholes, M., & Zhu, N. (2005). Is there information in the local portfolio choices of individuals? Working Paper. UC Davis.

Stromberg, D. (2004). Radio's impact on public spending. *Quarterly Journal of Economics, 119(1),* 189–221.

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

Tetlock, P., Saar-Tsechansky, M., & Macskassy, S. (2008). More than words: quantifying language to measure firms' fundamentals, *Journal of Finance*, 63, 1437-1467.

Tetlock, P. (2009). All the News That's Fit to Reprint: Do Investors React to Stale Information? Working Paper. Columbia University.

Veldkamp, L. (2006). Media frenzies in markets for financial information. *American Economic Review*, *96(3)*, 577-601.

Zhu, N. (2002). The local bias of individual investors. Working Paper. UC Davis.

Table 1: Summary Statistics

Account and trading data are taken from the large discount brokerage database with demographic information in Barber and Odean (2000). Accounts in Area is the number of accounts within 100 kilometers of the city's local paper headquarters. Local stocks are stocks of firms within 100 kilometers of the city's local paper headquarters. A paper "covers" an earnings announcement in our sample if a story about that firm appears on day 0, 1 or 2 after the firm's earnings announcement in the ProQuest database. Earnings announcement dates are taken from COMPUSTAT and confirmed with I/B/E/S. Pr(covering a non-local EA) is the frequency in which that paper covered a non-local earnings announcement. Pr(covering a local EA) is the frequency in which that paper covered a local earnings announcement.

	Accounts				Paper	
	Households in Area	Avg # of stocks held	Local Stocks Held / Local Stocks Avail	Non-Local Stocks Held / Non-Local Stocks Avail	Pr(covering a local EA)	Pr(covering a non-local EA)
Atlanta (JOURNAL-CONSTITUTION)	398	2.35	0.135	0.035	0.716	0.019
Boston (GLOBE)	635	2.12	0.137	0.050	0.225	0.015
Denver (POST)	488	2.08	0.109	0.041	0.263	0.015
Detroit (NEWS)	317	2.63	0.130	0.036	0.317	0.006
Houston (CHRONICLE)	607	2.18	0.164	0.047	0.372	0.011
Las Vegas (REVIEW-JOURNAL)	167	2.20	0.087	0.018	0.282	0.010
Los Angeles (TIMES)	1913	2.13	0.218	0.112	0.279	0.047
Minneapolis (STAR-TRIBUNE)	445	2.08	0.218	0.036	0.305	0.012
New Orleans (TIMES-PICAYUNE)	134	2.41	0.086	0.016	0.657	0.006
New York (TIMES)	2808	2.28	0.211	0.152	0.300	0.154
Pittsburgh (POST-GAZETTE)	318	2.28	0.151	0.031	0.365	0.038
Sacramento (BEE)	500	2.20	0.128	0.077	0.011	0.002
San Antonio (EXPRESS NEWS)	142	2.31	0.091	0.016	0.401	0.035
San Diego (UNION-TRIBUNE)	517	2.14	0.161	0.047	0.500	0.029
San Francisco (CHRONICLE)	4076	2.17	0.444	0.151	0.260	0.008
Seattle (POST-INTELLIGENCER)	1019	2.11	0.301	0.060	0.461	0.013
St. Louis (POST-DISPATCH)	241	2.06	0.143	0.021	0.628	0.008
St. Petersburg (TIMES)	243	2.52	0.069	0.028	0.243	0.015
Washington D.C. (POST)	983	2.42	0.188	0.079	0.513	0.029
USA TODAY	--	--	--	--	0.023	0.041
WALL STREET JOURNAL	--	--	--	--	0.438	0.318

	Mean	Standard Deviation	1st Percentile	10th Percentile	Median	90th Percentile	99th Percentile
Market Capitalization (in millions)	5660	10676	62	466	2513	12825	54479
SUE	0.001	0.065	-0.120	-0.014	0.002	0.013	0.113
Stocks Held (per household)	2.263	2.406	1.000	1.000	1.000	4.000	12.00
Stocks Traded (per month per household)	1.75	1.78	1.00	1.00	1.00	3.00	9.00
Number of papers that cover EA	0.879	1.534	0.000	0.000	0.000	3.000	7.00
Local Media Coverage (Dummy)	0.028	0.166	0.000	0.000	0.000	0.000	1.000

Table 2: Media Effects and the Trading of Households

Every firm earnings announcement in our sample corresponds to nineteen distinct observations representing trading in each of nineteen major U.S. cities which we call "trading areas." The dependent variable is the natural logarithm of one plus the dollar trading volume in each trading area for the firm that makes the earnings announcement where local trading volume is available between January 1991 and November 1996 and is taken from the discount brokerage database in Barber and Odean (2000). Trading volume is considered on day 0, 1 or 2 following the earnings announcement date as identified by COMPUSTAT and I/B/E/S. *Local Media Coverage* takes the value one if the local newspaper (i.e. the newspaper that corresponds to the trading area) wrote a story about the firm on day +0, +1 or +2 following the earnings announcement. Controls include the natural logarithm of market capitalization (size), dummy variables for varying quintiles of earnings surprise (SUE quintiles), dummy variables for coverage in each of our local and national newspapers, dummy variables for each of the nineteen trading areas and date fixed effects. Robust standard errors clustered by firm are in parentheses. *, **, and *** represent significance at the 10%, 5% and 1% levels, respectively.

Dependent Variable: Log Dollar Trading Volume					
Local Media Coverage	0.746***	0.658***	0.648***	0.372***	0.477***
	(0.0830)	(0.0739)	(0.0739)	(0.0466)	(0.0593)
Firm Size		0.0700***	0.0761***	0.0516***	0.0577***
		(0.0103)	(0.0108)	(0.00999)	(0.0113)
SUE Quintile = Lowest			0.0758***	0.0418***	0.0526***
			(0.0158)	(0.0154)	(0.0171)
SUE Quintile = 2			0.0128	0.00409	0.00988
			(0.00950)	(0.00853)	(0.00921)
SUE Quintile = 4			0.0308**	0.00856	0.0235**
			(0.0119)	(0.00943)	(0.0107)
SUE Quintile = Highest			0.0845***	0.0529***	0.0661***
			(0.0185)	(0.0177)	(0.0197)
Coverage in Boston Globe				0.0426	0.0305
				(0.0409)	(0.0426)
Coverage in Denver Post				0.0161	-0.00496
				(0.0467)	(0.0501)
Coverage in Detroit News				0.105	0.0814
				(0.0639)	(0.0641)
Coverage in Houston Chronicle				0.0391	0.0257
				(0.0448)	(0.0431)
Coverage in Las Vegas Review-Journal				-0.0622	-0.0671
				(0.0447)	(0.0474)
Coverage in Los Angeles Times				0.0479	0.0618*
				(0.0303)	(0.0356)
Coverage in New York Times				0.0362**	0.0341**
				(0.0146)	(0.0154)
Coverage in Pittsburgh Post-Gazette				0.0870**	0.0891*
				(0.0421)	(0.0457)
Coverage in San Antonio Express News				-0.0193	-0.0298
				(0.0379)	(0.0419)
Coverage in San Diego Union Tribune				0.0431	0.0531
				(0.0567)	(0.0612)
Coverage in San Francisco Chronicle				0.165**	0.261***
				(0.0640)	(0.0759)
Coverage in Seattle Post-Intelligencer				0.103	0.168
				(0.159)	(0.187)
Coverage in St. Louis Post-Dispatch				-0.0426	-0.0794**
				(0.0333)	(0.0381)
Coverage in St. Petersburg Times				0.0447	0.0519
				(0.118)	(0.140)
Coverage in Minneapolis Star-Tribune				0.0652*	0.0580
				(0.0376)	(0.0401)
Coverage in Atlanta Journal-Constitution				-0.0565*	-0.0780**
				(0.0320)	(0.0353)
Coverage in Sacramento Bee				0.0690	-0.152
				(0.256)	(0.259)
Coverage in Washington Post				0.0586**	0.0587**
				(0.0256)	(0.0288)
Coverage in New Orleans Times-Picayune				-0.0855***	-0.104***
				(0.0311)	(0.0326)
Coverage in USA Today				0.165***	0.173***
				(0.0353)	(0.0382)
Coverage in Wall Street Journal				-0.00506	-0.00561
				(0.0129)	(0.0145)
Industry Fixed Effects	NO	NO	YES	YES	YES
City Fixed Effects	NO	NO	NO	NO	YES
Date Fixed Effects	YES	YES	YES	YES	YES
Observations	276982	276982	273999	265928	273999
Adjusted R^2	0.026	0.039	0.041	0.036	0.058

Table 3: Media Effects among Buys and Sells

Every firm earnings announcement in our sample corresponds to nineteen distinct observations representing trading in each of nineteen major U.S. cities which we call "trading areas." The dependent variable is the natural logarithm of one plus the dollar trading volume in each trading area for the firm that makes the earnings announcement where local trading volume is available between January 1991 and November 1996 and is taken from the discount brokerage database in Barber and Odean (2000). Trading volume is considered on day 0, 1 or 2 following the earnings announcement date as identified by COMPUSTAT and I/B/E/S. Local Covered takes the value one if the local newspaper (i.e. the newspaper that corresponds to the trading area) wrote a story about the firm on day 0, 1 or 2 following the earnings announcement. The first five columns consider only buy volume while the second five columns only consider sell volume. Controls include the natural logarithm of market capitalization (size), dummy variables for varying quintiles of earnings surprise (SUE quintiles), dummy variables for coverage in each of our local and national newspapers, dummy variables for each of the nineteen trading areas and date fixed effects. Robust standard errors clustered by firm are in parentheses. *, **, and *** represent significance at the 10%, 5% and 1% levels, respectively.

	Dependent Variable: Log Dollar Trading Volume (Only Buys)					Dependent Variable: Log Dollar Trading Volume (Only Sells)				
Local Media Coverage	0.462***	0.408***	0.402***	0.339***	0.291***	0.381***	0.341***	0.335***	0.307***	0.260***
	(0.0588)	(0.0524)	(0.0521)	(0.0429)	(0.0406)	(0.0464)	(0.0426)	(0.0430)	(0.0392)	(0.0374)
Firm Size		0.0446***	0.0480***	0.0347***	0.0347***		0.0312***	0.0342***	0.0272***	0.0272***
		(0.00684)	(0.00714)	(0.00693)	(0.00693)		(0.00515)	(0.00540)	(0.00580)	(0.00580)
SUE Quintile = Lowest			0.0388***	0.0224**	0.0224**			0.0438***	0.0345***	0.0345***
			(0.0104)	(0.0108)	(0.0108)			(0.00864)	(0.00922)	(0.00922)
SUE Quintile = 2			0.000173	-0.00178	-0.00178			0.0150***	0.0135**	0.0135**
			(0.00716)	(0.00688)	(0.00688)			(0.00564)	(0.00558)	(0.00559)
SUE Quintile = 4			0.00898	0.00475	0.00475			0.0261***	0.0221***	0.0221***
			(0.00885)	(0.00807)	(0.00807)			(0.00690)	(0.00610)	(0.00610)
SUE Quintile = Highest			0.0457***	0.0328***	0.0328***			0.0469***	0.0395***	0.0395***
			(0.0120)	(0.0118)	(0.0118)			(0.0113)	(0.0122)	(0.0122)
Industry Fixed Effects	NO	NO	YES	YES	YES	NO	NO	YES	YES	YES
Paper Fixed Effects	NO	NO	NO	YES	YES	NO	NO	NO	YES	YES
City Fixed Effects	NO	NO	NO	NO	YES	NO	NO	NO	NO	YES
Date Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Observations	276982	276982	273999	273999	273999	276982	276982	273999	273999	273999
Adjusted R2	0.022	0.031	0.032	0.036	0.044	0.015	0.021	0.022	0.024	0.032

Table 4: Media Effects and Alternative Definitions of Household Trading

Every firm earnings announcement in our sample corresponds to nineteen distinct observations representing trading in each of nineteen major U.S. cities which we call "trading areas." The dependent variable is some measure of trading volume in each trading area for the firm that makes the earnings announcement where local trading volume is available between January 1991 and November 1996 and is taken from the discount brokerage database in Barber and Odean (2000). Trading volume is considered on day 0, 1 or 2 following the earnings announcement date as identified by COMPUSTAT and I/B/E/S. Local Covered takes the value one if the local newspaper (i.e. the newspaper that corresponds to the trading area) wrote a story about the firm on day 0, 1 or 2 following the earnings announcement. In the first three columns, the dependent variable is a dummy variable that takes the value one if there was any trading in the firm making the earnings announcement. In the second three columns, the dependent variable is the natural logarithm of one plus the number of accounts that traded in the firm making the earnings announcement. Controls include the natural logarithm of market capitalization (size), dummy variables for varying quintiles of earnings surprise (SUE quintiles), dummy variables for coverage in each of our local and national newspapers, dummy variables for each of the nineteen trading areas and date fixed effects. Robust standard errors clustered by firm are in parentheses. *, **, and *** represent significance at the 10%, 5% and 1% levels, respectively.

	Dependent Variable: Local Trading Dummy			Dependent Variable: Log Number of Accounts Traded		
	ALL TRADES	BUYS ONLY	SELLS ONLY	ALL TRADES	BUYS ONLY	SELLS ONLY
Local Media Coverage	0.0467***	0.0273***	0.0230***	0.0501***	0.0311***	0.0278***
	(0.00674)	(0.00428)	(0.00369)	(0.00592)	(0.00421)	(0.00390)
Firm Size	0.00495***	0.00286***	0.00227***	0.00615***	0.00375***	0.00290***
	(0.00104)	(0.000587)	(0.000529)	(0.00118)	(0.000736)	(0.000611)
SUE Quintile = Lowest	0.00469***	0.00189**	0.00297***	0.00594***	0.00257**	0.00385***
	(0.00156)	(0.000931)	(0.000824)	(0.00183)	(0.00117)	(0.00100)
SUE Quintile = 2	0.00101	-2.93e-06	0.00111**	0.00115	-0.000161	0.00152**
	(0.000824)	(0.000615)	(0.000456)	(0.00101)	(0.000754)	(0.000614)
SUE Quintile = 4	0.00221**	0.000504	0.00187***	0.00235**	0.000377	0.00233***
	(0.000980)	(0.000752)	(0.000503)	(0.00112)	(0.000857)	(0.000646)
SUE Quintile = Highest	0.00566***	0.00257**	0.00333***	0.00697***	0.00347***	0.00419***
	(0.00183)	(0.00104)	(0.00111)	(0.00205)	(0.00124)	(0.00129)
Industry Fixed Effects	YES	YES	YES	YES	YES	YES
Paper Fixed Effects	YES	YES	YES	YES	YES	YES
City Fixed Effects	YES	YES	YES	YES	YES	YES
Date Fixed Effects	YES	YES	YES	YES	YES	YES
Observations	273999	273999	273999	273999	273999	273999
Adjusted R2	0.064	0.048	0.033	0.057	0.043	0.031

Table 5: Media Effects and Determinants of Local Trading

Every firm earnings announcement in our sample corresponds to nineteen distinct observations representing trading in each of nineteen major U.S. cities which we call "trading areas." The dependent variable is the natural logarithm of one plus the dollar trading volume in each trading area for the firm that makes the earnings announcement where local trading volume is taken from the discount brokerage database in Barber and Odean (2000). Trading volume is considered on day 0, 1 or 2 following the earnings announcement date as identified by COMPUSTAT and I/B/E/S. Local Covered takes the value one if the local newspaper (i.e. the newspaper that corresponds to the trading area) wrote a story about the firm on day 0, 1 or 2 following the earnings announcement. Local Firm Dummy takes the value one if the firm is local to the trading market. Fraction of Accounts Holding is the fraction of accounts in the trading area holding the stock as of the first day of the month. Fraction of Accounts Trading is the fraction of accounts in the trading area that traded the stock in the prior month. Controls include the natural logarithm of market capitalization (size), dummy variables for varying quintiles of earnings surprise (SUE quintiles), dummy variables for coverage in each of our local and national newspapers, dummy variables for each of the nineteen trading areas and date fixed effects. Robust standard errors clustered by firm are in parentheses. *, **, and *** represent significance at the 10%, 5% and 1% levels, respectively.

	Dependent Variable: Log Dollar Trading Volume				
	ONLY NON-LOCAL FIRMS				
Local Media Coverage	0.286*** (0.0420)	0.328*** (0.0528)	0.381*** (0.0568)	0.453*** (0.0568)	0.281*** (0.0509)
Firm Size	0.0517*** (0.0100)	0.0576*** (0.0112)	0.0234*** (0.00543)	0.0484*** (0.00920)	0.0229*** (0.00504)
SUE Quintile = Lowest	0.0421*** (0.0154)	0.0521*** (0.0170)	0.0315*** (0.0101)	0.0434*** (0.0141)	0.0283*** (0.00942)
SUE Quintile = 2	0.00461 (0.00852)	0.00969 (0.00920)	0.00266 (0.00798)	0.00711 (0.00863)	0.00184 (0.00782)
SUE Quintile = 4	0.00899 (0.00944)	0.0231** (0.0106)	0.0247*** (0.00886)	0.0214** (0.00963)	0.0226*** (0.00849)
SUE Quintile = Highest	0.0536*** (0.0177)	0.0649*** (0.0196)	0.0493*** (0.0123)	0.0590*** (0.0172)	0.0464*** (0.0118)
Local Firm Dummy		0.492*** (0.0597)			0.334*** (0.0517)
Fraction of Accounts Holding			28.72*** (3.045)		23.47*** (2.755)
Fraction of Accounts Trading				71.23*** (7.245)	51.78*** (5.210)
Industry Fixed Effects	YES	YES	YES	YES	YES
Paper Fixed Effects	YES	YES	YES	YES	YES
City Fixed Effects	YES	YES	YES	YES	YES
Date Fixed Effects	YES	YES	YES	YES	YES
Observations	265928	273999	273999	273999	273999
Adjusted R2	0.040	0.057	0.072	0.064	0.081

Table 6: Media Effects, Newspaper Delivery and Local Weather

Every firm earnings announcement in our sample corresponds to nineteen distinct observations representing trading in each of nineteen major U.S. cities which we call "trading areas." The dependent variable is the natural logarithm of one plus the dollar trading volume in each trading area for the firm that makes the earnings announcement and is taken from the discount brokerage database in Barber and Odean (2000). Local Covered takes the value one if the local newspaper (i.e. the newspaper that corresponds to the trading area) wrote a story about the firm on day 0, 1 or 2 following the earnings announcement date as identified by COMPUSTAT and I/B/E/S. Weather data are taken from the National Climatic Data Center at ftp://ftp.ncdc.noaa.gov/pub/data/gsod. For each trading area, we use weather from the station closest to the local newspaper's zip code. Local Hail is a dummy variable that takes the value one if there was hail during the two days after the earnings announcement. Local Snow is a dummy variable that takes the value one if there was at least 8 inches of new snowfall during the two days after the earnings announcement. Extreme Weather is a dummy variable takes the value of one if either Local Hail or Local Snow take the value of one. Controls include the natural logarithm of market capitalization (size), dummy variables for varying quintiles of earnings surprise (SUE quintiles), dummy variables for coverage in each of our local and national newspapers, dummy variables for each of the nineteen trading areas and date fixed effects. Robust standard errors clustered by firm are in parentheses. *, **, and *** represent significance at the 10%, 5% and 1% levels, respectively.

	Dependent Variable: Log Dollar Trading Volume		
Local Media Coverage	0.478*** (0.0594)	0.478*** (0.0593)	0.479*** (0.0594)
Firm Size	0.0577*** (0.0113)	0.0576*** (0.0113)	0.0576*** (0.0113)
SUE Quintile = Lowest	0.0526*** (0.0171)	0.0526*** (0.0171)	0.0526*** (0.0171)
SUE Quintile = 2	0.00985 (0.00921)	0.00986 (0.00921)	0.00984 (0.00921)
SUE Quintile = 4	0.0235** (0.0107)	0.0235** (0.0107)	0.0235** (0.0107)
SUE Quintile = Highest	0.0660*** (0.0197)	0.0660*** (0.0197)	0.0660*** (0.0197)
Local Hail	0.0485 (0.0351)		
Local Hail * Local Media Coverage	-0.767*** (0.132)		
Local Snow		0.0981 (0.0745)	
Local Snow * Local Media Coverage		-0.871*** (0.126)	
Local Extreme Weather			0.0646* (0.0344)
Local Extreme Weather * Local Media Coverage			-0.809*** (0.105)
Industry Fixed Effects	YES	YES	YES
Paper Fixed Effects	YES	YES	YES
City Fixed Effects	YES	YES	YES
Date Fixed Effects	YES	YES	YES
Observations	273999	273999	273999
Adjusted R2	0.058	0.058	0.058

Table 7: Media Effects and the Timing of Stories

Every firm earnings announcement in our sample corresponds to nineteen distinct observations representing trading in each of nineteen major U.S. cities which we call "trading areas." The dependent variable is the natural logarithm of one plus the dollar trading volume in each trading area for the firm that makes the earnings announcement where local trading volume is available between January 1991 and November 1996 and is taken from the discount brokerage database in Barber and Odean (2000). In column 1 only trading volume on day 2 following the earnings announcement is considered whereas in column 2 only trading volume on day 1 following the earnings announcement is considered. In the third column, the observations in column 1 and 2 are pooled. Local Coverage on Matched Day takes the value 1 if there was a news article in the local newspaper on the same day than the trading day (1 day after or 2 days after). Local Coverage on Non-Matched Day takes the value 1 if there was a news article on a different day than the trading day (1 day after or 2 days after). Controls include the natural logarithm of market capitalization (size), dummy variables for varying quintiles of earnings surprise (SUE quintiles), dummy variables for coverage in each of our local and national newspapers, dummy variables for each of the nineteen trading areas and date fixed effects. Robust standard errors clustered by firm are in parentheses. *, **, and *** represent significance at the 10%, 5% and 1% levels, respectively.

	Dependent Variable: Local Trading Dummy		
	1 DAY AFTER	2 DAYS AFTER	BOTH DAYS
Local Coverage & Local Trading Match	0.280***	0.102**	0.245***
	(0.0490)	(0.0461)	(0.0446)
Local Coverage & Local Trading No Match	0.144***	0.112***	0.110***
	(0.0437)	(0.0205)	(0.0237)
Firm Size	0.0226***	0.00976***	0.0160***
	(0.00385)	(0.00181)	(0.00267)
SUE Quintile = Lowest	0.0198***	0.0115***	0.0161***
	(0.00645)	(0.00408)	(0.00455)
SUE Quintile = 2	0.0113**	0.00124	0.00653*
	(0.00464)	(0.00377)	(0.00345)
SUE Quintile = 4	0.0141***	0.00138	0.00847**
	(0.00511)	(0.00336)	(0.00365)
SUE Quintile = Highest	0.0275***	0.00751*	0.0179***
	(0.00807)	(0.00441)	(0.00555)
Industry Fixed Effects	YES	YES	YES
Paper Fixed Effects	YES	YES	YES
City Fixed Effects	YES	YES	YES
Date Fixed Effects	YES	YES	YES
Observations	273999	273999	547998
Adjusted R2	0.036	0.023	0.026