

WHICH NEWS MOVES STOCK PRICES? A TEXTUAL ANALYSIS¹

Jacob Boudoukh², Ronen Feldman³, Shimon Kogan⁴ and Matthew Richardson⁵

December 11, 2012

Abstract:

A basic tenet of financial economics is that asset prices change in response to unexpected fundamental information. Since Roll's (1988) provocative presidential address that showed little relation between stock prices and news, however, the finance literature has had limited success reversing this finding. This paper revisits this topic in a novel way. Using advancements in the area of textual analysis, we are better able to identify relevant news, both by type and by tone. Once news is correctly identified in this manner, there is considerably more evidence of a strong relationship between stock price changes and information. For example, market model R^2 s are no longer the same on news versus no news days (i.e., Roll's (1988) infamous result), but now are 16% versus 33%; variance ratios of returns on identified news versus no news days are 120% higher versus only 20% for unidentified news versus no news; and, conditional on extreme moves, stock price reversals occur on no news days, while identified news days show an opposite effect, namely a strong degree of continuation. A number of these results are strengthened further when the tone of the news is taken into account by measuring the positive/negative sentiment of the news story.

¹ Corresponding author: Shimon Kogan, GSB 5.159, McCombs School of Business, University of Texas at Austin, 1 University Station, B6600, Austin, TX 78712, Tel: +1 (512) 232-6839, email shimon.kogan@austin.utexas.edu. We would like to thank John Griffin and seminar participants at the University of Texas, Austin, and Stern NYU for their comments and suggestions.

² Arison School of Business, IDC, Israel

³ School of Business Administration, The Hebrew University

⁴ McCombs School of Business, University of Texas at Austin

⁵ Stern School of Business, New York University and NBER

I. Introduction

A basic tenet of financial economics is that asset prices change in response to unexpected fundamental information. Early work, primarily through event studies, seemed to confirm this hypothesis. (See, for example, Ball and Brown (1968) on earning announcements, Fama, Fisher, Jensen and Roll (1969) on stock splits, Mandelker (1974) on mergers, Aharony and Swary (1980) on dividend changes, and Asquith and Mullins (1986) on common stock issuance, among many others.) However, since Roll's (1988) provocative presidential address that showed little relation between stock prices and news (used as a proxy for information), the finance literature has had limited success at showing a strong relationship between prices and news, e.g., also see Shiller (1981), Cutler, Poterba and Summers (1989), Campbell (1991), Berry and Howe (1994), Mitchell and Mulherin (1994), and Tetlock (2007), to name a few. The basic conclusion from this literature is that stock price movements are largely described by irrational noise trading or through the revelation of private information through trading.

In this paper, we posit an alternative explanation, namely that the finance literature has simply been doing a poor job of identifying true and relevant news. In particular, common news sources for companies such as those in the Wall Street Journal stories and Dow Jones News Service, *et cetera*, contain many stories which are not relevant for information about company fundamentals. The problem of course is for the researcher to be able to parse through which news stories are relevant and which are not. Given that there are hundreds of thousands, possibly millions, of news stories to work through, this presents a massive computational problem for the researcher. Fortunately, advances in the area of textual analysis allow for better identification of relevant news, both by type and tone. This paper employs one such approach based on an information extraction platform (Feldman, Rosenfeld, Bar-Haim and Fresko (2011), denote Feldman et al. (2011)).

There is a growing literature in finance that uses textual analysis to try and convert qualitative information contained in news stories and corporate announcements into a quantifiable measure by analyzing the positive or negative tone of the information. One of the earliest papers is Tetlock (2007) who employs the General Inquirer, a well-known

textual analysis program, alongside the Harvard-IV-4 dictionary (denote IV-4) to calculate the fraction of negative words in the *Abreast of the Market* Wall Street Journal column. Numerous papers have produced similar analyses to measure a document's tone in a variety of financial and accounting contexts, including Davis, Piger, and Sedor (2006), Engelberg (2008), Tetlock, Saar-Tsechansky and Macskassy (2008), Demers and Vega (2010), Feldman, Govindaraj, Livnat and Segal (2010), and Loughran and McDonald (2011), among others. While all these papers support the idea that news, transformed into a sentiment measure, have important information for stock prices, none represent a significant shift in thinking about the overall relation between stock prices and information. Part of the reason is that, other than refinements of IV-4 for financial applications (e.g., Engelberg (2008) and Loughran and McDonald (2011)), the textual analysis methodology is similar.⁶

The aforementioned textual analysis methodology (Feldman et al. (2011)) employed in this paper is quite different. It combines not only a dictionary-based sentiment measure as in Tetlock (2007) and Loughran and McDonald (2011), but also an analysis of phrase-level patterns to further break down the tone of the article and a methodology for identifying relevant events for companies (broken down into 14 categories and 56 subcategories). While the methodology is for the most part based on sets of rules (as opposed to say machine learning),⁷ the implementation employs the commonly used technique of running and refining these rules on a subset of training articles. This procedure greatly improves the accuracy. In terms of relating stock prices to news, the methodology provides a number of advantages over existing approaches. In particular, over the sample period 2000-2009 for all S&P500 companies, the Dow Jones Newswire produces over 1.9M stories, only 50% of which we identify as relevant events. As discussed shortly, this breakdown into identified and unidentified news makes a massive difference in terms of our understanding of stock price changes and news. Moreover, employing a more sophisticated textual analysis methodology than one based on a simple count of positive versus negative words further improves the results. In other words, when we can identify the news, and more accurately

⁶ Some exceptions include Li (2010), Hanley and Hoberg (2011), and Grob-Klubmann and Hautsch (2011) who all use some type of machine learning-based application.

⁷ Some parts of the implementation, such as locating names of companies and individuals, employ machine-learning technology, that is, the use of statistical patterns to infer context.

evaluate its tone, there is considerably more evidence of a strong relationship between stock price changes and information.

This paper documents several new results. First, and foremost, using the aforementioned methodology that allows us to automatically and objectively classify articles into topics (such as analyst recommendations, financial information, acquisitions and mergers, etc.), we compare days with no-news, unidentified news, and identified news on several dimensions. In particular, we show that stock-level volatility is similar on no-news days and unidentified news days, consistent with the idea that the intensity and importance of information arrival is the same across these days. In contrast, on identified news days, the volatility of stock prices is over double that of other days. This evidence is provided further support by noting that identified news days are 31-34% more likely to be associated with extreme returns (defined by the bottom and top 10% of the return distribution) while unidentified and no news days are slightly more likely to be associated with moderate day returns (in the middle 30-70% range of the returns distribution). A major finding is that when we revisit Roll's (1988) R^2 methodology and estimate the R^2 from a market model regression for all days and for unidentified news days, consistent with his results, R^2 levels are the same for all days and for unidentified news days. However, when we estimate the same model over just identified news days, the R^2 drops dramatically from an overall median of 28% to 16%, the precise result that Roll (1988) was originally looking for in his work.

Second, beyond the parsing of news into identified events and unidentified news, the methodology provides a measure of article tone (that is, positive versus negative) that builds on Tetlock (2007) and others. As mentioned above, we perform both an analysis of phrase-level patterns (e.g., by narrowing down to the relevant body of text, taking into account phrases and negation, etc.) and employ a dictionary of positive and negative words more appropriate for a financial context. Using this more advanced methodology, in contrast to a simple word count, we show that our measure of tone can substantially increase R^2 on identified news days, but not on unidentified news days, again consistent with the idea that identified news days contain price-relevant information. Another finding is that tone variation across topics and within topics is consistent with one's intuition. For example,

deals and partnership announcements tend to be very positive while legal announcements tend to be negative. Analyst recommendations and financial information, on average, tend to be more neutral, but tend to have greater variation within the topic. Moreover, some of these topics are much more likely to appear on extreme return days (e.g., analyst recommendations, financials) while others are not (e.g., partnership). This suggests that different topics may have different price impact. Finally, the results are generally consistent with a positive association between daily returns and daily tone, with this relationship being more pronounced using the methodology presented here than of the more standard simple word count.

Third, the above discussion contemporaneously relates relevant news to stock price changes. An interesting issue is whether the differentiation between identified and unidentified news has forecast power for stock price changes. There is now a long literature, motivated through work in behavioral finance and limits of arbitrage, that stock prices tend to underreact or overreact to news, depending on the circumstances (see, for example, Hirshleifer (2000), Chan (2003), Vega (2006), Gutierrez and Kelley (2008), Tetlock, Saar-Tsechansky, and Macskassy (2008), and Tetlock (2010)). This paper documents an interesting result in the context of the breakdown of Dow Jones news into identified and unidentified news. Specifically, conditional on extreme moves, stock price reversals occur on no news and unidentified news days, while identified news days show an opposite effect, namely a small degree of continuation. That news days tend to be associated with future continuation patterns while no news days see reversals is consistent with (1) our methodology correctly parsing out relevant news, and (2) a natural partition between underreaction and overreaction predictions in a behavioral context. As an additional test, we perform an out-of-sample exercise based on a simple portfolio strategy. The resulting gross Sharpe ratio of 1.7 illustrates the strength of these results.

While our paper falls into the area of the literature that focuses on using textual analysis to address the question of how prices are related to information, the two most closely related papers to ours, Griffin, Hirschey and Kelly (2011) and Engle, Hansen and Lunde (2011), actually lie outside this textual analysis research area. Griffin, Hirschey and Kelly (2011) cross-check global news stories against earnings announcements to try and uncover relevant

events. Engle, Hansen and Lunde (2011) utilize the Dow Jones Intelligent Indexing product to match news and event types for a small set of (albeit large) firms. While the focus of each of these papers is different (e.g., Griffin, Hischey and Kelly (2011) stress cross-country differences and Engle, Hansen and Lunde (2011) emphasizing the dynamics of volatility based on information arrival), both papers provide some evidence that better information processing by researchers will lead to higher R^2 s between prices and news.

This paper is organized as follows. Section II describes the data employed throughout the study. Of special interest, we describe in detail the textual analysis methodology for inferring content and tone from news stories. Section III provides the main results of the paper, showing a strong relationship between prices and news, once the news is appropriately identified. In section IV, we reexamine a number of results related to the existing literature measuring the relationship between stock sentiment and stock returns. Section V discusses and analyzes the forecasting power of the textual analysis methodology for future stock prices, focusing on continuations and reversals after large stock price moves. Section VI concludes.

II. Data Description and Textual Analysis Methodology

A. Textual Analysis

With the large increase in the amount of daily news content on companies over the past decade, it should be no surprise that the finance literature has turned to textual analysis as one way to understand how information both arrives to the marketplace and relates to stock prices of the relevant companies. Pre mainstream finance, early work centered on document-level sentiment classification of news articles by employing pre-defined sentiment lexicons.⁸ The earliest paper in finance that explores textual analysis is Antweiler and Frank (2005) who employ language algorithms to analyze internet stock message boards posted on “Yahoo Finance”. Much of the finance literature, however, has focused on word counts based on dictionary-defined positive versus negative words.

⁸ See, for example, Lavrenko, Schmill, Lawrie, Ogilvie, Jensen, and Allan (2000), Das and Chen (2007) and Devitt and Ahmad (2007), among others. Feldman and Sanger (2006) provide an overview.

For example, one of the best known papers is Tetlock (2007). Tetlock (2007) employs the General Inquirer, a well-known textual analysis program, alongside the Harvard-IV-4 dictionary to calculate the fraction of negative words in the *Abreast of the Market* Wall Street Journal column. A plethora of papers, post Tetlock (2007), apply a similar methodology to measure the positive versus negative tone of news across a wide variety of finance and accounting applications.⁹ Loughran and McDonald (2011), in particular, is interesting because they refine IV-4 to more finance-centric definitions of positive and negative words.¹⁰

More recently, an alternative approach to textual analysis in finance and accounting has been offered by Li (2010), Hanley and Hoberg (2011), and Grob-Klubmann and Hautsch (2011). These authors employ machine learning-based applications to decipher the tone and therefore the sentiment of news articles.¹¹ The basic approach of machine learning is not to rely on written rules per se, but instead allow the computer to apply statistical methods to the documents in question. In particular, supervised machine learning uses a set of training documents (that are already classified into a set of predefined categories) to generate a statistical model that can then be used to classify any number of new unclassified documents. The features that represent each document are typically the words that are inside the document (bag of words approach).¹² While machine learning has generally come to dominate rules-based classification approaches (that rely solely on human-generated rules), there are disadvantages, especially to the extent that machine learning classifies documents in a non transparent fashion that can lead to greater misspecification.

In this paper, in contrast, classification is not used at all. Instead, a rule based information extraction approach is employed, appealing to recent advances in the area of textual analysis (Feldman et al. (2011)). That is, we extract event instances out of the text based on a set of

⁹ See, for example, Davis, Piger, and Sedor (2006), Engelberg (2008), Tetlock, Saar-Tsechansky and Macskassy (2008), Kothari, Li and Short (2009), Demers and Vega (2010), Feldman, Govindaraj, Livnat and Segal (2010), and Loughran and McDonald (2011), among others.

¹⁰ For a description and list of the relevant words, see http://nd.edu/~mcdonald/Word_Lists.html.

¹¹ Other papers, e.g., Kogan et. al. (2011), use machine learning to link features in the text to firm risk.

¹² See Manning and Schutze (1999) for a detailed description and analysis of machine learning methods.

predefined rules. For instance, when we extract an instance of an *Acquisition* event, we find who is the acquirer, who is the acquiree, optionally what was the amount of money paid for the acquisition, and so forth. Feldman et al. (2011) employ a proprietary information extraction platform specific to financial companies, which they denote *The Stock Sonar* (TSS), and which is available on commercial platforms like Dow Jones. This textual analysis methodology differs from current rules-based applications in finance in three important ways.

First, TSS also adheres to a dictionary-based sentiment analysis. In particular, the method uses as a starting point the dictionaries used by Tetlock (2007) and Loughran and McDonald (2011), but then augments it by adding and subtracting from these dictionaries. Beyond the usual suspects of positive and negative words, a particular weight is placed on sentiment modifiers such as “highly”, “incredible”, “huge”, *et cetera* versus lower emphasis modifiers such as “mostly” and “quite” versus opposite modifiers such as “far from”. For example, amongst the modifiers, the most commonly used word in the context of the S&P 500 companies over the sample decade is “highly”, appearing over 6,000 times. A typical usage is:

By the end of 2005 Altria is highly likely to take advantage of the provisions of the American Jobs Creation Act of 2004. (Dow Jones Newswire, at 18:16:25 on 03-15-2005.)

These words were adjusted to the domain of financial news by adding and removing many terms, depending on the content of thousands of news articles. Specifically, for developing these lexicons and rules (to be discussed in further detail below), a benchmark consisting of thousands of news articles was manually tagged. The benchmark was divided into a training set (providing examples) and a test set (kept blind and used for evaluating the progress of the methodology). The rulebook was run repeatedly on the system on thousands of articles, each time revised and iterated upon until the precision was satisfactory (e.g., >90%).

Second, this same approach was used to create a set of rules to capture phrase-level sentiments. Current systems employed in finance so far have operated for the most part at the word level, but compositional expressions are known to be very important in textual analysis. For example, one of the best known illustrations involve double negatives such as

“reducing losses” which of course has a positive meaning, yet would likely yield a negative word count in most schemes. For example, combination phrases with “reducing” appear over 1,200 times for the S&P 500 companies in our sample, such as:

Mr. Dillon said the successful execution of Kroger's strategy produced strong cash flow, enabling the Company to continue its "financial triple play" of reducing total debt by nearly \$400 million, repurchasing \$318.7 million in stock, and investing \$1.6 billion in capital projects. (Dow Jones Newswire, at 13:19:20 on 03-08-2005.)

Other examples include words like “despite” which tend to connect both positive and negative information. For example, the word “despite” appears over 3,600 times across our S&P 500 sample. A typical sentence is:

Wells Fargo & Co.'s (WFC) fourth-quarter profit improved 10% despite a continued slowdown in the banking giant's once-booming home mortgage business. (Dow Jones Newswire, at 12:04:21 on 01-18-2005.)

A large number of expressions of this sort are considered jointly with the word dictionary to help better uncover the sentiment of the article.

Third, and most important, TSS sorts through the document and parses out the meaning of the document in the context of possible events relevant to companies, such as new product launches, lawsuits, analyst coverage, financial news, mergers, *et cetera*. The initial list of events were chosen to match commercial providers such as CapitalIQ but were augmented by events likely to impact stock prices. This process led to a total of 14 event categories and 56 subcategories within events. For example, the events fall into one of the following categories: *Analyst Recommendations*, *Financial*, *Financial Pattern*, *Acquisition*, *Deals*, *Employment*, *Product*, *Partnerships*, *Inside Purchase*, *Facilities*, *Legal*, *Award*, *Stock Price Change* and *Stock Price Change Pattern*. Consider the *Analyst Recommendation* category.¹³ In terms of subcategories, it contains nine subcategories, including *analyst expectation*, *analyst opinion*, *analyst rating*, *analyst recommendation*, *credit - debt rating*, *fundamental analysis*, *price target*, etc.¹⁴

¹³ In practice, the categories, defined in terms of *Pattern*, represent cases in which an event was identified but the reference entity was ambiguous.

¹⁴ For a complete list of the categories and subcategories, see <http://shimonkogan.tumblr.com>.

Because events are complex objects to capture in the context of textual analysis of documents, considerable effort was applied to write rules that can take any news story and then link the name of a company to both the identified event and sentiment surrounding the event. For example, a total of 4,411 rules were written to identify companies with the various event categories and subcategories. Because every event is phrased in different ways, the process of matching companies to identified events is quite hard. For example, consider the following three sentences in the “Deals” category for different companies in the early January, 2005 period:

1. Northrop Grumman Wins Contract to Provide Navy Public Safety. (Dow Jones Newswire, at 17:02:21 on 01-03-2005.)
2. A deal between UBS and Constantia could make sense, Christian Stark, banks analyst at Cheuxvreux wrote in a note to investors. (Dow Jones Newswire, at 10:17:26 on 01-03-2005.)
3. Jacobs Engineering Group Inc. (NYSE:JEC) announced today that a subsidiary company received a contract to provide engineering and science services to NASA's Johnson Space Center (JSC) in Houston, Texas. (Dow Jones Newswire, at 12:45:03 on 01-04-2005.)

The methodology behind TSS managed to get a recall of above 85% by first identifying candidate sentences that may contain events (based on the automatic classification of the sentences) and then marking these sentences as either positive or negative for each event type (through quality assurance (QA) engineers). The tagged sentences were then used as updated training data for the sentence classifier and the QA cycle was repeated.

An additional difficulty is that sentences which identify the events may not mention the specific name of the company which is the subject of the sentence. The methodology underlying TSS is able to resolve these indirect references by analyzing the flow of the article. Examples of typical sentences are

1. For Fiscal Year 2006, the company announced that it is targeting pro forma earnings per share growth of 22 to 28 percent or \$0.76 to \$0.80 per share. (Dow Jones Newswire, at 12:06:01 on 01-26-2005.)
2. Based on results from November and December periods, the retailer expects fourth-quarter earnings to come in towards the end of previous guidance. (Dow Jones Newswire, at 13:13:15 on 01-06-2005.)

In the former case, the article referred to Oracle, while in the latter case the article referred to J.C. Penney. The TSS methodology was able to determine that the company mentioned in

the previous sentence was also the subject of this sentence and hence J.C. Penney could be tied to this event with negative sentiment. More generally, for each company, TSS tries to identify the exact body of text within the document that refers to that company so that the sentiment calculations will be based only on words and phrase that are directly associated with that company. For example, one technique is to consider only words within a range of the mention of the main company in the document. Another is to avoid historical events cited in documents by capturing past versus present tense. Like the document sentiment analysis, a training set of documents were used to refine the rulebook for events and then evaluated against a test set.

B. Data Description and Summary

The primary dataset used in this paper consists of all documents that pass through the Dow Jones Newswire from January 1, 2000 to December 31, 2009. For computational reasons, we limit ourselves to the S&P500 companies with at least 20 trading days at the time the news stories are released. Over the sample period, the dataset therefore includes at some time or another 791 companies. To avoid survivorship bias, we include in the analysis all stocks in the index as of the first trading day of each year. We obtain total daily returns from CRSP.

TSS methodology described in II.A processes each article separately and generates an output file in which each article/stock/day is represented as an observation. For each of these observations, TSS reports the total number of words in the article, the number of relevant words in the article, the event (and sub-event) identified, and the number of positive and negative features as identified by TSS. For the same set of articles we also count the number of positive and negative words using IV-4 (see, for example, Tetlock (2007)).¹⁵ In terms of sentiment score, after parsing out only relevant sentences, and

¹⁵ It should be pointed out that Tetlock (2007), and others that followed, do not apply a word count blindly to IV-4. For example, Tetlock (2007) counts words in each of the 77 categories in IV-4 and then collapses this word count into a single weighted count based on a principal components analysis across the 77 categories.

determining the appropriate context of words at the phrase-level, the sentiment score is analyzed through the standard method of summing up over positive and negative words, e.g., $S = \frac{P - N}{P + N + 1}$, where P and N stand for the number of positive and negative words, respectively.

A key feature of our methodology is its ability to differentiate between relevant news for companies (defined in our context as those related to specific firm events) as opposed to unidentified firm events. For each news story, therefore, our application of TSS produces a list of relevant events connected to this company and to this particular piece of news. It is possible that multiple events may be connected to a given story. In our analysis we ignore the *Stock Price Change* and *Stock Price Change Pattern* categories as these categories do not, on their own, represent fundamental news events. We also ignore *Award*, *Facilities*, and *Inside Purchase*, since these categories do not contain a sufficient number of observations. We are therefore left with eight main categories.

To be more precise, our goal is to analyze the difference in return patterns based on the type of information arrival. We therefore classify each stock/day into one of three categories:

1. *No news* – observations without news coverage.
2. *Unidentified news* – observations for which none of the news coverage is identified.
3. *Identified news* – observations for which at least some of the news coverage is identified as being at least one of the above events.

Moreover, we define “new” news versus “old” news by whether the news identifies the same event that had been identified in similar recent news stories of that company.¹⁶ Specifically, a given event coverage is considered “new” if coverage of the same event type (and the same stock) is not identified during the previous five trading days.

¹⁶ See Tetlock (2011) for a different procedure for parsing out new and stale news.

Since our goal is to relate information arrival to stock returns, which are observed at the stock/day level, we rearrange the data to follow the same stock/day structure. To that end, we consolidate all events of the same type for a given stock/day into a single event by averaging their scores. The resulting dataset is structured such that for each stock/day we have a set of indicators denoting which events were observed, and when observed, the relevant score for each of the event types. We also compute a *daily score* by adding the number of positive and negative features across all relevant articles.

In order to ensure that the analysis does not suffer from a look-ahead bias, we use the article timestamp and line it up with the trading day. Specifically, we consider date t articles those that were released between 15:31 on date $t-1$ and 15:30 on date t . Date t returns are computed using closing prices on dates $t-1$ and t . Articles released on non-trading days (weekends and holidays) are matched with the next available trading day.

Table 1 provides an overview of the data. The first column in panel A reports the number of observations under each of the day classifications. First, we see that most days have no news coverage, i.e., 696,985 of 1,229,359 stock/day observations contain no news reported on the Dow Jones Newswire. Second, and most important, the vast majority of the days with news coverage, 374,194 of 532,374, do not have a single topic-identified news event. As shown in columns 2-4 of Panel A, most identified news days contain only a single-identified event (124,158 of 158,180). We also observe that identified news days contain a larger number of articles compared with unidentified news days (6.1 vs. 2.6 per stock/day). While the number of words per article does not seem to vary much by day type, the number of relevant words (as identified by TSS) is much larger on identified news days (81 vs. 49). The bottom part of Panel A reports the same set of statistics by event type. For example, the row labeled *Acquisition* contains all day/stock observations in which an acquisition event type was observed. Note that this sorting is not mutually exclusive as there may be day/stock observations with multiple event types. The largest event type is *Financials*, with 69,205 observations. Outside of financials, the other event types contain between 10,047 observations (*Partnerships*) and 30,101 (*Deals*).

Panel B of Table 1 reports the average firm returns, market returns, and factor characteristics (size, book-to-market, and momentum) of observations across stock/day types. Consistent with the prior literature, we find that firm size is correlated with media, even if this effect is small for our sample of S&P500 firms -- quintile assignment of 4.48 for no news vs. 4.71 for unidentified news and 4.76 for identified news. Importantly, return and factor characteristics are very similar for identified and unidentified news days. In unreported results we considered a fourth category, stock/days with both identified and unidentified news. The results were unaffected by merging these categories.

A key finding of this paper is that when we can identify news, the news matters. As a first pass at the data, Table 2 provides a breakdown of news stories by the distribution of returns. In brief, the main result is that identified news days are more likely than unidentified news to lie in the negative and positive tails of the return distribution. On the surface, this is consistent with rational models, which would suggest that information arrival should be associated with increases in volatility.

In particular, if identified news days proxy for information arrival, then we should find that news arrival would be concentrated among days with large return movements, positive or negative. To relate news arrival intensity with returns, we assign daily returns into percentiles separately for each stock and year: bottom 10%, next 20%, middle 40%, next 20%, and top 10%. We perform the assignment for each stock separately to control for cross-sectional variation in total return volatility, and perform the assignment for each year separately to control for large time-series variations in average return volatility, e.g., 2008-9. The columns in Table 2 group observations according to this split. The first three rows of the table show that extreme day returns are associated with somewhat larger number of articles (for each stock appearing in the news) and on these days, there is a larger total number of words used in the articles.

Next, we compare the observed intensity of different day types to the intensity predicted under the null that these distributions are independent. For example, the null would suggest that of the 700 thousand no news days, 70 thousand would coincide with returns at the bottom 10%, 140 thousand would coincide with returns at the following 20%, and so forth.

The results in rows five through fourteen report the difference between the observed intensity and the null in percentage terms.

Several observations are in order. First, we find that no news days are less concentrated among days with large price changes: -6.6% (-6.5%) for the bottom (top) 10% of days. This is consistent with the notion that news coverage proxies for information arrival. Interestingly though, we observe a very similar pattern for unidentified news days: 2.2% (1.1%) for the bottom (top) 10% of days. Second, in sharp contrast to these results, we find that identified news days are 30.8% (34.2%) more likely to coincide with the bottom (top) 10% of return days. Thus, while we might expect under independence to have 15,818 identified news stories in the lower tail, we actually document 20,690 news stories. That is, identified news days, but not unidentified news days, are much more likely to be extreme return days.

Third, this last pattern is also observed when we examine the frequency of individual event types, one at a time. The bottom part of Table 2 shows a U-shaped pattern suggesting that each of the event types is more likely to coincide with extreme return days compared with moderate return days. It should be noted that for some event types, the pattern is not symmetric. For example, *Deals* are more likely to appear on extreme positive days, compared with extreme negative days. This is consistent with the intuition that deals would generally be regarded as a positive event for the firm. At the same time, *Legal* events are more likely to coincide with extreme negative days compared with extreme positive days. The news categories with the greatest concentration of events in the tails – *Analyst Recommendations* and *Financial* – are not surprisingly dispersed in a much more symmetric way.

III. R^2

A seminal paper on the question of whether stock prices reflect fundamental information is Roll (1988). In that paper, Roll (1988) argues that once aggregate effects have been removed from a given stock, the finance paradigm would imply that the remaining variation

of firm returns would be idiosyncratic to that firm. As a proxy for this firm specific information, Roll (1988) uses news stories generated in the financial press. His argument is that, on days without news, idiosyncratic information is low, and the R^2 s from aggregate level regressions should be much higher. Roll (1988) finds little discernible difference. Thus, his conclusion is that it is difficult to understand the level of stock return variation. Working off this result, a number of other papers reach similar conclusions with respect to prices and news, in particular, Cutler, Poterba and Summers (1989), and Mitchell and Mulherin (1994).

The evidence that asset prices do not reflect seemingly relevant information is not just found with equity returns. For example, Roll (1984)'s finding that, in the frozen concentrated orange juice (FCOJ) futures market, weather surprises explain only a small amount of variability of futures returns has been a beacon for the behavioral finance and economics literature. Given that weather has theoretically the most important impact on FCOJ supply, and is the focus of the majority of news stories, Roll (1984) concludes, like in his 1988 paper, that there are large amounts of "inexplicable price volatility". In contrast, Boudoukh, Richardson, Shen and Whitelaw (2007) show that when the fundamental is identified, in this case temperatures close to or below freezing, and when relevant path dependencies are taken into consideration, e.g., first freeze versus second, third etc., there is a close relationship between prices and weather surprises. In this section, we make a similar argument to Boudoukh, Richardson, Shen and Whitelaw (2007). We parse out news stories into identified versus unidentified events and reevaluate Roll's (1988) finding and conclusion.

In a different context, and using a different methodology, Griffin, Hirschey and Kelly (2011) and Engle, Hansen and Lunde (2011) also provide evidence that price volatility can be partially explained by news. For example, by cross-checking global news stories against earnings announcements to try and uncover relevant events, Griffin, Hirschey and Kelly (2011) document better information extraction can lead to higher R^2 s between prices and news. Engle, Hansen and Lunde (2011) utilize the *Dow Jones Intelligent Indexing* product to match news and event types for a small set of (albeit large) firms, and show that the arrival of this public information has explanatory power for the dynamics of volatility.

The results of Table 2 suggest that our textual analysis methodology will have similar success at linking identified events to stock return variation.¹⁷ Therefore, as a more formal look at the data, we study the link between news arrival and volatility by computing daily return variations on no news days, unidentified news days, and identified news days. Specifically, for each stock we compute the average of squared daily returns on these day types. We then calculate the ratio of squared deviations on unidentified news days to no news days, and the ratio of squared deviations on identified news days to no news days.¹⁸ If both unidentified and identified news days have no additional effect on stock volatility, then we should find that these ratios are distributed around one.

Table 3 reports the distribution of these variance ratios. Consistent with Table 2 results, we find that the median variance ratio of unidentified news days is close to one (i.e., 1.2) while the variance ratio of identified news days exceeds two. That is, the median stock exhibits return variance on identified news days that is 2.2 times the variance of no news days. The result appears quite robust with over 90% of stocks exhibiting variance ratios exceeding one on identified news days.

Figure 1 depicts the distribution of these ratios across the 672 stocks for which these ratios are available (out of 791), winsorized at 10.¹⁹ As evident, the ratios are not distributed around one for neither unidentified nor identified news days. However, the difference in distributions between unidentified and identified news days' ratios is clear: the variance ratio is much higher on identified news days compared with unidentified news days. These results clearly demonstrate that our day classification has power to distinguish between days on which price-relevant information arrives and days on which information may or may not arrive, but if it does, it is not price-relevant.

¹⁷ Note that, while most researchers focus on Roll's (1988) R^2 result, Roll (1988) also provided evidence that kurtosis was higher on news versus no news days, a result similar to that provided in Table 2.

¹⁸ We include only stocks with at least 20 observations for all day classifications.

¹⁹ We eliminate stocks for which we do not have at least twenty trading days of under each of the day categories.

The middle part of Table 3 reports variance ratios for each of the event types (Acquisition, Analyst Recommendations, etc.). The event-level analysis reveals similar patterns with median variance ratios exceeding 1 for all event types and exceeding two for two of the eight event categories, in particular, Analyst Recommendations and Financial. Most striking is that for 25% of the firms, five of the event types exceed variance ratios of two. In general, consistent with a priori intuition, Acquisitions, Legal, Financial and Analyst Recommendations appear to be the most informative.

As an additional measure of the informative of news, Section 2 defined “new” news versus “old” news by whether the news identifies the same event that had been identified in similar recent news stories of that company. One might expect that new news would have more information and thus greater price impact. Indeed, we find that among identified news days there is a substantial difference between the variance of old news days, with a median variance ratio of 1.5, and new news days, with the corresponding statistics of 2.2.

This fact, that variances are higher on days in which we can identify important events and on days with “new” news, supports a relation between prices and fundamentals. As a more formal analysis, we reproduce the aforementioned Roll (1988) analysis for our setting. Table 4 reports results for a reinvestigation of the R^2 analysis of Roll (1988). Specifically, we estimate a one-factor pricing model and a four-factor pricing model separately for each firm and for each day classification: all, no news, unidentified news, and identified news.²⁰ We repeat the same analysis at the 2-digit SIC industry classification thereby imposing a single beta for all firms within a given industry and utilizing weighted least squared regressions. All R^2 are adjusted for the number of degrees of freedom.

The results in the top part of Table 4 report the mean and median R^2 across firms (columns 2 and 3) and industries (columns 5 and 6). Consider the median calculations for the CAPM model at the firm level. The R^2 s are similar on no news and unidentified news days (i.e., 33% vs. 30%). The magnitude of the R^2 s and similarity of these numbers between no news and news days (albeit unidentified) are consistent with Roll’s puzzling results. However,

²⁰ We impose a minimum of 40 observations to estimate the regressions.

R^2 s are much lower on identified news day, i.e., 15.9%. The difference in R^2 between identified news and no-news days is striking – the ratio of median R^2 between identified news and no-news days is 2.1, in sharp contrast to Roll’s results.

Roll’s original theory-based conjecture, dramatically refuted empirically by his 1988 work, was that the performance of a market model, as measured by R^2 , should be much worse during days on which firm-specific information arrives, compared with days when no such information arrives. In contrast to Roll’s results, our results do lend support to this conjecture, since we are able to better proxy for firm-specific information arrival days using event identification.

Our results appear to be robust to the pricing model and firm/industry specification. For example, the results are analogous for the four factor model that, along with the market, includes the book-to-market, size and momentum factors. In particular, the ratio of median R^2 between no-news and identified news days is still greater than two, and the R^2 s between no-news and unidentified days is again similar. All these results change only barely when we perform the analysis at an industry level in which we constrain the betas against the 1- or 4-factor models to be the same within industry. Constraining the betas allows greater degrees of freedom for subsequent analysis when we try and understand the source of the differences between the R^2 s of no-news versus unidentified days. Specifically, in the next section, we ask the question whether our estimate of news sentiment/tone, coupled with the exact event identifier, can help explain these R^2 s. As a brief preview, we find that, even in a simple regression framework using the score S defined in section II.B, there is a strong link between this information and the unexplained variation from factor model regressions.

IV. Measuring Sentiment

One of the main applications of textual analysis in finance has been to link sentiment scores to both contemporaneous and future stock returns. The evidence is statistically significant albeit weak in magnitude. For example, Tetlock (2007) and Tetlock, Saar-Tsechansky and Macskassy (2008), show that negative word counts of news stories about firms based on IV-4 have contemporaneous and forecast power for the firms’ stock returns, though the R^2 s

are low. Loughran and McDonald (2011) argue that for a finance context the Harvard dictionary is not appropriate and build a sentiment score using a more finance-centric dictionary. Their application focuses on creating a dictionary appropriate for understanding the sentiment contained in 10-K reports. For their 10-K application, sentiment scores based on word counts from this alternative dictionary generally provide a better fit.

In this section, we first extend the analysis of Section III on news versus no news R^2 s to include sentiment scores. In the above analysis, we showed that identified news days are a good proxy for information arrival. Below, we show that the sentiment of these articles, i.e., the directional content of this information, has explanatory power for returns. As a preview, consider Table 4. Table 4 shows that market model regressions on news days have low R^2 , that is, most of the variation of stock returns is idiosyncratic in nature. A reasonable hypothesis is that the R^2 s should increase if idiosyncratic information is incorporated directly. We use the sentiment score as our proxy for this direct information, and we compare the score based on TSS and that using IV-4.

Recall that for each day and event type (within the day) we compute a sentiment score using the number of positive and negative features identified by TSS. For comparison purposes, we also compute a score using IV-4, similar to Tetlock (2007). We refer to these scores as “IV4”. Table 5 provides a set of summary statistics with respect to sentiment scores.

The first column in the table reports the number of observations classified as unidentified and identified news days (first two rows), followed by the number of observations falling into each of the event types.²¹ The set of columns under “TSS” report score statistics for each of the classifications. For example, of the 374,194 unidentified news days, TSS is able to compute a sentiment score for only 158,180. In contrast, virtually all identified news days are matched with sentiment output from TSS. The remaining columns in the column block report the mean, percentiles (10%, 50%, 90%), and spread between the top and bottom 10% of observations within each category. The next block of columns, under “IV4”, reports the

²¹ Recall that the sum of observations under all event types exceeds the number of observations under “identified days” since they are, on average, multiple events for each identified news day.

same set of statistics using the IV-4 based dictionary. The last column in the table reports the correlation between the TSS and IV4 scores.

First, for virtually every category, the number of observations with available TSS scores is smaller than the number of observations with available IV4 scores available. This is consistent with the set of negative and positive words in the IV4 dictionary being generally larger than the set of positive and negative features in TSS. The average score for unidentified and identified news days is on average positive, demonstrating the tendency of media coverage to have a positive tone. This bias is similar in magnitude for TSS and IV4.

Second, TSS appears to produce more discerning sentiment scores compared with IV4. For both unidentified and identified days, the spread of TSS scores is much larger than the spread of IV4 scores; the difference between the top and bottom 10% of identified news days is 1.23 under TSS but only 0.50 under IV4.²² This holds across many of the event types. Examining variations across event types, we find that TSS scores vary much more than IV4 scores. Also, the variation in average TSS scores is consistent with one's priors about these event types. For example, the average scores of *Analyst Recommendations* is close to neutral (0.06) consistent with the idea that analysts revisions are equally likely to be positive as they are to be negative. On the other hand, legal events are on average negative and correspond to negative TSS scores (-0.21), while partnership events are on average positive and correspond to positive TSS scores (0.63).

These differences between TSS and IV4 scores are not merely an artifact of rescaling. The last column in Table 5 reports the correlation between TSS and IV4 scores. While the correlations are positive, they range between 0.17 and 0.38 -- far from one. In fact, for three of the eight event types, event-specific scores correlations are lower than 0.20.

To see the additional explanatory power of event-specific scores, consider the results of Table 6. The R^2 's reported in the table are adjusted R^2 's derived from industry regressions. We augment the one-factor or four-factor models with event-level scores obtained from

²² Recall that the score ranges from -1 to 1.

TSS or IV4, utilizing weighted least squared regressions estimated on the 2-digit SIC industry level (with at least 80 observations). That is, we assume that all firms within the industry have the same return response magnitude to a given event type but we allow this magnitude to vary across events and industries. Focusing on identified event days, we see that at the firm level, daily scores obtained from TSS increase R^2 from a median of 16.2% to 17.2% under the one-factor model, and from 17.1% to 20.2% under the four-factor model (while essentially unimproved using IV-4 scores). Most important, these increases are attained only for identified news days. In contrast, for unidentified news days, there is no increase in R^2 's when sentiment scores are taken into account. In other words, to link stock prices to information, it is necessary to measure both the news event and the tone (i.e., sentiment) of this news.

In order to investigate this further, we also report R^2 from a weighted least squares pooled industry regression while separating observations by event types. Consider the CAPM-like model. The results show a large degree of variation across events. For example, *Acquisitions* (12.8%) and *Financial* (12.6%) are lower than the 16% cited above for identified news days, and substantially lower than the 33% on no news days and 30% on unidentified news days. In contrast, *Analyst Recommendations* (14.8%), *Deals* (19.9%), *Employment* (16.4%), *Partnerships* (23.0%) and *Product* (28.6%) produce much higher R^2 's. Of particular interest, the increase in adjusted R^2 's are all positive once the news' sentiment is taken into account, with the percent increase in the ratio of R^2 's ranging from 11% to 62%, the latter being *Analyst Recommendations*. Sample size aside, to the extent these categories can be further broken down and the sentiment of each event be incorporated, one would expect an even greater bifurcation of the R^2 's between unidentified/no news days and further refined identified news days. In conclusion, the TSS sentiment score, when allowing for event-specific scores, increases the explanatory power significantly for all event types.

For comparison purposes, Table 6 also provides the same type of analysis for IV-4 scores. While IV-4 scores also add to the explanatory power of stock market returns on news and specific event days, the gains are of a significantly smaller magnitude. In fact, for all event types, TSS dominates the IV-4 score methodology. This result is loosely consistent with

previous analyses using IV-4 scores which show statistical (albeit economically weak) significance (e.g., see Tetlock (2007), and Loughran and McDonald (2010), among others).

V. News Type, Reversals and Continuations

Though the results of Sections III and IV are supportive of one of the main hypotheses from efficient markets, namely, that prices respond to fundamental information, the growing literature in the area of behavioral finance also has implications for our research. There are a number of papers that describe conditions under which stock prices might under- or overreact based on well-documented behavioral biases. (See, for example, Daniel, Hirshleifer and Subrahmanyam (1998), Barberis, Shleifer and Vishny (1998), Hong and Stein (1999), Hirshleifer (2002), and Barber and Odean (2008), among others.) Essential findings from this literature based on behavioral theory are that (i) investors only partially adjust to real information, leading to a continuation of the price response to this information, and (ii) investors overreact to shocks to prices (i.e., unreal information), leading to higher trading volume and reversals of these shocks.

Indeed, there are a number of studies that provide some empirical support for these hypotheses. For example, Stickel and Verrecchia (1994) and Pritamani and Singal (2001) report stock price momentum after earnings announcements. Tetlock, Saar-Tsechansky and Macskassy (2008) report similar underreaction to news events focused on negative words (as measured through a word count based on a textual analysis). The closest papers to ours, however, are Chan (2003) and Tetlock (2010, 2011) who focus on days with and without news. Specifically, Chan (2003) separates out companies hit by price shocks into those with public news versus no news. Chan (2003) finds that after bad news stock prices continue to drift down while after no news stock prices reverse. Tetlock (2010, 2011) generally finds that public news, and especially new as opposed to stale news, reduce the well-known short-term reversals of stock returns. In contrast, Gutierrez and Kelley (2008) do not find any difference.

In this section, we extend the above analyses to our dataset, in particular, to our differentiation of public news into identified news events versus unidentified news. To the

extent the behavioral literature tries to explain the theories of under- and overreaction in terms of stock price responses to *real* news versus *false* news, our methodology provides an effective way to study this issue further.

As mentioned above, the results so far suggest a strong contemporaneous response of stocks to their media coverage on identified news days but not on unidentified news days. One interpretation is that identified news days are days on which price-relevant information arrives. To examine this, we measure return autocorrelation on different day types (i.e., no-news, unidentified news, and identified news). Table 7 reports the results of a weighted least squared regression in which the dependent variable is day $t+1$ returns. In the first column of the table, the independent variables are time t returns and day classification dummies (no-news days dummy is dropped), along with the TSS daily sentiment score.

Consistent with the aforementioned literature (e.g., Chan (2003)), Table 7 suggests a reversal following no-news days. For example, the day t return coefficient of -0.037 implies a negative daily autocorrelation of 3.7%. While this negative autocorrelation is consistent with microstructure effects such as bid-ask spreads and state prices, the reversal is sizable considering the S&P500 universe of stocks in our sample and their average bid-ask spreads. Unidentified news days are characterized by reversals too, while the magnitude of the negative autocorrelation is smaller compared with no-news days (i.e., -2.9%). This reduction in the magnitude of reversals on news days is also consistent with findings in the prior literature (e.g., Tetlock (2010, 2011)).

The more interesting and novel finding is that the well-known reversal result disappears when we condition on identified news days. The marginal coefficient of 4.4% implies that identified news days are followed by continuations (i.e., positive autocorrelation of 0.7%). Furthermore, we find that on identified news day, the continuation follows the direction of the day t sentiment – positive sentiment days are followed by higher than average returns on subsequent days, controlling for date t return. For example, using Table 5, consider the 90% quantile of TSS scores, i.e., 0.83, times the 0.068 coefficient, which adds an additional 5.6bp to the continuation. While these numbers are arguably small, the coefficients are all statistically significant at the 1% level.

Columns 2-9 of Table 7 study these patterns for each of the event types separately. In these regressions, we set the event dummy to be equal to one if the event occurred on date t and zero otherwise. This specification contrasts days on which a specific event took place with all other days. The results suggest that many event types exhibit continuations, with the largest ones following *Analyst Recommendations*, *Deals*, *Employment*, and *Financials*. Together, the results in the table suggest that the contemporaneous price response to identified news days is unlikely to be due to irrational over-reaction to news coverage of the events underling our study. If anything, it suggests that the price response is insufficiently strong for many of the event types.

The analysis underlying Table 7 has resulted from a pooled time series regression of all S&P 500 stocks over the 10-year sample period. The regression model imposes the same coefficients across all stocks, making it difficult to gauge the economic significance of the results. To further evaluate the economic magnitude of the impact the type of news has on stock prices, we consider three separate zero-cost strategies, each implemented on a subset of the day classifications – no-news, unidentified news, and identified news.

Specifically, following no-news and unidentified news days we follow a *reversal* strategy which goes long one unit of capital across all stocks with time t excess returns (returns minus CRSP value-weighted returns) falling below minus one standard deviation based on lagged 20 day excess returns, and go short one unit of capital across all stocks with time t returns exceeding plus one standard deviation.²³ In contrast, following identified news days, we follow a *continuation* strategy which goes long one unit of capital across all stocks with time t excess returns (returns minus CRSP value-weighted returns) exceeding plus one standard deviation based on lagged 20 day excess returns, and go short one unit of capital across all stocks with time t returns falling below minus one standard deviation. The return of the zero cost strategy is equal to the returns on the longs minus the returns on the shorts, provided that we have at least five stocks on the long side and at least five stocks on the short side. Thus, at any time, the strategy is long $1/N_L$ units of capital across N_L positive

²³ The results are robust to changes in the threshold.

extreme moves and short $1/N_S$ units of capital across N_S negative extreme moves. In all cases we hold the stocks for one day.

Table 8 reports results for these three strategies separately. The top panel of the table reports time series regressions for a four-factor model. The alpha of the strategy formed using no-news stocks is 0.16% (per day), suggesting a return pattern consistent with reversals. In sharp contrast, the strategy formed using identified-news stocks exhibits large continuation and daily alpha of 0.20%; finally, the strategy formed using unidentified news days produces an alpha that is indistinguishable from zero. These results, differentiating between identified and unidentified news, further highlight the importance of parsing out the content of the news. Of some note, the strategies load significantly only on one of the aggregate factors, resulting in negative albeit small betas. Apparently, going long and short stocks subject to extreme moves is on average short the market.

The second panel of the table reports the average daily return of the each of the strategies. The continuation strategy using identified news stocks generates an average daily return of 19.6bp per day and was profitable in every single year in our sample, including the crash of internet stocks and the financial crisis period. The reversal strategy using no news days was also profitable in every year, achieving an average daily return of 15.6bp per day. The bottom panel of the table reports the mean and standard deviation of daily returns along with the strategies' annualized alphas. While the reversals strategy produces a Sharpe ratio of 1.8, the continuations strategy produces a Sharpe ratio of 1.7 – economically large given that our sample includes only S&P500 stocks.²⁴ It should be pointed out that unlike many standard long-short strategies analyzed in the academic literature (e.g., value or momentum), the strategy evaluated here exhibits daily variations in the number of stocks held. The bottom part of the table reports the average number of stocks held by each of the strategies on the long and the short sides. Not surprisingly, the number of stocks in the reversal strategy is much larger than the number of stocks in the continuation strategy as we

²⁴ The strategy turns over the stocks after one day of trading, resulting in significant roundtrip trading costs, even for liquid stocks such as those in the S&P500. The use of the trading strategy is illustrative and is intended to demonstrate the economic significance through aggregation across S&P500 stocks each period.

have over four times more no-news observations compared with identified news observations.

Figure 2 facilitates a further comparison between the strategies' relative performance over the 10-year sample period. The figure, reported in log scale, shows what would have happened to the value of a \$1 invested in each of the strategies on Feb 1st, 2000.²⁵ In addition to the three strategies, we include the cumulative value of investing in the market, in excess of the risk-free rate. Consistent with the Sharpe ratios, and transaction costs aside, both the reversals and continuations strategies are very profitable. Of course, as mentioned in footnote 23, an individual trading on a daily basis would generate significant transactions costs, resulting in much lower profits.

The results provided in Table 8 ignore the information contained in the sentiment scores associated with identified news events. To rectify this, we evaluate the economic significance of the daily sentiment score by including it in the trading strategy described above. Specifically, with the strategy described above, we focus on identified news, going long (short) one unit of capital of stocks with large positive (negative) excess returns (exceeding one standard deviation). However, we now break this strategy into two strategies by sorting additionally on the sentiment scores across each stock each day. That is, the first strategy, labeled "Long Score", goes long stocks with above median sentiment scores among all stocks with large positive returns (that day) and goes short stocks with below median score among all stocks with large negative returns (that day). The strategy labeled "Short Score" still goes long stocks with large positive returns but now with below median scores, and goes short stocks with large negative returns but with above median scores. Thus, effectively, we have split the holdings of the previous strategy into two portfolios of roughly equal size. If the continuation pattern observed following identified news days is unrelated to the informational content of the news that day, then the two strategies would yield similar results.

²⁵ Recall that we use the first 20 trading days to compute lagged volatility.

The top part of Table 9 reports time-series alphas of these zero-cost strategies controlling for the four factors. The findings are clear: the informational content of the news appears to affect the continuation strategy dramatically. The “Long Score” strategy generates a daily alpha of 34bp while the short-score strategy produces an alpha that is indistinguishable from zero. Clearly, this strategy is quite volatile since it relies on a relatively smaller universe of stocks by halving the existing portfolio. That said, the “Long Score” strategy produces positive returns in 9 of the 10 years in our sample compared to the “Short Score” strategy which splits the period by returning positive profits in 7 of the 10 years.

As a final comment, the strategy focuses on one-day holding periods. It seems worthwhile commenting on the price patterns following no-news days and identified news days on subsequent days. Figure 3 plots the cumulative abnormal returns of the three strategies described in Table 8 based on no news, unidentified news and identified news for days 1 through 10 following extreme moves of S&P 500 stocks. First, consider no-news days. Stocks exhibit a reversal that appears to continue on, reaching around 40bps at the end of the two-week period. The results also now show a reversal, albeit of 20bps, for unidentified news stocks. These results are again consistent with the broad theme of this paper that no news and unidentified news days display similar characteristics. In contrast, following identified news days, stocks exhibit a one-day continuation with no clear subsequent price movement. This suggests that, whatever under-reaction to “real” news that takes place, it is short lived.

VI. Conclusions

The bottom line from this paper is in stark contrast to the last 25 years of literature on stock prices and news. We find that, when information can be identified and that the tone (i.e., positive versus negative) of this information can be determined, there is a much closer link between stock prices and information. Examples of results include market model R^2 s that are no longer the same on news versus no news days (i.e., Roll’s (1988) infamous result), but now are 16% versus 33%; variance ratios of returns on identified news days double than those on no news and unidentified news days; and, conditional on extreme moves, stock

price reversals occur on no news unidentified news days, while identified news days show continuation.

The methodology described in this paper may be useful for a deeper analysis of the relation between stock prices and information, especially on the behavioral side (e.g., as pertaining to the reversals/continuation analysis of Section V). There is a vast literature in the behavioral finance area arguing that economic agents, one by one, and even in the aggregate, cannot digest the full economic impact of news quickly. Given our database of identified events, it is possible to measure and investigate “complexity”, and its effect on the speed of information processing by the market. For example, “complexity” can be broken down into whether more than one economic event occurs at a given point in time, how news (even similar news) gets accumulated through time, and cross-firm effects of news. We hope to explore some of these ideas in future research.

References

- Aharony, J., and I. Swary, 1980, Quarterly Dividend And Earnings Announcements And Stockholders' Returns: An Empirical Analysis, *Journal of Finance* 35, 1–12.
- Antweiler, W., and M. Z. Frank, 2004, Is All That Talk Just Noise? The Information Content Of Internet Stock Message Boards, *Journal of Finance* 59, 1259–1293.
- Asquith, P., and D. W. Mullins, 1986, Equity Issues And Offering Dilution, *Journal of Financial Economics* 15, 61–89.
- Ball, R., and P. Brown, 1968, An Empirical Evaluation Of Accounting Income Numbers, *Journal of Accounting Research* 6, 159–178.
- Berry, T. D., and K. M. Howe, 1994, Public Information Arrival, *Journal of Finance* 49, 1331–1346.
- Campbell, J. Y., 1991, A Variance Decomposition For Stock Returns, *Economic Journal* 101, 157–179.
- Chan, W. S., 2003, Stock Price Reaction To News And No-News: Drift And Reversal After Headlines, *Journal of Financial Economics* 70, 223–260.
- Cutler, D. M., J. M. Poterba, and L. H. Summers, 1989, What Moves Stock Prices?, *Journal of Portfolio Management* 15, 4–12.
- Das, S. R., and M. Y. Chen, 2007, Yahoo! For Amazon: Sentiment Extraction From Small Talk On The Web, *Management Science* 53, 1375–1388.
- Davis, A. K., J. Piger, and L. M. Sedor, 2012, Beyond The Numbers: An Analysis Of Optimistic And Pessimistic Language In Earnings Press Releases, *Contemporary Accounting Research* 29, 845–868.
- Demers, E., and C. Vega, 2010, Soft Information In Earnings Announcements: News Or Noise?, *Working Paper, INSEAD*.
- Devitt, A., and K. Ahmad, 2007, Sentiment Polarity Identification In Financial News: A Cohesion-

- Based Approach, *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, 984–991.
- Engelberg, J., 2008, Costly Information Processing: Evidence From Earnings Announcements, *Working Paper, University of North Carolina*.
- Engle, R. F, M. Hansen, and A. Lunde, 2011, And Now, The Rest Of The News: Volatility And Firm Specific News Arrival, *Working Paper*.
- Fama, E. F., L. Fisher, M. C. Jensen, and R. Roll, 1969, The Adjustment Of Stock Prices To New Information, *International Economic Review* 10, 1–21.
- Feldman, R., S. Govindaraj, J. Livnat, and B. Segal, 2010, ManagementS Tone Change, Post Earnings Announcement Drift And Accruals, *Review of Accounting Studies* 15, 915–953.
- Feldman, R., B. Rosenfeld, R. Bar-Haim, and M. Fresko, 2011, The Stock Sonar Sentiment Analysis Of Stocks Based On A Hybrid Approach, *Proceedings of the Twenty-Third Innovative Applications of Artificial Intelligence Conference*, 1642–1647.
- Feldman, R., and J. Sanger, 2006, The Text Mining Handbook, *Cambridge University Press*.
- Griffin, J. M., N. H. Hirschey, and P. J. Kelly, 2011, How Important Is The Financial Media In Global Markets?, *Review of Financial Studies* 24, 3941–3992.
- Grob-Klubmann, A., and N. Hautsch, 2011, When Machines Read The News: Using Automated Text Analytics To Quantify High Frequency News-Implied Market Reactions, *Journal of Empirical Finance* 18, 321–340.
- Gutierrez, R. C., and E. K. Kelley, 2008, The Long-Lasting Momentum In Weekly Returns, *Journal of Finance* 63, 415–447.
- Hanley, K. W., and G. Hoberg, 2012, Litigation Risk, Strategic Disclosure And The Underpricing Of Initial Public Offerings, *Journal of Financial Economics* 103, 235–254.
- Kogan, S., Routledge, B. R., Sagi, J. S., and N. A. Smith, 2011, Information Content of Public

Firm Disclosures and the Sarbanes-Oxley Act, *Working Paper*.

Kothari, S. P., X. Li, and J. E. Short, 2009, The Effect Of Disclosures By Management, Analysts, And Financial Press On The Equity Cost Of Capital: A Study Using Content Analysis, *Accounting Review* 84, 1639–1670.

Lavrenko, V., M. Schmill, D. Lawrie, P. Ogilvie, D. Jensen, and J. Allan, 2000, Mining Of Concurrent Text And Time Series, *Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 37–44.

Li, F., 2010, The Information Content Of Forward-Looking Statements In Corporate Filings A Nave Bayesian Machine Learning Approach, *Journal of Accounting Research* 48, 1049–1102.

Loughran, T., and B. McDonald, 2011, When Is A Liability Not A Liability? Textual Analysis, Dictionaries, And 10-Ks, *Journal of Finance* 66, 35–65.

Mandelker, G., 1974, Risk And Return: The Case Of Merging Firms, *Journal of Financial Economics* 1, 303–335.

Manning, C. D., and H. Schutze, 1999, Foundations Of Statistical Natural Language Processing, *Cambridge Massachusetts: MIT Press*.

Mitchell, M. L., and J. H. Mulherin, 1994, The Impact Of Public Information On The Stock Market, *Journal of Finance* 49, 923–950.

Roll, R., 1984, Orange Juice and Weather, *American Economic Review*, 74, 5, 861–880.

Roll, R., 1988, R2, *Journal of Finance* 43, 541–566.

Shiller, R. J., 1981, Do Stock Prices Move Too Much To Be Justified By Subsequent Changes In Dividends?, *American Economic Review* 71, 421–436.

Tetlock, P. C., 2007, Giving Content To Investor Sentiment: The Role Of Media In The Stock Market, *Journal of Finance* 62, 1139–1168.

Tetlock, P. C., 2010, Does Public Financial News Resolve Asymmetric Information?, *Review of*

Financial Studies 23, 3520–3557.

Tetlock, P. C., M. Saar-Tsechansky, and S. Macskassy, 2008, More Than Words: Quantifying Language To Measure Firms Fundamentals, *Journal of Finance* 63, 1437–1467.

Vega, C., 2006, Stock Price Reaction To Public And Private Information, *Journal of Financial Economics* 82, 103–133.

1 Tables

Table 1: Summary Statistics
Panel A

	# Obs.	# Obs. with event count			# Tickers	# Articles (daily)	# Words (per art.)	# Relv. Words (per art.)
		=1	=2	>3				
Total	1,229,359	124,158	26,600	7,422	791	3.6	325	59
No News	696,985	NA	NA	NA	790	NA	NA	NA
Unid News	374,194	NA	NA	NA	791	2.6	329	49
Iden News	158,180	124,158	26,600	7,422	790	6.1	316	81
Acquisition	22,270	11,811	7,068	3,391	724	8.6	302	76
Analyst Rec	12,411	6,665	3,920	1,826	680	8.5	335	66
Deals	30,101	17,001	8,790	4,310	718	6.8	315	93
Employment	21,489	14,024	5,269	2,196	741	6.3	283	87
Financial	69,205	48,873	14,989	5,343	783	7.6	309	71
Legal	10,764	6,244	2,881	1,639	581	8.6	291	71
Partnerships	10,047	4,765	3,347	1,935	587	7.3	371	110
Product	25,181	14,775	6,936	3,470	652	7.1	366	108

Panel B

	Stock Return (daily)	Market Ret (daily)	SIZE	BM	MOM
No News	4.6bp	1.4bp	4.48	2.92	2.85
Unid News	2.9bp	0.6bp	4.71	2.92	2.84
Iden News	3.3bp	0.8bp	4.76	2.88	2.81
Total	3.9bp	1.1bp	4.59	2.91	2.84

The table reports summary statistics for observations (stock/day) classified as first as having no news, unidentified news (i.e., containing news all with unidentified events), or identified news (i.e., containing news with *some* identified events), and then by event type. Panel A reports the total number of observations and their distribution by event type count, the number of unique tickers, the average number of articles per day, the average number of words per article (day), and the average number of relevant words (as identified by TSS) per article (day). Panel B reports the average daily stock return, average daily market return, and the average size, book-to-market, and momentum quintile assignments.

Table 2: Event Frequency Across Return Ranks

Return rank	<i>p0-p10</i>	<i>p10-p30</i>	<i>p30-p70</i>	<i>p70-p90</i>	<i>p90-p100</i>
# of articles	4.5	3.5	3.3	3.5	4.4
# of words	1,445	1,140	1,093	1,133	1,423
% of rel. words	17.5%	16.6%	16.5%	16.6%	17.3%
No News	-6.6%	0.5%	2.8%	0.5%	-6.5%
Unid News	2.2%	-0.4%	-0.4%	-0.4%	1.1%
Iden News	30.8%	-5.7%	-10.7%	-5.4%	34.2%
Acquisition	18.8%	-4.2%	-6.2%	-5.4%	25.3%
Analyst Rec	71.1%	-11.9%	-22.1%	-10.3%	61.9%
Deals	6.5%	-4.0%	-3.1%	-1.9%	17.6%
Employment	20.5%	-2.1%	-5.8%	-2.8%	12.3%
Financial	66.2%	-11.1%	-23.0%	-10.3%	68.2%
Legal	16.8%	0.2%	-5.1%	-4.4%	11.9%
Partnerships	3.3%	-2.3%	-1.0%	-2.9%	11.2%
Product	10.4%	-2.2%	-4.1%	-2.6%	15.6%

The table reports summary statistics of all observations based on return rank sorts. For each stock and every year separately, we assign each day based on its percentile return rank – bottom 10%, following 20%, middle 40%, following 20%, and top 10%. The statics reported are the average number of article per observation, the average number of words, the fraction of all words identified as relevant (per TSS). Next, we report the difference between the observed distribution and the distribution that would obtain under independence based on observations’ classification as having no news, unidentified news (i.e., containing news all with unidentified events), or identified news (i.e., containing news with some identified events). For example, out of a total of 700K no news observations, 70K should fall under the bottom 10% of returns, but only 65K do resulting in a -6.6% difference. The bottom panel of the table groups observations into non-mutually exclusive event types and reports the results of the same comparison described above.

Table 3: Variance Ratios by Day and Event Type

	$p10$	$p25$	Median	$p75$	$p90$	N
Unid News	0.75	0.95	1.2	1.5	2.0	754
Iden News	1.05	1.49	2.2	3.3	5.1	672
Acquisition	0.64	1.01	1.5	2.7	6.3	294
Analyst Rec	1.05	1.70	2.7	5.1	13.5	190
Deals	0.65	0.94	1.4	2.1	4.0	273
Employment	0.65	0.96	1.4	2.4	4.1	329
Financial	1.29	1.95	2.9	4.7	7.6	581
Legal	0.46	0.71	1.3	2.3	4.9	116
Partnerships	0.49	0.69	1.1	1.5	2.0	110
Product	0.62	0.86	1.3	2.1	3.1	223
Old news	0.74	1.08	1.5	2.2	3.6	432
New news	1.13	1.59	2.2	3.5	5.8	654

The table reports daily return variations (daily returns squared) ratios where no-news days serve as the denominator in all calculations. Specifically, we compute squared returns for each stock under each of the day classifications, provided at least 20 observations were available. Next, we computed the ratio for each day and event types for each of the stocks. The table reports the distribution of these ratios. Observations (stock/day) are classified as having no news, unidentified news (i.e., containing news all with identified events), or identified news (i.e., containing news with some identified events). Identified news days are classified by event types (not mutually exclusive). “New news” classification denotes observations for which at least one of the events types did not appeared in the previous five trading days and “Old news” denotes the complementary set of identified news days.

Table 4: R^2 s – Firm and Industry-level Regressions (adjusted R^2)

	Firm Level			Industry Level		
	Mean R^2	Med R^2	N	Mean R^2	Med R^2	N
1 Factor						
All	28.6%	27.8%	791	28.5%	27.6%	60
No News	32.1%	33.3%	774	32.3%	34.1%	60
Unid News	30.8%	30.3%	721	30.9%	30.1%	58
Iden News	18.5%	15.9%	597	17.3%	16.2%	55
$\frac{NoNews}{IdenNews}$	1.73	2.10		1.87	2.10	
4 Factors						
All	33.6%	32.7%	791	31.8%	31.0%	60
No News	38.0%	38.6%	774	36.3%	36.9%	60
Unid News	35.8%	35.9%	721	34.0%	33.6%	58
Iden News	22.3%	19.6%	597	19.2%	17.1%	55
$\frac{NoNews}{IdenNews}$	1.71	1.96		1.89	2.16	

Panel A of the table reports daily return regressions with one factor (total market, value weighted) in the first panel, and with four factors in the second panel (market, value, size and momentum). Regressions are run separately for each day category – on all days, no news days, unidentified news days (i.e., containing news all with unidentified events), and identified news days (i.e., containing news with some identified events). Firm level regressions estimate firm-level betas and R^2 s while industry level regressions estimate 2-digit SIC industry level betas and R^2 s. All R^2 's are adjusted for the number of degrees of freedom and regressions use WLS. Firms (industries) with fewer than 40 (80) observations are excluded from the firm-level (industry-level) regressions.

Table 5: TSS and IV4 Scores – Summary Statistics

	Event count	N	TSS					IV4					IV4-TSS	
			mean	p10	p50	p90	p90-p10	N	mean	p10	p50	p90	p90-p10	Corr.
Unid News	374,194	135,643	0.25	-0.50	0.50	0.75	1.25	364,969	0.33	0.01	0.32	0.69	0.68	0.33
Iden News	158,180	158,097	0.38	-0.40	0.50	0.83	1.23	158,150	0.32	0.08	0.31	0.58	0.50	0.34
Acquisition	22,270	22,228	0.54	0.25	0.50	0.80	0.55	22,238	0.36	0.06	0.36	0.67	0.61	0.19
Analyst Rec	12,411	12,371	0.06	-0.67	0.00	0.75	1.42	12,401	0.23	-0.06	0.22	0.54	0.60	0.32
Deals	30,101	30,082	0.60	0.50	0.67	0.83	0.33	30,084	0.42	0.12	0.43	0.69	0.57	0.19
Employment	21,489	21,482	0.42	-0.41	0.50	0.80	1.21	21,475	0.30	-0.07	0.33	0.64	0.71	0.17
Financial	69,205	69,124	0.29	-0.54	0.50	0.83	1.37	69,156	0.29	-0.01	0.29	0.62	0.63	0.38
Legal	10,764	10,761	-0.21	-0.80	-0.40	0.57	1.37	10,763	0.16	-0.16	0.17	0.50	0.66	0.31
Partnerships	10,047	10,040	0.63	0.50	0.67	0.83	0.33	10,040	0.44	0.16	0.46	0.69	0.53	0.24
Product	25,181	25,166	0.55	0.18	0.67	0.83	0.65	25,155	0.34	0.00	0.37	0.65	0.65	0.31

The table reports summary statistics of scores computed based on TSS and IV4 identification of positive (P) and negative (N) features. Scores in all cases equal to $\frac{P-N}{P+N+1}$. Observations are sorted by day type: unidentified news (i.e., containing news all with unidentified events), or identified news (i.e., containing news with some identified events), and event types. The first column reports the total number of observations under each day type. The column block titled “TSS” (“IV4”) reports results obtained from TSS(IV4). N corresponds to the number of observations with non-missing scores, and the remaining columns report the distribution of scores. The last column in the table reports the correlation between TSS and IV4 scores for each of the days types.

Table 6: Industry-level R^2 's with Event-Level Scores (adjusted R^2)

	Factor(s) Only Mean R^2	with TSS Scores Mean R^2	with IV4 Scores Mean R^2	N	ΔR^2 (TSS)
1 Factor					
Unid News	30.9%	30.9%	31.0%	58	1.00
Idea News	17.3%	18.9%	17.5%	55	1.09
New News	15.7%	17.9%	16.0%	55	1.14
Acquisition	13.1%	14.7%	14.3%	31	1.12
Analyst Rec	17.3%	25.3%	21.5%	32	1.46
Deals	20.6%	25.7%	26.3%	34	1.25
Employment	19.4%	24.7%	23.5%	37	1.27
Financial	14.1%	17.0%	15.0%	52	1.20
Legal	16.9%	23.5%	20.8%	23	1.39
Partnerships	25.0%	29.7%	28.4%	17	1.19
Product	28.2%	34.3%	34.0%	30	1.22
4 Factors					
Unid News	34.0%	34.0%	34.1%	58	1.00
Idea News	19.2%	20.8%	19.4%	55	1.09
New News	17.7%	20.0%	18.0%	55	1.13
Acquisition	15.1%	16.5%	16.3%	31	1.09
Analyst Rec	19.8%	27.8%	24.0%	32	1.40
Deals	23.0%	28.1%	28.5%	34	1.22
Employment	22.4%	27.7%	26.6%	37	1.24
Financial	15.8%	18.7%	16.8%	52	1.18
Legal	20.8%	27.6%	24.5%	23	1.33
Partnerships	28.5%	33.3%	31.6%	17	1.17
Product	30.5%	37.0%	36.3%	30	1.21

Panel A of the table reports daily return regressions with one factor (total market, value weighted) in the first panel, and with four factors in the second panel (market, value, size and momentum) – with and without article sentiment scores. Regressions are run separately for each day category – on all days, no news days, unidentified news days (i.e., containing news all with unidentified events), and identified news days (i.e., containing news with some identified events). Sentiment score include a separate dummy for each event type (e.g., “acquisition” dummy) and event-level scores (e.g., “acquisition” score). The column titled “TSS” (“IV4”) reports results using TSS(IV4) scores. Industry level regressions estimate 2-digit SIC industry level betas and R^2 s. All R^2 's are adjusted for the number of degrees of freedom and regressions use WLS. Industries with fewer than 80 observations are excluded from the industry-level regressions.

Table 7: Return Reversals and Continuations

(Ret_1)	By day type	Acquisition	AnalystRec	Deals	Employment	Financial	Legal	Partner	Product
Ret_0	-0.037 [0.014]***	-0.024 [0.012]**	-0.026 [0.012]**	-0.026 [0.012]**	-0.026 [0.012]**	-0.031 [0.013]**	-0.025 [0.012]**	-0.025 [0.012]**	-0.025 [0.012]**
$I_{unid} \times Ret_0$	0.008 [0.006]								
$I_{iden} \times Ret_0$	0.044 [0.009]***								
$I_{unid} \times Score_0$	0.002 [0.015]								
$I_{iden} \times Score_0$	0.068 [0.025]***								
$I_{event} \times Ret_0$		-0.005 [0.017]	0.042 [0.024]*	0.04 [0.014]***	0.048 [0.016]***	0.047 [0.010]***	0.038 [0.034]	0.013 [0.027]	-0.002 [0.011]
$I_{event} \times EventScore_0$		0.075 [0.072]	0.167 [0.073]**	0.117 [0.073]	0.104 [0.045]**	0.076 [0.030]**	-0.016 [0.048]	0.058 [0.129]	-0.026 [0.045]
I_{unid}	-0.007 [0.009]								
I_{iden}	-0.045 [0.016]***								
I_{event}		-0.03 [0.046]	-0.052 [0.037]	-0.077 [0.048]	-0.075 [0.029]**	-0.035 [0.022]	-0.034 [0.027]	-0.071 [0.087]	-0.011 [0.030]
Constant	0.052 [0.024]**	0.048 [0.024]**	0.048 [0.024]**	0.048 [0.024]**	0.048 [0.024]**	0.049 [0.024]**	0.048 [0.024]**	0.048 [0.024]**	0.048 [0.024]**
Observations	1,228,568	1,228,568	1,228,568	1,228,568	1,228,568	1,228,568	1,228,568	1,228,568	1,228,568
R^2	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001

The dependent variable in all regressions is day $t + 1$ stock returns. The dependent variables include day t stock returns, in all specifications, and day t classification dummy: I_{unid} equals one for days with unidentified news, I_{iden} equals one for days with identified news, and I_{event} equals one for days with the column event type. "Score" denotes daily scores and "EventScore" denotes even-level score. All regressions use WLS with time clustered standard errors.

Table 8: Reversals and Continuations Strategies

	No News	Unid News	Iden News
Alpha	0.161 [0.027]***	0.062 [0.032]*	0.197 [0.041]***
Mkt-rf	0.197 [0.035]***	0.164 [0.040]***	-0.203 [0.064]***
SMB	-0.067 [0.070]	-0.132 [0.081]	0.057 [0.106]
HTM	-0.152 [0.075]**	-0.06 [0.097]	-0.006 [0.111]
UMD	0.064 [0.049]	0.072 [0.060]	-0.012 [0.064]
Observations	2487	2417	2030
R^2	0.046	0.021	0.022

Daily Statistics			
Mean	0.156%	0.059%	0.196%
s.d.	1.340%	1.545%	1.857%
SR	1.837	0.604	1.666
Drawdown	-18.2%	-46.5%	-34.5%
N Longs	35.8	21.7	11.2
N Shorts	38.0	20.6	10.5

Year	Mean Strategy Returns		
2000	0.351%	0.108%	0.087%
2001	0.094%	-0.039%	0.522%
2002	0.207%	0.117%	0.519%
2003	0.126%	0.050%	0.040%
2004	0.067%	0.032%	0.215%
2005	0.086%	0.061%	0.141%
2006	0.048%	0.023%	0.121%
2007	0.043%	-0.020%	0.205%
2008	0.235%	0.102%	0.105%
2009	0.312%	0.164%	0.026%

The table reports zero-cost trading strategy returns based on day $t - 1$ classification: no-news, unidentified news (i.e., containing news all with unidentified events) and identified news (i.e., containing news with some identified events). The first two columns' strategy goes long (short) stocks classified on day $t - 1$ as having at least one standard deviation negative (positive) returns, i.e., it is a reversal trading strategy. The last column strategy goes long (short) stocks classified on day $t - 1$ as having at least one standard deviation positive (negative) returns, i.e., it is a continuation trading strategy. Standard deviations are computed on a rolling basis, for each stock separately, using the prior 20 days. Holding period for all strategies is one day and we assume an investment of one unit of capital in the long strategy and one unit of capital in the short strategy. Each day strategy requires at least 5 long position and at least short positions. The top panel reports four factor time-series regressions. The middle panel reports the average daily return statistics and the average number of stocks held in the long and short side of the strategies, and the bottom panel reports the average strategy daily return for each of the years in the sample.


Table 9: Reversals and Continuations Strategies – Score Based

	Long Score	Short Score
Alpha	0.342	0.083
	[0.087]***	[0.082]
Mkt-rf	-0.369	-0.048
	[0.163]**	[0.135]
SMB	0.266	0.143
	[0.230]	[0.197]
HTM	-0.458	0.21
	[0.253]*	[0.225]
UMD	0.156	-0.171
	[0.144]	[0.114]
Observations	751	740
R^2	0.106	0.015

Year	Mean Strategy Returns	
2000	0.363%	0.411%
2001	0.638%	0.240%
2002	0.611%	0.843%
2003	0.309%	-0.326%
2004	0.294%	0.045%
2005	0.253%	-0.016%
2006	0.416%	-0.045%
2007	0.382%	0.217%
2008	-0.031%	0.068%
2009	0.296%	0.092%
Total	0.322%	0.106%


The table reports zero-cost trading strategy returns based on day $t-1$ classification into identified news (i.e., containing news with some identified events) and day $t-1$ sentiment score. The “Long Score” strategy goes long (short) stocks classified on day $t-1$ as having at least one standard deviation positive (negative) returns and above (below) median scores. The “Short Score” strategy goes long (short) stocks classified on day $t-1$ as having at least one standard deviation positive (negative) returns and below (above) median scores. Standard deviations are computed on a rolling basis, for each stock separately, using the prior 20 days. Holding period for all strategies is one day and we assume an investment of one unit of capital in the long strategy and one unit of capital in the short strategy. The top panel reports four factor time-series regressions. The bottom panel reports the average strategy daily return for each of the years in the sample. The “Long Score” strategy goes long half of the stocks

Figure 1: Variance Ratios




The figure reports the distribution of variance ratios unidentified news days and identified news days, compared with no-news days. We estimate variances by squaring daily returns and compute a separate ratio for each stock/day type. Stocks with less than 20 observations in each of the day types are excluded. Ratios are winsorized at 10.

Figure 2: Strategies' Cumulative Payoffs



The figure plot cumulative returns over the sample period of zero-cost trading strategies based on day $t - 1$ classification: no-news, unidentified news (i.e., containing news all with unidentified events) and identified news (i.e., containing news with some identified events). No-news and identified news day strategies go long (short) stocks classified on day $t - 1$ as having at least one standard deviation negative (positive) returns, while identified news strategy goes long (short) stocks classified on day $t - 1$ as having at least one standard deviation positive (negative) returns. Standard deviations are computed on a rolling basis, for each stock separately, using the prior 20 days. Holding period for all strategies are one day and we assume an investment of one unit of capital in the long strategy and one unit of capital in the short strategy, with 1 initial investment. For comparison, we include the cumulative market excess returns.

Figure 3: Strategies' CAR



The figure plot cumulative returns for zero-cost trading strategies based on day $t - 1$ classification: no-news, unidentified news (i.e., containing news all with unidentified events) and identified news (i.e., containing news with some identified events). Each strategy goes long (short) stocks classified on day $t - 1$ as having at least one standard deviation positive (negative) returns. Standard deviations are computed on a rolling basis, for each stock separately, using the prior 20 days. Holding period for all strategies varies from one day to ten days and we assume an investment of one unit of capital in the long strategy and one unit of capital in the short strategy.