

e-Information: A Clinical Study of Investor Discussion and Sentiment

by Sanjiv Das, Asís Martínez-Jerez and Peter Tufano

Abstract

In this clinical study, we examine the information flow for four stocks over seven months to trace the relationship between on-line discussion, news activity and stock price movements. On-line discussions are characterized by numerous unsubstantiated rumors, substantial on-point exchanges and quick dissemination of imminent and recently-released information. An interview with a frequent poster reveals the benefits and costs he perceived from his participation in on-line discussions. Using various language processing routines, we create sentiment and disagreement measures based on the comments posted on the message boards. We combine this “e-Information” with other components of the traditional information set to characterize public information flow. We analyze the determinants of sentiment and disagreement, and trace links between news, e-Information and stock returns. This intensive clinical study of on-line discussions suggests mechanisms whereby individual investors and groups analyze and digest company information.

Sanjiv R. Das
srdas@scu.edu
Santa Clara University
500 El Camino Real
Santa Clara, CA 95053

Asís Martínez-Jerez
asis@hbs.edu
Harvard Business School
Soldiers Field Road
Boston, MA 02163

Peter Tufano*
ptufano@hbs.edu
Harvard Business School
Soldiers Field Road
Boston, MA 02163

J.E.L.: G12, G14

Keywords: Internet Message Boards, Information dissemination, Market efficiency, Volatility.

Drafts: September 2001; January 2002; July 12, 2002; October 2003; November 2004

We would like to thank executives at The Motley Fool, Yahoo!, Silicon Investor and Raging Bull for their cooperation; Sarah Woolverton, Sarah Erikson, John Kosmidis and Jose Camacho for assistance with data collection; Mike Chen and John Sheridan for assistance with developing some of the computer code; and David Leinweber, Jacob Sisk, anonymous referees, and seminar participants at the HBS Brown Bag, the University of Texas at Austin, Yale, Ohio State, the Atlanta Finance Forum, the American Finance Association Meetings, and Northwestern University for their comments. We also thank Glenn, who candidly shared his experiences as a frequent poster. This research does not represent the views of any of the firms from which we obtained chat board information nor the views of the firms studied. This work has been funded in part by the Dean Witter Foundation and a Breetwor Fellowship (Das), the Bank of Spain (Martínez-Jerez) and the Division of Research at the Harvard Business School (Tufano and Martínez-Jerez). For additional resources on this topic, see <http://www.people.hbs.edu/ptufano/einfo>.

* Corresponding author.

1. Introduction

In light of the large body of research on informationally-efficient markets, there seems little left to learn from continued empirical examinations of information and markets.¹ It would seem similarly pointless for individual investors to try to compete with professional analysts. However, individual investors persist, and understanding their investing decisions has been one of the most vibrant research streams in recent years. Scholars have studied the alleged overconfidence of male investors, on-line traders and executives by examining the decisions that they take.²

Technology now gives us a new way to study investor reasoning. In the past few years, technology has enabled people to share information, opinions and analyses via stock message boards (or chat rooms) posted in cyberspace. These boards provide a new real-time window into some discussions of individual investors. We believe that to understand investor behavior, it is instructive to peek through this window to observe how information is digested, how sentiment evolves and how perceptions are related to prices.³

The methodology we adopt in this paper is to employ a clinical, i.e., small sample, approach to understanding investor behavior. We believe that there is an important methodological virtue to close inspection of a small sample. Before one can begin to frame hypotheses or construct tests, it is important to establish a base level of understanding in an area. Our paper is therefore decidedly descriptive. This inductive type of research has a long and varied role in the history of thought in economics (Blaug 1992). We are not utilizing it to affirm or falsify theory, as much as to suggest a series of working conjectures (or hypotheses) about individual perceptions and group dynamics that can be developed through subsequent model building and large-scale empirical study.

We have three goals in this paper. First, we closely analyze the people who share their opinions (posters) and their discussions surrounding a few stocks. While we could obtain little publicly available information on the identity of posters, we augmented this using phenomenological inquiry by interviewing an extensive poster to understand why someone would spend substantial amounts of time posting messages to one of the boards

¹ For at least thirty-five years, financial economists have formally considered the notion of informationally-efficient markets, which instantaneously process all available information and impound it into security prices (Samuelson (1965), Fama (1965 and 1975)). Extensive literatures examine the degree to which markets are informationally-efficient, including studies of the performance of filter-rules, information reactions to events, informational anomalies, etc. Accounting scholars have studied how quickly and fully information (e.g. earnings releases) is incorporated into security prices (Bernard and Thomas (1990), Kothari (2000)). Market-microstructure scholars have modeled how new information, the dispersion of beliefs, and the intensity of private beliefs should affect the volumes, volatilities and spreads of stock prices (O'Hara (1995), Hasbrouck (1996)). Hundreds of articles on the informational efficiency of financial markets have been published over the past quarter century in major finance and accounting journals.

² For examples of this literature, see the work of Roll (1986), Odean (1998), Gervais and Odean (2001), and Barber and Odean (2001).

³ In this sense, our work is in the spirit of research that examines how technology affects financial markets; for instance, Garbade and Silber (1978) demonstrated that the opening of a Trans-Atlantic telegraph cable in the 19th century led to immediate adjustments of prices between London and New York.

we study. We also focus attention on the discussions themselves.⁴ While there is a perception that postings are “garbage,” to the contrary, discussions sustain on-point exchanges, generate possibly non-public information, quickly disseminate public information from news stories, and serve as forums where investors digest information. They are also sources of numerous unsubstantiated rumors, adding noise to the information flow. Nevertheless, the fact that even some non-public information may be released on the boards—and the observation that posters use the boards to test their own analyses and obtain those of others—may explain why posters and surfers continue to frequent these chat board sites.

Second, using language-processing algorithms, we measure the intensity and dispersion of sentiment (which we dub e-Information) for over 170,000 messages posted about four stocks. We then analyze the determinants of the level of sentiment and disagreement among posters. Perhaps not surprisingly, we find that there is a close relationship between sentiment levels, historical stock prices and news. In addition, disagreement is related to the intensity of discussion.

Finally, we explore the usefulness of expressed investor sentiment (e-Information) to predict stock returns. Our clinical study confirms other studies that fail to find predictive power forecasting returns (Antweiler and Frank 2004, Das and Chen 2003).

Our work complements the emerging research on word of mouth in the form of discussions by online posters. Scholars are studying how on-line discussions predict product acceptance in various markets. Godes and Mayzlin (2002) examine how metrics of online discussion activity are correlated with the long-run performance of 44 television shows and find that the breadth of discussion is a predictor of short-run ratings success. Chevalier and Mayzlin (2003) find a relationship between on-line consumer-produced recommendations (at sites like Amazon.com and BarnesandNoble.com) and book sales.

Just as on-line discussions may be informative about TV ratings or book sales, they may enlighten our understanding of how investors think about stocks. Various papers—all using empirical samples—study different aspects of this relationship. Wysocki (1999) shows that the number of messages is strongly related to various balance-sheet characteristics of a cross-section of companies. Das and Chen (2003) develop a methodological approach to classifying information posted on electronic forums. Their algorithms are used to analyze the information sets considered in this paper. Bagnoli, Beneish and Watts (1999) look at earnings whispers, i.e. small investor point estimates of market sentiment, and find that this information could be used to advantage in trading strategies. Tumarkin and Whitelaw (2001) examine buy-sell message identifiers on the Raging Bull message board for 73 stocks, and find no evidence that these point-estimate opinions are predictive of equity returns. Antweiler and Frank (2004) use off-the-shelf classifier algorithms to study 45 stocks in the major indexes. They conclude that chat room information does not predict returns, though there may be some predictive content for the volatility of returns and trading volume. Antweiler and Frank (2002) find that stocks with very high posting levels tend to have lower returns and

⁴ This interpretive and post-positivist approach to research has been useful in many other settings, especially in the area of Marketing.

higher risk than other stocks after controlling for various asset pricing factors, and they conclude that posting volume may be a priced factor itself.

Our paper is similar to these in that we also look at the relationship between online discussions and returns, but we have a slightly different agenda, as we also are trying to understand the evolution of sentiment and disagreement themselves. In addition, our clinical study—which involves a detailed reading of messages—can help to characterize the nature and content of the discussions. To the best of our knowledge, our study is also the first to analyze the content of news postings for market sentiment using automated tools.

The rest of the paper proceeds as follows. Section 2 deals with our clinical design. Section 3 deals with demographics and describes the identity of posters, detailing one interview with an especially active investor-discussant. Section 4 reports a clinical examination of the nature of the discussions and the quality of information in those discussions. Section 5 describes our computer-generated measures of sentiment and disagreement (e-Information) that are extracted using language-processing algorithms. Section 6 analyzes the determinants of our e-Information measures. Section 7 examines the relationship of e-Information to the price formation process. Finally, in Section 8 we summarize the hypotheses that emerge from this clinical investigation. The Appendix provides a short review of the technical approach and algorithms used to extract e-Information from Internet sources.

2. Sample Design

Many existing studies of information and security markets study a single information source across a large number of companies, e.g., reactions to press announcements of a particular type. In contrast, we examine a wide range of information sources about a small number of companies. In particular, we study four firms over a period of seven months, using these four firms as archetypes for different information environments where traditional information and e-Information flows vary. As befits a clinical study, we attempt to dig deeply into these four firms, using our observations to drive hypotheses for large-scale studies. We have consciously not selected pathological examples where stock message boards were used to explicitly manipulate prices (Leinweber and Madhavan (2001)).

To select the four firms for our preliminary study, we first collected information on the 3724 firms that had at least one posting on The Motley Fool (TMF) stock message board in the period July 1, 1998 through January 31, 1999.⁵ Then we classified each by the number of TMF messages. For the 504 firms in the TMF list that had at least 25 posts in the period July 1, 1998 through January 31, 1999, we collected the number of major news stories from Factiva. We defined a “major news story” as one in which the name of the company was either in the headline or was mentioned in the lead paragraph and appeared at least three times in the body of the article. We stratified the 504 firms into quintiles along two dimensions (number of posts and news stories) and selected one firm

⁵ The Motley Fool graciously provided us with the data to perform this screening, but subsequent posting information for this and other boards was collected through a proprietary web crawler program.

from each of the four extreme categories of the joint distribution. See **Figure I** for a plot of the joint distribution and for the positioning of our four firms. The four firms are shown below:

		Traditional Information Environment (number of news stories)	
		Rich (High)	Poor (Low)
E-Information Environment (number of posts)	Rich (High)	Amazon.com	General Magic
	Poor (Low)	Delta Air Lines	Geoworks

These four stocks were not selected to be representative of the average stock, but rather to help us to understand the *extremes* of information flow. **Table I** provides summary statistics on the four firms. Delta Airlines was an old-economy company with a large work force, substantial institutional ownership, and positive earnings. Amazon was considered a flagship new-economy company. General Magic and Geoworks were small, not very profitable firms attempting to serve the new economy, but General Magic—founded by former Apple Computer executives—had an extremely high level of posting activity for a firm its size. Each of these firms was quite unique, characterized by quite different information flows, which is the dimension along which we sought to stratify this sample. **Table I** shows that these firms are different along other dimensions as well, particularly in levels of trading activity. The two firms with substantial discussion were ones with active trading (Amazon and General Magic with 9.4% and 3.5% daily turnover) and the firms with less substantial discussion evidenced less active daily trading (Delta’s trading volume in shares and value exceeds General Magic but its share turnover is a quarter of General Magic’s).⁶

The time frame we study (seven months in the late 1990s) is admittedly short. We sought to look at a period that was long enough so that firms could experience different types of news, but not so long as to prevent meaningful data collection and manual analysis of some of the data. *Ex post*, this was a time of high valuations and trading activity. It was also a period that could justifiably be characterized as one where there was a great deal of uncertainty about the future, and where sentiment and perception drove values. Thus, it is a particularly appropriate time to study investor sentiment. If we compare the level of postings during this period to those before and after it, as shown in **Figure II**, we see that perhaps no period of time is “representative” in this new and evolving phenomenon. While the period we study was a period of relatively high postings for Amazon, it is approximately equal to the posting levels we observe today. In contrast, the period was part of a general upward trend in posting activity for Delta Airlines.

⁶ The NYSE fact book shows that the annual turnover for 1998-99 was 77% across all NYSE stocks. Therefore, assuming 250 trading days in a year, the daily turnover is 0.31%. Given that there are many stocks on the NYSE that do not trade at all, the turnover for more active stocks such as those in our study is likely to be somewhat higher.

3. The Nature of the Posters

Posting a message is a quasi-anonymous act. Posters select a screen-name (or multiple screen names) and select how much information to reveal. At the outset, it is important to remember that posters may—or may not—be “representative” of the average or marginal investor in stocks. For Delta, where institutional investors hold 75% of the shares, posters—who by all accounts are individual investors—surely do *not* represent the average investor. For the smaller stocks (General Magic and Geoworks) where individuals hold 85 to 90% of all shares, they are likely to be more representative. As we shall see later, the posters on the boards we study appear representative at least along one dimension—their sentiment about the stocks seems to contemporaneously track the stock price.

In this section, we provide some information on the characteristics of the people who generated the messages we study. We report (a) the distribution of the number of unique posters and the posting activity by posters for our four stocks, and (b) our interview with one of the most prolific posters in our sample.

Number of posters and posting activity per screen name: **Table II** reports data on the posters for our four boards for the stocks we study. For example, over our study period, 12,169 unique screen-names posted 102,820 messages on the four Amazon boards. While some individuals might have posted under multiple names, we suspect that most of these represent unique individuals. For the other three stocks, the number of unique posters ranged from 404 (for Delta) to 3,208 (for General Magic.) These numbers can be compared with the number of holders of record for each of these stocks in our sample period. If all of the posters were investors, they would represent 2% of the registered holders of Delta but 528% of the shareholders of Amazon!⁷ These figures are not reliable however because (a) we know that some of the posters were short sellers or those thinking about buying the stocks, so the numerator is overestimated; and (b) the denominator is underestimated as shares held in omnibus accounts (like 401k plans or under some brokerage relations) show up as one shareholder of record, even though a large number of individuals would consider themselves investors.

At first glance, the boards do not appear concentrated, when we calculate share of messages by poster. In **Table II, Panel A**, we calculated Herfindahl indices (see Table II for details) for the sixteen stock boards (four stocks x four board vendors). In only four cases was the share of message concentration about the level that the Department of Justice would consider “mildly concentrated” (i.e., 1000), were this to be product market data, and in all four cases the number of postings is small. The average poster posts 3.4 to 20.2 messages per stock over the window we study, and the median poster on all of the boards posts 1 to 3 messages. These are relatively large and active discussion groups.

However, the Herfindahl index analysis hides important concentration, which is captured by the large difference between the mean and median, which reflects a skewed distribution of posting activity. **Table II, Panel B** reports the number of postings by

⁷ As of fiscal year end 1998 the number of registered shareholders for the four sample firms were: AMZN (2,304); DAL (21,672); GMGC (725); GWRX (7,800). Source: Compustat.

screen name. **Figure III, Panel A** is a histogram of this data. 2,552 screen-names posted only one message in the seven-month period we studied. Another 5,193 submitted two to five posts over this seven-month period. In sharp contrast, the most prolific poster on the Amazon boards shared 4,907 messages over seven months on the Silicon Investor site, or over 25 messages a day. In aggregate, the top five posters in our four stocks accounted for 9 to 18% of all postings. The picture that emerges is of a relatively small and vigorous core of frequent posters, surrounded by a substantially large number of occasional posters and by an even larger group of “lurkers,” people who read the postings but do not contribute messages (Nielsen/NetRatings reported that the number of unique visitors per month to Yahoo Finance was 10.6 million in March 2004, up from 8.8 million in March 2003; biz.yahoo.com estimated that in 2002, the average daily number of visitors to Yahoo! Finance was 20.3 million).

Business analysts report that the typical visitor of stock chat boards was a middle-aged man. Women represented only between 20-40% of the visitors to the four boards analyzed in this paper, and more than half of the visitors were between 25-40 years old. However, the population of visitors was not completely homogeneous across boards. Yahoo! Finance also shows an above-average percentage of users with income over \$100,000 per year. Roughly 30 percent of its users fit that description, a figure that approaches that of brokerages, and is almost 10 points higher than AOL, MSN, or the digital media universe average. Motley Fool had a higher proportion of visitors over 55 years old and Yahoo! had a larger proportion of women in its visitors.⁸ The unavailability of comparable data across message boards makes it difficult to exploit these differences in a full-blown empirical analysis. However, we have provided the available details of the demographics of posters to the various boards in **Figure III, panel B**.

While we might understand why someone might post a few messages and then lose interest, it is less clear what motivates someone to post over 5000 messages about a single stock in a bit more than half a year. The time and effort expended by this person was considerable. Why?

Profile of an Active Poster: In a large empirical study, it is impossible and unwise to spend too much time on an individual data point; rather it is normal to discard anomalous observations. In contrast, in a clinical study, one is able—and indeed encouraged—to understand outliers. It is in this spirit that we interviewed Glenn R., the most prolific poster on the Amazon boards. We emailed Glenn and asked if he would be willing to talk to us about his posting activity, and we set up a time for a telephone interview.

As we have presented earlier drafts of this paper at seminars, participants have made a number of smug remarks about posters. Our discussion with Glenn failed to uncover any of the “lunacy” that some academics had predicted. To the contrary, he gave an interesting interpretation of membership in a posting group. We report our interview findings to give readers a first-hand look at an active poster.

Glenn was in his late-40s when he was posting the messages we studied. He has an undergraduate degree in engineering from a large Midwestern university, and

⁸ Nielsen Net-ratings 2002. The Yahoo! figures should be taken with some reserve, because Nielsen does not report Yahoo! Finance independent from the main Yahoo! website.

completed some of the requirements for a business degree. He owned a small chain of jewelry stores, including an on-line jewelry store and was self-employed. He did most of his postings on nights and weekends, when he was not otherwise busy at work. He estimated that he spent approximately 30 hours a week interacting on the boards. He was a client of a large brokerage firm and read the professional analyst reports he received. A few years later, when we interviewed him, he was still able to cite analysts by name. He was interested in stocks and in technology, so he gravitated toward tech stocks. He also actively searched the web for news stories about these stocks. He provided four explanations for his activity.

1. “I wanted to learn.” A number of times in our conversation, Glenn emphasized that he lived in a small town of 15,000 people and there were no investment clubs in his town. He reports, there were “not many people in town that he could talk to about investing.” His activity in the Silicon Investor board was equivalent to membership in an investing club. Glenn was interested in technology and technology stocks, and was keen on learning from “people who had more experience than (he) had.” In particular, he felt the boards were quite good in providing information on market microstructure details and technical analysis, especially the nuances of shorting stocks and the daily fluctuations in the outstanding float of the stock. Glenn approached stocks from the perspective of fundamental analysis, but was intrigued by the approaches of technical analysts that seemed to give them “a better batting average.”
2. He felt that the professional analysts missed many of the details about firms, and he used the discussion boards to test out his analyses. Glenn did not believe that he, nor any of the active members of the Amazon board, had any proprietary or inside information. “I don’t think there was any truly inside information...the whole group had no better idea than the next person.” However, they did have the time, experience and inclination to analyze the fundamental data on Amazon in a careful way. As he explains, “I was perceiving this firm as a retailer and I was in the retail business. There was no question that the cost of fulfillment was higher than in regular stores. Others didn’t understand issues of costs.” While much of this information was in public disclosures, it was buried in footnotes and was labor intensive to pull out. This information was “missed by a lot of the analysts.”
3. The boards provided Glenn with colleagues that he enjoyed. While we calculated that there were 925 posters on Silicon Investor during our study period, Glenn estimated that there were a much smaller number (50 or 60) that were relatively active. Of these he came to know five or six personally, through phone calls or in-person meetings. Unlike the “cheerleaders,” these people helped each other “see through” the news stories. They discussed stock picks and non-investing business advice off line.
4. The boards provided Glenn with a venue to engage in enjoyable debate and to earn the respect of others. Academics enjoy discourse, but sometimes forget that others may value the exchange of ideas as much as the profits from trading. Glenn called this the “entertainment value” of the boards; the ability to interact with others and enter into a sustained discussion with them. Moreover, the

discussion was self-reinforcing. “I enjoyed putting forth an opinion and then having to justify it.” According to Glenn, people who earned positive reputations were those who were able to more accurately predict the short-run stock price or next earnings numbers, and those who provided superior insights. Glenn was proud to develop a reputation for the latter. “People wanted to know what I thought...it was a feeling of accomplishment.”

In retrospect, Glenn felt that he *lost* more money as a result of participating on the boards than if he had not. He explained that by virtue of having to stake out and argue a position in public, he felt that he probably became more “stubborn” about his opinions, and therefore he held onto his positions longer than he might otherwise have. He reports that while he lost money on his Amazon position, he profited on a few other positions that he followed regularly.

This interview gives a new dimension on discussion boards and the investing process. It reminds us that investing is not necessarily a solitary activity, but can be a *communal* activity (Das and Sisk, 2003). People voluntarily join communities because they perceive some benefits. For Glenn, the benefits included the enjoyment of coming to know other like-minded people, the ability to share ideas, and the ability to develop a reputation for clear thinking.⁹ An on-line community focused on a particular stock is no less valuable to its participants than one that revolves around a television show or game.

To the extent that the online community serves as a social group or debating society, its economic impact is probably secondary. However, to the extent it serves as a vehicle for testing ideas and analyses, it frames a few interesting questions that could form the basis for subsequent research. One could ask, what are the relative returns from communal vs. individual analysis? Glenn believed that his analysis would be improved by testing it with others, but this it is an untested assertion. More narrowly, can discussion, even among well meaning investors, have the impact of producing even more severe biases, like the hardening of Glenn’s investment bias? Ex post, Glenn reached this self-critical conclusion, but it is possibly a more general phenomenon.

4. The Substance of the Discussions

In the spirit of exploiting the clinical nature of our research design, we turned our attention to the content of the postings. In particular, we sought to understand what subjects were discussed, whether discussions stayed on point, and whether the

⁹ These observations provide motivations for voluntary postings on stock chat message boards. The theoretical justification for posting might be related to models of information disclosure. Suppose each investor receives a noisy signal about future stock price (e.g., their opinion as to the importance of a new product announcement). By sharing their signals with others, they can verify the information *before* trading, or can share the signal with others *after* trading with the hopes that their interpretation will lead to the desired movement in share prices. On a more mundane level, stock chat boards can be locations where disgruntled shareholders, customers, employees and former employees can share their experiences with others. Modeling the many reasons why people may care to share their opinions with others, and why readers might respond to these opinions is a complex question that is beyond the scope of this paper, but we feel these comments offer some insights to modelers. Admati and Pfleiderer (2001) and Bacharach and Board (2002) have begun to address these questions. Mayzlin (2003) provides a theoretical model of how online conversations can be used by retailers to induce certain buying patterns and the equilibria that result from this activity.

discussions revealed meaningful information. To do this, we conduct two experiments. In the first, we begin our analysis with *actual news* released by firms, and look before and after the news release to understand any foreshadowing of the news prior to the release and the subsequent digestion of the news after the release. In the second, we examine a set of *rumors* on the boards and track them through time. From the first analysis we find that the discussion boards seem to play an important role in rapidly disseminating news, sometimes “breaking stories” before they are covered widely. From the second analysis, we conclude that discussion boards are rumor mills for many unsubstantiated claims and a poor source of inside information.

In the first experiment, we selected 16 seemingly newsworthy¹⁰ press releases by the four companies, and for each traced how the “news” was communicated to investors through traditional media as well as through postings. (**Table III** lists the events). For each, we sought to understand how the press or message boards provided advance information of the event, and how they responded to the event. With respect to response, we measured the speed from the press release to the first discussion on the message boards, the time series of subsequent discussion, and the nature of the discussion. The evidence in **Table III** is anecdotal, but nonetheless helps us to understand the nature of the manner in which information is shared. We found the following patterns, which can be thought of as empirically driven “hypotheses” about the different functions of the boards:

1. Message boards provided factual foreshadowing of subsequent press releases. In quite a few instances, posters provided readers with advance warning of subsequent news events. For example, one of our tracer events is the spin-off of the DataRover division by General Magic. Nine days before the DataRover spin-off, someone reported on the Yahoo! message board that they had found a new DataRover website that did *not* mention General Magic. This site was apparently taken off-line in a few hours by General Magic, which was probably testing the URL for the imminent spin-off. Readers of the board could not only read the message, but also confirm it by going to the site. In another example, one day before General Magic announced an agreement with Microsoft, someone posted that the two looked like they would share a booth at the Consumer Electronics Show, a tip-off to some closer relationship. A third General Magic poster alerted readers to a local radio broadcast, which suggested that the firm would enter into an agreement with Intuit, which was not publicly announced until a few hours later. In other instances, posters speculated about upcoming splits and bond issues. In two others, we see advance discussion of upcoming earnings numbers, or so-called “whisper numbers” as studied by Bagnoli et al. (1999). Put together, this anecdotal evidence suggests that posters provide active surveillance, especially of smaller companies, well before the traditional press picks up news events. In these instances, the posters seemed to “stumble across” non-public information, rather than being privy to inside information. Having stumbled across it, they then shared it, possibly to get others to validate it or to make sense of it.

¹⁰ The economic materiality of these announcements is not substantial ex post. We performed event studies on the 16 announcements and found that even though ten of the events have abnormal 1-day or 3-day returns over 5%, these returns are statistically significant in only four cases.

Felton and Kim (2002) examine the content of the message boards for Enron Corporation from 1997 to 2001 and find a similar pattern. In particular, they report a series of posts from alleged insiders that foretold of the subsequent events at Enron.

2. *Rapid postings disseminate company information quickly.* Companies tend to issue press releases either before markets open (sometimes in the middle of the night) or after markets close. **Table III** shows the number of minutes between the time-stamps on each press release and the first posting of the news on one of the stock chat boards. For many, but not all announcements, the first post is incredibly soon after the news is posted, and in a few cases, prior to the time stamp of the first major news wire story. First posts tend to contain a short notice of the news, often with a URL directing readers to the press release. The boards are apparently serving to disseminate information to interested investors quite quickly. This observation is consistent with those of the poster, Glenn, who mentioned that it was his routine to scan the press about Amazon and post links to new and important stories.

3. *Extensive on-point discussion is sustained for eight hours after news releases.* **Figure IV** characterizes the postings that followed these 16 events. Panel A shows that in the few hours immediately after a news event, posting volume rises, but then tails off over time. Panel B displays the composition of the posts over the first eight post-news hours. The nature of the discussion was measured by categorizing each subsequent post into one or more of five possible categories (asks question, offers alleged fact, shares opinion, comments unrelated to news event and spam/garbage). The first three categories of posts can be thought of as on-point postings, i.e., ones that relate to the news at hand. We see that for the first four hours after a news event, over two-thirds of all posts are on-point, and even eight hours later, about half are still discussing the news (as opposed to other issues or spam).

4. *The on-line discussion is a mix of questions, answers and opinions.* We categorize on-point posts as asking a question, offering an alleged fact, or proposing an opinion about the meaning of the news. For the first hour, we see more of a question-and-answer pattern, with a quarter of all posts and a third of the on-point posts either asking a question or supplying a fact. Over time, the discussion tends toward more analysis: What does this news event mean for the company and its stock price? Much of this later discussion does not focus on communicating facts, but rather *interpreting* them, consistent with our understanding of the primary function of the board. This observation is consistent with our discussions with Glenn, the prolific Amazon poster.

5. *The boards are a source of unsubstantiated rumors.* The research design in this first experiment is biased in that it is conditioned on validated news, i.e., all of the stories we studied were real. In the second experiment, we began our analysis with rumors on the boards and tracked them subsequently. In particular, we searched our sample postings for messages related to mergers and acquisitions, which are material corporate events. Our goal was to identify events with enough potential materiality that they might be ultimately reported. We searched for news related to the posting rumors in Factiva

for the period January 1998 to August 1999 (i.e. six months before and after our sample period).¹¹

Table IV reports the results of our analysis. In total, we identified 54 merger and acquisition rumors on the discussion boards, with seven of these meriting at least five posts. However, the rumors are almost entirely ungrounded in fact. Just one of the 54 rumors preceded an announcement by the company of an actual merger. Nine others preceded similar rumors in the business press, but not an actual transaction. One preceded a press denial by the company.¹² The vast majority (43 of 54) did not result in a transaction or even in a rumor in the press.¹³

Thus, this second experiment produces a different picture than the first. The boards are quite good at quickly disseminating information and providing investors with a forum to digest it. They also are good at sharing not-quite-yet-public information. However, we see no evidence that the boards are privy to truly material inside information. This is consistent with the observation of our prolific poster, and also consistent with the empirical results, which we confirm later in this article, that postings cannot predict returns.

There is a less obvious conclusion from this analysis. While the boards are places where many rumors are suggested, we did not see evidence that they were “rumor mills,” where these rumors themselves were the source of sustained discussion. In 47 of 54 instances, the rumors generated fewer than four subsequent posts each, and unsubstantiated rumors generated less discussion.

If one were to construct a “wheat and chaff” measure for the boards, they would probably perform poorly. On the positive side, of the sixteen actual news events in **Table III**, half were foreshadowed on the board discussions. A .500 batting average sounds quite good, but this is conditional on knowing that something actually happened. In contrast, an avid reader of the boards continually scanning for merger announcements would have enjoyed a .019 batting average, with the remainder of the stories being unsubstantiated rumors.¹⁴

Both our interview with a “hyperposter” and this analysis of rumors suggests that boards can speed up the process of information dissemination and possibly even analysis, but do not reliably produce new information.

5. The Concept and Measurement of e-Information

While our subject is technology-enabled discussions, the research tools we used up to this point have been decidedly low-tech, involving reading messages. In the

¹¹ The keywords used for the search were: merge, merger, hostile, acquisition, acquire, takeover, target, tender, offer, and stock swap.

¹² In two instances, the press story announcing the rumor cited the Internet as the source of the rumors.

¹³ The one “missing” observation is a rumor on the boards that followed, rather than preceded, some press rumors.

¹⁴ While it is impossible to compare these without knowing the gains and losses of trading on this information, it seems that the gains from being right 2% of the time would be more than offset by being wrong 98% of the time.

remainder of the study, we use technology to classify the 170,953 messages and to relate our measures of sentiment and disagreement to information sources as well as to stock prices.

5.1 *Definition and Motivation of e-Information*

The simplest characterizations of the flow of information are activity measures: simple counts of the numbers of news stories or postings, or the length of news story or posting. These metrics are used by Mitchell and Mulherin (1994; number of news stories) and Wysocki (1999; number of postings). These activity measures indicate the level of interest, excitement, puzzlement or “buzz” about the information set, similar to the measure of the decibels of noise in trading pits used by Coval and Shumway (2001). Activity measures are based on the notion that discussion (whether in person, by electronic posting, or news stories) is correlated with the salience and newness of information releases.

Wysocki’s (1999) study of posting activity on the web provided the first activity measure of chat room activity. In particular, he studied the *number* of postings on 3,478 of the 8,011 firms on the Yahoo! Message boards. He found that certain types of firms are more likely to generate high posting activity. These firms had high short-seller activity, high market values relative to fundamentals, low institutional holdings, high trading volume, extreme performance, or extensive analyst following. In addition, he studied a panel of 50 firms to understand the time series properties of the number of postings. He found that posting volume varied by day of week, increased near earnings announcements, and was related to the volume of trading activity and abnormal stock returns.

Capturing the *content* of the information is a more complicated matter. In this paper, we extract a subjective measure of the meaning of the information using five computer algorithms that read and categorize the content of each individual message. The algorithms parse from each message a single measure: the degree to which the message conveys a buy, sell or neutral sentiment about a stock. By aggregating these messages over some time period, we can gauge the average sentiment as well as the distribution of “posting sentiment” manifested by the stock message board information flow. We also use this technology to classify the “news sentiment” of news stories. We call the combination of activity measures and content measures (distribution of sentiment indices) “e-Information.”¹⁵

We electronically “read” the messages to uncover sentiment because this intangible quality is critical to many of models used in financial economics. In behavioral finance, investor sentiment (or noise trader sentiment) is used to explain deviations in prices from “rational” levels (see DeLong, Shleifer, Summers and Waldman (1990)). To measure sentiment, academics have used the closed-end fund discount (Lee, Shleifer and Thaler (1991)), flows into mutual funds (Goetzman, Massa and Rouwenhorst (2000)),

¹⁵ Tumarkin and Whitelaw (2001) study a subset of postings from one message board, which permits posters to voluntarily classify their short-term opinion about each stock. While these voluntary disclosures are convenient for study, only less than a quarter of posters choose to reveal a “short term opinion,” and the board that permits this disclosure accounted for less than 5% of the total postings in our sample.

and subjective determinations based on reading various news stories (Gay et al. (1994)). We create new measures of sentiment by examining posting or news stories.

We use a variety of procedures to provide the classification scheme, and benchmark the results against results of human interpretation. (**Appendix A** contains a description of the methodology we use to classify messages, summarized from Das and Chen, 2003). We also use this methodology to extract sentiment measures from news stories as well.

We are hopeful that these new measures of sentiment will be useful for a variety of reasons:

(a) Posting sentiment reflects the widely available opinions of a set of retail investors. The information is freely available on the Internet and the web sites that offer this information are widely visited. Das and Chen (2003) find that the contemporaneous link between message board sentiment and stock returns is a strong one. This complements the results of Antweiler and Frank (2002) that suggest that this relationship holds after many controls, suggesting a price factor interpretation for postings.

(b) News stories have always had the power to affect many people, but technology makes more news available to more potential investors more quickly. We have the capacity to “sign” the news to determine its likely impact on investor sentiment. Recent studies, such as Busse and Greene (2001) and Fleming and Remolona (1999), show rapid responses to alleged news events in the equity and bond markets.

(c) The disaggregated observation of sentiment allows us to calculate various distribution measures. We can not only calculate an average measure (net bulls less bears), but we can also calculate measures of the degree to which investors disagree, which is a variable that theoretical research suggests will correlate with both trading volumes and volatility (Kim and Verrecchia (1991)).

(d) We believe the detailed nature of the data allows for new and more detailed understanding of sentiment. The stock-by-stock observation of sentiment is more granular than broad market-wide sentiment indices. The high frequency of the data allows us to plot changes in sentiment over time.

5.2 *The Calculation of e-Information Measures*

The four firms we studied together received over 170,000 posts over seven months. However, our primary interest is in trying to extract some meaning from the messages, in particular the “bullishness and bearishness” of the posts and the extent to which posters seem to agree or disagree.

The **Appendix** describes the methodology we use to classify messages (and news stories). Briefly, we adopt a voting system.¹⁶ The raw message is read by the program, and the five language processing routines use different rules to classify whether the message is a buy, a sell or neutral. We then count the number of “votes” across the five

¹⁶ Antweiler and Frank (2004) use two classifiers similar to the five we use. To improve the overall signal to noise ratio, we adopt a voting rule to determine content. They analyze the two separately, rather than using a voting rule.

different measures. Messages that receive at least 3 bullish (bearish or neutral) votes are categorized in that category; otherwise they are not categorized (nc). This approach enhances the signal to noise ratio of the classification procedure.¹⁷

While the unit of observation for classification is the message, we also create *daily* measures of the e-Information (as well as measures for other time intervals). Our primary measure is a *sentiment index*, defined as the number of buy messages less the number of sell messages (excluding null and not classified messages). The sentiment index picks up the net bullish sentiment and is an “absolute” or unscaled measure of sentiment. In addition, we calculate a number of other related measures:

- *Sentiment sign*: 1 if sentiment > 0; -1 if sentiment < 0, 0 otherwise.
- *Sentiment percentage*: Sentiment index divided by all messages for the day. This is essentially a scaled measure of sentiment. It assesses the extent of the day’s discussion that comprises bullish or bearish opinion. We include null messages as well as ambiguous messages in the total count in the denominator of this measure because they are both symptomatic of the absence of strong sentiment.
- *Opinion index*: Fraction of all messages that are classified as either buy or sell. One interpretation of the complement of the opinion index (i.e., 1-index) is that it represents the extent of questions, non-directional comments and noise in the discussion.
- *Disagreement index*: This index is defined as

$$\left| \frac{|BUYS - SELLS|}{BUYS + SELLS} - 1 \right|$$

(or n.a. if Buys+Sells = 0).

This measure is intended to capture whether the opinionated posters have the same view, or whether there is dispersion of belief. If everyone is on the buy or sell side of the market, this index is 0%, but the index can rise to 100% (or 1) if the opinions are split equally into buys and sells.¹⁸

Figure V graphs e-Information variables for the four firms over our sample period and **Table V** reports the overall categorization of messages for the four stocks over the entire sample period. For the four firms, we can classify about 40-50% of the messages as either buys or sells. About 6-8% of the total posts are net buys. This low level of positive sentiment reflects the fact that there tend to be large numbers of both buy and sell messages, shown by the disagreement index of 80% or more (100% would mean that the buys and sells are equally split.) This high disagreement level is not surprising, given

¹⁷ We have other information that is not currently incorporated in our analysis, including the length of the post and the stated identity of the poster. As posters use self-chosen nicknames that can be changed at will, we cannot be assured that two posts by different named posters represent two different individuals.

¹⁸ To the extent that total messages are highly correlated with signed messages, the disagreement index is like an unsigned version of the sentiment percentage.

that discussion takes place when there are differences of opinion, and given the fact that longs and shorts can both participate on the boards.¹⁹

5.3 *Confirming our Measures of Sentiment and Dispersion*

Before we analyze the determinants of sentiment and dispersion, or their relationship to the price formation process, it is useful to confirm that these measures seem reliable. We describe some tests used to validate our procedure. Furthermore, in the spirit of transparency, in the Appendix, we provide readers with a sample of messages and the machine coding so they can get a sense of the output of the classification. Detailed description of the technology is carried out in the methodological paper of Das and Chen (2003) and we refer to some of those results here. In addition, we undertook some additional validations and these are also provided below.

5.3.1 *Accuracy of Machine Classification Voting Scheme*

First, we compared the accuracy of the machine classification with a human-coded training sample of 64 messages. Our classifier employs a voting scheme, and requires agreement by a simple majority of votes to classify a message (i.e. 3 out of 5 algorithms need to agree on the classification). However, varying the minimum number of votes needed for agreement changes the percentage of “attempted classification” (i.e. number of messages on which an opinion was generated), and also the “accuracy” of classification (i.e. the extent of agreement with humans on the messages that attained a voting majority). These two measures are inversely related to each other. These measures are reported in the following table (data from Das and Chen, 2003):

Number of votes required for a classifier opinion	Percent of messages attempted for classification	Percentage of attempted messages correctly classified
2	100%	53.13%
3	81.25	61.54
4	31.25	85.00
5	7.81	100.00

1. Our tests showed that this approach provided an accuracy of attempted classification of around 62%, i.e. the methodology was correct about 2 out of 3 times, compared with random choice accuracy levels of 33% given the three way selection of bullish, bearish or neutral
2. The voting scheme results in discarding messages that do not attain a simple majority amongst classifiers. While this is a sensible way to remove noise and retain the signals, it may result in throwing away too many messages in case the

¹⁹ We separately conducted numerical simulations which show that if buy and sell messages arrived independently and uniformly, the average value of disagreement is about 0.65. Therefore, the level of disagreement of 0.8 appears high, but in comparison to the average of 0.65, it is reflective of a positive correlation of arrivals of buy and sell messages.

level of disagreement among classifiers is too high. However, we found this not to be the case. A simple majority-voting scheme results in a loss of approximately 20% of the messages.

3. Of the messages classified incorrectly, we are more concerned with messages that were actually buys and were classified as sells, and vice versa. There is less concern that neutral messages were classified as buys or sells and vice versa. We found that the level of error in classification where buys and sells were mistaken for each other (the beta error rate) was lower than the average error, and is about 20%, compared to about 33% on average. This is reassuring, since it means that the sentiment index is less likely to be erroneous in an extreme way.

Second, in Das and Chen (2003), a much more extensive test was carried out comparing the machine-coded messages with sentiment measures that were self-described by posters. In July 2001, Yahoo! began to allow posters to identify the sentiment of their posts, using a five-point scale, ranging chosen from a set {Strong Sell, Sell, Hold, Buy, Strong Buy} - we collapsed this scale to {Sell, Null, Buy}. While this information was not available during the time of our study, it provides a truly out-of-sample test for our algorithm. The results showed that the algorithm was able to classify messages consistently in agreement with the revealed sentiment of the posters. The chi-square statistic in Das and Chen (2003) based on about 66,000 messages, is highly significant at the 0.1% level, strongly rejecting the fact that there is random assignment in the classifier as opposed to accurate classification. These results also reconfirmed the beta error rate (classifying buys as sells and vice-versa) to be less than 20% as stated above.

5.3.2 Accuracy of Human Classification

The high level of textual ambiguity of the messages implies that even human agreement on the messages may be reasonably less than 100%. Hence, we ran some tests to see what the level of human-human agreement was. On average, we found this to be 72%, not substantially higher than that of machine-human agreement (62%). The following table presents the results of an analysis of 374 messages by two human subjects.

Human1 \ Human 2->	Buy	Sell	Null	Total
Buy	62	7	6	75
Sell	13	93	14	120
Null	32	30	117	179
Total	117	120	137	374

In this table, the rows represent the classification by Human 1 and the columns the classification by Human 2. The cross tabulation presents the intersection of their message interpretations. The F-stat (chi-square = 268.30, degrees of freedom = 4) rejects

the independence of rows and columns at the 0.1% level. From the table we notice that there is a heavy diagonal, which implies a reasonable level of agreement, yet there are mismatches even when comparing one person's classification to the others. Overall, the beta error (buys classified as sells, and vice versa) lies in the 10-15% range (as opposed to 20% in machine classifier case).

6. The Analysis of Sentiment and Disagreement

Investor opinion is likely shaped by a variety of forces. In this section, we report on the relationship of various variables to our two measures of interest - sentiment and disagreement.

First, we would expect that opinions change slowly, so that there would be persistence in the sentiment and disagreement time series. This will be evidenced by autocorrelation.

Second, by their revealed preferences, posters have an obvious interest in stock returns and volatility. If the posters' impressions were formed by the level of prices, we expect to see a positive relationship between returns and sentiment; but equally plausibly, we could see a negative relationship if the posters tended to be contrarians. If volatility represents uncertainty, then high levels of volatility might be associated with more disagreement given the overall level of uncertainty. We measure returns with data from CRSP. We measure forward-looking uncertainty with implied volatilities on short-dated options reported on Bloomberg.

Third, we may see evidence of "*reinforced persistence*." Investor's views are likely to be more persistent when they are reinforced by data. This would suggest that bullish sentiment is likely to be more autocorrelated when returns are positive.

Fourth, we suspect that people turn from "lurkers" to "posters" when they either have questions or strong opinions. Changes in the level of sentiment and increasing levels of disagreement might be related to the level of "new posters" in the group.

Fifth, discussion is more likely when various parties disagree or when the level of sentiment is on average high. We would therefore anticipate that the level of posting activity (controlling for day of the week effects)²⁰ would be related to the levels of sentiment and disagreement, with higher posting activity related to higher absolute levels of sentiment (either positive or negative) and greater disagreement.

Sixth, our discussion with "Glenn" and our analysis of the content of postings suggests that these on-line discussions take place in the context of a rich environment in which posters are collecting and studying news, filings, analyst reports and non-financial

²⁰ It is important to control for day of week effects because there are noticeable differences in these patterns for the various variables we study. Unlike company filings, press releases and analyst revisions, which are virtually absent on weekends, press and posting activity continues through the weekend, although at lower levels of activity. About 16% of weekly posting activity taking place on Saturdays and Sundays (or 23% from close of market on Friday to market open on Monday.) Company filings are most likely on Fridays, but press releases are more likely made in the middle of the week. Analyst revisions tend to be issued in the second half of the week. Since both company announcements and professional announcements are relatively infrequent, these patterns probably reflect conscious public relations strategies. News stories and postings also peak mid-week.

data. We expect a positive relationship between news sentiment (extracted from news stories using the same algorithm) and posting sentiment. We also expect a positive relationship between news disagreement and posting disagreement. To measure news sentiment and disagreement, we applied the algorithm described above to the major news stories on Factiva.²¹ In addition, we collected Press releases and filing information from Factiva, EDGAR, and Global Access, as well as analyst reports and earnings revisions from Investext and from IBES.

Our sample includes information that can be easily obtained by a *retail* investor *without real-time monitoring*. However, we do not have information that is not stored, such as TV and radio broadcasts. In the current analysis, we also exclude information that might be available only to large institutional investors (e.g., conference call proceedings prior to web-broadcasting or private communications with management prior to Regulation FD; see Bushee, Matsumoto and Miller (2001)).

Table VI breaks down the information events for each of the four companies. Over the seven month period, there were 168 press releases, 58 filings, 207 analyst forecast revisions, 1,667 major news stories and 170,953 stock chat posts. The dispersion in the information releases is intentionally large as we were attempting to capture four different types of firms. For example, there are 73 times more postings at Amazon than at Delta, but 15 times more stories about Delta than about either General Magic or Geoworks. **Table VII** reports the univariate statistics and definitions of the variables we use in our analyses in the remainder of the paper.

6.2 Empirical Evidence on the Determinants of Sentiment and Disagreement

One of our goals is to understand the drivers of investor sentiment. Our clinical methodology encourages us to frame our findings in terms of generated hypotheses. In this section, we consider the evidence provided by our sentiment measure on the four boards we study. These results complement the picture of poster behavior emerging from our interview with a frequent poster.

We begin with simple regressions of sentiment and disagreement measures on explanatory variables such as lagged values of sentiment, the current and lagged values of stock returns, posting volume, trading volume, lagged market return, and current and lagged sentiment derived from news sources using our algorithm. **Table VIII**, panel A examines the determinants of posting volume. Panel B provides an empirical analysis of the levels of sentiment, and panel C shows an analysis of the levels of disagreement.

Panel A shows the determinants of posting volume. There are more messages posted when the stock's trading volume is high and when there are new posters in the prior seven days. There is also evidence of persistence in posting volume for DAL and GMGC (while not significant for AMZN and GWRX). For two stocks (AMZN, GMGC) there is a positive relationship between the number of news stories and the level of message posting. This is consistent with the hypothesis that news is most salient for stocks with high message posting volume. More active boards may use discussion to interpret and digest news releases, as we discuss earlier. For all four stocks, but

²¹ As in section 2.2 we define "major news story" as one in which the name of the company is either in the headline or is mentioned in the lead paragraph and is mentioned at least three times in the body of the article.

significantly only for the two low traditional information stocks, there is a negative relationship between contemporaneous returns and postings, suggesting that message volume picks up when the stocks do poorly. This may be consistent with loss aversion—perhaps losses are more salient to message posters than gains. When bad news arrives, the impact on increased posting volume is observed for all stocks except DAL, albeit not statistically significant, suggesting that e-information from message board discussion may be required to compensate for lower news content. Disagreement is also found to relate to message volume; for all stocks but AMZN, contemporaneous disagreement is significant, though lagged disagreement is not (except for DAL, which is negative). Disagreement and discussion go hand in hand. The high adjusted R-square for three of four stocks and significant F-statistics suggest that we are able to model posting volume reasonably well.²²

Panel B suggests a high level of persistence in the sentiment level, as predicted, shown by the significance of lagged sentiment in three of the four regressions. In addition, sentiment is also positively related to the total volume of postings across all four boards. Sentiment is also positively related to trading volume (although only significant for GWRX), and to current stock returns, which implies that small investors are attentive to market activity and influenced by it. This is consistent with other studies that find the same result [see Das and Chen (2003), and Antweiler and Frank (2004)]. It appears from the data that the market takes time to digest the impact of stock returns into sentiment. This “time to digest” hypothesis is revealed in **Table VIII**, panel B, where there is strong affirmative evidence for this conjecture in the two stocks that have active posting levels (AMZN, GMGC), and it is insignificant on the boards with lower posting activity (DAL, GWRX). The results show that sentiment is driven by contemporaneous and lagged stock returns, which is consistent with the information digestion idea.²³ Except for GMGC, overall market return did not influence stock-specific sentiment.

We tested whether sentiment was more persistent when the market return reinforced the previous day’s sentiment, but this analysis yielded non-significant results. In regressions not reported in the paper we added a dummy variable that took the value of 1 if lagged stock return and lagged posting sentiment had the same sign and 0 otherwise. This variable never resulted significant and the results on other variables were not noticeably affected. Thus, we do not observe that sentiment persistence is affected by its accuracy.

In univariate results, **Table IX**, investor sentiment and news sentiment have strong positive correlations for all firms but GMGC. However, in this multivariate setting, news sentiment is unrelated to message board sentiment (with the curious exception of GMGC). It is likely that other more fundamental factors, i.e., current returns, are more fundamental drivers so that news per se adds limited incremental content.

²² We tested for endogeneity in a system of two equations for the number of postings and disagreement (the two dependent variables in Tables VIII A and VIII C respectively using a Hausman test against a 2SLS alternative, and failed to reject the OLS specification for each of the 4 stocks, even at a very high level of significance.

²³ We are grateful to an anonymous referee for suggesting this analysis.

Our empirical analysis of disagreement in panel C shows that it is primarily related to message posting volume. As anticipated, disagreement and discussion go hand-in-hand. Disagreement is significantly persistent for 2 of the stocks, those with high posting volumes. Surprisingly, few of the other variables are correlated with disagreement. Implied volatility is unrelated to disagreement, even we thought it might capture the future uncertainty of returns for a stock. From the intercept term, we can get some idea of the range of base level disagreement, which seems to lie in the 0.35-0.70 range, which means that the difference between bullish and bearish messages is about 35-70% of total signed messages.

Overall, we find that that sentiment changes slowly and is related to stock returns and trading volume. Higher message volume and disagreement are related to one another, however our daily evidence does not suggest that disagreement leads to discussion, as we had anticipated.

7. e-Information and the Price Formation Process

Market efficiency implies that public information is immediately embedded in the stock price (Fama (1965 and 1975)). If message boards just rehash old information, then e-Information would not make markets any more efficient nor predict stock returns. However, message boards and e-Information would be useful if the boards provide new information about the firm (whether new facts or new analysis of public information) or about the state of mind of traders whose actions might affect prices. We use simple specifications to explore whether stock returns are related to our measure of sentiment and other information state variables. We would expect that this sentiment index would be most salient for firms where the traditional information environment is poor and the chat board information environment is rich (General Magic). Our null hypothesis is that stock returns are *not* predictable using these measures.

7.1 Correlations.

In **Table IX** we report the contemporaneous correlations between our information variables and market variables (returns, excess returns, share turnover, implied volatility and intraday volatility and average bid-ask spreads) as well as the autocorrelations of market variables.

For the two small firms, and, to a lesser extent for Amazon, there is a significant contemporaneous positive correlation between the number of news stories and market returns, which is consistent with the press writing about high performing stocks. For those same firms, there is contemporaneous positive correlation between news *sentiment* and returns, which is indicative of the press writing *favorably* about high performing stocks. For three of the four stocks (excluding the low-news/low-posting Geoworks), there is contemporaneous significant correlation between message board sentiment measures and returns. For the two firms with substantial chat activity, there is a negative correlation between disagreement and returns; when people disagree, returns tend to be lower. Or conversely, when stocks fall, there tends to be greater discussion and disagreement.

While the contemporaneous correlation between information variables and returns is modest, the contemporaneous correlation between the e-Information variables and turnover, volatility, bid-ask spreads and jumps is more robust, especially for the most actively discussed firms in our sample (Amazon and General Magic). These non-return aspects of the financial markets are also correlated with one another, as shown on the right hand part of Panel A, and highly autocorrelated, as shown in Panel B.

7.2 Are Returns for these Four Stocks Explicable Using e-Information?

For the four stocks we study, we examine whether the e-Information and other information variables help to explain returns. If e-Information variables are meaningful, at a minimum we should observe contemporaneous correlations between them and returns. If they contain truly new and novel information, we might observe e-Information predicting returns.

For stock returns, we use not only close-to-close returns, but also open-to-close returns. We use this latter measure because we seek to understand whether reading posts from the day before and the night/early morning prior to market opening would permit a trader to predict subsequent returns. Because a bleary-eyed trader could at best transact at the open price (not the prior night's close), the traditional close-to-close return could not be executed. We use a simple specification, similar to the ones used by Wysocki (1999) and Mitchell and Mulherin (1994). **Table VII** defines the variables we use, which include the following:

- Announcements, including:
 - i. Press releases by the firm
 - ii. SEC filings by the firm
 - iii. Analyst revisions
- News story information
 - i. Abnormal measures of news stories are created by calculating the residual when the measures are regressed against their lagged values as well as day of week and month dummies.
 - ii. Sentiment indices from news stories.
- Posting information, including:
 - i. Abnormal measures of postings are created by calculating the residual when the measures are regressed against their lagged values as well as day of week and month dummies. The time periods for aggregating posting information are the same as for return information. Thus, for the close-to-close analysis we consider the number of posts in the period 4:00 pm prior day to 4:00 pm trading day as the number of posts at time t . For the open-to-close analysis, we consider the period 9:30 am to 4:00 pm on a trading day, as well as the abnormal number of “premarket” posts from market close the prior trading day through 9:30 am the current trading day as the number of posts at times t and $t-1$ respectively.
 - ii. Sentiment using the same time conventions.

- Interactions: The first three events (press releases, filings and analyst reports) are probably most salient if they generate news, chat or change in sentiment. Therefore we create interaction variables by multiplying them by the abnormal news levels, news sentiment, abnormal posting levels and posting sentiment measures.²⁴

We include both contemporaneous and lagged information variables. We want to see if there is a relationship with contemporaneous variables to verify if our constructed measures are sheer noise or whether they are picking up signals that confirm the current state of the market. We use lagged information sources to determine whether the information seems to be useful in predicting returns.

Table X provides the results of this inquiry. Panel A looks at the relationship between contemporaneous information measures and close-to-close market-adjusted returns. Panel B repeats this analysis using open-to-close market adjusted returns. Both of these panels show whether returns are related to the *same day's* announcements, news activity, and posting activity. Evidence of significant relationships would not suggest market inefficiency; to the contrary, it would be consistent with information production being impounded quickly into prices.

Panels C and D repeat the analysis in the first two panels, except that all of the independent variables are lagged. Announcements and news levels are from the prior day, postings follow the convention described above. Relationships here would suggest that a trader could observe data today (until the open of the market) and profit from it.

Panels A and B, which examine contemporaneous information flows and returns, hint that the information environment has a rich set of interactions with the price formation process. For these four stocks, there are indications that stock returns may be higher on days when:

- There are more news stories (AMZN, GMGC) – notice that this is statistically significant for the two high e-information boards, implying a possible catalyst relationship from message board activity, when e-information crosses a threshold level;
- This news conveys more positive sentiment, as measured by our sentiment algorithm applied to the text of news stories (GMGC, GWRX) – here the effect is statistically significant for the two boards that have low news volume. The observation that returns and news are more related for low-news stocks may suggest a certain diminishing marginal impact of news stories;
- The message board postings reflect greater positive sentiment (especially for stocks with more active postings; AMZN, GMGC) – this is consistent with the findings of Antweiler and Frank (2002) who find that stocks with high posting volumes tend to have a significant relationship with contemporaneous returns, suggestive of a price factor in postings; and
- There is a more positive posting sentiment (AMZN, GMGC, GWRX). Separately, stock returns seem to be lower on days when low traditional

²⁴ The specifications we show only include the press release interaction terms, as the others tended not to be material.

information firms issue press releases (GMGC, GWRX) and when the company press release is combined with more news (especially GWRX).

Our posting sentiment index is more closely related to contemporaneous prices than is the sheer numbers of postings, the measure used by Wysocki (2000). These results are encouraging, in that they suggest that our sentiment index, applied to either short messages or longer news stories, captures the tone of the text. Furthermore, news and sentiment measures can help us to “sign” various news events, in this case, company press releases. Not surprisingly, press releases seem to gain importance as they are interpreted by the news and by board posters.

Panels C and D are encouraging, at least for those who are proponents of informationally-efficient markets. Using information that arrives prior to the opening of the market, including overnight posting activity, we find that no “information” variable is consistently informative to a trader who will transact over the course of the day. Of the 43 firm-coefficients in each panel (4 stocks x 10 or 11 variables), 3 or 4 are significant at the 10% level, just as chance would predict.²⁵

Our results are consistent with those found by Tumarkin and Whitelaw (2001) and by Antweiler and Frank (2004). Our three papers were independently produced and use different samples and different methodologies for coding the information content of the message boards.²⁶ Nevertheless, all three of the papers show no predictive power for the message boards to explain subsequent stock returns. In a word, all three of these small sample papers suggest that people trade first and talk later, with returns preceding postings, rather than the other way around.²⁷

8. Discussion and Summary

Our goal in this paper is to clinically study the process of investor discussion and sentiment formation, using stock message boards as a window into investor behavior. Using various language processing routines, we create sentiment and disagreement measures based on the comments posted on the message boards. This clinical study can provide a sense of the phenomenon and suggest a number of *hypotheses* that can be developed in subsequent theory and empirical tests.

We find that a small core of members of online communities carry out an extended discussion that has a number of positive attributes. Our investigation suggests the hypotheses that the boards provide readers with a community of like-minded

²⁵ As a sensitivity check, we performed all these analyses using changes, as opposed to levels, of sentiment. We defined changes as differences in sentiment, percentage differences in sentiment, and residuals from a regression on lagged sentiment and calendar variables. The results are directionally similar.

²⁶ Antweiler and Frank develop a “bullishness index” using Bayesian classifiers and support vector methods. They focus on a broad collection of stocks, whereas our paper examines a few stocks in depth, using a clinical approach. Tumarkin and Whitelaw (2001) use messages from the Raging Bull board, only when the message poster indicates her sentiment. They do not parse the sentiment from the text of the messages. Rather than discard messages which did not reveal their position on the stock, we used five algorithms to parse the opinion of the poster. See Das and Chen (2003) for complete details of the sentiment extraction methodology.

²⁷ We would say that based on the limited samples and time frames of all of the existing papers, this remains a hypothesis, not a proven fact.

investors, and the group delivers on-point discussions, quick dissemination of new public information, and in some instances provide foreshadowing of subsequent news releases. However, these benefits come at the cost of a large number of “false positives” in the form of unsubstantiated rumors. The calculus that leads someone to spend a great deal of time investing in these discussions has more to do with sharing opinions than sharing (or collecting) any private information. We posit this hypothesis on the basis of our interview with an extensive poster, from our inspection of the specific content of the boards, and from our econometric analysis of the predictability of returns using sentiment. While some posters apparently value the benefits of testing their ideas with others, a cost may be to harden the opinions of some.

We extract a time series of sentiment and disagreement measures from the message boards. We find that there is a close relationship between sentiment levels and lagged sentiment, posting activity, stock returns, and lagged stock returns. Once we control for these factors, posting sentiment seems unrelated to news sentiment. In addition, disagreement is related to the intensity of discussion. The discussion on the boards provides one mechanism by which sentiment is created and firmed up. Thus, the paper provides an in-depth study of the mechanism driving investor sentiment.

While sentiment does not apparently predict returns, returns drive sentiment, suggesting a hypothesis that members of the on-line community are more likely to extrapolate past returns, rather than to be contrarian, leading to behavior consistent with the representativeness heuristic (Lakonishok, Shleifer and Vishny, 1994). Our extensive documentation of the environment on the message boards is also consistent with the idea that people will engage in social interaction to mitigate the costs of bounded rationality and for opportunistic reasons (Baker, 1984). This paper suggests the value of sociological inquiry amongst small investor groups.

Appendix: Classifier Algorithms

The classification algorithms in this paper were developed from several different ideas in the field of linear algebra and statistical theory. The algorithms in this paper are described in Das and Chen (2003). Earlier work in a different text classification domain comes from the work of Koller and Sahami (1997) and Dom, Chakrabarti, Agrawal and Raghavan (1998).

Instead of a single classification algorithm, the sentiment index comes from the majority vote of five distinct algorithms. Messages that are ambiguous tend to be characterized by low agreement amongst diverse classifiers, and hence are better discarded as noise, a feature that is implemented by the voting structure. Keeping messages with high agreement amongst classifiers results in better signal detection. The voting algorithm discards on average about 20% of the data. The accuracy of the algorithms tends to be in the 62% range. Human agreement on message classification is slightly higher at about 72%.

The algorithms classify messages into 3 categories: (a) bullish, (b) bearish and (c) neutral/null content. The algorithms are based on a set of keywords carried in a “lexicon”. The lexicon contains words that have meaningful finance implications, and are hand picked or chosen via the use of a discriminant function. Each word in the lexicon is tagged with a bullish or bearish flag.

The classification process also uses a training data set or “grammar.” We pick a set of messages from internet boards and hand classify them. These messages form the basis for further classification of messages. We may also think of these messages as being a base set which reflects the types of messages that we may see. For example, we may assess the “closeness” (using some metric) of messages to this base grammar, and use this measure as a basis for assigning a classification.

In the following paragraphs, we describe the five algorithms in brief. Technical discussion is avoided and the reader is referred to the original papers for more details.

Naive Classifier (NC): This algorithm is the simplest of the ones we used. It consists of simply undertaking a count of the words in the message that appear in the lexicon. Each word that appears is assigned a value of -1 if it is signed as bearish in the lexicon and is assigned a value of $+1$ if it is signed as bullish. The net count for the message is made, and if it exceeds a given positive threshold T , we assign the message a buy classification. If the message count is less than a lower threshold $-T$, it is assigned a sell classification.

Vector Distance Classifier (VDC): Each classified message in the training set is coded as a vector of word counts. If the training set is of size M , then the number of word vectors is M too. The total length of each word vector is the size of the lexicon, denoted L . The matrix of training messages is denoted X , and is of dimension $(L \times M)$. Every incoming message is also coded as a vector of words, which may be denoted Y , which is of dimension $(L \times 1)$. The vector distance between any message in the training set, say X_i and the message vector Y is given by the angle between the two vectors, given by the

standard formula for vectors, i.e. $\cos(\theta) = X \cdot Y / (|X| |Y|)$, which is a number between 0 and 1. ($|X|$ stands for the norm of vector X , and $X \cdot Y$ is the dot product of vectors X and Y). We classify the message Y as having the same class as the message from the training set to which it is closest in the training set, i.e. the one with the highest value of $\cos(\theta)$.

Discriminant Classifier (DC): In this algorithm, we also undertake a word count, but weight each word in the message by its discriminant value from the training set. The discriminant for a word w is given by the ratio of the across-class variance to its within-class variance. This is a simple way to find words that give better classificatory performance on the training data set. To fix intuition, let's suppose that the word "bullish" always appears exactly twice in every buy message and never appears in a message in any other category. The discriminant function for this word would be infinity as its numerator, the across class variance, would be a finite number, and the denominator, the within class variation, would be zero. Once a weighted average word count is taken, the classifier scores the message and accordingly determines its buy, sell or null category.

Adjective-Adverb Phrase Classifier (AAPC): This classifier is based on the simple notion that the parts of a message that contain more emphasis are those that have adverbs and adjectives. To identify the parts of speech of each word in the message, we use a dictionary for parsing each word. Whenever we come across an adjective or adverb, we take triplets of words around the message and submit them for lexical analysis. The score for the message determines its classification.

Bayesian Classifier (BC): This classifier uses the training set as a prior and finds the posterior probabilities for each message to undertake classification in a Bayesian setting. The classifier has three elements: words (w), messages (m) and categories (c). Our final goal is to obtain the conditional probability of the category given a message, i.e. $P(c|m)$. Using Bayes' theorem, this is equal to:

$$P(c | m) = \frac{P(m | c)P(c)}{\sum P(m | c)P(c)}$$

The probability $P(c)$ is simply the proportion of messages in category c . The probability $P(m|c)$ is multinomial over the words in the message. This is given by:

$$P(m | c) = \binom{n(m)}{n(m, w)} \prod_{k=1}^w \theta(c, w)^{n(m, w)}$$

where $\theta(c, w)$ is the probability of seeing word w in category c . $n(m, w)$ is the number of words w in message m . These basic statistics are developed via word counts on the training set.

Implementation: We used these five classifiers, and the results from these classifiers was fed into a voting scheme where we choose to undertake classification of a message when 3 or more of the 5 classifiers agreed on a classification. If we do not get this absolute majority, the message is not classified.

In the following paragraphs we examine a few messages and report their classifications. These messages were posted to the Yahoo! message board for Delta

Airlines (ticker: DAL). The message and its classification by each algorithm as well as the final vote are provided below. The messages are reported verbatim, including spelling and grammatical errors; 3 represents a buy message, 1 is a sell and 0 is a neutral..

Message	N C	V D C	D I S C R	A P C	B C	V o t e
Curious communication: does the management of this company often use a public message board to communicate matters of employment issues? when has any manager of this company ever responded to sniper attacks on this silly board? there have certainly been many arguments in the 700+ messages to this board about mechanics vs pilots, ground workers vs management and so on. i do not believe you are the real donburkett. i think you play with these people.	0	0	0	0	1	0
contact if you are who you say you are contact one of the schedulers by company e-mail for any meeting you want to call, not some suspected ringleader you want to crucify.	0	0	0	0	0	0
Skippy's informed: it looks like im going to have to buy a lap to top to keep hawkster off my a___.i was in orlando this weekend and was told that dal express is kicking butt, exceeding all expectations. the only problem is that they can not get enough airplanes fast enough. does anybody know what dal got for the sale of the greenbriar center? i hope the guys out in dfw get their problems taken care of. as of today june 29 the system load factor is near 80 percent, looks like a great quarter to me.	3	3	3	3	1	3
Month loads how do loads look compared to june 97? from what else i see in the industry, i am guessing flat. july bookings, however, look terrific. i'd also bet that june yields were strong.	3	1	3	3	0	3
So are you guys cutting commissions to rumor, now reported by holly hegegan, says dal and ual are cutting agency commission to 4% tomorrow. i am skeptical of this report, but i am glad i am not a travel agent.	1	1	1	1	1	1
Res Commissions: if it is true, then that is part of the full meaning of the proposed "alliances"--reduce cost to enhance share holder value..and do not forget, deregulation made all this possible.. maybe all the ramrats should vote for a union, so they would have an official mouth piece, rather than boring us with their personal bitching--this is an investment board--buy low and sell high or short the stuff and leave your personal bitching at home..that is what wife/husbands/friends are for... maybe a stock split is in the offering to further dilute the esopholdings... cheers..	0	3	0	0	3	0
Employee Gripping employee griping remember one thing about this board and employee griping, employee moral does play a big part in company earnings. if employees are not happy then that reflects attitude towards the customer. i agree, i do not like the bitchin either to an extent but i am glad to know all the facts about the company not just somebody telling me how great delta airlines is, (which by the way is not true anymore) because they let go some of the greatest people they ever had with 7.5.	1	1	0	1	1	1

References

- Admati, A. R. and P. C. Pfleiderer, 2001, "Disclosing Information on the Internet: Is It Noise or Is It News?" Unpublished manuscript.
- Antweiler, W. and M. Frank, 2002, "Internet Stock Message Boards and Stock Returns," working paper, University of British Columbia.
- Antweiler, W. and M. Frank, 2004, "Is all that talk just noise? The information content of internet stock message boards," *Journal of Finance*, v59(3), 1259-1295.
- Bacharach, M. and O. Board, 2002, "The quality of information in electronic groups," *Netnomics* 4, 73-97.
- Bagnoli, M., M. Beneish, and S. G. Watts, 1999, "Whisper Forecasts of Quarterly Earnings per Share," *Journal of Accounting and Economics*, v28(1), 27-50.
- Baker, W. E., 1984, "The Social Structure of a National Securities Market," *American Journal of Sociology*, v89(4), 775-811.
- Barber, B. and T. Odean, 2001, "Boys will be Boys: Gender, Overconfidence, and Common Stock Investment," *Quarterly Journal of Economics*, v116(1), 261-292.
- Bernard, V. L. and J. K. Thomas, 1990, "Evidence that stock prices do not fully reflect the implications of current earnings for future earnings," *Journal of Accounting and Economics*, v13, 305-340.
- Blaug, M., 1992, *The Methodology of Economics*, Cambridge: Cambridge University Press.
- Brooks, C., 1998, "Predicting stock market volatility: Can market volume help?" *Journal of Forecasting*, v17(1), 59-80.
- Bushee, B., D. Matsumoto and G. Miller, 2001, "Open versus closed conference calls: The determinants and effects of broadening access to disclosure," unpublished manuscript.
- Busse, J. and T. C. Green, 2001, "Market efficiency in real time," unpublished Emory University working paper.
- Chakrabarti, S., B. Dom, R. Agrawal and P. Raghavan, 1998, "Scalable feature selection, classification and signature generation for organizing large text databases into hierarchical topic taxonomies," *The VLDB Journal*, Springer-Verlag.

- Chevalier, J. and D. Mayzlin, 2003, "The effect of word of mouth on sales: Online book reviews," Yale University working paper ES#28 & MK#15 (dated Aug. 6, 2003).
- Coval, J. and T. Shumway, 2001, "Is sound just noise?" *Journal of Finance*, v56(5), 1887-1910.
- Das, S. and M. Y. Chen, 2003, "Yahoo! For Amazon: Sentiment extraction from small talk on the Web," working paper (original May 2000), Santa Clara University.
- Das, S. and J. Sisk, 2003, "Financial Communities," working paper, Santa Clara University.
- DeLong, J. B., A. Shleifer, L. Summers and R. J. Waldman, 1990, "Positive feedback investment strategies and destabilizing rational speculation," *Journal of Finance*, v45(2), 379-395.
- Fama, E., 1965, "The behavior of stock market prices," *Journal of Business*, v38, 34-105.
- Fama, Eugene, 1975, *Foundations of Finance*, Basic Books, New York.
- Felton, J. and J. Kim, 2002, "Warnings from the Enron Message Board," *The Journal of Investing*, v11(3).
- Fleming, M. and E. M. Remolona, 1999, "Price formation and liquidity in the U.S. treasury market: The response to public information," *Journal of Finance* v54(5), 1901-15.
- Garbade, K. and W. Silber, 1978, "Technology, Communication and the Performance of Financial Markets: 1840-1975," *Journal of Finance*, v33, 819-832.
- Gay, G., J. R. Kale, R. W. Kolb, and T. H. Noe, 1994, "(Micro) Fads in asset prices: Evidence from the futures market," *Journal of Futures Markets*, v14, 637-59
- Gervais, S. and T. Odean, 2001, "Learning to be Overconfident," *Review of Financial Studies*, v14(1), 1-27.
- Glosten, L. R., 1987, "Components of the bid-ask spread and the statistical properties of transaction prices," *Journal of Finance*, v42, 1293-1307.
- Glosten, L. R. and P. R. Milgrom, 1985, "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders," *Journal of Financial Economics*, v14, 71-100.
- Goetzman, W., M. Massa and G. Rouwenhorst, 2000, "Behavioral Factors in Mutual Fund Flows," unpublished manuscript, Yale University.

- Godes, D. and D. Mayzlin, 2003, "Using online conversations to study word of mouth communication," Unpublished manuscript, Harvard Business School.
- Harris, M., and A. Raviv, 1993, "Differences of Opinion Make a Horse Race," *Review of Financial Studies*, 6, 473-506.
- Hasbrouck, J., 1996, "Modeling market microstructure time series," unpublished manuscript, New York University.
- Ho, T. and H. R. Stoll, 1981, "Optimal dealer pricing under transactions and return uncertainty," *Journal of Financial Economics*, v9, 47-73.
- Holthausen, R. W. and R. E. Verrecchia, 1988, "The effect of sequential information releases on the variance of price changes in an intertemporal multi-asset market," *Journal of Accounting Research*, v26(1), 82-106.
- Karpoff, J. M., 1987, "The relation between price changes and trading volume: a survey," *Journal of Financial and Quantitative Analysis*, v22(1), 109-126.
- Kim, O. and R. E. Verrecchia, 1991, "Trading volume and price reactions to public announcements," *Journal of Accounting Research*, v29(2), 302-321.
- Koller, D. and M. Sahami, 1997, "Hierarchically classifying documents using very few words," International Conference on Machine Learning 14, Morgan-Kaufman, San Mateo, CA.
- Kothari, S.P., 2001, "Capital markets research in accounting," *Journal of Accounting and Economics*, v31(1-3), 105-231.
- Lakonishok, J., A. Shleifer, and R.W. Vishny, 1994. "Contrarian Investment, Extrapolation, and Risk," *Journal of Finance*. v49(5), 1541-1578.
- Lee, C., A. Shleifer and R. Thaler, 1991, "Investor Sentiment and the Closed-end fund puzzle," *Journal of Finance*, v46(1), 75-109.
- Leftwich, R., 1981, "Evidence of the impact of mandatory changes in accounting principles on corporate loan agreements," *Journal of Accounting and Economics* v3, 3-36.
- Leinweber, D. and A. Madhavan, 2001, "Three hundred years of stock market manipulation," *The Journal of Investing*, v10(7), 7-16.
- Lo, A., and C. MacKinlay, 1990, "When are contrarian profits due to stock market overreaction?" *Review of Financial Studies*, v3, 175-208.

- MacKinlay, A. C., 1997, "Event studies in economics and finance," *Journal of Economic Literature*, v35, 13-39.
- Mendenhall, R. R., 1991, "Evidence on the possible underweighting of earnings-related information," *Journal of Accounting Research*, v29(1), 170-179.
- Mitchell, M. and J. H. Mulherin, 1994, "The impact of public information on the stock market," *Journal of Finance*, v49, 923-950.
- Odean, T., 1998. "Are Investors Reluctant to Realize Their Losses?" *Journal of Finance*, v53(5), 1775-1798.
- O'Hara, M., 1995, *Market Microstructure Theory*, Blackwell Publishers, Cambridge, Mass.
- Roll, R., 1986, "The Hubris Hypothesis of Corporate Takeovers," *Journal of Business*, v59(2), 197-216.
- Samuelson, P., 1965, "Proof that properly anticipated prices fluctuate randomly," *Industrial Management Review*, v6, 41-49.
- Stickel, S. E. and R. E. Verrecchia, 1994, "Evidence that trading volume sustains stock price changes," *Financial Analysts Journal*, November-December, 57-67.
- Stoll, H. R., 1978, "The supply of dealer services in security markets," *Journal of Finance*, v33, 1133-1151.
- Tumarkin, R. and R. Whitelaw, 2001, "News or Noise? Internet message board activity and stock prices," *Financial Analysts Journal*, v57, 41-51.
- Wysocki, P., 1999, "Cheap talk on the Web: The determinants of postings on internet stock message boards," unpublished manuscript, University of Michigan.

Figure I: Distribution of Posting Volume and Number of News Stories, 1998 Q3/4.

This figure plots for each of 504 firms that received at least 25 posts on The Motley Fool over the period 7/1/98-1/31/99 against the number of major news stories recorded in Factiva. A “major news story” is defined as one in which the firm’s name is either in the title of the story or in the lead paragraph (plus three additional mentions.) The horizontal and vertical bands represent the top and bottom quintiles along each dimension and the four sample firms are identified (DAL = Delta Air Lines, AMZN = Amazon.com, GWRX = Geoworks, and GMGC = General Magic).

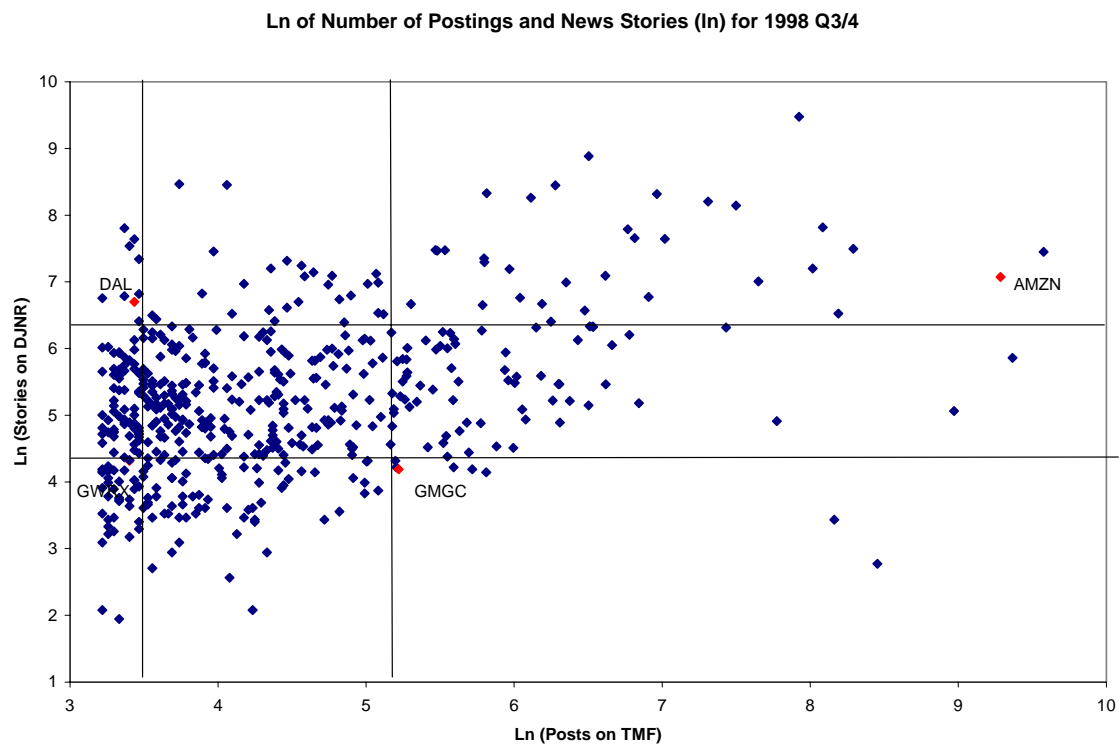


Figure II: Posting Activity, 1997-2003.

This figure presents a plot of trend in the number of messages per quarter from the inception of the earliest message board on Yahoo! (26 November, 1997), until 26 August 2003. We provide plots for two tickers in our study, AMZN and DAL. (GWRX and GMGC failed in the dotcom bust and subsequently removed from the Yahoo! Website, so we cannot track the same time series.) The total number of messages is as follows: AMZN = 666,820, DAL = 79,126. The spike after for DAL after our study follows September 11, 2001, and reflects increased interest in DAL an airline stock.




Figure III: Message Poster Demographics

Panel A: Histogram of Poster Message Distribution

The figure below shows the number of messages posted by screen name for posters on all message boards over the seven month period. In aggregate, there was one screen-name which posted close to 5000 messages, and 7552 screen-names which posted only one message.




Figure III (continued)

Panel B: Message Posters and Their Use of Electronic Financial Media


We do not have independent data on the demographics of message posters. Yahoo! Finance reveals certain demographic characteristics of their users, which is reported below.

- Posters are predominantly male, high-income, medium-to-high net worth, with a midterm to long-term trading orientation.
- Forty-five percent have portfolios over \$100,000.
- About 85 percent are conducting fewer than 10 trades a month.
- Yahoo! Finance commanded one-third of top financial content site traffic.
- Yahoo! Finance engaged users most loyal to transaction-based relationships.
- More than half of all Yahoo! Finance engaged users were online account holders.
- Nearly one in 10 Yahoo! Finance engaged users were account holders with E*Trade or American Express.
- Nearly half of Yahoo! Finance engaged users visited CBS MarketWatch.
- Yahoo! Finance engaged users seek value: More than one in 10 shopped at BizRate.com.
- Men made up a majority of engaged users, and their degree of engagement exceeded that of women.
- Middle-aged adults made up the core of Yahoo! Finance engaged users.
- High-income households are the heaviest Yahoo! Finance engaged users.
- The percentage of online users that use each of the major financial portals is:
 - Yahoo! Finance: 10%
 - CNBC: 9%
 - Motley Fool: 5%
 - CBS Marketwatch: 5%
 - Bloomberg: 4%
 - Nasdaq: 4%
 - CNNFN: 4%
 - Morningstar: 3%
 - Forbes: 3%
 - Wall Street Journal: 2%

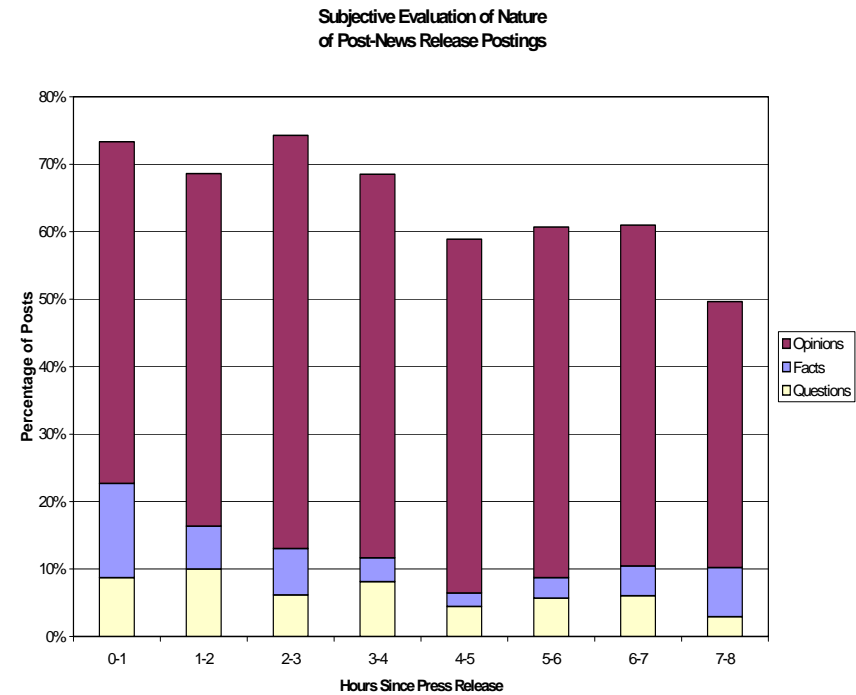
Source: Jupiter Research, 3/14/00 "Consumer-Created Content: Creating and Valuing User-Generated Programming"; 1/5/01 "Yahoo! Finance Engaged Users and Online Affiliates"; 4/5/02 "Portal Channels: Mapping the Competitive Landscape"; 9/18/02 "Jupiter Consumer Survey. Brokerage & Wealth Management, 2002"

Figure IV: Discussion Activity around 16 News Events

We report on sixteen press releases made by the four companies which seemed to indicate a material change in the firm’s business or financing. For each event, we summarize the speed of information dissemination by traditional and chat room sources. Panel A shows the aggregate number of posts after the time-stamp of the press release. For each news release, we manually categorized the postings in the subsequent eight hours into on-point (relating to the news event) or off-point (unrelated to the news event or spam), and we also categorized the on-point postings into ones that asked a question, offered an alleged fact, and proposed an opinion. Panel B shows the distribution of the on-point postings. Table III provides detailed information on the nature of the events, the first posting activity, and the first press activity.



Panel A: Number of postings by hour after 16 selected corporate press releases.



Panel B: Distribution of type of posting by hour after 16 selected corporate press releases. Postings are classified as on-point if related to the news story, and off-point otherwise. The histogram shows the percentage of on-point posts (the height of each bar) and the nature of the on-point posts (asks question, provides alleged fact, proposes opinion.)

Figure Va: Numbers of Postings by Day for the Four Sample Companies

The figures below show the total number of messages posted on the boards for each of the stocks under analysis. The graphs measure the daily number of messages by stock for the period July 1st, 1998 to January 31st, 1999. The graphs' scales differ as a function of the maximum number of daily postings for each stock.




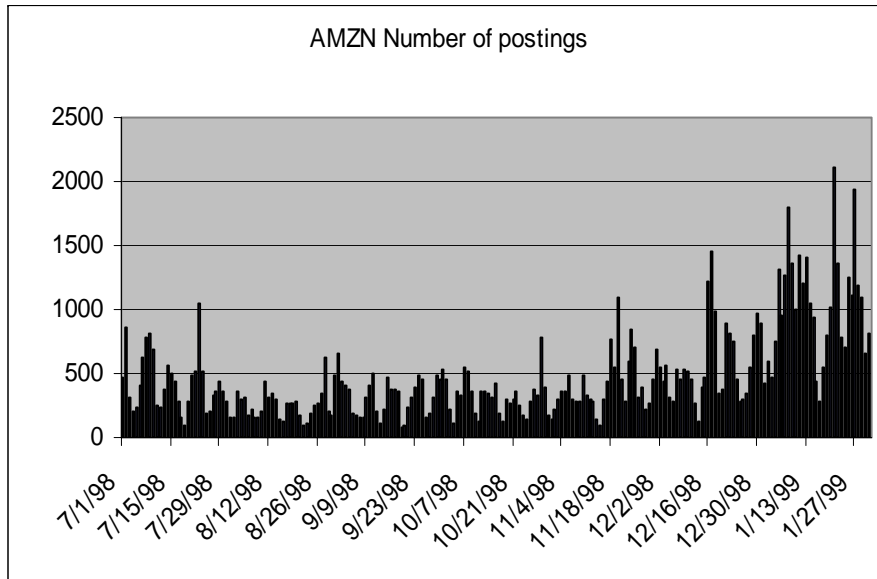


Figure Vb: Sentiment Percentage by Day for the Four Sample Companies

The figures below show the sentiment percentages for the messages posted on the boards for each of the four stocks. Sentiment percentage is calculated as (Buy-Sell Messages)/Total number of postings. The graphs measure the daily sentiment percentage for the period July 1st, 1998 to January 31st, 1999.

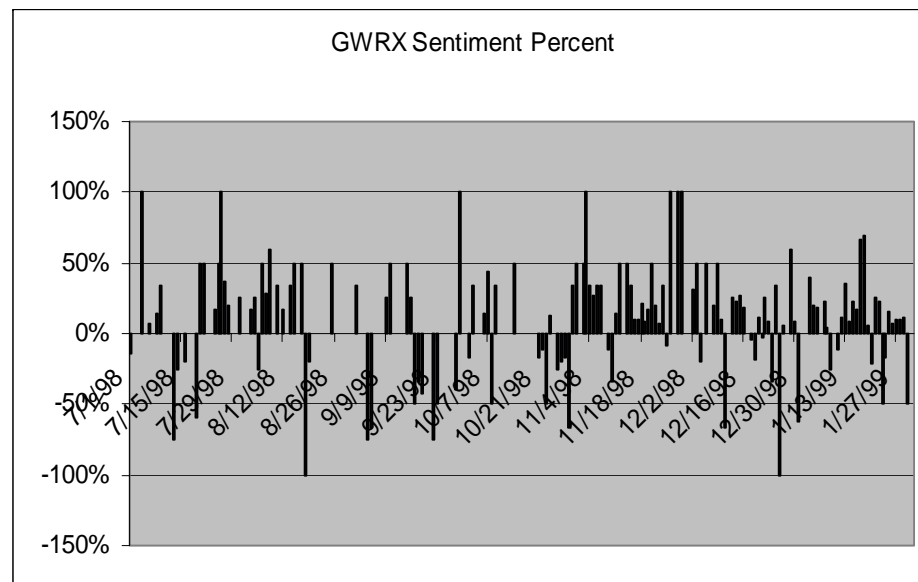
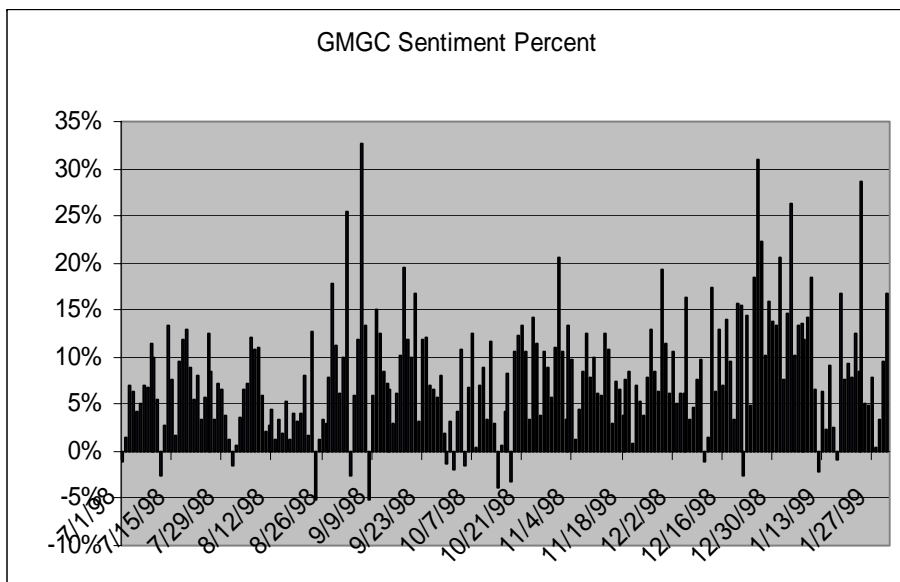
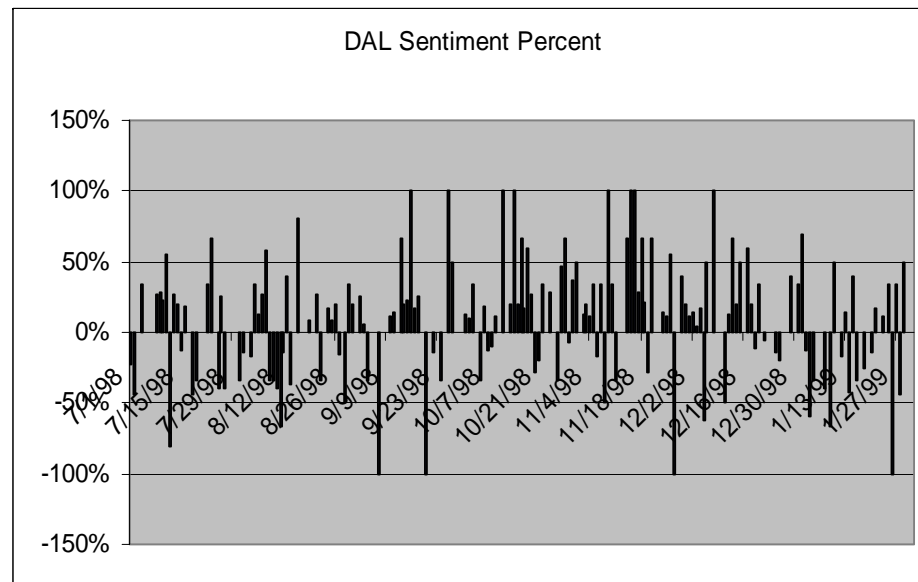



Table I
Characteristics of the Four Companies Studied

The table below provides basic business and financial information for the four companies we study. The data sources and dates are given in the table. The information for this table comes from onsource.com, Hoovers, Bloomberg and public filings. All \$ in millions. Financial figures are as of the end of each company's fiscal year: 12/31/98 for Amazon.com and General Magic, 6/30/99 for Delta Air Lines, and 3/31/1999 for Geoworks.

Firm	Amazon.com, Inc.	Delta Air Lines, Inc	General Magic, Inc.	Geoworks Corporation
Business	On-line retailer	Major air carrier	Voice Application Service Provider to telecom and Internet companies	Provider of wireless software solutions
Industry	Retail (specialty; non apparel)	Airlines	Software and programming	Communications services
Stock listing (ticker)	NASDAQ (AMZN)	NYSE (DAL)	NASDAQ (GMGC)	NASDAQ (GWRX)
Market value (Year end)	17054	7984	168	56
Year founded	1995	1924	1990	1983
Total assets	2471.6	16750.0	36.3	18.2
Total sales	1639.8	14597.0	2.3	8.8
Net income	-720	1101	-38.9	-15.8
Institutional ownership	30%	75%	10%	15%
Number of institutions	442	776	61	54
Bond rating	B	BBB	not rated	not rated
Number of analysts (equity + fixed income)	26 + 2	14 + 11	4+ 0	5+0
Number of employees	2100	74000	169	110
Avg trading volume (M shares / day)	27.75	1.19	1.04	0.28
Avg Volume (% Outstanding)	9.37%	0.83%	3.50%	1.73%
Average \$ value of trades/day (\$ mil)	\$555	\$65	\$8	\$1

Table II
Posting Activity by Screen-Name and Poster Concentration

The table below provides information on posting activity per poster on the four major stock message boards (Yahoo!, The Motley Fool, Silicon Investor and Raging Bull) for the period July 1, 1998 through January 31, 1999 for the four stocks. A poster is defined by a unique screen name. Panel A uses the Herfindahl measure: $\sum (\text{Marketshare} \cdot 100)^2$; where market-share of poster i is the share of messages over the period we study: $\text{Numberofpostings}_i / \sum_{j=1}^n \text{Numberofpostings}_j$.

Panel A: Message share by boards: Herfindahl indices

STOCK	board	Number of Posters	Herfindahl index	Number of postings
AMZN	RagingBull	162	286	517
	SiliconInvestor	866	478	29543
	TMF	1031	175	10854
	Yahoo	<u>10110</u>	<u>25</u>	<u>61906</u>
	Average	3042	241	25705
DAL	RagingBull	0	n/a	0
	SiliconInvestor	10	2648	47
	TMF	23	510	31
	Yahoo	<u>371</u>	<u>171</u>	<u>1313</u>
	Average	101	1110	348
GMGC	RagingBull	43	842	238
	SiliconInvestor	189	214	2297
	TMF	62	455	185
	Yahoo	<u>2914</u>	<u>60</u>	<u>62164</u>
	Average	802	393	16221
GWRX	RagingBull	6	1837	7
	SiliconInvestor	29	1514	172
	TMF	10	2152	29
	Yahoo	<u>359</u>	<u>274</u>	<u>1764</u>
	Average	101	1444	493

**Table II (cont.)
Posting Activity by Screen-Name and Poster Concentration**

Panel B: Number of messages posted over the sample period by screen name and stock.

Number of postings	AMZN		DAL		GMGC		GWRX		TOTAL	
	Number of posters	Percentage	Number of posters	Percentage	Number of posters	Percentage	Number of posters	Percentage	Number of posters	Percentage
>4,000	1	0%							1	0%
1001-4000	7	0%			6	0%			13	0%
501-1000	10	0%			14	0%			24	0%
101-500	114	1%			101	3%	2	0%	217	1%
51-100	135	1%	3	1%	112	3%	5	1%	255	2%
26-50	281	2%	5	1%	150	5%	6	1%	442	3%
11-25	730	6%	17	4%	349	11%	14	3%	1110	7%
6-10	949	8%	27	7%	365	11%	37	9%	1378	9%
2-5	3964	33%	111	27%	992	31%	126	31%	5193	32%
1	5978	49%	241	60%	1119	35%	214	53%	7552	47%
	12169	100%	404	100%	3208	100%	404	100%	16185	100%

Table III
Tracer Analysis: Backwards Analysis of Subsequently Released News Stories

We report on sixteen press releases made by the four companies which seemed to indicate a newsworthy change in the firm's business or financing. For each event, we summarize the speed of information dissemination by traditional and chat room sources.

Company	Event	Release Date and Time	First Press Notice	Minutes to First Press Activity	Subsequent Press Activity	First Posting	Minutes to First Post	Subsequent Related Postings	Foreshadowing
Geoworks	Poor 1st Quarter Results	January 26, 1999; 8:02 AM EST	Dow Jones NewsWire 1/26; 9:34 AM EST	92	Warning of Price drop 1/25/99	1/26/1999 8:45 AM (Yahoo!)	43	15 Posts before market closes; 8 more before market opens next day	Discussion of poor earnings possibilities given large price drop 1/25/99
Geoworks	GWRX announces Cooperation with Optimay	January 20, 1999; 6:00 AM EST	Dow Jones NewsWire 1/21; 3:53 AM EST	1313	Warning of Price Jump 1/19/99	1/20/1999 7:07 AM (Yahoo!)	67	18 Posts before market closes, 9 more before market opens next day; 4 before close.	
Geoworks	Dave Grannan Named as CEO	January 11, 1999; 8:05 AM EST	Dow Jones NewsWire 1/11; 8:03 AM EST (preliminary information)	-2	1/11: 2 Dow Jones Articles 1/12: Wall Street Journal	1/11/99 9:15 AM (Yahoo!)	70	16 posts before market closes; 4 more before opening next day; 6 before close	
Geoworks	Debuts Enhanced Phone with Mitsubishi	November 16, 1998; 6:00 AM EST	Dow Jones NewsWire 11/16; 11:10 AM EST (preliminary)	310	11/16: 7 Dow Jones Articles, 4 Price Alerts; 1/17 Wall Street Journal	11/16/99 10:06 AM (Yahoo!)	246	35 posts before market closes; over 80 before next day market opening	
Delta	Alliance With Korean Air	August 6, 1998; 8:00 PM EST	Dow Jones NewsWire 8/6 8:25 PM EST	25	8/6: 1 Dow Jones Article; 8/10: M2 Presswire release	N/A	NEVER	N/A	
Delta	Finances Beat Consensus	July 16, 1998; 9:14 AM EST	Dow Jones NewsWire 7/16 8:59 AM EST (preliminary)	-15	7/16: 4 Dow Jones Articles; 7/17: 5 (4 Newspaper, 1 BusinessWire)	7/16/99 9:25 AM (Yahoo!)	11	2 before market closes	
Delta	Announces Strategic Mgmt Reorganization in Int'l Markets	December 9, 1998; 7:01 AM EST	Dow Jones NewsWire 12/9 7:03 AM EST (preliminary)	2	12/9: 3 Dow Jones Articles; 12/10: M2 Presswire	N/A	NEVER	N/A	
Delta	2 for 1 Stock Split Approved	October 22, 1998 1:54 PM EST	Dow Jones NewsWire 10/22 1:55 PM EST (preliminary)	1	10/22: 2 Dow Jones Articles; 10/23: M2 PressWire	10/22/99 4:01 PM (Yahoo!)	127	1 after market closes 10/22, 3 on 10/27	Questions about split timing 10:46am

Table III (continued)
Tracer Analysis: Backwards Analysis of Subsequently Released News Stories

Company	Event	Release Date and Time	First Press Notice	Minutes to First Press Activity	Subsequent Press Activity	First Posting	Minutes to First Post	Subsequent Related Postings	Foreshadowing
General Magic	Agreement with Intuit for voice access to financial info	November 9, 1998; 4:04 AM EST	Dow Jones NewsWire 11/9 4:09 AM EST (preliminary)	5	11/9: 4 Dow Jones Articles; 11/10: 1 Dow Jones Article	11/9/98 12:39 AM (Yahoo!)	515	54 before market opens; 112 more before close; over 90 before open next day	Info of radio broadcast of Quicken news before press release.
General Magic	2nd Quarter Results	July 29, 1998; 4:02 PM EST	Dow Jones NewsWire 7/29 4:06 PM EST (preliminary)	4	7/29: 2 (Dow Jones, BusinessWire); 7/30: Wall Street Journal	7/29/98 4:11 PM (Yahoo!)	9	130 between release and 9pm that night	Discussion of potential results beforehand
General Magic	Alliance With Microsoft for Auto-Enabled PC	January 7, 1999; 7:34 AM EST	Dow Jones NewsWire 1/7 2:23 PM EST (preliminary)	409	1/7: 2 Dow Jones Articles	1/7/99 7:53 AM (Yahoo!)	19	59 before market opens; 179 more before market close	Rumor of shared booth at CES between MSFT and GMGC (1/6 10:12AM)
General Magic	Spin off of DataRover Division as an independent company	October 28, 1998; 8:10 PM EST	LA Times and New York Times, 10/30	>1500	10/30: 3 (Bloomberg, LA Times, ComputerWire) 11/2: 1 Electronics News (two subsequent mentions in articles referencing Q3 Results)	10/28/1998 8:41 PM (Yahoo!)	31	32 Additional Posts on boards before market opens the following morning	Posting of suspicious independent DataRover URL nine days before announcement.
Amazon	Acquisition of Jungle	August 4, 1998; 7:30 AM EST Note: Filing on 8/8/98 (8-K)	Dow Jones NewsWire 8/4; 9:05 AM EST	95	8/8: 17 8/9: 69 8/10: 3 8/11-18: 27	8/4/98 7:47 AM (Silicon Investor). 7:58 AM (Yahoo!)	17	5 additional related posts before the opening of the market (9:30 am); 21 more by the close of the market (4:00 pm); 22 additional before the next morning's market open	
Amazon	\$500 million debt issue	January 28, 1999; 7:25 AM EST	Dow Jones Newswire 1/28 7:41 AM EST	16	1/29: 1 Dow Jones Article; 1/29: 6; 2/1: 4	1/28/99 7:40 AM (Yahoo!)	15	34 before market open; over 300 more before market closes	Bond Discussion in days before (no facts, though)
Amazon	Enters European Book Market	October 15, 1998; 5:00 AM EST	Dow Jones Newswire 10/15 5:51 AM EST	51	10/17: 3 Newspaper Articles	10/15/99 6:10 AM (Motley Fool)	70	6 before market opens, 6 more before closes	
Amazon	Amazon Announces 3 for 1 Split	November 19, 1998; 5:45 PM EST	Dow Jones NewsWire 11/19 5:47 PM EST	2	11/20: 9; 11/21: 3. 11/22: 1; 11/23: 1	11/19/98 5:51 PM (Yahoo!)	6	Over 300 more posts before midnight 11/19	Split speculation in days before announcement

Table IV
Analysis of Outcomes of Board Initiated Rumors

The table below displays the number of merger and acquisitions rumors about the four sample firms posted on four bulletin boards between 7/1/98-1/31/99 and their relationships with news stories published in the period 1/1/98-8/31/99. Panel A reports which posting rumors received news coverage and the tone of the hypothetical news story (rumor, denial or confirmatory report). 11 of 54 received some press coverage. For those 11, Panel B reports the timing of the postings when there was some press coverage.

Panel A: Posting Rumors and Related Press Coverage

	Related Press Coverage				TOTAL
	Event Transpired	Press Rumor Surfaced	Firm Issues Denial	No Press Coverage	
Rumor with 4 or fewer posts	1	4	1	41	47
Rumor with five or more posts ^a	0	5	0	2	7

^a A rumor that is mentioned in five or more different postings.

Panel B: Number and Timing of Relevant Postings Relative to Press Coverage

Firm involved	Delta	Delta	Amazon	Amazon	Amazon	Amazon	Amazon	Delta	Delta	General Magic	General Magic
Press Outcome	Firm Denial	Merger Transpired	Press Rumor	Press Rumor	Press Rumor	Press Rumor	Press Rumor	Press Rumor	Press Rumor	Press Rumor ^a	Press Rumor
>1 mo. before story	1	1	1	-	19	11	5 ^b	3	-	-	6 ^c
Month before story	1	-	-	1	6	-	-	-	-	4	-
Month after story	-	-	-	-	3	-	4	-	2	-	-
>1 mo. after story	-	-	-	-	2	1	-	2	2	-	-

- a. The press stories mentioned the Internet as a source of the rumor.
- b. All posting two months or more before the news story.
- c. All postings four months or more before the news story. Press cites the Internet as a source of the rumor.

Table V
e-Information Variables for the Four Companies Studied

The table below shows the classification of the messages posted on the four major stock message boards (Yahoo!, The Motley Fool, Silicon Investor and Raging Bull) for the period July 1, 1998 through January 31, 1999 for the four stocks. The top panel shows the message classification for the entire period, using the algorithm described in the Appendix. "Opinion" is defined as the percentage of all messages which are either buys or sells. "Sentiment" is defined as the net number of buy minus sell messages. "Sentiment %" divides the sentiment measure by the total number of messages. "Disagreement" is defined as $|\text{Sentiment}|/(\text{Buy} + \text{Sell Messages}) - 1$.

Company	AMZN		DAL		GMGC		GWRX	
TOTAL PERIOD								
MESSAGES								
Buy	29367	29%	404	29%	15276	24%	557	28%
Sell	23017	22%	293	21%	10949	17%	372	19%
Neutral	36916	36%	535	38%	31854	49%	813	41%
Nonclassified	13363	13%	166	12%	6835	11%	236	12%
Total	102663	100%	1398	100%	64914	100%	1978	100%
OPINION	51%		50%		40%		47%	
SENTIMENT	6350		111		4327		185	
SENTIMENT %	6%		8%		7%		9%	
DISAGREEMENT	88%		84%		84%		80%	
DAILY AVERAGES								
No message days	0		10		0		20	
OPINION								
Mean	52%		45%		41%		45%	
Median	52%		50%		41%		50%	
Std Deviation	5%		25%		5%		27%	
SENTIMENT								
Mean	30		1		20		1	
Median	20		0		17		0	
Std Deviation	43		2		17		3	
DISAGREEMENT								
Mean	85%		43%		80%		43%	
Median	87%		50%		82%		50%	
Std Deviation	10%		41%		13%		40%	

Table VI
Information Events for the Firms in the Sample

The table below shows the total number of information events (press releases, filings, analyst reports/revisions, major news stories and posts) for the four sample firms over the time period July 1, 1998 through January 31, 1999.

Firm	Amazon	Delta	General Magic	Geoworks	Total
News quintile	High	High	Low	Low	
Chat quintile	High	Low	High	Low	
Press releases	22	109	20	17	168
Filings	26	10	12	10	58
Analyst reports/revisions	135	68	-	4	207
Major news stories	987	549	66	65	1,667
Postings	<u>102,663</u>	<u>1,398</u>	<u>64,914</u>	<u>1,978</u>	<u>170,953</u>
Total	103,833	2,134	65,012	2,074	173,053

Table VII
Variable Definitions

Variable	Definition	SUMMARY STATISTICS: Mean (Std. Deviation)				
		All firm- periods	AMZN	DAL	GMGC	GWRX
Implied volatility	Implied volatility on at-the-money call options, taken from Bloomberg	0.880 (0.354)	0.948 (0.168)	0.450 (0.074)	1.243 (0.140)	n/a n/a
Share turnover	Number of shares traded that day divided by number of shares outstanding	0.040 (0.054)	0.096 (0.051)	0.009 (0.004)	0.036 (0.038)	0.020 (0.057)
<u>Announcements</u>						
Press release	Dummy variable. Equals 1 if the company has made a press release on that day	0.209 (0.374)	0.142 (0.350)	0.459 (0.500)	0.135 (0.343)	0.101 (0.303)
Filing	Dummy variable. Equals 1 if the company has made a SEC filing on that day	0.090 (0.281)	0.149 (0.357)	0.061 (0.240)	0.081 (0.274)	0.068 (0.252)
Analyst revision	Dummy variable. Equals 1 if any analyst has issued some sort of earnings revision on that day	0.122 (0.327)	0.250 (0.434)	0.216 (0.413)	0.000 (0.000)	0.020 (0.141)
<u>News activity</u>						
Abnormal news stories	Residual from regression of news stories on prior day news stories, day of week, and month of year dummy variables.	0.000 (2.996)	0.000 (4.765)	0.000 (5.445)	0.000 (0.822)	0.000 (0.951)
Lagged abnormal stories	Note: lag here determined by last calendar day, where the lag of Monday includes news published in the three previous days.					

**Table VII (continued)
Variable Definitions**

Variable	Definition	SUMMARY STATISTICS: Mean (Std. Deviation)				
		All firm- periods	AMZN	DAL	GMGC	GWRX
Posting activity						
Abnormal number of posts (close-to-close)	Residual from regression of posts on the 4:00 pm prior day to 4:00 pm trading day period, on its lag, and day of week, and month of year dummy variables.	0.000 (105.9)	0.000 (181.0)	0.000 (2.43)	0.000 (109.2)	0.000 (14.42)
Abnormal number of market posts	Residual from regression of posts on trading day from 9:30 am to 4 pm on its lag, and day of week, and month of year dummy variables.	0.000 (52.20)	0.000 (82.28)	0.000 (2.29)	0.000 (63.86)	0.000 (9.26)
Abnormal number of pre-market posts	Residual from regression of posts from 4 p.m. prior day to 9:30 am on trading day on its lag, and day of week, and month of year dummy variables.	0.000 (77.16)	0.000 (133.7)	0.000 (3.03)	0.000 (76.96)	0.000 (9.08)
Sentiment level (close-to-close)	Number of buy messages - number of sell messages from 4:00 pm prior day to 4:00 pm trading day.	12.682 (24.42)	30.723 (48.42)	0.581 (1.96)	23.034 (18.20)	0.973 (3.50)
Sentiment level during market hours	Number of buy messages - number of sell messages from 9:30 am to 4:00 pm	4.973 (13.46)	9.378 (23.25)	0.284 (1.19)	9.655 (10.01)	0.574 (1.88)
Sentiment level during pre-market hours	Number of buy messages - number of sell messages from 4:00 pm prior day to 9:30 am trading day. For Monday it includes all posting activity from 4:00 pm Friday to 9:30 am Monday	8.855 (19.20)	21.345 (31.22)	0.297 (1.46)	13.378 (13.22)	0.399 (2.86)
Interaction terms	Note: the information events are dummy variables indicating a press release (etc.) on the current day. The four interaction terms indicate the abnormal level of news stories or posts, or the level of news sentiment or posting sentiment.					
Lags	Note: lags are determined by previous trading day, not by previous calendar day					

Table VIII
Analysis of Posting Activity Characteristics

The dependent variables in this table are the different dimensions of e-Information: Posting Volume, Sentiment and Disagreement. Data from the four boards is pooled for each of the four firms during the period 7/1/98-1/31/99.

Panel A: Analysis of Posting Volume

Dependent variable: Number of postings	AMZN		DAL		GMGC		GWRX	
	Coeff	p-value	Coeff	p-value	Coeff	p-value	Coeff	p-value
Constant	-112.390	(0.551)	3.006	(0.456)	-75.038	(0.531)	-7.213	(0.022)
<u>Posting Activity</u>								
Lagged number of postings	0.015	(0.827)	0.319	(0.001)	0.443	(0.000)	0.059	(0.372)
Disagreement	48.375	(0.774)	4.645	(0.000)	172.958	(0.035)	8.777	(0.005)
Lagged disagreement	87.219	(0.610)	-2.742	(0.018)	38.120	(0.631)	3.313	(0.286)
<u>Stock Market Activity</u>								
Trading volume	0.000	(0.000)	0.000	(0.147)	0.000	(0.000)	0.000	(0.000)
Stock return	-276.077	(0.159)	-18.428	(0.246)	-448.101	(0.001)	-28.416	(0.017)
Lagged stock return	-293.856	(0.092)	31.059	(0.031)	-136.846	(0.229)	30.783	(0.000)
Equal weighted stock market return	339.337	(0.756)	-43.126	(0.319)	712.895	(0.332)	141.738	(0.239)
Implied volatility	37.369	(0.687)	-4.232	(0.562)	-85.237	(0.188)		
<u>News Activity</u>								
Average news sentiment	-8.835	(0.133)	0.389	(0.423)	-2.171	(0.935)	-3.957	(0.356)
Abnormal news stories	6.466	(0.004)	0.137	(0.337)	26.168	(0.004)	-2.078	(0.158)
New posters in the last seven days	0.593	(0.000)	0.022	(0.817)	0.625	(0.000)	0.252	(0.000)
Number of observations		143		114		144		100
Adjusted R-squared		0.870		0.311		0.772		0.870
F-statistic		56.51		4.40		33.22		48.23

Table VIII (cont.)
Analysis of Posting Activity Characteristics

Panel B: Analysis of Sentiment

Dependent variable: Sentiment level	AMZN		DAL		GMGC		GWRX	
	Coeff	p-value	Coeff	p-value	Coeff	p-value	Coeff	p-value
Constant	-30.262	(0.040)	-0.468	(0.745)	-5.927	(0.633)	-0.550	(0.209)
<u>Posting Activity</u>								
Lagged sentiment	0.271	(0.002)	-0.087	(0.337)	0.134	(0.080)	0.189	(0.017)
Number of postings	0.034	(0.035)	0.125	(0.000)	0.034	(0.004)	0.061	(0.001)
<u>Stock Market Activity</u>								
Trading volume	0.000	(0.928)	0.000	(0.203)	0.000	(0.190)	0.000	(0.076)
Stock return	104.340	(0.004)	12.635	(0.040)	41.548	(0.045)	-0.734	(0.689)
Lagged stock return	122.188	(0.000)	5.438	(0.319)	58.101	(0.001)	1.660	(0.282)
Equal weighted stock market return	208.854	(0.294)	15.257	(0.331)	195.921	(0.074)	-21.798	(0.165)
Implied volatility	13.120	(0.432)	0.073	(0.979)	2.323	(0.813)		
<u>News Activity</u>								
Average news sentiment	1.387	(0.194)	-0.043	(0.817)	6.908	(0.095)	0.239	(0.726)
Abnormal news stories	0.285	(0.483)	0.065	(0.252)	0.241	(0.862)	-0.371	(0.139)
New posters in the last seven days	0.034	(0.051)	0.017	(0.641)	0.030	(0.220)	0.014	(0.247)
Number of observations	143		144		144		144	
Adjusted R-squared	0.733		0.096		0.425		0.645	
F-statistic	28.890		2.080		8.550		20.950	

Table VIII (cont.)
Analysis of Posting Activity Characteristics

Panel C: Analysis of Disagreement

Dependent variable: Disagreement in posting activity	AMZN		DAL		GMGC		GWRX	
	Coeff	p-value	Coeff	p-value	Coeff	p-value	Coeff	p-value
Constant	0.560	(0.000)	0.434	(0.211)	0.716	(0.000)	0.341	(0.001)
<u>Posting Activity</u>								
Lagged disagreement	0.377	(0.000)	0.106	(0.295)	0.169	(0.046)	0.086	(0.416)
Number of postings	0.000	(0.774)	0.035	(0.000)	0.000	(0.035)	0.010	(0.005)
Lagged number of postings	0.000	(0.819)	0.000	(0.980)	0.000	(0.615)	-0.001	(0.665)
<u>Stock Market Activity</u>								
Trading volume	0.000	(0.682)	0.000	(0.709)	0.000	(0.029)	0.000	(0.128)
Stock return	-0.125	(0.228)	0.016	(0.991)	-0.205	(0.157)	0.363	(0.379)
Lagged stock return	-0.090	(0.328)	-2.205	(0.078)	-0.230	(0.058)	-0.464	(0.132)
Equal weighted stock market return	0.138	(0.811)	1.835	(0.624)	-1.470	(0.061)	-2.083	(0.613)
Implied volatility	0.001	(0.984)	-0.445	(0.480)	-0.026	(0.704)		
<u>News Activity</u>								
Average news sentiment	-0.003	(0.271)	-0.035	(0.410)	-0.033	(0.245)	0.098	(0.504)
Abnormal news stories	-0.001	(0.482)	-0.007	(0.575)	0.003	(0.774)	0.083	(0.096)
New posters in the last seven days	0.000	(0.399)	0.002	(0.822)	0.000	(0.543)	-0.001	(0.647)
Number of observations	143		114		144		100	
Adjusted R-squared	0.262		0.13		0.228		0.057	
F-statistic	4.35		2.13		3.81		143	

Table IX
Correlations Between Information and Financial Markets

Panel A reports the same-day contemporaneous correlations between information variables and financial market variables (return, excess return, volume, implied volatility, intraday volatility, and bid-ask spreads.) Panel B reports the autocorrelations for the financial market variables. Each cell represents the median correlation among the four firms, the median p-value among the four firms, and the identity of the firms for which the p-value is .05 or better where a=Amazon, d=Delta, m= General Magic and x = Geoworks.

Panel A: Contemporaneous Correlations: Information and Financial Markets

	Press Release	Filing	Analyst Revision	# Major news stories	News Sentiment	Number of Postings	New Posters last 7days	Analyst Sentiment	Sentiment	Disagreement	Open to close return	Close to close return	Equal Weighted Mkt.Return	Share turnover	Intraday volatility	Implied Volatility	Bid-Ask %	
Filing	0.024 0.768	1.000																
Analyst Revision	-0.043 0.352	-0.032 0.660 a	1.000															
No of major news stories	0.601 0.000 admx	-0.019 0.821	0.072 0.383 a	1.000														
News Sentiment	0.327 0.000 dmx	-0.004 0.779	0.013 0.879	0.390 0.000 admx	1.000													
Number of Postings	0.058 0.485 m	-0.013 0.519	0.150 0.070 a	0.284 0.010 amx	-0.021 0.209 ax	1.000												
New posters in last 7 days	0.040 0.647 m	-0.095 0.302	0.111 0.179 x	0.222 0.047 ax	0.010 0.384 a	0.728 0.000 admx	1.000											
Analyst Sentiment	-0.089 0.284 a	0.000 0.437 a	-0.518 0.000 ad	-0.101 0.224 a	-0.012 0.890 a	0.108 0.067 a	-0.128 0.012 ax	1.000										
Sentiment	0.055 0.571 m	0.014 0.364	0.026 0.753	0.251 0.006 amx	0.096 0.014 adx	0.562 0.000 admx	0.500 0.000 amx	0.080 0.069 a	1.000									
Disagreement	0.023 0.785	0.019 0.367	0.100 0.254	-0.075 0.248 a	-0.107 0.252 a	0.269 0.001 admx	0.111 0.047 ax	-0.028 0.110	-0.341 0.120 am	1.000								
Open-to-close return	0.030 0.520 x	-0.010 0.563	-0.040 0.631	0.176 0.054 mx	0.147 0.296 mx	-0.044 0.601	-0.002 0.652	0.119 0.152	0.103 0.217 d	-0.048 0.577 m	1.000							
Close-to-close	0.088 0.375 mx	-0.018 0.727	0.022 0.793	0.248 0.012 amx	0.265 0.023 amx	-0.066 0.521	0.038 0.623	0.029 0.705	0.210 0.011 adm	-0.093 0.340 am	0.902 0.000 admx	1.000						
Equal weighted market return	-0.026 0.726	0.017 0.760	0.090 0.275 x	-0.031 0.708	-0.028 0.745	-0.040 0.383 m	0.051 0.409	0.030 0.715 a	0.133 0.146 ad	-0.070 0.410 m	0.243 0.003 adm	0.371 0.000 adm	1.000					
Share turnover	0.196 0.041 mx	0.035 0.446 m	0.001 0.988	0.407 0.000 admx	0.209 0.152 am	0.495 0.000 amx	0.336 0.001 amx	-0.028 0.534	0.321 0.000 amx	0.018 0.172 a	0.184 0.050 ax	0.217 0.017 amx	-0.043 0.615	1.000				
Intraday volatility	0.028 0.737	0.008 0.561	0.073 0.457	0.171 0.040 am	0.119 0.249 am	0.120 0.186 am	-0.051 0.175 ad	-0.073 0.380	0.028 0.523 a	-0.010 0.597	0.127 0.207 m	0.089 0.114 am	-0.169 0.116 am	0.436 0.000 adm	1.000			
Implied volatility	0.017 0.080	-0.023 0.785	0.069 0.451	0.049 0.555 a	-0.023 0.633 a	0.030 0.227 a	-0.164 0.000 adm	-0.215 0.073 a	-0.031 0.332 a	-0.060 0.495	-0.151 0.069 m	-0.190 0.022 am	-0.024 0.769 m	0.185 0.025 ad	0.450 0.000 adm	1.000		
Bid-Ask %	-0.088 0.181 m	-0.094 0.258 x	-0.014 0.865	-0.259 0.002 amx	-0.063 0.435 a	-0.367 0.000 admx	-0.395 0.000 amx	-0.010 0.902	-0.287 0.001 amx	0.056 0.168 ax	-0.101 0.226 a	-0.149 0.079 am	-0.038 0.633 a	-0.276 0.001 amx	0.360 0.001 adm	0.257 0.002 dm	1.000	
Number of jumps	0.085 0.311 x	-0.014 0.533	0.071 0.389	0.136 0.113 x	0.128 0.181 am	0.166 0.083 mx	0.206 0.015 amx	-0.154 0.062	0.145 0.338 ax	0.015 0.386 x	0.042 0.404 x	0.072 0.394 x	0.047 0.296	0.186 0.031 mx	-0.103 0.073 dm	0.057 0.214	-0.185 0.028 admx	

Table IX (continued)
Correlations Between Information and Financial Markets

Panel B: Autocorrelations for Financial Markets Variables

Autocorrelations	AMAZON		DELTA		GENERAL MAGIC		GEOWORKS	
	Coeff.	p-value	Coeff.	p-value	Coeff.	p-value	Coeff.	p-value
Open/close return	-0.045	0.624	0.249	0.005	-0.098	0.368	-0.033	0.728
Close/close return	0.102	0.279	0.223	0.008	0.026	0.831	-0.019	0.820
Share turnover	0.544	0.000	0.412	0.000	0.542	0.000	0.515	0.000
Intraday volatility	0.602	0.000	0.508	0.000	0.486	0.000	0.265	0.042
Implied volatility	0.702	0.000	0.866	0.000	0.529	0.000	n/a	n/a
Bid/ask spread	0.820	0.000	0.795	0.000	0.724	0.000	0.687	0.000
Number of jumps	0.049	0.571	0.249	0.037	0.215	0.040	0.315	0.019

Table X
Explaining Returns Using Information Variables

The dependent variables in this table are the market adjusted return (raw return less S&P 500) earned by each of the four firms during the period 7/1/98-1/31/99. For Panel A and C, this return is calculated from the close of the market the prior trading day to close of the market the current day, while for Panels B and D, it represents the return from the open of the market to the close of the market. Panels A and B use contemporaneous information, i.e., information from the same day as the return. This includes all filings, press releases and news stories with that date and postings for the same period for which returns are calculated. Panels C and D use lagged information, i.e., filings, press releases and news stories from the prior day, and postings that preceded the return computation period.

Panel A

Dependent variable: market adjusted returns, close-to-close All independent variables measured the same "day."	AMZN		DAL		GMGC		GWRX	
	Coeff	p-value	Coeff	p-value	Coeff	p-value	Coeff	p-value
Constant	0.142	(0.000)	0.023	(0.066)	0.175	(0.000)	0.009	(0.646)
Volatility	-0.152	(0.000)	-0.054	(0.043)	-0.153	(0.000)	0.000	(0.991)
<u>Announcements</u>								
Press Release	0.037	(0.136)	-0.001	(0.891)	-0.115	(0.000)	-0.342	(0.000)
Filing	-0.002	(0.877)	-0.011	(0.110)	0.003	(0.890)	-0.012	(0.733)
Revision	-0.002	(0.815)	0.001	(0.899)			0.004	(0.856)
<u>News Activity</u>								
Abnormal number of stories	0.002	(0.113)	0.001	(0.548)	0.028	(0.008)	0.016	(0.381)
News sentiment	-0.001	(0.537)	-0.002	(0.576)	0.020	(0.076)	0.056	(0.006)
<u>Posting Activity</u>								
Abnormal number of posts	-3.0E-06	(0.934)	-0.001	(0.365)	-1.0E-04	(0.067)	6.5E-05	(0.889)
Sentiment level	5.0E-04	(0.006)	0.002	(0.174)	6.2E-04	(0.071)	-0.004	(0.249)
<u>Interaction terms: Press release dummy interacted with</u>								
...abnormal number of stories	-4.3E-04	(0.906)	-0.002	(0.447)	0.018	(0.299)	0.142	(0.000)
...sentiment of stories	-0.016	(0.007)	0.003	(0.392)	0.017	(0.249)	-0.030	(0.398)
...abnormal number of posts	2.0E-04	(0.145)	-0.001	(0.474)	-9.0E-05	(0.336)	0.008	(0.194)
...sentiment of posts	0.001	(0.008)	0.001	(0.581)	0.001	(0.061)	0.022	(0.087)
Number of observations	146		147		147		105	
Adjusted R-squared	0.1892		0.1009		0.4425		0.7541	
F-statistic	3.42		1.82		8.23		16.13	

Panel B

Dependent variable: market adjusted returns, open-to-close
 All independent variables measured the same "day."

	AMZN		DAL		GMGC		GWRX	
	Coeff	p-value	Coeff	p-value	Coeff	p-value	Coeff	p-value
Constant	0.111	(0.000)	0.008	(0.576)	0.158	(0.000)	-0.009	(0.660)
Volatility	-0.107	(0.002)	-0.018	(0.530)	-0.143	(0.000)	0.004	(0.598)
<u>Announcements</u>								
Press Release	0.007	(0.710)	0.000	(0.961)	-0.101	(0.000)	-0.326	(0.000)
Filing	-0.010	(0.405)	-0.011	(0.093)	0.000	(0.997)	0.066	(0.065)
Revision	-0.008	(0.403)	-0.004	(0.478)			0.026	(0.575)
<u>News Activity</u>								
Abnormal number of stories	0.003	(0.025)	0.001	(0.543)	0.023	(0.015)	0.025	(0.196)
News sentiment	-0.003	(0.030)	-0.001	(0.738)	0.006	(0.591)	0.052	(0.001)
<u>Posting Activity</u>								
Abnormal number of posts	0.000	(0.238)	0.001	(0.510)	0.000	(0.375)	-0.001	(0.043)
Sentiment level	0.001	(0.001)	0.003	(0.403)	0.001	(0.021)	-0.007	(0.326)
<u>Interaction terms: Press release dummy interacted with</u>								
...abnormal number of stories	-0.001	(0.694)	-0.002	(0.407)	0.010	(0.472)	0.113	(0.002)
...sentiment of stories	-0.003	(0.484)	0.002	(0.566)	0.018	(0.184)	0.004	(0.855)
...abnormal number of posts	0.001	(0.057)	-0.002	(0.175)	0.000	(0.767)	0.005	(0.386)
...sentiment of posts	0.001	(0.188)	0.001	(0.785)	0.002	(0.188)	0.039	(0.011)
Number of observations	146		147		147		105	
Adjusted R-squared	0.229		0.070		0.371		0.745	
F-statistic	7.02		1.12		9.85		23.14	

Panel C

Dependent variable: market adjusted returns, close-to-close
All independent variables measured the prior "day".

	AMZN		DAL		GMGC		GWRX	
	Coeff	p-value	Coeff	p-value	Coeff	p-value	Coeff	p-value
Constant	0.048	(0.695)	-0.010	(0.480)	-0.057	(0.240)	-0.020	(0.555)
Volatility	-0.043	(0.215)	0.021	(0.454)	0.047	(0.229)	0.017	(0.319)
<u>Announcements</u>								
Press Release	0.005	(0.844)	0.002	(0.791)	-0.028	(0.355)	0.042	(0.342)
Filing	-0.010	(0.510)	0.001	(0.923)	-0.010	(0.548)	-0.064	(0.377)
Revision	-0.009	(0.558)	-0.001	(0.837)			0.035	(0.164)
<u>News Activity</u>								
Abnormal number of stories	0.000	(0.723)	0.000	(0.965)	0.006	(0.674)	-0.010	(0.685)
News sentiment	0.000	(0.914)	-0.001	(0.852)	0.008	(0.636)	0.016	(0.238)
<u>Posting Activity</u>								
Abnormal number of posts	0.000	(0.450)	-0.001	(0.169)	0.000	(0.813)	-0.001	(0.196)
Sentiment level	0.000	(0.550)	0.002	(0.208)	0.000	(0.550)	-0.003	(0.367)
<u>Interaction terms: Press release dummy interacted with</u>								
...abnormal number of stories	-0.001	(0.724)	-0.001	(0.728)	-0.020	(0.224)	-0.048	(0.065)
...sentiment of stories	0.003	(0.781)	-0.002	(0.602)	0.000	(0.986)	0.039	(0.092)
...abnormal number of posts	0.000	(0.981)	0.001	(0.252)	0.000	(0.326)	0.005	(0.105)
...sentiment of posts	0.000	(0.142)	-0.004	(0.142)	0.000	(0.901)	-0.008	(0.438)
Number of observations	145		146		146		104	
Adjusted R-squared	0.050		0.072		0.118		0.075	
F-statistic	0.890		1.200		2.130		17.210	

Panel D

Dependent variable: market adjusted returns, open-to-close
All independent variables measured the prior "day".

	AMZN		DAL		GMGC		GWRX	
	Coeff	p-value	Coeff	p-value	Coeff	p-value	Coeff	p-value
Constant	0.036	(0.366)	-0.019	(0.164)	-0.078	(0.064)	-0.029	(0.342)
Volatility	-0.027	(0.570)	0.041	(0.120)	0.058	(0.071)	0.017	(0.300)
<u>Announcements</u>								
Press Release	0.005	(0.791)	-0.001	(0.808)	0.003	(0.898)	0.020	(0.686)
Filing	-0.019	(0.167)	0.002	(0.756)	-0.010	(0.477)	-0.106	(0.108)
Revision	-0.012	(0.359)	-0.002	(0.710)			-0.019	(0.488)
<u>News Activity</u>								
Abnormal number of stories	0.001	(0.596)	-0.001	(0.566)	-4.5E-04	(0.966)	-0.020	(0.388)
News sentiment	0.001	(0.634)	0.000	(0.975)	0.008	(0.633)	0.041	(0.008)
<u>Posting Activity</u>								
Abnormal number of posts	4.2E-05	(0.479)	0.001	(0.465)	1.0E-05	(0.920)	-0.003	(0.571)
Sentiment level	9.9E-06	(0.967)	0.002	(0.367)	7.1E-05	(0.889)	-0.005	(0.349)
<u>Interaction terms: Press release dummy interacted with</u>								
...abnormal number of stories	0.003	(0.354)	0.000	(0.814)	-0.025	(0.045)	-0.014	(0.735)
...sentiment of stories	-0.001	(0.849)	-0.002	(0.431)	-0.009	(0.608)	-0.003	(0.929)
...abnormal number of posts	-3.0E-05	(0.664)	-0.001	(0.378)	0.000	(0.400)	0.003	(0.642)
...sentiment of posts	-0.001	(0.224)	0.002	(0.563)	0.000	(0.995)	3.6E-04	(0.981)
Number of observations	145		146		146		104	
Adjusted R-squared	0.069		0.082		0.118		0.083	
F-statistic	1.950		1.290		4.170		37.840	