Front Page News: The Effect of News Consumption on Financial Markets

Anastassia Fedyk*

March 15, 2017

Abstract

This paper estimates the effect of attention to news on financial markets, using quasi-random variation in positioning of news articles on the Bloomberg terminal. I focus on a category of news articles, some of which are pinned to the prominent "front page" positions at the top of the news screen in a process independent of their content. This offers a natural experiment in positioning between the articles that are pinned and those that are not. The front page and non-front page articles are indistinguishable by either algorithmic analysis, or by the target audience of active finance professionals. I find that pinning a news article to the front page leads to 280% higher trading volumes and 180% larger price changes within the first ten minutes after publication. These articles also see a much stronger short-term price drift, with 21% higher continuation in returns at the five-minute level. The effect is strongest during the first 30-45 minutes after the news, and partially reverses over subsequent hours. A comparison against differential reactions following news articles of varying editorial importance indicates that news positioning plays an even stronger role in driving market activity than news importance.

Keywords: information diffusion, news positioning, asset pricing, trading volume

---

*Harvard University, Department of Economics and Harvard Business School. Mail: Baker Library 244C, 25 Harvard Way, Boston, MA 02163. Email: afedyk@hbs.edu. I am particularly indebted to John Campbell, Lauren Cohen, David Laibson, Chris Malloy, Andrei Shleifer, and Jeremy Stein for their invaluable guidance and advice. I am also grateful to Malcolm Baker, John Beshears, Vitaly Bord, Kirill Borusyak, Joshua Coval, Kent Daniel, Stefano DellaVigna, William Diamond, Tatiana Fedyk, Yuriy Fedyk, Xavier Gabaix, Tom Glocer, Robin Greenwood, Sam Hanson, James Hodson, John Horton, Alex Imas, Marc Kaufmann, Owen Lamont, Andres Liberman, Marina Niessner, Matthew Rabin, Martin Schmalz, Stas Sokolinski, Adi Sunderam, Carmen Wang, Vladimir Zdorovtsov, Florian Zimmermann, and seminar participants at Harvard University and State Street Global Advisers for insightful comments. This paper was a finalist for the Hillcrest Behavioral Finance Award. I thank Lauren Cohen and Chris Malloy for sharing the QuantQuote data and Katelyn Haruki for providing outstanding research assistance. Funding for the survey of financial experts was generously provided by the Harvard Business School.

1 Introduction

How does information get incorporated into asset prices? A number of theoretical models propose potential frictions in information consumption that may prevent even publicly available information from being instantaneously reflected in prices.[1] Multiple empirical studies lend suggestive evidence to this view.[2] However, isolating the causal effect of information consumption has remained difficult without a more detailed understanding of the variation in attention to individual pieces of information.

In this paper, I capture the causal effect of persistent attention to news on market dynamics using a natural experiment in the way news articles are pinned to the top of the Bloomberg terminal news screen, in the "front page" positions. I show that positioning a news article about a security on the front page has a substantial effect on the security's short-term market dynamics. News articles that get pinned to the front page induce 280% higher trading volumes and 180% larger price changes within ten minutes of publication, and 21% more continuation in returns over consecutive five-minute intervals. The effects last for approximately 30-45 minutes after publication, and partially reverse during the following one to two hours. Interestingly, differences in news positioning have an even stronger effect on market dynamics than differences in news importance.

My empirical design exploits a natural experiment based around a category of Bloomberg news articles whose placement depends on the contemporaneous volume of other articles, rather than on their own content. The news articles in my hand-collected sample fall into three categories: "primary important," "secondary important," and "all other" news. News articles marked as "primary important" are always pinned in the prominent front page positions, displacing the previous front page news and remaining on the front page for up to several hours. News articles marked as "all other" are never placed in the front page positions. News articles marked as "secondary important" constitute the category of interesting variation. Any particular news article in this category is given a front page slot if and only if, at the precise moment when the article is released, there is at least one such slot remaining from the "primary important" news. As a result, "secondary important" news

[1] See, for example, Peng and Xiong (2006), DellaVigna and Pollet (2009), and Andrei and Hasler (2014) on limited attention to publicly available information; and Harris and Raviv (1993), Kandel and Pearson (1995), Cao and Ou-Yang (2009), and Banerjee and Kremer (2010) on differential interpretations of public information.

[2] See, for example, Foster, Olsen, and Shevlin (1984), Bernard and Thomas (1989), and Peress (2008) on the post-earnings announcement drift; Loh (2010), Da, Engelberg, and Gao (2011, 2015), Drake, Roulstone, and Thornock (2012), Schmidt (2013), and Ben-Rephael, Da, and Israelsen (2017) on predictability of market dynamics from proxies of attention; Huberman and Regev (2001), Tetlock (2011), Gilbert, Kogan, Lochstoer, and Ozyildirim (2012), and Fedyk and Hodson (2014) on reactions to stale news; and Carvalho, Klagge, and Moench (2011) and Marshall, Visaltanachoti, and Cooper (2014) on reactions to false news.

articles that make it to the front page position and those that do not are marked as equally significant. Their positions vary due to contemporaneous numbers of "primary important" articles, rather than their own underlying content.

I structure the empirical analysis of the market dynamics following front page versus non-front page "secondary important" articles using a theoretical framework that reflects standard models of limited attention and gradual information diffusion.[3] The three-period model considers a news signal published by the main news source of interest and also reported by alternative news sources. The framework incorporates two standard features from models of gradual information diffusion: (1) only a fraction of investors are attentive to the news signal from each source in each period; and (2) investors update their beliefs in a naïve Bayesian manner, incorporating their own information but not rational expectations of the information that may have been obtained by other investors. Front page positioning is represented by more prominent and longer-lasting reporting by the main news source. A larger incidence of investors are attentive to the news signal from the main source when it is published on the front page, and this persists into the second period.

The model makes several predictions. First, the front page news articles are accompanied by larger immediate trading volumes and absolute price changes. Second, the larger returns accompanying front page news articles are more likely to continue in the short-term. Third, the front page news articles are followed by *lower* longer-term price continuation.

My empirical results confirm these predictions. There are striking differences in market dynamics following "secondary important" news articles that are pinned to the front page and those that are not. Consistent with the first prediction, front page news articles are accompanied by substantially higher trading volumes and absolute price changes for the tagged securities during the minutes after publication. For example, these articles are, on average, accompanied by 280% larger trading volumes and 180% larger absolute price changes during the ten minutes immediately following article publication.

Since pinning a news article to a front page position makes it visible for a longer period of time, the front page positioning also induces more persistent short-term market reactions, confirming the second prediction of the gradual information diffusion model. Front page news articles are accompanied by significantly higher serial correlation in price changes over a variety of horizons on the order of minutes. For example, these articles are, on average, followed by 21% larger serial correlations in price changes across consecutive five-minute intervals, and 19% larger serial correlations in price changes across consecutive ten-minute intervals.

[3]For models of limited attention and gradual information diffusion, see, for example, Hong and Stein (1999), DellaVigna and Pollet (2009), and Andrei and Hasler (2014).

The effects of front page positioning are relatively short-term. Most articles are displaced from the front page by new incoming "primary important" articles within hours, and occasionally within minutes. Those that are not displaced get removed after at most a couple of hours. Correspondingly, the differential price drift following the exogenous front page positioning lasts for approximately 30-45 minutes. Over longer horizons of one to two hours, the larger returns from the first 30-45 minutes are actually subject to 15-30% more *reversal* if the news article is printed on the front page, consistent with the third prediction of the model. Thus, the initial differential reaction to equally important news in different positions at least partially corrects within hours of publication.

I compare the market effects of news positioning against the effects of news importance. In particular, I estimate market dynamics following two sets of news articles that receive equally prominent positions but that differ in their importance, as marked by the editorial staff. These are: (1) "secondary important" articles that make it to the front page; and (2) "primary important" articles, all of which make it to the front page by default. Articles in both of these categories are prominently positioned, but the articles in the second category are marked by the editorial staff, ex ante, to be more important than those in the first category.

I find that news importance is not as significant in driving short-term market activity as news positioning. Trading volumes following news publication are not statistically different for securities mentioned in more ("primary important") versus less important ("secondary important") news articles printed on the front page. Absolute price changes are 78% (62%) larger during the first five (ten) minutes following the more important news articles, but the relative difference is smaller than that induced by the news positioning. The short term price drift is statistically indistinguishable for more and less important news articles, holding front page position constant. Overall, the results indicate that news positioning plays an even larger role for short-term market dynamics than the importance of the underlying news.

I perform a number of additional analyses to confirm that the results are not driven by systematic differences between "secondary important" articles that receive a front page slot and those that do not. First, I consider the possibility that, in general, news articles published during quieter times are accompanied by larger market reactions, and that the larger reactions to the front page news articles are driven by them coming out during quieter times. To address this possibility, I hold position constant and compare non-front page "secondary important" articles released during times with different amounts of contemporaneous news activity. I document that the non-front page articles released during quiet times are, if anything, accompanied by *less* substantial reactions than the non-front page articles released during busy times.

Second, using techniques from machine learning and a representative corpus of financial news from Reuters, I learn the mixtures of topics generally discussed in financial news, such as earnings announcements, technology, and litigation. I then use the trained model to compare the distributions of identified topics appearing in the text of the individual Bloomberg news articles in my hand-collected samples. I find no systematic differences between the distributions of topics discussed in the front page versus non-front page "secondary important" news articles. The distribution of topics covered by the "primary important" news articles, by comparison, does differ slightly from the distribution of topics in the "secondary important" news.

Third, a survey of active financial professionals and MBA students from top business school programs indicates that absent salient positioning, market participants find front page "secondary important" headlines to be indistinguishable from non-front page ones. The participants consist of key decision makers at a broad range of financial institutions, including broker dealers such as Bank of America and Goldman Sachs, investment management firms such as BlackRock and PIMCO, hedge funds such a Bridgewater and Tudor, and private equity firms such as Blackstone and Warburg Pincus. These individuals consistently identify the "primary important" news articles as, on average, more impactful than "secondary important" news articles ("primary important" headlines are chosen as more impactful 63% of the time, significantly different from 50% at the 1% level). By contrast, the survey participants identify the front page "secondary important" news articles as more impactful than their non-front page counterparts only 48% of the time, not significantly different from 50%.

My findings build on the growing literature on the impact of media on financial markets.[4] Prior empirical strategies for estimating the causal impact of media use exogenous variation in news arrival through weather-related disruptions (see Engelberg and Parsons (2011)), newspaper strikes (see Peress (2014)), and disruptions to boat routes (see Koudijs (2016)), and variation in headline complexity (see Umar (2017)). Klibanoff, Lamont, and Wizman (1998) find that for closed-end country funds, the incidence of news on the front page of the *New York Times* is correlated with a higher elasticity of price with respect to asset value. Huberman and Regev (2001) further highlight the importance of prominent news positioning by comparing market reactions after similar (albeit not identical) information reported in *Nature* and non-front page articles in the media such as *New York Times* in November 1997, and then on the front page of the *New York Times* in May 1998.

The contribution of the present paper is threefold. First, exogenous variation in position-

[4]See Busse and Green (2002), Chan (2003), Fehle, Tsyplakov, and Zdorovtsov (2005), Fang and Peress (2009), Engelberg, Sasseville, and Williams (2012), Solomon (2012), Dougal, Engelberg, Garcia, and Parsons (2012), Rogers, Skinner, and Zechman (2013), Hillert, Jacobs, and Müller (2014), Ahern and Sosyura (2014), Liu, Sherman, and Zhang (2014), Yuan (2015), and Boulland, Degeorge, and Ginglinger (2016).

ing at the article-level resolution allows me to compare individual information events and to tie them to precise market dynamics for specific firms at specific times.

Second, the front page position in my setting not only induces more attention overall, but also changes the *dynamics* of attention. By construction, all news articles passing through the Bloomberg news screen begin at the top of the screen; however, the front page news articles remain at the top for longer, corresponding to more continued attention. This allows me to observe the causal effect of *persistence* in attention on market dynamics, and to directly test models of gradual information diffusion.

Third, while prior studies of the causal impact of media focus on exogenous variation in availability, form, or complexity of information, I restrict my attention to news articles that are all publicly available and indistinguishable by the target audience. My analysis offers the first systematic evidence that it is not enough to make financially-relevant information easily accessible: how saliently the information is presented plays an important role in determining whether the information is immediately and efficiently incorporated into asset prices.

The remainder of the paper proceeds as follows. Section 2 describes the data and the natural experiment in news positioning. Section 3 outlines the conceptual framework of market dynamics following more and less prominently positioned news. Section 4 presents the key empirical findings on the differential market dynamics following front page and non-front page news articles. Section 5 explores the effect of news importance, holding position constant, by comparing "secondary important" news articles that are positioned on the front page against "primary important" front page news articles. Section 6 presents additional analyses of news content, confirming that the front page "secondary important" articles in the sample are indistinguishable from their non-front page counterparts by both algorithmic analysis and the target market participants. Section 7 concludes.

2 Data Sources and Empirical Strategy

In order to capture the casual effect of the dynamics of news consumption on market dynamics, I use quasi-random variation in positioning of news articles on the Bloomberg terminal. Two key features of these data make them especially well-suited to the current analysis. First, Bloomberg is one of the largest financial news providers and a main source of news for financial professionals, making it an ideal setting to estimate the effect of attention to news on financial markets. Second, the data include a natural experiment of quasi-random positioning for a subset of news articles. The news data are merged with market data to relate news consumption to trading volume and price formation.

2.1 Natural Experiment in News Positioning

In this subsection, I describe the quasi-random variation in news positioning that I use in my research design. In particular, I concentrate on a subset of news articles that are sometimes prominently positioned, and sometimes not, depending on the volume of other articles released around the same time and not on the characteristics of the news articles themselves.

The full sample of news passing through the Bloomberg terminal is aggregated from a variety of sources in real-time. The sources of news include key national and international news wires from a comprehensive set of news organizations, company filings, press releases, and content from web sources, including blogs and social media. The news articles are disseminated electronically to over 300,000 finance professionals through the subscription-based terminal. Overall, there are over 3.5 million articles tagged with U.S. equity securities during the sample period of March 22, 2014 - March 2, 2015.

There are differences in how Bloomberg presents individual news articles on the terminal. Generally, the news screen features a scrolling list of news articles, where newly published articles replace the older ones at the top of the screen. However, some of the news articles written directly by Bloomberg News get pinned to the top of the screen, in what I term "front page" positions. At any given point in time, there are at most three such pinned articles. Figure 1 features a screenshot of a Bloomberg news screen. The top three articles are pinned and remain at the top, while the articles below continually move down as new publications arrive.

[FIGURE 1 AROUND HERE]

Effectively, there are three categories of news articles in the sample: "primary important" (PI) articles; "secondary important" (SI) articles, and "all other" articles. The assignment of individual articles to these categories reflects the journalistic and editorial opinion regarding the importance of a given piece of news. Each of the two important categories, PI and SI, comprises roughly 0.1-0.5% of all news, so both of these categories of articles capture news of fairly rare perceived importance. I hand-collect all articles that are tagged with at least one publicly traded U.S. equity security, that are published between 8AM and 5PM EST during the sample period, and that are either in the PI category (a total of 1,250 PI news articles) or the SI category (a total of 2,210 SI news articles).

For the most part, PI articles represent significant company news, such as earnings reports and M&A decisions. A few representative examples of the sample of PI news are provided in Panel 1 of Table 1.

SI articles likewise include significant events, such as changes in regulation and drug approvals. However, this sample also features news articles that are likely to capture the readers' curiosity, but that are less immediately relevant to financial markets, such as moves of top well-known traders and perks in financial firms. A few representative examples of SI articles are presented in Panel 2 of Table 1.

[TABLE 1 AROUND HERE]

The classification of the articles into categories of relative importance plays a role in how prominently the articles are positioned. When an article from the PI category is released, it is immediately placed in a prominent front page position, displacing whichever news article was in that position previously. Once on the front page, a news article remains there until the earlier of two things occurs: either a new PI article comes out and displaces the old article, or a predefined amount of time (on the order of hours) elapses. Occasionally, there are not enough PI articles at a given point in time to fill all of the front page slots. In this case, the next SI article to be published, upon its release, takes the available front page position. The process of article positioning is depicted in Figure 2.

[FIGURE 2 AROUND HERE]

As a result, there are two categories of news articles deemed equally important but having different positions: the SI articles that come out at a time when there are no available front page slots and the SI articles that come out at a time when front page slots are available. I hand-collect the positions of the SI articles in my sample. This subset of the news sample – category SI articles in various positions – forms the basis for my causal analysis.

Screening of the articles confirms that there are no systematic differences in content between SI news articles that are placed on the front page and those that are not. Both include significant events, such as:

- "T-Mobile Said to Plan to Turn Down Iliad's $15 Billion Offer" (not front page)
- "Chipotle Probed for New Outbreak of Different E. Coli Strain" (front page)

But both front page and non-front page SI news articles also feature news events that carry less immediately relevant impact for financial markets. For example:

- “Morgan Stanley Gets 90,000 Applications for Summer Program” (not front page)
- “Pimco Said to Have Discussed Firing Gross Before Exit to Janus” (front page)

In Section 6.2, I compare the texts of the front page and non-front page SI news articles formally using machine learning techniques. I find no systematic differences between the two categories of news. Similarly, a survey of active finance professionals and MBA students from top business school programs indicates that human financial experts cannot reliable distinguish the front page SI news articles from the non-front page ones.

Table 2 presents the distribution over time of PI and SI news articles published between the hours of 8AM and 5PM. All numbers are cited in ticker-articles, so that articles tagged with more than one ticker are included one time for each tagged U.S. security. Overall, there are 1,364 PI article-tickers in the sample and 4,055 SI article-tickers, of which 627 are given a front page position. The articles are roughly evenly distributed across the months of the year, with fewer articles in March since the sample covers only twelve days in March (March 22-31, 2014 and March 1-2, 2015). Over hours of the day, PI news articles peak at the start and end of the business day, during 8-9AM and especially during 4-5PM, while the SI news articles are more evenly distributed during the day, except a lower incidence during the 4-5PM window. Consistent with the SI articles’ positioning being determined by the concurrent volume of PI articles, a lower percentage of SI articles makes it to the front page during the hours that see a higher volume of PI articles. The correlation between the hourly numbers of PI articles and the hourly likelihoods of SI articles receiving front page positions is -56%.

[TABLE 2 AROUND HERE]

Examining the timing of news releases in the sample, I find no evidence of strategic release timing of the SI news articles. In particular, given the patterns in the data, it is unlikely that the editorial staff is faced with multiple SI articles to be released at the same time and strategically releases the more important ones first. Of all the front page SI news articles, only 1.9% have a non-front page article released up to one minute before or after the front page article’s publication. Similarly, only 1.1% of the front page SI articles are accompanied by non-front page SI articles within 30 seconds before or after. A mere 0.2% of the articles in the front page SI sample have a non-front page article released within 10 seconds of their publication. Overall, the volume of SI news articles is sufficiently low as to leave little scope for influencing article position by strategically timing the exact seconds of when the articles are released.

I also find that the article volume does not appear to be driven by editorial targets. I observe the distribution of articles across days and find that the volumes of PI and SI news articles vary dramatically from day to day. The number of PI news articles ranges from 0 to 28 per day, while SI articles can number anywhere between 0 and 67 per day. There is also little relationship between the numbers of PI and SI articles on any given day. The daily numbers of the two types of articles display a low correlation of 27%. As shown in Figure 3, any given day can see a large number of PI articles accompanied by few SI articles, and vice versa. Overall, the distribution of PI and SI articles across days indicates that the editorial staff is not targeting particular numbers of high-importance articles. Instead, the patterns are more consistent with the evaluation of each article's importance being based on its own merit, independently of the volume of other news.

[FIGURE 3 AROUND HERE]

2.2 Market Data

I use the security ticker tags to merge the news position data with market data from several sources. Industry classification, market capitalization, and shares outstanding come from Compustat. High frequency price and trading data come from QuantQuote. The second-resolution QuantQuote data include all tickers listed on NYSE and NASDAQ exchanges, and provide prices and numbers of shares traded for each second during the market open. The data are adjusted for splits, dividends, and symbol changes.

The high frequency tests are run using news articles tagged with all firms for which there are pricing data in QuantQuote, and shares outstanding and NAICS industry codes in Compustat. The merged sample includes 459 front page SI news articles, 2,631 non-front page SI news articles, and 495 PI news articles. Merging with market data reduces the PI sample more substantially than the sample of SI news articles due to the fact that PI news articles are more likely to come out during the hours of 8-9AM EST and especially 4-5PM EST. As a result, it is more common for these news articles not to be accompanied by any quotes or trades in the QuantQuote data within an hour of publication.

3 Conceptual Framework

Before estimating the market impact of news positioning, I derive the key empirical predictions from a conceptual framework of investor attention to front page and non-front page news.

The conceptual framework follows the setups in Hirshleifer and Teoh (2003) and DellaVigna and Pollet (2009). There is a risk-free asset with a zero rate of return and a single risky security with a stochastic payoff R normally distributed with mean $\overline{R}$ and variable σ_R^2, realized in an unmodeled final period T. In the relatively short-term empirical settings that I consider, the realized value R can be taken to denote, for example, the end of day price for day-traders trading on individual news releases, or the price on which an asset settles in the days following an earnings announcement. The risky asset is in fixed supply X. For expositional simplicity, I fix $X = 0$, so that the asset is in zero net supply; this simplifies the notation without affecting the results.

There is a continuum of investors with total mass equal to 1, who maximize mean-variance utility. In particular, let $W^{(i)}$ denote investor i's final wealth at the end of the game at time T. Then at any point in time t, investor i maximizes expected utility of the form

$$\mathbb{E}_{i,t}\{W^{(i)}\} - \frac{A^{(i)}}{2} Var_{i,t}\{W^{(i)}\} \tag{1}$$

with respect to his current holdings. For simplicity, I take the risk-aversion coefficient to be identical across investors: $\forall i, A^{(i)} = A$. Each investor i is initially endowed with wealth $W_0^{(i)}$. There are no liquidity constraints.

Information in this framework is modeled as a signal arriving at a particular point in time and gradually diffusing across the population of investors. In particular, there are four periods in the model. In period 0, investors form prior expectations regarding the distribution of R. In period 1, a noisy signal (news) is released, and investors update their expectations accordingly. In periods 2 and 3, investors continue to update their beliefs following the news signal. At the end of the game, in the unmodeled period T, the true value of R is realized and the investors consume their final wealth. I assume the following form for the news signal: $N = R + \epsilon$, where ϵ is a normally distributed noise term, independent of R, with mean μ and variance σ_ϵ^2.

The news signal is not immediately observed by all investors. Instead, the main news source, S, reports the news signal N for some number of periods. Mass γ of investors are attentive to the main source S in each period t. Thus, in each period t that S reports the news signal N, a fraction γ of investors who had not observed the news signal prior to t now become aware of N.

I model the difference between front page and non-front page news with two key features. First, front page news articles induce more attention overall, so that the fraction of investors attentive to the news signal is higher: $\gamma = \overline{\gamma}$ in the case of front page news and $\gamma = \underline{\gamma} < \overline{\gamma}$

in the case of non-front page news. Second, front page news articles correspond to signals being reported by S for longer. Thus, for non-front page news, investors can observe the signal N from the main source S only in period 1. For front page news, however, investors can also observe the signal from the main source S in period 2.

Investors may also learn the news from alternative sources, albeit at a lower rate. In particular, in any period when the news is not being reported by S, a fraction ξ of uninformed investors still observe the news signal. This additional information channel can be interpreted as investors finding the news through filters or active searches once it scrolls off the top of the Bloomberg terminal screen, or reading the news from other publications. This channel is a minor one in the model; in particular, I assume that:

$$\xi < \frac{1-\overline{\gamma}}{1-\underline{\gamma}}\underline{\gamma} \tag{2}$$

This condition ensures that once the main news source stops actively reporting the news (i.e., when the news is not on the front page), the fraction of informed investors does not increase faster than when the source continues to report (front page news). Consistent with the information disseminating relatively slowly over the short horizons considered in my empirical analysis, I also assume that both γ and ξ are small: $\gamma, \xi << 1/2$.

The model timeline is depicted in Figure 4. In each period t, let I_t denote the set of the informed investors who observe the news signal either during or prior to t, and let $F_t = |I_t|$ be the share of informed investors. I denote the remaining uninformed investors by U_t. Let F_t^{NFP} and F_t^{FP} denote the values of F_t in the cases of front page and non-front page news, respectively. Figure 4 illustrates the arrival of information and the evolution of the share of informed investors for both front page and non-front page news.

[FIGURE 4 AROUND HERE]

The key frictions in the model are that (1) some investors are inattentive; and (2) the investors update their beliefs in a naïve Bayesian manner. Namely, some of the investors do not observe the public signal, and all investors update their beliefs with respect to only their own information, without taking into account the information sets and actions of others. These assumptions are common modeling devices in models of gradual information diffusion (see Hong and Stein (1999), Hirshleifer and Teoh (2003), or Peng and Xiong (2006)).

I characterize the price path and trading volume following a news signal as a function of the fraction of attentive investors F_t (Section 4.1). The empirical predictions for the differences in market dynamics following front page and non-front page news are then derived in Section 4.2.

3.1 Evolution of Prices and Trading Volumes

I begin by characterizing the price levels and trading volumes in terms of the fraction of attentive investors F_t, without distinguishing whether the news signal is reported on the front page or not.

Price levels. First, note that the uninformed investors hold the prior beliefs that the return R is normally distributed with mean $\overline{R}$ and variance σ_R^2. The informed investors update their beliefs in a naïve Bayesian manner, incorporating the news signal. Hence, their beliefs are given by:

$$\forall t \in \{0,1,2,3\}, i \in I_t : \mathbb{E}_t^{(i)}\{R\} = \frac{\sigma_\epsilon^2 \overline{R} + \sigma_R^2 (N-\mu)}{\sigma_R^2 + \sigma_\epsilon^2}; Var_t^{(i)}\{R\} = \frac{\sigma_R^2 \sigma_\epsilon^2}{\sigma_R^2 + \sigma_\epsilon^2} \tag{3}$$

Next, note that optimization of the mean-variance preferences given by (1) with the above beliefs results in the following demand functions by the two groups of investors during any period t:

$$\forall t \in \{0,1,2,3\}, i \in I_t : x_t^{(i)} = \frac{\sigma_\epsilon^2 \overline{R} + \sigma_R^2 (N-\mu) - (\sigma_R^2 + \sigma_\epsilon^2) P_t}{A \sigma_R^2 \sigma_\epsilon^2} \tag{4}$$

$$\forall t \in \{0,1,2,3\}, i \in U_t : x_t^{(i)} = \frac{\overline{R} - P_t}{A \sigma_R^2} \tag{5}$$

where P_t denotes the price of the risky asset in period t.

The market clearing condition each period is that the total demand from the informed and uninformed investors must equal the zero net supply. Hence, in each period t, the price of the asset P_t must satisfy:

$$\forall t \in \{0,1,2,3\} : F_t \frac{\sigma_\epsilon^2 \overline{R} + \sigma_R^2 (N-\mu) - (\sigma_R^2 + \sigma_\epsilon^2) P_t}{A \sigma_R^2 \sigma_\epsilon^2} + (1-F_t) \frac{\overline{R} - P_t}{A \sigma_R^2} = 0 \tag{6}$$

Solving this equation gives the following expression for the price of the asset during each period t:

$$\forall t \in \{0,1,2,3\} : P_t = \frac{\sigma_\epsilon^2 \overline{R}}{\sigma_\epsilon^2 + F_t \sigma_R^2} + \frac{F_t \sigma_R^2}{\sigma_\epsilon^2 + F_t \sigma_R^2} (N-\mu) \tag{7}$$

Absolute price changes. Taking the first differences yields the absolute price changes between any two consecutive periods:

$$\forall t \in \{1,2,3\} : \Delta P_t = |P_t - P_{t-1}| = \frac{(F_t - F_{t-1})\sigma_R^2 \sigma_\epsilon^2 |N - \mu - \overline{R}|}{(\sigma_\epsilon^2 + F_t\sigma_R^2)(\sigma_\epsilon^2 + F_{t-1}\sigma_R^2)} \tag{8}$$

Price continuation. In order to calculate the continuation in the price path, recall that the news signal has the form $N = R + \epsilon$, where R and ϵ are independent normal variables with $R \sim \mathcal{N}(\overline{R}, \sigma_R^2)$ and $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$. Hence, price continuation, measured as the slope in a regression predicting next period's price change from the current period's price change, is given by:

$$\forall t \in \{1,2\} : Cont(t, t+1) = \frac{Cov(\Delta P_t, \Delta P_{t+1})}{Var(\Delta P_t)} = \left(\frac{F_{t+1} - F_t}{F_t - F_{t-1}}\right)\left(\frac{\sigma_\epsilon^2 + F_{t-1}\sigma_R^2}{\sigma_\epsilon^2 + F_{t+1}\sigma_R^2}\right) \tag{9}$$

Trading volumes. Trading volume in each period t is defined as half of the sum of absolute changes in holdings, across all investors, from the prior period $t-1$ to t. In each period, the newly informed investors, i.e. investors $i \in I_t \cap U_{t-1}$ change their demand following receipt of the news signal, inducing a change in the equilibrium price and the other investors' equilibrium holdings. Let $x_t^{(I)}$ denote the equilibrium holdings, in period t, of an investor $i \in I_t$; similarly, let $x_t^{(U)}$ denote the equilibrium holdings of an investor $u \in U_t$. Trading volume in each period can be expressed as a function of the newly informed investors' holdings as follows:

$$\forall t \in \{1,2,3\} : TV_t = (F_t - F_{t-1})|x_t^{(I)} - x_{t-1}^{(U)}| \tag{10}$$

Taking the holdings from (4)-(5) and the equilibrium price levels from (7) then gives the following expression for each period's trading volume:

$$\forall t \in \{1,2,3\} : TV_t = (F_t - F_{t-1})\frac{(1 - F_t)\sigma_\epsilon^2 + F_{t-1}(\sigma_\epsilon^2 + \sigma_R^2)}{A(\sigma_\epsilon^2 + F_t\sigma_R^2)(\sigma_\epsilon^2 + F_{t-1}\sigma_R^2)}|N - \mu - \overline{R}| \tag{11}$$

3.2 Empirical Predictions

I now compare the expressions for price changes, trading volumes, and price continuation for front page and non-front page news, and derive empirical predictions for differential market dynamics following different article positions.

Before proceeding, I note the evolution of the share of informed investors, F_t, in the cases of front page and non-front page news. In the first period, $F_0^{NFP} = F_0^{FP} = 0$. After that, the share of informed investors after non-front page news evolves as follows:

$$F_t^{NFP} = \begin{cases} \underline{\gamma} & \text{for t} = 1 \\ \underline{\gamma} + (1 - \underline{\gamma})\xi & \text{for t} = 2 \\ \xi + (1 - \xi)(\underline{\gamma} + (1 - \underline{\gamma})\xi) & \text{for t} = 3 \end{cases} \tag{12}$$

For front page news, meanwhile, the share of informed investors evolves as follows:

$$F_t^{FP} = \begin{cases} \overline{\gamma} & \text{for t} = 1 \\ \overline{\gamma} + (1 - \overline{\gamma})\overline{\gamma} & \text{for t} = 2 \\ \xi + (1 - \xi)(\overline{\gamma} + (1 - \overline{\gamma})\overline{\gamma}) & \text{for t} = 3 \end{cases} \tag{13}$$

Combining the shares of informed investors in (12)-(13) with the price changes in (8) gives the immediate absolute price changes after front page and non-front page news:

$$\Delta P_1^{NFP} = \frac{\underline{\gamma}}{\sigma_\epsilon^2 + \underline{\gamma}\sigma_R^2}\sigma_R^2\sigma_\epsilon^2|N - \mu - \overline{R}|; \quad \Delta P_1^{FP} = \frac{\overline{\gamma}}{\sigma_\epsilon^2 + \overline{\gamma}\sigma_R^2}\sigma_R^2\sigma_\epsilon^2|N - \mu - \overline{R}| \tag{14}$$

Given that $\overline{\gamma} > \underline{\gamma}$, the absolute price change is larger following front page news than following non-front page news.

Similarly, the trading volumes at the news release in period $t = 1$ are given by:

$$TV_1^{NFP} = \frac{\underline{\gamma}(1 - \underline{\gamma})}{A(\sigma_\epsilon^2 + \underline{\gamma}\sigma_R^2)}|N - \mu - \overline{R}|; \quad TV_1^{FP} = \frac{\overline{\gamma}(1 - \overline{\gamma})}{A(\sigma_\epsilon^2 + \overline{\gamma}\sigma_R^2)}|N - \mu - \overline{R}| \tag{15}$$

The relationship between immediate trading volume around the news signal and the percentage of immediately informed investors is non-monotonic. Trading volume is low if either all or none of the investors see the news immediately, and trading volume is maximized when the split between immediately attentive and inattentive investors is roughly even. In the setting I consider, the proportion of the population who see any news article immediately (within the first few minutes of publication) is relatively low (substantially less than half of the population) even for front page news. As a result, the split of attentive versus inattentive investors is more even and the immediate trading volume is higher when the news is printed on the front page.

Together, the price and volume expressions give the first empirical prediction regarding the immediate market response to front page and non-front page news.

Prediction 1 (Immediate Market Response) *Front page news articles are followed by larger trading volumes and absolute price moves immediately (within minutes) after the news.*

How does the price response play out outside of the immediate window? To see this, I turn to the continuation in the price path. I begin with the short-term continuation:

$$Cont^{NFP}(\Delta P_1, \Delta P_2) = \frac{(1-\underline{\gamma})\xi}{\underline{\gamma}} \times \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2 + [\underline{\gamma} + (1-\underline{\gamma})\xi]\sigma_R^2} \tag{16}$$

$$Cont^{FP}(\Delta P_1, \Delta P_2) = (1-\overline{\gamma}) \times \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2 + [\overline{\gamma} + (1-\overline{\gamma})\overline{\gamma}]\sigma_R^2} \tag{17}$$

Note that from condition (2), the first term of $Cont^{FP}(\Delta P_1, \Delta P_2)$ is larger than the first term of $Cont^{NFP}(\Delta P_1, \Delta P_2)$. The second term is larger in $Cont^{NFP}(\Delta P_1, \Delta P_2)$, since $\overline{\gamma} > \underline{\gamma} > \xi$. However, for sufficiently low levels of immediate attention $\overline{\gamma}$ and $\underline{\gamma}$, the former effect dominates. This results in the following empirical prediction.

Prediction 2 (Immediate Return Continuation) *Front page news articles are accompanied by higher continuation in the short-term price changes.*

While front page news articles are followed by a larger immediate reaction that continues in the short-term, the longer term dynamics are quite different. To see this, note that the continuation in returns from the second to the third period for front page and non-front page news is given by:

$$Cont^{NFP}(\Delta P_1, \Delta P_2) = (1-\xi) \times \frac{\sigma_\epsilon^2 + \underline{\gamma}\sigma_R^2}{\sigma_\epsilon^2 + [\xi + (1-\xi)(\underline{\gamma} + (1-\underline{\gamma})\xi)]\sigma_R^2} \tag{18}$$

$$Cont^{FP}(\Delta P_1, \Delta P_2) = \frac{(1-\overline{\gamma})\xi}{\overline{\gamma}} \times \frac{\sigma_\epsilon^2 + \overline{\gamma}\sigma_R^2}{\sigma_\epsilon^2 + [\xi + (1-\xi)(\overline{\gamma} + (1-\overline{\gamma})\overline{\gamma})]\sigma_R^2} \tag{19}$$

Note that with $\xi < \underline{\gamma} < \overline{\gamma} << 1/2$, expressions (18)-(19) imply that the continuation from the second period to the third is actually lower for front page news compared to non-front page news. This yields the third empirical prediction of the gradual information diffusion framework.

Prediction 3 (Delayed Return Continuation) *Front page news articles induce lower continuation in the long-term price changes.*

In the next section, I test Predictions 1, 2, and 3 by observing the market dynamics following front page and non-front page Bloomberg news articles in my hand-collected sample. For the immediate news release window, $t = 1$, I look at the 5-10 minutes after each individual news article is published. As the short-term subsequent window, $t = 2$, I consider 30-45 minutes following the news, as the front page news articles tend to remain prominently positioned for no more than a couple of hours and sometimes as little as a few minutes. For the longer horizon, $t = 3$, I consider windows of 60, 90, and 120 minutes following the news release.

4 News Positioning and Market Dynamics

Using the natural experiment in news positioning, in this section, I present the empirical analysis of the causal effect of the dynamics of news consumption on financial markets.

4.1 News Positioning and Short-Term Market Dynamics

I begin the analysis of differential activity following comparable front page and non-front page news articles by observing the short-term trading volume surges and price dynamics following the two types of SI news. Placing a piece of news on the front page is associated with substantially larger trading volumes and absolute price changes within minutes of publication, as well as higher continuation in the short-term price paths.

Consistent with Prediction 1, the more saliently positioned front page news articles induce significantly higher trading volumes. The median 15-second trading volume, computed as the percentage of shares turned over during the ten minutes before and after SI news articles is displayed in Panel 1 of Figure 5. The median non-front page SI news article is accompanied by virtually no increase in trading volume (plotted in light blue in the figure) relative to the pre-news baseline. There is, however, a pronounced increase in the trading volumes following SI news articles that appear on the front page (displayed in dark blue). The difference in averages is even starker. Over the ten minutes after a news release, the average non-front page SI news article is accompanied by at total of 0.05% turnover. The average front page article is followed by a 280% higher ten-minute trading volume of 0.19%. The difference is statistically significant at the 1% level, with a t-statistic of 4.12, as reported in Panel 1 of Table 3. The estimated difference remains identical when controlling for month and hour fixed effects, log market capitalization, and industry fixed effects.

[FIGURE 5 AROUND HERE]

[TABLE 3 AROUND HERE]

Does the increased market activity reflected in trading volume correspond to increased price volatility? Panel 2 of Figure 5 presents the average absolute price changes following front page and non-front page SI news articles. The absolute price changes are calculated separately for each firm every five seconds. The graph averages the price changes in event time over the cross-section of firms.

Two patterns emerge from a visual inspection of the absolute price changes. First, the overall price change from the time of news publication to ten minutes later is much larger for SI news articles that are positioned on the front page than for those that are not. Second, corresponding to the more persistent attention garnered by the front page news articles being saliently positioned for longer, price changes after these news articles are more persistent. These two effects are investigated in greater detail below.

I begin the analysis of price changes by looking at the differential immediate price reactions to front page and non-front page SI news articles. Lending further support to Prediction 1, the average absolute price change within the first ten minutes after front page SI news articles is 60 basis points, compared to 21 basis points for non-front page SI news. The difference of 39 basis points is statistically significant at the 1% level, with a t-statistic of 5.06, as displayed in Panel 2 of Table 3. The result is robust to the inclusion of controls: the estimated difference is 40 basis points when accounting for month and hour fixed effects, and 36 basis points when also controlling for log market capitalization and industry fixed effects. The results are similar at a shorter horizon of five minutes following the news, with an average absolute price change of 45 basis points accompanying front page news articles, compared to 16 basis points for non-front page news articles (t-statistic on the difference of 4.92). The contrast is less stark, but still highly significant when the window is extended to one hour following the news. The average absolute price change over the hour following front page SI news articles is 1.09%, whereas the average absolute price change over the hour following non-front page SI news articles is 0.51% (t-statistic on the difference of 3.95)

Having established empirical support for the first prediction of the gradual information diffusion framework, I now turn to Prediction 2. As indicated by Panel 2 of Figure 5, the price paths following the front page SI news articles display more continuation, reflecting the more persistent attention garnered by news articles that stay at the top of the terminal screen for longer. I test the extent to which the front page positioning induces higher return continuation formally by running the following specification:

$$Ret_{s,i,[t+t_1,t+t_2]} = \alpha + \beta_1 Ret_{s,i,[t,t+t_1]} + \beta_2 FP_s + \beta_3 Ret_{s,i,[t,t+t_1]} \times FP_s \qquad (20)$$

$$+\gamma X_{i,t} + \epsilon_{s,i,[t+t_1,t+t_2]},$$

where $Ret_{s,i,[t,t+t_1]}$ denotes the return on security i during the immediate period $[t, t+t_1]$ after publication of the news article s, and $Ret_{s,i,[t+1,t+t_2]}$ is the return during the delayed period $[t+t_1, t+t_2]$. FP_s is an indicator variable equal to one for SI news articles that are pinned to the front page and zero for SI news articles not on the front page. The controls $X_{i,t}$ include month and hour of day fixed effects, as well as log firm size and industry fixed effects. The tests are run over the following time windows: $(t_1, t_2) \in \{(3\,min, 5\,min), (5\,min, 10\,min),$ $(5\,min, 15\,min)$, $(5\,min, 20\,min)$, $(10\,min, 20\,min)$, $(10\,min, 30\,min)\}$.

Confirming Prediction 2, front page news articles are followed by higher serial correlation in returns at all considered short-term horizons, except for the shortest horizon of $(t_1, t_2) = (3\,min, 5\,min)$. The coefficient of interest, β_3, is positive and highly statistically significant across the other time specifications, as displayed in Table 4. For example, relative to non-front page SI news articles, front page SI news articles induce 21% more continuation in returns from the first five minutes after publication to the next five minutes. This result is economically sizable. For every 1% price move within the first five minutes of a front page SI news articles, there is an additional 21 basis points move in the same direction during the following 5 minutes, compared to non-front page SI news articles. The effect is also precisely estimated, with t-statistics of 5.89 without controls, 5.88 with month and hour fixed effects, and 5.86 with the full set of controls including log firm size and industry fixed effects. Results over other windows are qualitatively similar, and quantitatively stronger for most specifications.

[TABLE 4 AROUND HERE]

Interestingly, the non-front page SI news articles are actually followed by short-term return reversal from the first five minutes to the next five to ten minutes, consistent with the literature on short-term price reversals.[5] The coefficient on $Ret_{s,i,[t,t+t_1]}$ not interacted with the front page indicator is negative and statistically significant for $(t_1, t_2) \in \{(5\,min, 10\,min), (5\,min, 15\,min)\}$. Effectively, these news articles, which are prominently positioned at the top of the terminal screen only for short periods of time, see the initial five-minute price reactions partially reverse within the following minutes. On the other hand, front page SI news articles, which are prominently positioned for longer, are followed by a strong price drift over the short term.

[5] See, for example, Ederington and Lee (1995), Chordia, Roll, and Subrahmanyam (2002), Zawadowski, Andor, and Kertész (2006), and Heston, Korajczyk, and Sadka (2010).

4.2 News Positioning and Longer-Term Price Dynamics

Placing a piece of news on the front page induces short-term continuation in returns, but for how long does this effect persist? As the front page news articles get removed from their prominent positions, the differences in diffusion of information contained in these articles and the non-front page articles gradually diminish. The conceptual framework predicts that at longer horizons, front page news articles should see *less* continuation in returns. I find evidence that some of the higher price continuation accompanying front page news at short horizons reverses over the longer horizons of one to two hours after the news.

To evaluate the longer-term price dynamics, I estimate the specification (20) over the following windows: $t_1 \in \{5\,min, 10\,min\}$ and $t_2 \in \{45\,min, 60\,min, 90\,min\}$. The results are reported in Panel 1 of Table 5.

[TABLE 5 AROUND HERE]

The results reveal an interesting pattern of dynamics: the immediate returns over the first five minutes after news publication are more predictive of subsequent returns following front page news relative to non-front page news up to approximately forty-five minutes. But over longer horizons of sixty or ninety minutes, the effect is no longer present. Continuation in returns from the first five minutes to the remainder of the first hour or ninety minutes is statistically indistinguishable for front page versus non-front page SI news articles, indicating that the drift over the shorter horizons reverses during the hour following the news.

Similarly, the difference between the continuation of the first ten-minute return following front page versus non-front page news articles is greatest over the forty-five minute window. During the first forty-five minutes, front page news articles induce an additional drift of nearly 60% of the initial ten-minute return. Expanding the window to sixty or ninety minutes, only half of the effect remains present, with greatly reduced statistical significance.

As I shift the window even further, the results lend empirical support to Prediction 3. Panel 2 of Table 5 reports estimates of specification (20) over the following windows: $t_1 \in \{30\,min, 45\,min\}$ and $t_2 \in \{90\,min,\ 120\,min\}$. The non-front page news articles are followed by a 24% continuation in returns from the first half-hour to the remainder of the 90-120 minutes. Front page news articles, however, see 15-20% less continuation. The difference is highly statistically significant. Similarly, the returns from the first forty-five minutes are substantially more likely to reverse if the news article is pinned to the front page. Non-front page news articles see a continuation of 14% from the first 45 minutes to the remainder of the first 90-120 minutes. By contrast, front page news articles actually see a substantial reversal over the same time windows.

Coupled with the results in Table 4, the longer-term price dynamics suggest that the effect of news positioning on market dynamics is substantial but relatively short-term. Pinning a piece of news on the front page induces a stronger drift in returns up to forty-five minutes after the news, which then partially reverses over the remainder of the first couple of hours after news publication. Theoretically, these patterns are fully consistent with the gradual information diffusion framework outlined in Section 3. Practically, the results indicate that for news articles consumed by sophisticated finance professionals through a subscription-based platform such as Bloomberg, the market dynamics track the discretionary positioning in real time.

5 News Positioning versus News Importance

In this section, I compare the estimated effects of news positioning against the effects of news importance. I estimate the relationship between news importance and market dynamics by holding news position constant and varying news importance – i.e., by comparing front page news articles from the PI and SI categories. The difference in market reactions following these two types of news is smaller than the difference induced by the front page positioning.

I limit my attention only to news articles that are pinned to the front page, so that there is no variation in the prominence of the article positions. I include all front page news articles, regardless of whether they are marked as "primary important" or "secondary important," and estimate the difference in market reactions following the more ("primary important") and less important ("secondary important") news articles.

First, I note that the trading volumes immediately following front page PI news articles do not significantly differ from the trading volumes following front page SI news articles. As displayed in Panel 1 of Table 6, during the first five minutes after a front page news article, on average, an additional 0.09% of shares turn over when the article is from the PI category, but this difference is not statistically significant. Similarly, during the first ten minutes, PI front page news articles are followed by an additional 0.11% in trading volume compared to front page SI news articles, also not statistically significant. The pattern remains similar over longer horizons, with an average of 0.22% additional shares turned over during the hour following PI front page news articles, with the difference remaining statistically insignificant.

[TABLE 6 AROUND HERE]

Second, while the absolute price changes are larger after front page PI articles than after front page SI articles, the effect is less significant and less persistent than the difference in absolute price changes induced by front page positioning. As can be seen from Panel 2 of Table 6, in the first five minutes, front page PI news articles are followed by an additional 0.35% absolute price change, an increase of 78% over the front page SI articles; the difference is significant but statistically weaker than the difference between front page and non-front page SI articles. The difference in absolute price changes following front page PI news articles versus front page SI news articles remains similar in magnitude and declines in statistical significance as the window is extended to ten and then sixty minutes. Overall, front page PI news articles are followed by larger price reactions immediately in the first five minutes, but do not see further differences from the front page SI news articles over longer horizons. This contrasts with the difference between front page and non-front page SI news articles documented in Table 3, which continues to grow over the hour following the news.

This result is corroborated by a comparison of the continuation in the price paths following PI and SI news articles, which I estimate using the following specification:

$$Ret_{s,i,[t+t_1,t+t_2]} = \alpha + \beta_1 Ret_{s,i,[t,t+t_1]} + \beta_2 PI_s + \beta_3 Ret_{s,i,[t,t+t_1]} \times PI_s \qquad (21)$$

$$+\gamma X_{i,t} + \epsilon_{s,i,[t+t_1,t+t_2]},$$

where $Ret_{s,i,[t,t+t_1]}$ denotes the return on security i during the immediate period $[t, t+t_1]$ after publication of the news article s, and $Ret_{s,i,[t+1,t+t_2]}$ is the return during the delayed period $[t + t_1, t + t_2]$. PI_s is an indicator variable equal to one if the front page article comes from the "primary important" category and zero if the article is from the "secondary important" category. The controls $X_{i,t}$ include month and hour fixed effects, log firm size, and industry fixed effects. The considered time windows are $(t_1, t_2) \in \{(5\,min, 10\,min), (5\,min, 15\,min)\}$. The results are reported in Table 7.

[TABLE 7 AROUND HERE]

The estimated coefficient on $Ret_{s,i,[t,t+t_1]} \times PI_s$ indicates that a front page PI news article is not accompanied by any more short-term price drift compared to a front page SI news article. The difference is neither economically notable, nor statistically significant. Over the same time horizons, the difference in price drift following front page and non-front page SI news articles is 21-26% and highly statistically significant (see Table 4).

Recall that this analysis considers *only* news articles positioned on the front page, but of both categories: "primary important" and "secondary important." Whereas the results in

Section 4.2 keep article importance constant (only SI articles) and vary front page positioning, the analyses in this section keep the positioning constant but vary article importance. As can be seen from a comparison of Tables 6-7 against Tables 3-4, differences in article importance correspond to milder differences in market dynamics than differences in article positioning. These findings suggest that article positioning is even more instrumental in driving market reactions than differences in article importance, as perceived by Bloomberg's journalistic and editorial staff.

Overall, news articles that are prominently positioned on the front page are accompanied by higher trading volumes, larger returns, and a substantially stronger drift that partially reverses over the longer term. The effects induced by news positioning are even stronger than the effects of news importance. These patterns support the importance of gradual information diffusion and highlight news consumption as playing a significant role in causally driving market dynamics around information releases.

6 Additional Analyses

I present additional analyses confirming the exclusion restriction of my natural experiment design: that the SI news articles that are pinned to the front page do not systematically differ from those that are not. First, I show that, holding position constant, the news articles published during quiet times (when more front page slots are available) do not generally induce stronger market reactions than the articles published during busy times. Second, I use machine learning techniques to show that the distribution of topics discussed in the texts of the front page and non-front page SI news articles do not systematically differ. Lastly, I conduct a survey of active finance professionals and MBA students at top business school programs to highlight that, in absence of the differential positioning, the target audience finds the two sets of news to be indistinguishable in terms of the importance and expected market impact of their content.

6.1 Quiet Times vs. Busy Times

In this subsection, I address the potential concern that the differential reactions to front page and non-front page SI news articles are driven by the fact that the former are released during generally quieter times (when there are fewer PI news articles; see Panel 2 of Table 2), rather than by different amounts of attention to the two types of articles. A few points are worth noting here.

First, to the extent that increased market activity during quiet times reflects increased attention dedicated to the security due to few other contemporaneous events, the results would still capture the attention channel. In fact, various indicators of "quiet times" have been used as indirect proxies for attention in prior work (see DellaVigna and Pollet (2009) on earnings announcements released on Fridays; or Hirshleifer, Lim, and Teoh (2009) on earnings announcements released contemporaneously with other announcements). The analysis in Section 4 captures the variation in attention more precisely through the salience of news positioning.

Second, in my sample, the SI news articles published during quiet times, when little goes on in the markets, are likely to be accompanied, if anything, by *less* market activity than the SI news articles published during the busier times. This would push in the direction of finding less market activity after the front page SI news articles (i.e., SI news articles released during the quiet times), dampening my results.

I test this conjecture by comparing the non-front page SI news articles released during relatively quiet times with non-front page SI news articles released during relatively busy times. Thus, I hold news position (non-front page) and news importance (SI) constant, and vary only the numbers of contemporaneous news releases.

To differentiate busy times from quiet times, I consider the contemporaneous volumes of articles within three time intervals: (1) on the same day as a given non-front page SI article; (2) within five hours of a given non-front page SI news article; and (3) within two hours of the given article. Non-front page SI news articles for which the contemporaneous volumes of other news fall at or below the median form the "quiet times" sample. Non-front page SI news articles with above-median contemporaneous volumes of other news form the "busy times" sample.

As displayed in Table 8, holding editorial importance markings and position constant, SI news articles that come out during quieter times are **not** accompanied by larger trading volumes and absolute price changes than the SI news articles published during busier times. If anything, price changes and trading volumes are smaller following non-front page SI news articles published during quiet times. These patterns are qualitatively consistent across definitions of quiet and busy times based on contemporaneous news volumes over the one day, five hour, and two hour windows. Statistically, the difference in absolute price changes after non-front page articles that come out during quiet and busy times is only discernible when the volume of contemporaneous news is measured on a daily level. Similarly, the difference in trading volumes is significant at the 1% level only using the daily window to define busy times, and at the 5% level otherwise.

These results confirm that the differential market reactions following front page and non-front page SI news articles are not driven by the SI news articles that come out during quiet times (and are therefore more likely to take an available front page position) carrying more important content than the articles that come out during busy times.

[TABLE 8 AROUND HERE]

6.2 Distributions of Topics

To rule out systematic differences in the texts of front page and non-front page articles, I directly analyze the text of the news articles across different positions and levels of importance. The distribution of topics discussed in front page SI news articles is statistically indistinguishable from the distribution of topics covered by non-front page SI articles. By contrast, the distribution of topics discussed in PI news articles does differ somewhat from the SI news articles, with a larger focus on company operations and the healthcare industry, and lower coverage of regulations and the financial services industry.

Topic analysis provides an intuitive way to compare the content value of different news articles. The existing literature on the effect of news on financial markets considers textual characteristics such as sentiment,[6] grammatical structure,[7] and complexity.[8] The methodology in this section contributes to the literature by proposing an intuitive approach to identifying common topics in financial news and representing the news articles in terms of these prototypical topics.

The topic analysis proceeds in two steps. First, I use a large corpus from Reuters to analyze textual patterns in financial news by representing the articles in the space of meaningful features and identifying a set of broadly applicable topics. Second, I apply the trained topic model to the news articles in the PI, front page SI, and non-front page SI samples of hand-collected Bloomberg news articles.

For the first step of the process, I use the Latent Dirichlet Allocation algorithm proposed by Blei et al (2003), following similar methods employed in genetics (see, for example, Pritchard, Stephens, and Donnelly (2000)). The Latent Dirichlet Allocation approach is particularly well suited to the problem at hand, because it represents all documents as being generated from an underlying set of topics by a latent process. This admits modeling of out-of-sample documents as mixtures over the topics identified from the training data –

[6]See Tetlock (2007), Tetlock, Saar-Tsechansky, and Macskassy (2008), Loughran and McDonald (2011), Garcia (2013), and Uhl (2014).

[7]See Engelberg (2008).

[8]See Li (2008), You and Zhang (2008), Miller (2010), Lehavy, Li, and Merkley (2011), Loughran and McDonald (2014), and Umar (2007).

i.e., modeling the news articles from the various Bloomberg categories in terms of topics identified from the larger sample of Reuters news. For a description of the Latent Dirichlet Allocation methodology, please refer to Appendix A.1.

In order to train the topic model on a dataset that is similar yet distinct from the Bloomberg news articles that I ultimately classify and evaluate, I use the Thomson Reuters Text Research Collection 2 (TRC2), part of the Thomson Reuters Research Collection described in Lewis, Yang, Rose, and Li (2004). This training corpus includes approximately 1.8M news articles spanning the full spectrum of financial news reported by Reuters during the period of 2008-2009, and is available from the National Institute of Standards and Technology. Appendix A.2 describes the pre-processing of the news articles in order to represent them in terms of meaningful textual features ready to be inputted into the Latent Dirichlet Allocation algorithm.

The output of the Latent Dirichlet Allocation model provides an intuitive conceptualization of the identified topics in terms of the most frequently occurring words conditional on each topic, which can be used to interpret the topics. I estimate the topic model for $k = \{10, 15, 20, 25, 30, 35, 40\}$ topics, and observe that the specification with 15 topics performs best in terms of model log likelihood (see Appendix A.3 for details on the topic model estimation). The topics identified in this specification are presented in Table 9. For each topic, the table displays the fifteen terms in the vocabulary that are most likely to appear conditional on the topic. For each of the topics, the set of common terms forms a single coherent theme; for clarity of reference, each topic is labeled with a concise name capturing its theme. For example, the topic whose most common terms are "court," "case," "judge," "federal," etc. is labeled "Litigation;" while the topic whose most common terms are "deal," "offer," "price," "bid," etc. is labeled "Mergers & Acquisitions."

[TABLE 9 AROUND HERE]

The topics in Table 9 are listed in order of their estimated frequencies, and the last column of the table presents the estimated frequencies for all topics. The most common topic to appear in the training corpus of Reuters financial news relates to technology, followed by financial reports such as earnings, and then news regarding financial institutions such as hedge funds and banks. Other common topics include automobile and air transport industries, litigation, and management. Overall, the identified topics are generally applicable and representative of concepts discussed in financial news.

For the second part of the process, I take advantage of the Latent Dirichlet Allocation's ability to represent out-of-sample documents as mixtures over the identified topics. I apply this to characterize the distribution of topics in news articles from three categories: (1) PI news articles; (2) SI news articles that appear on the front page; and (3) non-front page SI news articles.

The results suggest that there are some distinct topic patterns for the select set of news articles marked as more important by the editorial staff (PI news), but no differences in the content of SI news articles that make it to the front page and those that do not. The distribution of topics for each category of news is displayed in Figure 6. All three categories of news overweight content regarding the financial services industry, regulations, the retail industry, and company employees. Coverage of technology, earnings reports, and the healthcare industry is also common, although technology is far less ubiquitous than in the training corpus. PI news articles are substantially more likely to cover news related to the healthcare industry and company operations; they have a lower focus on the financial services industry and regulations. Front page and non-front page SI news articles are very similar in terms of the distribution of topics, with only minor differences (non-front page SI news articles are marginally more likely to feature news about M&A deals and company employees, whereas front page articles contain marginally more discussions of litigation and the financial services industry).

[FIGURE 6 AROUND HERE]

For a formal comparison of the distributions of topics across the different categories of news, I perform a Pearson χ-square test of independence pairwise between any two categories (see Rao and Scott (1981)). The results are tabulated in Table 10. In the main specification with 15 topics, the distribution of topics in PI news articles is weakly statistically significantly different, at the 10% level, from the distribution of topics covered by the front page SI news (see Panel 2 of Table 10). This is robust to varying the number of topics, with the difference becoming significant at the 5% level in the specification with 25 topics but falling short of the 10% statistical significance threshold when the number of topics is reduced to 10. More importantly, the front page and non-front page SI articles are statistically indistinguishable in terms of their textual content, with a p-value of 87.76% in the primary specification with 15 topics. The similarity in the two distributions is robust to varying the topic model specification, with all p-values above 75% (the lowest is 78.01% for 25 topics).

[TABLE 10 AROUND HERE]

Overall, the results point to some distinction in the content of "primary important" news articles from the content of "secondary important" articles. But the distributions of topics are remarkably similar for front page and non-front page SI news articles. This supports the identifying assumption of independence of the prominence of the SI news articles' positions from their underlying content.

6.3 Market Participants' Perceptions of News

In order to directly assess the market's perceptions of the underlying news articles in my hand-collected sample, I survey the target audience of the news: active finance professionals and current MBA students at top business schools. Without the differential positioning, these individuals perceive no differences between the headlines of front page and non-front page SI news articles. They do, however, perceive the PI news articles to be more impactful, supporting Bloomberg editorial staff's decisions to mark these articles as more important.

For this part of the analysis, I survey 109 financial experts from a number of financial institutions and top MBA programs. The breakdown of these individuals across affiliations is presented in Table 11. The majority of the sample (95%) covers active professionals from a representative landscape of financial institutions. The remainder consists of current MBA students at the Harvard Business School, Columbia Graduate School of Business, the University of Chicago Booth School of Business, UVA Darden School of Business, and the McDonough School of Business at Georgetown University.

[TABLE 11 AROUND HERE]

The sample of active finance professionals is representative of the full landscape of the financial services industry. The bulk (79%) of the active professionals come from large banks and broker dealers such as JP Morgan and Morgan Stanley, investment management firms such as BlackRock and State Street, hedge funds such as Bridgewater Associates and Tudor Investment Corporation, and private equity firms such as the Blackstone Group and Warburg Pincus. The remainder of the sample spans consulting firms such as the Boston Consulting Group, government agencies such as the Federal Reserve Board, financial offices of corporations such as Nike and Walt Disney, pension funds such as North Carolina Retirement System, insurance companies such as Massachusetts Mutual Life Insurance, and other areas of the financial services industry.

The respondents largely constitute key decision makers within their respective firms. Many of the professionals from larger corporations such as banks, broker dealers, and large investment management firms are at the principal or managing director levels within their organizations, including heads of regional offices. The sample also includes partners and C-level executives of the smaller companies.

My sample is broadly reflective of the client base receiving Bloomberg news through the terminal. Approximately 89% of the professionals report having used a Bloomberg terminal at some point, with 75% using the terminal occasionally on an ongoing basis, and 35% reporting frequent use. Almost a quarter of all respondents, 22%, state that they use the Bloomberg terminal all the time, on a daily basis.

In the survey, each respondent is asked to answer a series of twenty-five questions about news headlines. The respondent is told that the headlines come from a news provider who chooses how prominently the headlines are displayed based in part on the importance and market impact of the underlying news. Each question presents two headlines, and asks the respondent to specify which headline the respondent thinks had larger market impact and deserves more prominence. An screenshot with an example question is displayed in Figure 7.

[FIGURE 7 AROUND HERE]

The survey questions span two sets of comparisons: (1) between front page SI news articles and PI news articles; and (2) between front page SI news articles and non-front page SI news articles. In particular, in each question, one of the two headlines (in random position – either on the left or on the right of the screen) is from the front page SI news category. The other headline is randomly selected, with equal likelihoods, from the categories of PI news (approximately 36% of the questions) and non-front page SI news (approximately 64% of the questions).

The respondents are incentivized to identify the news importance as accurately as they can. Each respondent receives a $10 gift (an Amazon.com gift card or a lunch voucher to a venue of the respondent's choice) for completing the survey. In addition, the five respondents whose answers most closely match actual differences in positioning by the news provider receive additional prizes of $90 each.

The results indicate that the financial experts in the sample do not distinguish between the front page and non-front page SI news articles. Panel 1 of Table 12 presents the incidence of front page SI news articles being chosen as more important than non-front page SI news articles. The full sample of finance professionals identifies the front page articles as more impactful 48.42% of the time, not significantly different from 50%. This finding is robust to splitting the sample into active finance professionals (who choose the front page articles 48.46% of the time) and MBA students (who choose the front page articles 47.83% of the time). Thus, absent the differential positioning, the target audience of financial professionals does not perceive the front page SI news articles as being any more important than their non-front page counterparts.

[TABLE 12 AROUND HERE]

The market participants do, however, identify the "primary important" articles as more impactful. The respondents choose a PI news article over a front page SI news article 63% of the time, significantly different from 50% at the 1% level. Slicing the sample into active finance professionals and current MBA students reveals that the finance professionals choose the PI articles 63.5% of the time (significantly different from 50% at the 1% level), while the MBA students choose the PI articles 55% of the time (not statistically significantly different from 50%, likely due to the sample currently containing only 5 MBA students). Overall, the results point to the Bloomberg editorial staff correctly identifying, on average, the news most relevant for the target demographic: the higher importance ranking assigned to the PI news articles is corroborated by the surveyed market participants.

Similar patterns hold at the individual level. For each respondent, I calculate the percentage of times that the respondent chooses a front page SI news article over a non-front page one, and the percentage of times that the respondent chooses a PI news article over an SI one. A histogram of these individual-level percentages is displayed in Figure 8. The incidence of choosing front page articles over non-front page ones is presented in blue; the distribution is centered around 50%, is symmetric, and resembles a normal distribution. Overall, this distribution is consistent with there being no distinction between the two sets of articles, and the differences between individuals' choices coming from noise and the variation in the questions that the different individuals see. The incidence of choosing PI news articles over SI ones, presented in gray, paints a different picture. Practically no respondents choose PI news articles less than 40% of the time, and the distribution is centered around 60%, with a number of respondents choosing the PI news articles as often as 90-100% of the time.

[FIGURE 8 AROUND HERE]

Overall, the target audience of the news – current active finance professionals and future professionals currently enrolled in top MBA programs – perceive no systematic differences between the SI news articles that get placed on the front page and those that do not. This is consistent with the quasi-random positioning of these news articles. There is a stark juxtaposition between the significantly different market dynamics following these two sets of news and the market participants' lack of distinction between them in the survey. This juxtaposition highlights the extent to which salient news positioning can induce different reactions to otherwise identically important content.

7 Conclusion

This paper takes advantage of a natural experiment in news positioning to directly estimate the effect of news consumption on financial markets. For two news articles of equal importance, pinning one to a prominent position induces a 280% higher trading volume during the ten-minute window after the news, 180% larger absolute price change, and substantially higher short-term return continuation. Interestingly, differences in news positioning play an even larger role for market dynamics than differences in the perceived importance of the underlying news articles' content.

The results of this paper highlight the importance of the way in which information is presented and consumed, in the increasingly prolific modern news environment. Despite the push for transparency that brings progressively more information into the public domain, informational advantages persist – only here, they take the form of speedy and comprehensive attention to public news rather than possession of private information. As a result, even when we restrict our attention to public information, how a piece of news is presented to the public, who consumes it, and when, play an important role in determining whether the information is immediately and efficiently incorporated into asset prices.

References

[1] Ahern, K. R., and D. Sosyura (2014). "Who writes the news? Corporate press releases during merger negotiations." *The Journal of Finance,* 69(1): 241-291.

[2] Andrei, D., and M. Hasler (2014). "Investor attention and stock market volatility." *Review of Financial Studies*, 28(1): 33-72.

[3] Atkins, A. B., and E. A. Dyl (1990). "Price reversals, bid-ask spreads, and market efficiency." *Journal of Financial and Quantitative Analysis*: 535-547.

[4] Banerjee, S., and I. Kremer (2010). “Disagreement and learning: Dynamic patterns of trade.” *The Journal of Finance*, 65(4): 1269-1302.

[5] Ben-Rephael, A., Z. Da, and R. D. Israelsen (2017). “It Depends on Where You Search: A Comparison of Institutional and Retail Attention.” *Review of Financial Studies,* forthcoming.

[6] Bernard, V. L., and J. K. Thomas (1989). “Post-earnings-announcement drift: delayed price response or risk premium?” *Journal of Accounting Research,* 27: 1-36.

[7] Blei, D. M., A. Y. Ng, and M. I. Jordan (2003). “Latent dirichlet allocation.” *Journal of Machine Learning Research*, 3 993-1022.

[8] Boulland, R., F. Degeorge, and E. Ginglinger (2016). “News dissemination and investor attention.” *Review of Finance:* 1-31.

[9] Busse, J. A., and T. C. Green (2012). “Market efficiency in real time.” *Journal of Financial Economics*, 65(3): 415-437.

[10] Cao, H. H. and H. Ou-Yang (2009). “Differences of opinion of public information and speculative trading in stocks and options.” *Review of Financial Studies*, 22(1): 299-335.

[11] Carvalho, C., N. Klagge, and E. Moench (2011). “The persistent effects of a false news shock.” *Journal of Empirical Finance*, 18(4): 597-615.

[12] Chan, W. (2003). “Stock price reaction to news and no-news: drift and reversal after head- lines.” *Journal of Financial Economics*, 70: 223-260.

[13] Chordia, T., R. Roll, and A. Subrahmanyam (2002). “Order imbalance, liquidity, and market returns.” Journal of Financial economics 65(1): 111-130.

[14] Da, Z., J. Engelberg, and P. Gao (2011). “In search of attention.” *The Journal of Finance*, 66(5): 1461-1499.

[15] Da, Z., J. Engelberg, and P. Gao (2015). “The sum of all fears investor sentiment and asset prices.” *Review of Financial Studies,* 28(1): 1-32.

[16] DellaVigna, S. and J. M. Pollet (2009). “Investor inattention and Friday earnings announcements.” *The Journal of Finance*, 64(2): 709-749.

[17] Dougal, C., J. Engelberg, D. Garcia, and C. A. Parsons (2012). “Journalists and the stock market.” *Review of Financial Studies,* 25(3): 639-679.

[18] Drake, M. D.T. Roulstone, and J. Thornock (2012). "Investor information demand: Evidence from google searches around earnings announcements investor information demand." *Journal of Accounting Research,* 50(4), 1001?1040.

[19] Ederington, L. H., and J. H. Lee (1995). "The short-run dynamics of the price adjustment to new information." *Journal of Financial and Quantitative Analysis,* 30(1): 117-134.

[20] Engelberg, J.E. (2008). "Costly information processing: Evidence from earnings announcements." Working Paper, University of California San Diego.

[21] Engelberg, J. E., and C.A. Parsons (2011). "The causal impact of media in financial markets." *The Journal of Finance*, 66(1): 67-97.

[22] Engelberg, J., C. Sasseville, and J. Williams (2012). "Market madness? The case of mad money." *Management* Science 58(2): 351-364.

[23] Fang, L., and J. Peress (2009). "Media coverage and the cross?section of stock returns." *The Journal of Finance*, 64(5): 2023-2052.

[24] Fedyk, A. and J. Hodson (2015). "When Can the Market Identify Stale News?" Working Paper, Harvard Business School.

[25] Fehle, F., S. Tsyplakov, and V. Zdorovtsov (2005). "Can companies influence investor behaviour through advertising? Super bowl commercials and stock returns." *European Financial Management*, 11(5): 625-647.

[26] Foster, G., C. Olsen, and T. Shevlin (1984). "Earnings releases, anomalies, and the behavior of security returns." *Accounting Review,* 59(4): 574-603.

[27] Fung, A.K.W., D.M.Y. Mok, and K. Lam (2000). "Intraday price reversals for index futures in the US and Hong Kong." *Journal of Banking & Finance,* 24(7): 1179-1201.

[28] Garcia, D. (2013). "Sentiment during recessions." *The Journal of Finance,* 68(3): 1267-1300.

[29] Gilbert, T., S. Kogan, L. Lochstoer, and A. Ozyildirim (2012). "Investor inattention and the market impact of summary statistics." *Management Science,* 58(2): 336-350.

[30] Griffiths, T. L. and M. Steyvers (2004). "Finding scientific topics." *Proceedings of the National Academy of Sciences*, 101(suppl 1): 5228–5235.

[31] Harris, M. and A. Raviv (1993). "Differences of opinion make a horse race." *Review of Financial Studies*, 6(3): 473-506.

[32] Heston, S. L., R. A. Korajczyk, and R. Sadka (2010). "Intraday Patterns in the Cross?section of Stock Returns." *The Journal of Finance,* 65(4): 1369-1407.

[33] Hillert, A., H. Jacobs, and S. Müller (2014). "Media makes momentum." *Review of Financial Studies,* 27(12): 3467-3501.

[34] Hirshleifer, D., S. S. Lim, and S. H. Teoh (2009). "Driven to distraction: Extraneous events and underreaction to earnings news." *The Journal of Finance*, 64(5): 2289-2325.

[35] Hirshleifer, D., and S. H. Teoh (2003). "Limited attention, information disclosure, and financial reporting." *Journal of Accounting and Economics,* 36.1: 337-386.

[36] Hofmann, T. (1999). "Probabilistic latent semantic indexing." *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*: 50–57.

[37] Hong, H., and J. C. Stein (1999). "A unified theory of underreaction, momentum trading, and overreaction in asset markets." *The Journal of Finance*, 54(6): 2143-2184.

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

[39] Kandel, E., and N. D. Pearson (1995). "Differential interpretation of public signals and trade in speculative markets." *Journal of Political Economy*, 103(4): 831-872.

[40] Klibanoff, P., O. Lamont, and T. A. Wizman (1998). "Investor reaction to salient news in closed-end country funds." *The Journal of Finance,* 53.2: 673-699.

[41] Koudijs, P. (2016). "The boats that did not sail: Asset price volatility in a natural experiment." *The Journal of Finance*, 71(3):

[42] Lehavy, R., F. Li, and K. Merkley (2011). "The effect of annual report readability on analyst following and the properties of their earnings forecasts." *The Accounting Review,* 86(3): 1087-1115.

[43] Lewis, D.D., Y. Yang, T. G. Rose, and F. Li (2004). "RCV1: A New Benchmark Collection for Text Categorization Research." *Journal of Machine Learning Research,* 5: 361-397.

[44] Li, Feng (2008). "Annual report readability, current earnings, and earnings persistence." *Journal of Accounting and economics,* 45(2): 221-247.

[45] Liu, L., A. Sherman, and Y. Zhang (2014). "The Long-Run Role of the Media: Evidence from Initial Public Offerings." *Management Science*, 60(8): 1945-1964.

[46] Loh, R. K (2010). "Investor inattention and the underreaction to stock recommendations." *Financial Management,* 39.3: 1223-1252.

[47] Loughran, T., and B. McDonald (2011). "When is a liability not a liability? Textual analysis, dictionaries, and 10'Ks." *The Journal of Finance,* 66(1): 35-65.

[48] Loughran, T., and B. McDonald (2014). "Measuring readability in financial disclosures." *The Journal of Finance,* 69(4): 1643-1671.

[49] Marshall, B. R., N. Visaltanachoti, and G. Cooper (2014). "Sell the rumour, buy the fact?" *Accounting & Finance,* 54(1): 237-249.

[50] Miller, B. P (2010). "The effects of reporting complexity on small and large investor trading." *The Accounting Review,* 85(6): 2107-2143.

[51] Peng, L., and W. Xiong (2006). "Investor attention, overconfidence and category learning." *Journal of Financial Economics* 80(3): 563-602.

[52] Peress, Joel (2008). "Media coverage and investors' attention to earnings announcements." Working Paper, INSEAD.

[53] Peress, J. (2014). "The media and the diffusion of information in financial markets: Evidence from newspaper strikes." *The Journal of Finance*, 69(5): 2007-2043.

[54] Pritchard, J. K., M. Stephens, and P. Donnelly (2000). "Inference of population structure using multilocus genotype data." *Genetics*, 155(2): 945–959.

[55] Rao, J. and A. J. Scott (1981). "The analysis of categorical data from complex sample surveys: chi-squared tests for goodness of fit and independence in two-way tables." *Journal of the American Statistical Association*, 76(374): 221-230.

[56] Rogers, J.L., D.J. Skinner, and S.L.C. Zechman (2013). "The role of the media in disseminating insider-trading news." *Review of Accounting Studies*, 21(3): 711-739.

[57] Schmidt, D. (2013). "Investors' Attention and Stock Covariation." Working paper, HEC Paris.

[58] Solomon, D. H. (2012). "Selective publicity and stock prices." *The Journal of Finance*, 67(2): 599-638.

[59] Tetlock, P. (2011). "All the news that's fit to reprint: Do investors react to stale information?" *Review of Financial Studies*, 24(5): 1481-1512.

[60] Tetlock, P.C., M. Saar-Tsechansky, and S. Macskassy (2008). "More than words: Quantifying language to measure firms' fundamentals." *The Journal of Finance*, 63(3): 1437-1467.

[61] Uhl, M. W. (2014). "Reuters sentiment and stock returns." *Journal of Behavioral Finance,* 15.4: 287-298.

[62] Umar, T. (2017). "Attention Grabbers when Seeking Alpha." Working Paper, University of Chicago Booth.

[63] Yuan, Y. (2015). "Market-wide attention, trading, and stock returns." *Journal of Financial Economics,* 116(3), 548?564.

[64] You, H., and X.Zhang (2009). "Financial reporting complexity and investor underreaction to 10-K information." *Review of Accounting Studies,* 14(4): 559-586.

[65] Zawadowski, A. G., G. Andor, and J. Kertész (2006). "Short-term market reaction after extreme price changes of liquid stocks." *Quantitative Finance,* 6(4): 283-295.

Figure 1: Bloomberg terminal screen displaying company news as of 3:01PM EST on December 15, 2016. The first three lines are the articles pinned to the "front page." All other articles scroll off the screen with the arrival of more recent news.

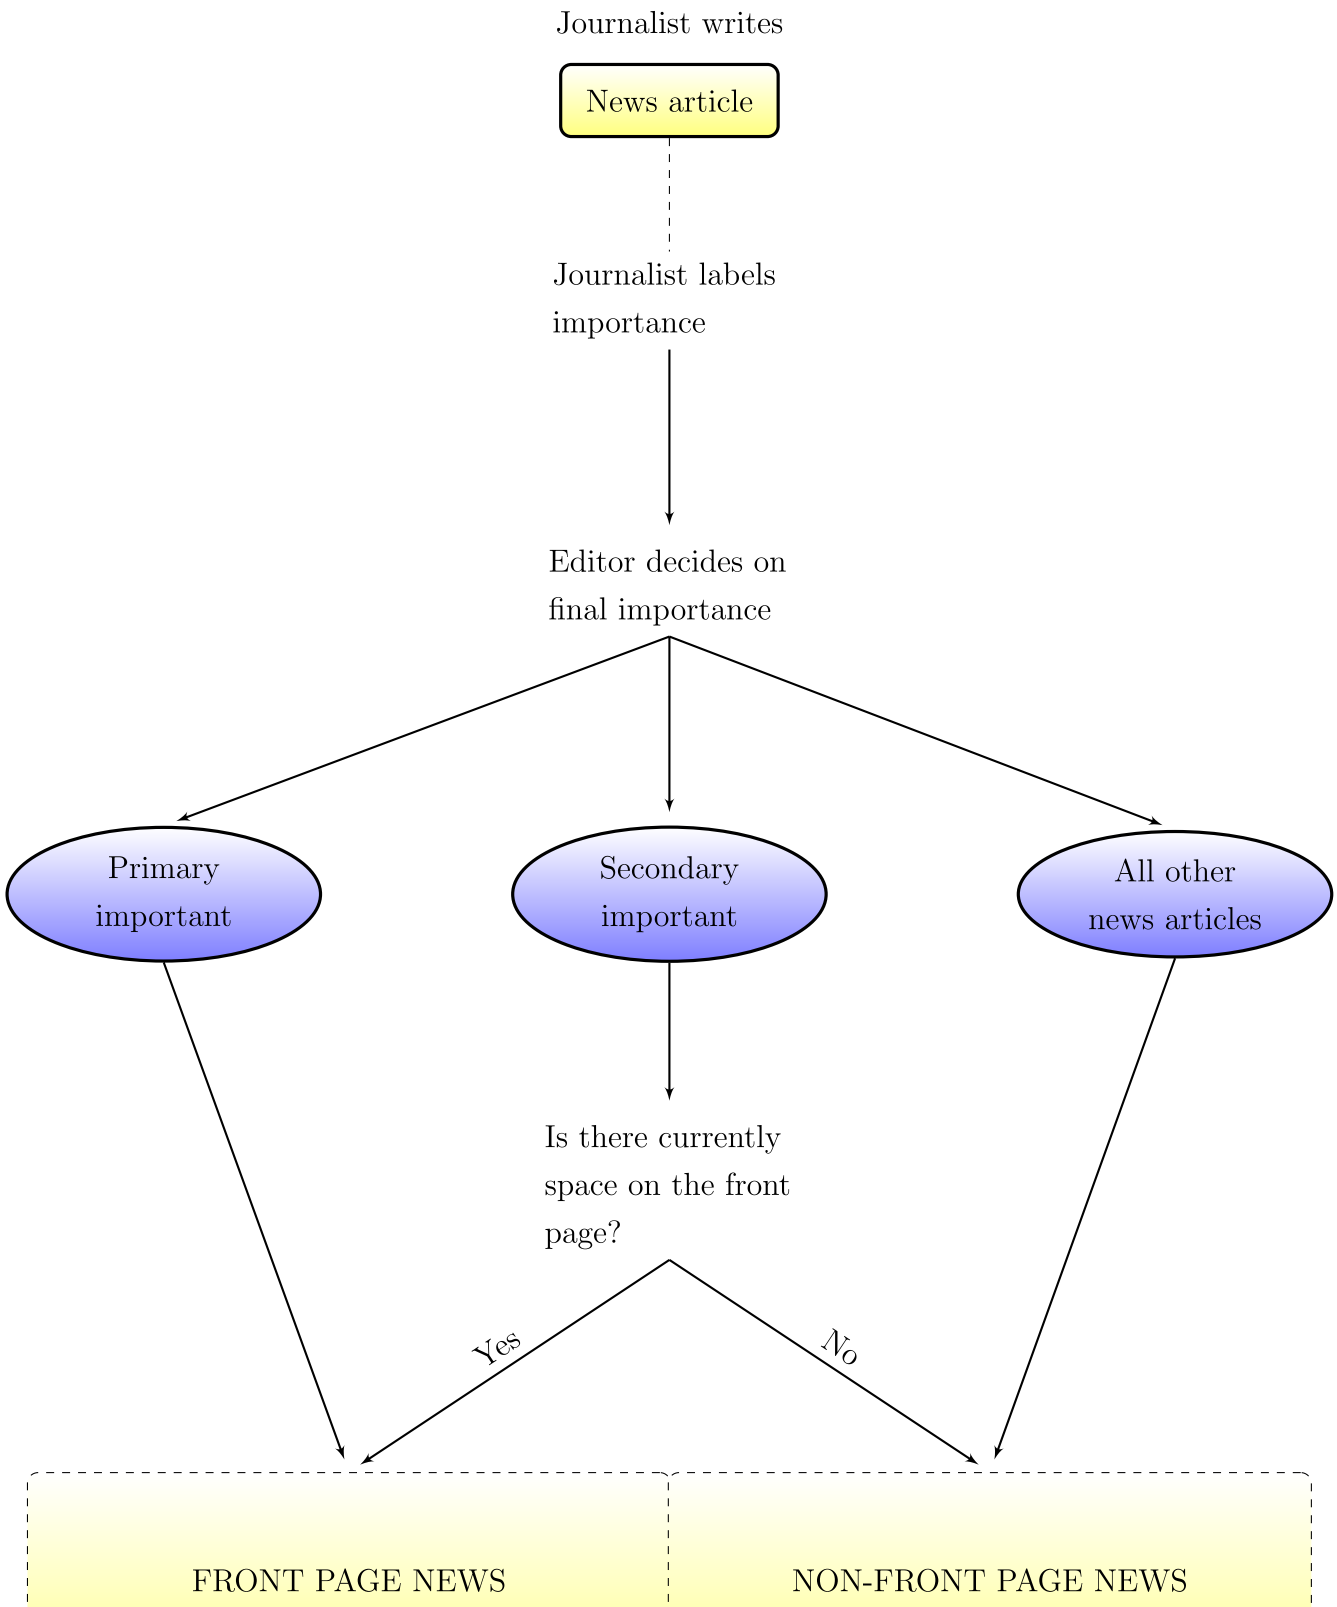

Figure 2: Process illustrating how Bloomberg news articles are pinned to the prominent front page positions at the top of the news screen.

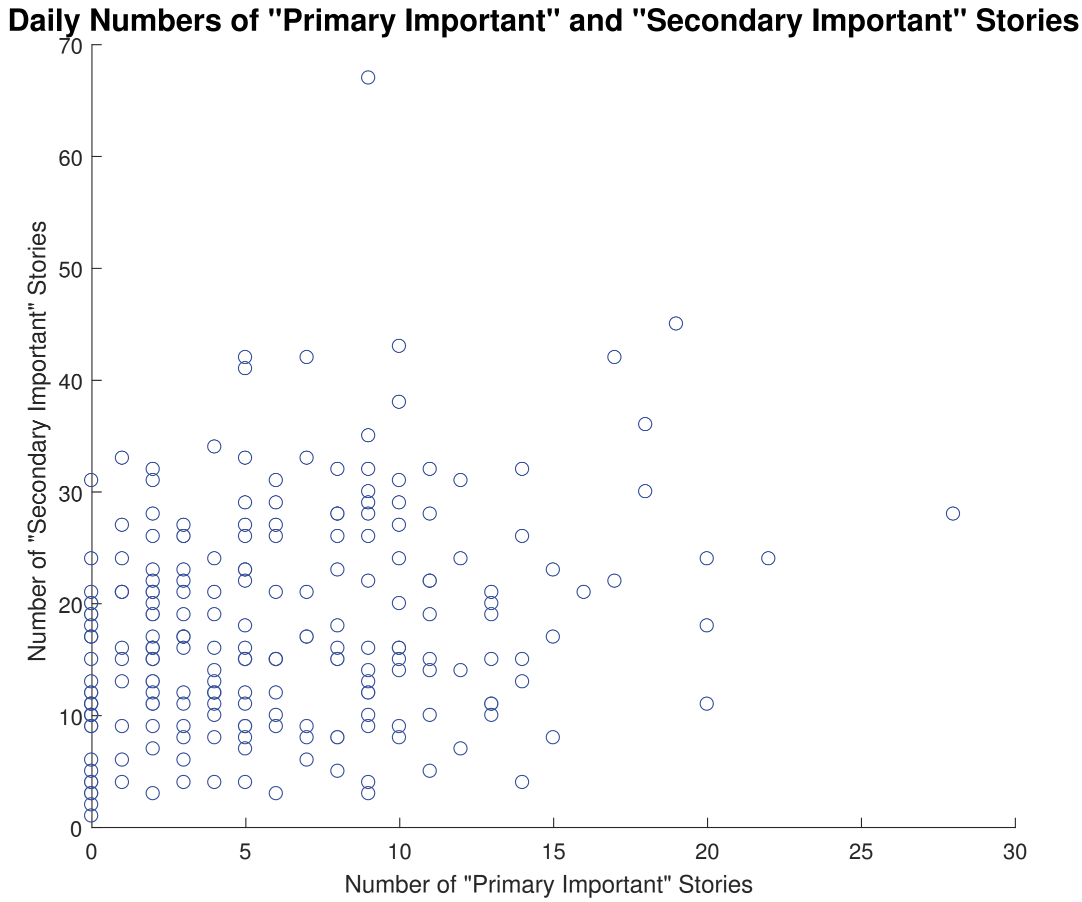

Figure 3: Distribution of the daily volumes of “primary important” (PI) and “secondary important” (SI) news articles across the days in the sample.

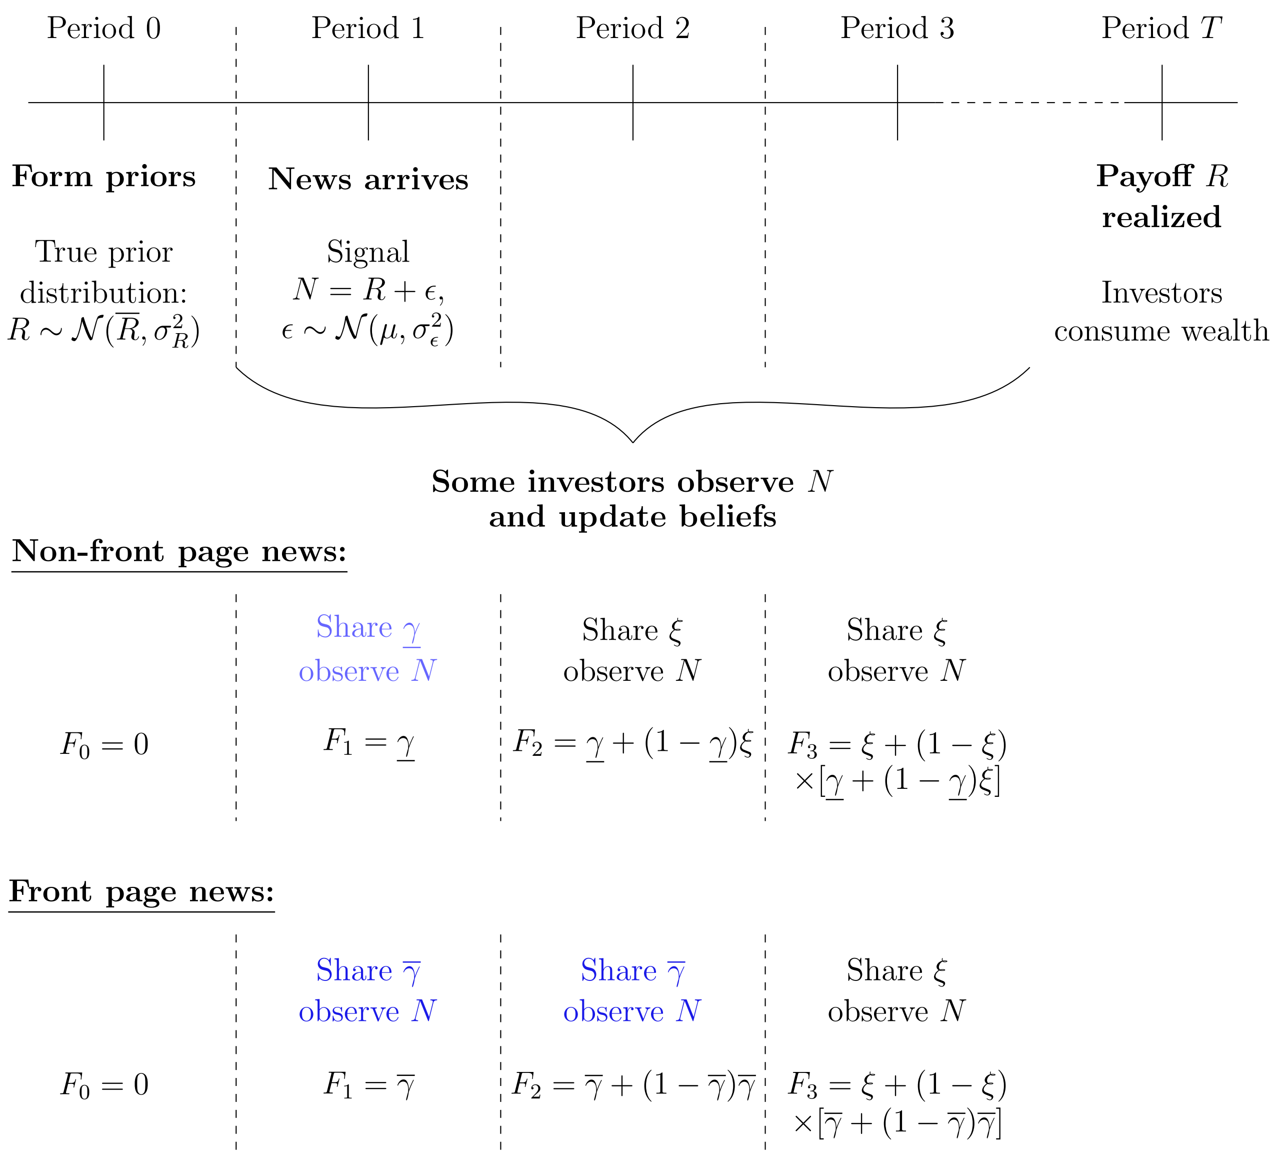

Figure 4: Model timeline.

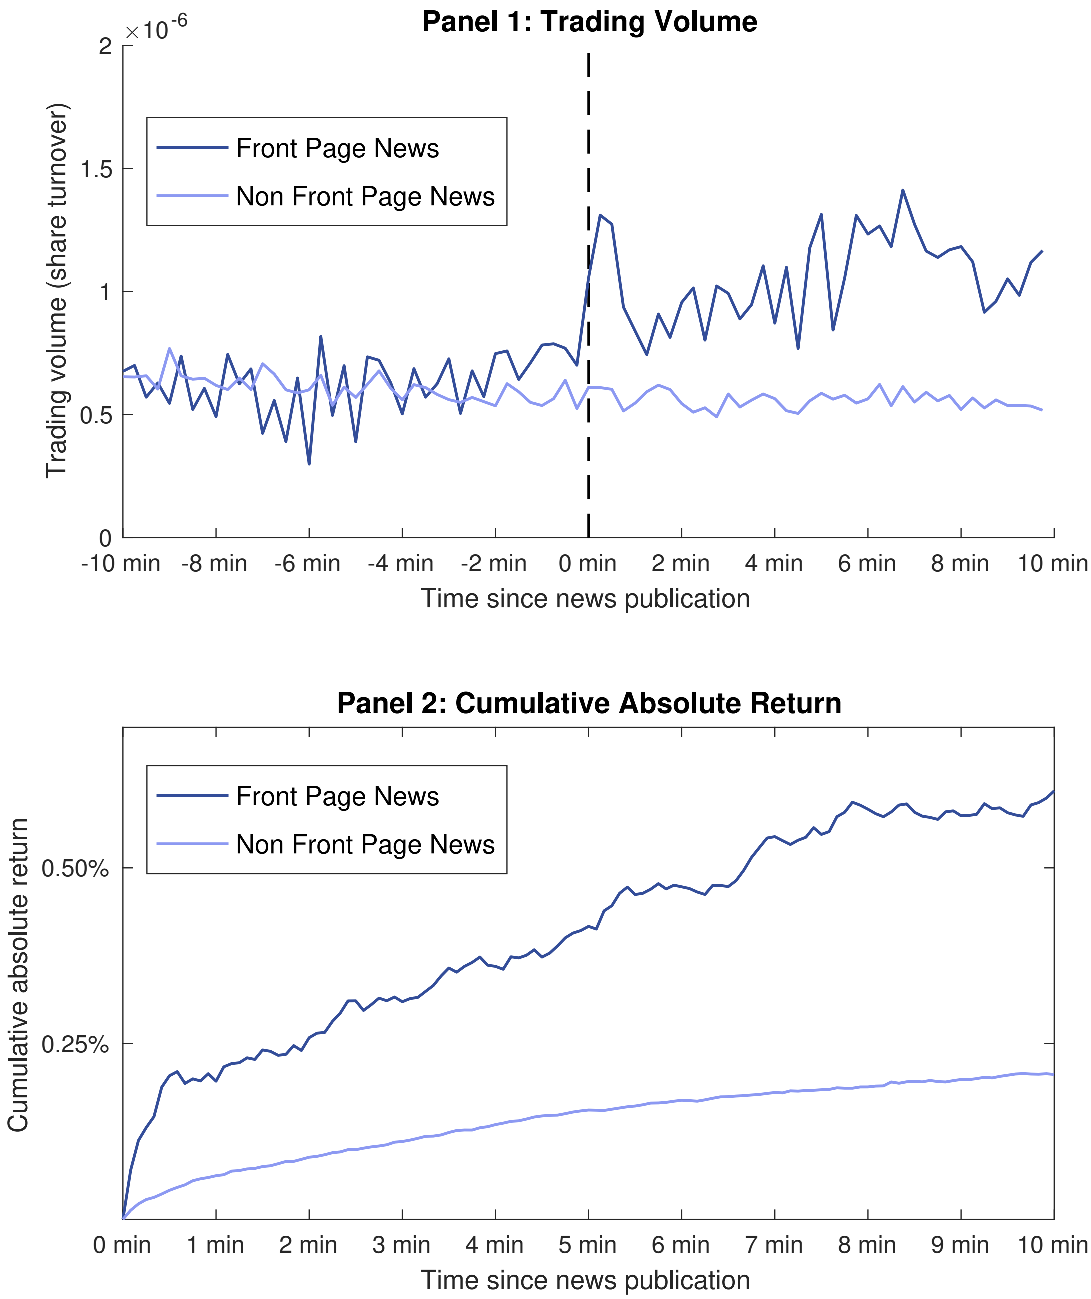

Figure 5: Market dynamics following front page and non-front page SI news articles. **Panel 1** displays the median trading volume by 15-second intervals during the ten minutes before and the ten minutes after news publication. The dark blue line plots trading volume around SI news articles that are prominently placed on the front page; trading volume around non-front page SI news articles is displayed in light blue. **Panel 2** presents the average absolute price changes from publication to up to ten minutes later, for front page SI news articles (in dark blue) and non-front page SI news articles (in light blue).

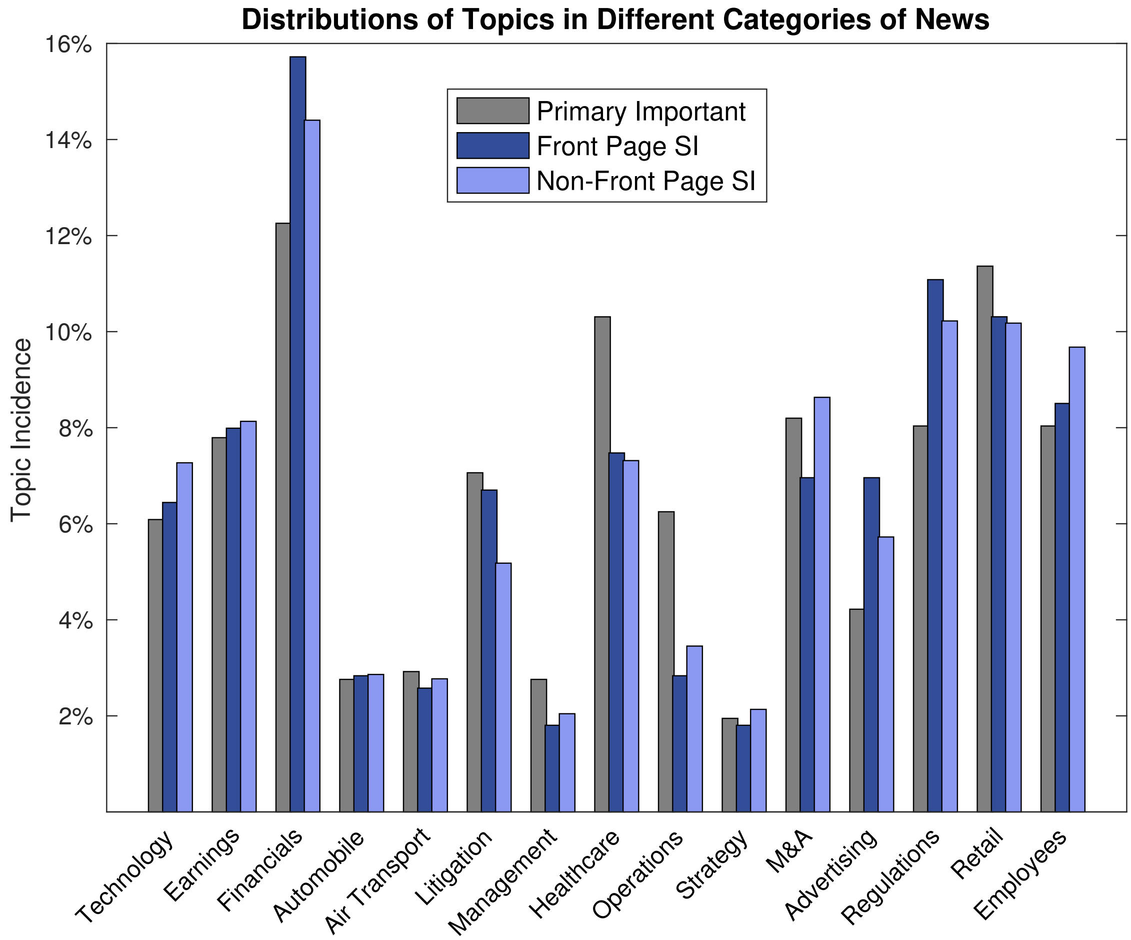

Figure 6: Distributions of topics across news articles from different categories. The figure presents the frequency of the identified topics in three types of news articles: PI news, SI news placed on the front page, and non-front page SI news.

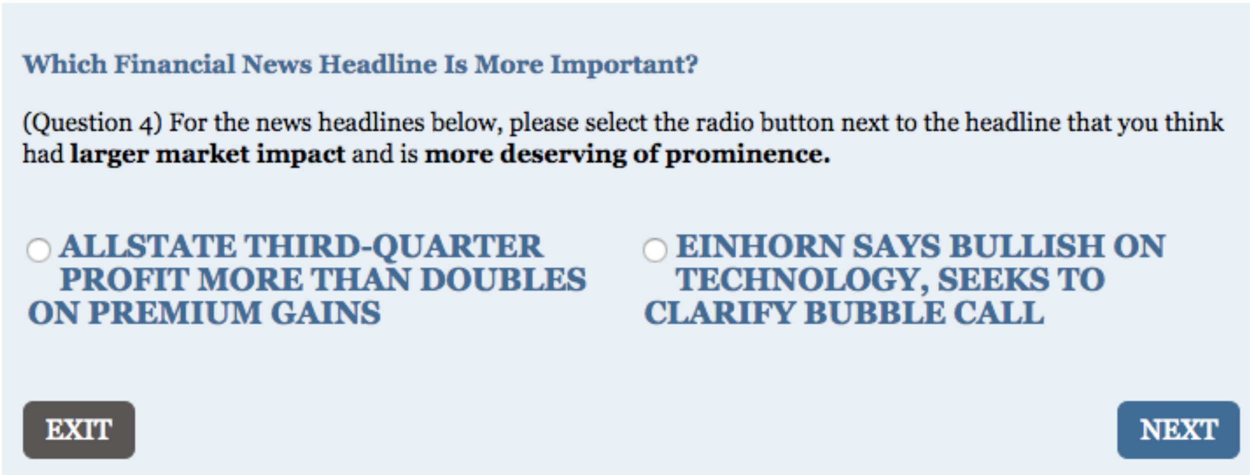

Figure 7: Example question from the survey administered to MBA students and active finance professionals. One of the presented headlines comes from the front page SI sample, while the other is either a non-front page SI headline or a PI headline.

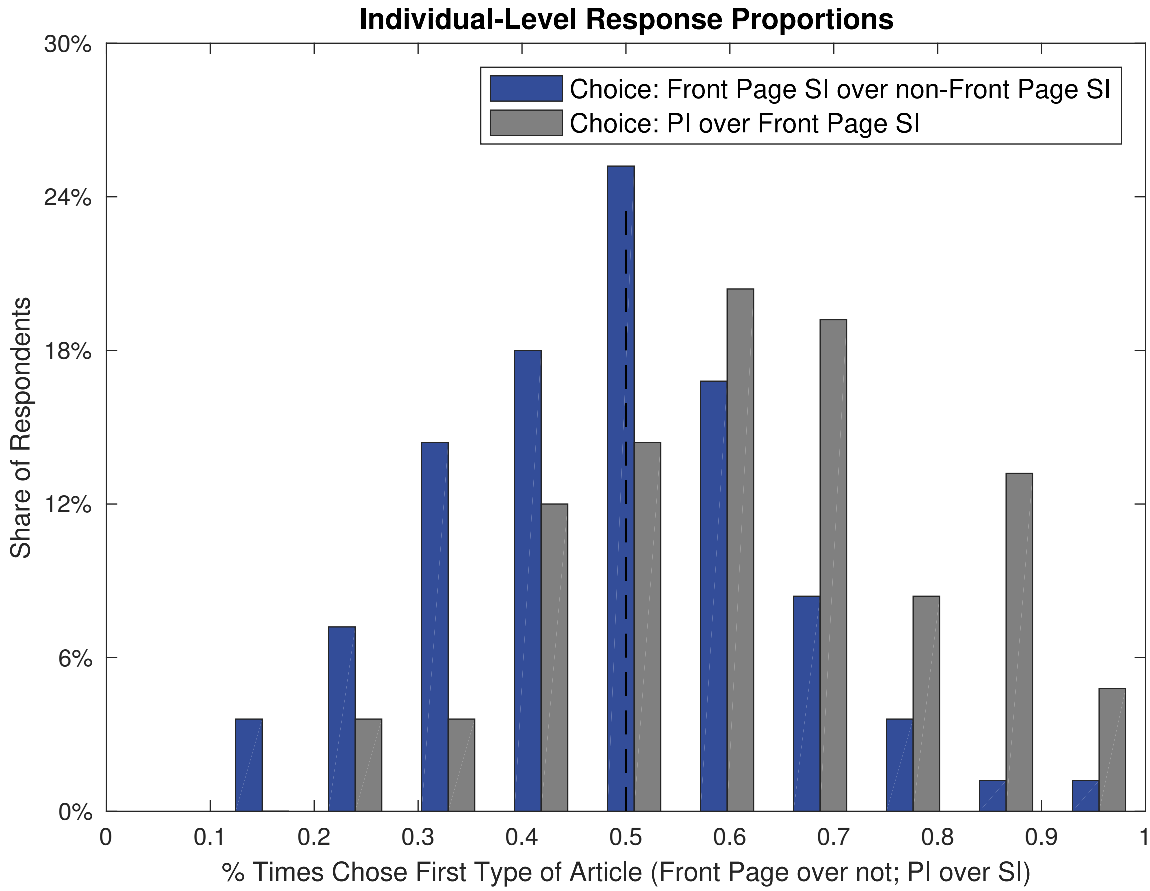

Figure 8: Individual-level responses from the financial experts. For each respondent, I calculate the percentage of times that respondent chose a front page SI headline over a non-front page one (plotted in blue), and the percentage of times the respondent chose a PI headline over a front page SI one. The figure displays the distribution of these percentages across respondents. [Note: The survey is currently running, and the results are updated on an ongoing basis.]

Table 1: Examples of articles in the “primary important” and “secondary important” categories. **Panel 1** presents randomly selected headlines in the “primary important” sample, all pinned to the front page. **Panel 2** lists randomly selected headlines in the “secondary important” sample, indicating whether each of the example articles was pinned to the front page.

Panel 1: “Primary important” news articles

Date	Headline
3/25/2014	Walgreen 2Q Adj. EPS Misses Est.
4/7/2014	Tekmira Says FDA Modifies TKM-Ebola Drug to Partial Hold
4/25/2014	United Technologies Reports SEC Formal Investigation, Subpoena
8/14/2014	Icahn Reports 6.63% Stake in Gannett, Urges Splitting Co.
10/28/2014	Amgen Restarts Buyback, Boosting Dividend; 2015 View Tops Ests.
12/23/2014	Stryker Said to Plan Smith & Nephew Takeover Bid Within Weeks
1/27/2015	Amgen 4Q Adj. EPS, Rev. Top Ests.; Ivabradine, T-vec Delayed
5/13/2015	Nissan Forecasts 6% Gain in Profit on U.S. Demand, Weak Yen
5/19/2015	Computer Sciences Corp. to Split Into Two Companies
7/30/2015	Sanofi Profit Beats Estimates as Multiple Sclerosis Drugs Gain
9/10/2015	Morrison Earnings Miss Analysts’ Estimate as Grocer Cut Prices
9/14/2015	Standard Chartered Said to Plan Cutting 250 Managing Directors
11/24/2015	Fed Says It’s Overhauling Standards for Large-Bank Examiners
1/15/2016	Wal-Mart to Close 269 Stores in U.S., Globally
2/25/2016	Apple Says U.S. Can’t Force It to Unlock Terrorist’s IPhone

Panel 2: “Secondary important” news articles

Date	Headline	FP
4/7/2014	Honeywell CEO Makes Biggest Executive Shift Naming Vice Chairmen	N
5/19/2014	AstraZeneca Chairman “Surprised” Pfizer Took Last Offer Public	N
6/3/2014	Robertson’s Stock Picker Singh Said to Become Newest Tiger Cub	Y
6/24/2014	Morgan Stanley Gets 90,000 Applications for Summer Program	N
7/10/2014	TRW Said to Receive Takeover Approach From ZF Friedrichshafen	Y
8/1/2014	Judge Grants Preliminary Approval to Apple e-Book Settlement	N
9/26/2014	Pimco Said to Have Discussed Firing Gross Before Exit to Janus	Y
12/5/2014	CNN’s Candy Crowley to Leave Cable News Network After 27 Years	N
1/20/2015	FXCM Plunges as Bailout Lets Leucadia Force Sale of Brokerage	Y
3/12/2015	Viacom Says Chairman Redstone Will Miss Monday’s Annual Meeting	N
4/28/2015	McDonald’s Axes Seven Sandwiches in Push to Get Its Menu Right	Y
6/3/2015	Pandora Internet Radio Wins U.S. Nod to Buy South Dakota Station	N
6/11/2015	Biotech Led by 29-Year-Old CEO Now Worth Billions With No Sales	Y
7/29/2015	High-Density Drone Flights Possible Within Decade, Google Says	N
9/21/2015	Clinton’s Tweet on High Drug Prices Sends Biotech Stocks Down	Y
10/22/2015	Amazon Sales Top Estimates on Prime Day Event, Cloud Computing	Y
12/21/2015	Chipotle Probed for New Outbreak of Different E. Coli Strain	Y
1/14/2015	Apple, Ericsson Sue Each Other Over Phone Patent Royalties	N
2/27/2016	Lenovo to Purge Adware From New PCs After Superfish Controversy	N

Table 2: Summary statistics of the hand-collected news sample. **Panel 1** presents the breakdown by month of publication of PI news articles, SI news articles, and SI news articles that are positioned on the front page. **Panel 2** presents the breakdown by hour of publication, and includes the percentage of SI articles that are positioned on the front page. The sample is restricted to the articles published between 8AM and 5PM EST.

Panel 1: News Articles By Month

Hour of Day	PI articles	SI articles	FP SI articles
January	106	346	78
February	100	175	42
March	51	105	28
April	138	418	41
May	104	422	47
June	141	343	42
July	126	470	68
August	126	407	58
September	172	373	79
October	140	407	54
November	82	301	37
December	78	288	53
Total	1,364	4,055	627

Panel 2: News Articles By Hour

Hour of Day	PI articles	SI articles	FP SI articles	% SI articles on FP
8AM - 9AM	199	426	77	18%
9AM - 10AM	162	547	59	11%
10AM - 11AM	119	614	96	16%
11AM - 12PM	104	522	98	19%
12PM - 1PM	81	491	73	15%
1PM - 2PM	111	461	58	13%
2PM - 3PM	143	403	80	20%
3PM - 4PM	76	420	71	17%
4PM - 5PM	369	171	15	9%
Total	1,364	4,055	627	15%

Table 3: Comparison of trading volumes and absolute price changes immediately following SI news articles that are pinned to the front page and those that are not. **Panel 1** looks at the trading volumes within five, ten, and sixty minutes of publication, while **Panel 2** considers the absolute price changes over the same time periods.

Panel 1: Trading Volume

	Front Page SI News	Non-Front Page SI News	Difference
First 5 min	0.10%	0.02%	0.07%**
Standard Error	*(0.00017)*	*(0.00001)*	*(0.00017)*
# Observations	408	2,296	–
First 10 min	0.19%	0.05%	0.14%**
Standard Error	*(0.00039)*	*(0.00003)*	*(0.00040)*
# Observations	415	2,345	–
First 1 hour	0.79%	0.26%	0.53%**
Standard Error	*(0.00177)*	*(0.00014)*	*(0.00177)*
# Observations	432	2,464	–

Panel 2: Absolute Price Changes

	Front Page SI News	Non-Front Page SI News	Difference
First 5 min	0.45%	0.16%	0.29%**
Standard Error	*(0.00059)*	*(0.00006)*	*(0.00059)*
# Observations	408	2,296	–
First 10 min	0.60%	0.21%	0.39%**
Standard Error	*(0.00077)*	*(0.00007)*	*(0.00077)*
# Observations	415	2,345	–
First 1 hour	1.09%	0.51%	0.58%**
Standard Error	*(0.00145)*	*(0.00027)*	*(0.00147)*
# Observations	432	2,464	–

** denotes significance at the 1% level.

Table 4: Short-term continuation in returns after front page and non-front page SI news articles. Each column runs the following specification:
$Ret_{s,i,[t+t_1,t+t_2]} = \alpha + \beta_1 Ret_{s,i,[t,t+t_1]} + \beta_2 FP_s + \beta_3 Ret_{s,i,[t,t+t_1]} \times FP_s + \gamma X_{i,t} + \epsilon_{s,i,[t+t_1,t+t_2]}$,
where FP_s is a dummy variable equal to 1 if the news article s is pinned to the front page; $Ret_{s,i,[t,t+t_1]}$ denotes the return on security i during the immediate period $[t, t+t_1]$ after publication of the news article s, and $Ret_{s,i,[t+1,t+t_2]}$ is the return during the delayed period $[t+t_1, t+t_2]$. The main coefficient of interest is β_3 on the interaction term $Ret_{s,i,[t,t+t_1]} \times FP_s$ (highlighted in blue). The tests are run over the following time windows: $(t_1, t_2) \in \{(5min, 10min), (3min, 5min), (5min, 15min), (5min, 20min), (10min, 30min), (10min, 30min)\}$. Columns marked with (1) do not include any controls. For columns marked with (2), the vector of controls $X_{i,t}$ consists of hour and month fixed effects. Columns marked with (3) also control for log firm size and industry fixed effects.

	$t_1 = 5\ min, t_2 = 10\ min$			$t_1 = 3\ min, t_2 = 5\ min$		
	(1)	(2)	(3)	(1)	(2)	(3)
$Ret_{s,i,[t,t+5\ min]}$	-0.149**	-0.150**	-0.146**	0.000	0.002	0.004
Standard error	*(0.030)*	*(0.030)*	*(0.030)*	*(0.030)*	*(0.030)*	*(0.030)*
$Ret_{s,i,[t+5\ min]} \times FP_s$	0.206**	0.206**	0.205**	0.052	0.051	0.041
Standard error	*(0.035)*	*(0.035)*	*(0.035)*	*(0.034)*	*(0.034)*	*(0.034)*
# FP SI articles	415	415	415	408	408	408
# Non-FP SI articles	2,345	2,345	2,345	2,296	2,296	2,296

	$t_1 = 5\ min, t_2 = 15\ min$			$t_1 = 5\ min, t_2 = 20\ min$		
	(1)	(2)	(3)	(1)	(2)	(3)
$Ret_{s,i,[t,t+5\ min]}$	-0.121**	-0.121**	-0.115**	-0.002	-0.003	-0.000
Standard error	*(0.038)*	*(0.038)*	*(0.038)*	*(0.049)*	*(0.048)*	*(0.049)*
$Ret_{s,i,[t+5\ min]} \times FP_s$	0.261**	0.260**	0.247**	0.310**	0.310**	0.310**
Standard error	*(0.044)*	*(0.044)*	*(0.045)*	*(0.056)*	*(0.056)*	*(0.057)*
# FP SI articles	422	422	422	425	425	425
# Non-FP SI articles	2,364	2,364	2,364	2,385	2,385	2,385

	$t_1 = 10\ min, t_2 = 20\ min$			$t_1 = 10\ min, t_2 = 30\ min$		
	(1)	(2)	(3)	(1)	(2)	(3)
$Ret_{s,i,[t,t+5\ min]}$	0.027	0.031	0.033	-0.017	-0.015	-0.017
Standard error	*(0.029)*	*(0.029)*	*(0.030)*	*(0.040)*	*(0.040)*	*(0.040)*
$Ret_{s,i,[t+5\ min]} \times FP_s$	0.192**	0.190**	0.192**	0.420**	0.419**	0.428**
Standard error	*(0.034)*	*(0.034)*	*(0.034)*	*(0.046)*	*(0.046)*	*(0.046)*
# FP SI articles	425	425	425	427	427	427
# Non-FP SI articles	2,385	2,385	2,385	2,414	2,414	2,414

**, *, and † denote significance at the 1%, 5%, and 10% levels, respectively.

Table 5: Continuation in returns over longer horizons following front page and non-front page SI news articles. Each column runs the following specification:
$Ret_{s,i,[t+t_1,t+t_2]} = \alpha + \beta_1 Ret_{s,i,[t,t+t_1]} + \beta_2 FP_s + \beta_3 Ret_{s,i,[t,t+t_1]} \times FP_s + \gamma X_{i,t} + \epsilon_{s,i,[t+t_1,t+t_2]}$,
where FP_s is a dummy variable equal to 1 if the news article s is positioned on the front page; $Ret_{s,i,[t,t+t_1]}$ denotes the return on security i during the immediate period $[t, t+t_1]$ after publication of the news article s, and $Ret_{s,i,[t+1,t+t_2]}$ is the return during the delayed period $[t+t_1, t+t_2]$. The vector of controls $X_{i,t}$ is empty in columns marked with (1), consists of month and hour fixed effects in columns marked with (2), and also includes log firm size and industry fixed effects in columns marked with (3). The main coefficient of interest is β_3 on the interaction term $Ret_{s,i,[t,t+t_1]} \times FP_s$ (highlighted in blue).
Panel 1 looks at the following time windows: $t_1 \in \{5\,min, 10\,min\}$, $t_2 \in \{45\,min, 60\,min, 90\,min\}$.
Panel 2 considers the following time windows: $t_1 \in \{30\,min, 45\,min\}$, $t_2 \in \{90\,min, 120\,min\}$.

Panel 1: Longer-term continuation from the first 5-10 minutes

	$t_1 = 5\,min, t_2 = 45\,min$			$t_1 = 10\,min, t_2 = 45\,min$		
	(1)	(2)	(3)	(1)	(2)	(3)
$Ret_{s,i,[t,t+5\,min]}$	0.105	0.102	0.104	0.103†	0.104†	0.096†
Standard error	*(0.084)*	*(0.084)*	*(0.085)*	*(0.057)*	*(0.057)*	*(0.058)*
$Ret_{s,i,[t+5\,min]} \times FP_s$	0.368**	0.374**	0.389**	0.585**	0.588**	0.606**
Standard error	*(0.098)*	*(0.099)*	*(0.100)*	*(0.065)*	*(0.065)*	*(0.066)*
# FP SI articles	430	430	430	430	430	430
# Non-FP SI articles	2,440	2,440	2,440	2,440	2,440	2,440

	$t_1 = 5\,min, t_2 = 60\,min$			$t_1 = 10\,min, t_2 = 60\,min$		
	(1)	(2)	(3)	(1)	(2)	(3)
$Ret_{s,i,[t,t+5\,min]}$	0.423**	0.465**	0.464**	0.240**	0.241**	0.235**
Standard error	*(0.103)*	*(0.104)*	*(0.105)*	*(0.077)*	*(0.077)*	*(0.078)*
$Ret_{s,i,[t+5\,min]} \times FP_s$	-0.181	-0.184	-0.170	0.334**	0.335**	0.352**
Standard error	*(0.121)*	*(0.122)*	*(0.124)*	*(0.088)*	*(0.089)*	*(0.090)*
# FP SI articles	432	432	432	432	432	432
# Non-FP SI articles	2,464	2,464	2,464	2,464	2,464	2,464

	$t_1 = 5\,min, t_2 = 90\,min$			$t_1 = 10\,min, t_2 = 90\,min$		
	(1)	(2)	(3)	(1)	(2)	(3)
$Ret_{s,i,[t,t+5\,min]}$	0.325**	0.324**	0.321**	0.112	0.110	0.109
Standard error	*(0.102)*	*(0.103)*	*(0.106)*	*(0.079)*	*(0.079)*	*(0.081)*
$Ret_{s,i,[t+5\,min]} \times FP_s$	-0.192	-0.189	-0.192	0.287**	0.292**	0.296**
Standard error	*(0.122)*	*(0.122)*	*(0.124)*	*(0.091)*	*(0.091)*	*(0.092)*
# FP SI articles	434	434	434	434	434	434
# Non-FP SI articles	2,481	2,481	2,481	2,481	2,481	2,481

**, *, and † denote significance at the 1%, 5%, and 10% levels, respectively.

Panel 2: Longer-term continuation from the first 30-45 minutes

	$t_1 = 30\,min, t_2 = 90\,min$			$t_1 = 45\,min, t_2 = 90\,min$		
	(1)	(2)	(3)	(1)	(2)	(3)
$Ret_{s,i,[t,t+5\,min]}$	0.245**	0.251**	0.253**	0.142**	0.143**	0.147**
Standard error	*(0.034)*	*(0.034)*	*(0.035)*	*(0.018)*	*(0.018)*	*(0.018)*
$Ret_{s,i,[t+5\,min]} \times FP_s$	-0.139**	-0.142**	-0.145**	-0.256**	-0.256**	-0.253**
Standard error	*(0.040)*	*(0.040)*	*(0.040)*	*(0.022)*	*(0.023)*	*(0.023)*
# FP SI articles	434	434	434	434	434	434
# Non-FP SI articles	2,481	2,481	2,481	2,481	2,481	2,481

	$t_1 = 30\,min, t_2 = 120\,min$			$t_1 = 45\,min, t_2 = 120\,min$		
	(1)	(2)	(3)	(1)	(2)	(3)
$Ret_{s,i,[t,t+5\,min]}$	0.243**	0.249**	0.247**	0.140**	0.140**	0.138**
Standard error	*(0.035)*	*(0.035)*	*(0.036)*	*(0.020)*	*(0.020)*	*(0.021)*
$Ret_{s,i,[t+5\,min]} \times FP_s$	-0.185**	-0.191**	-0.192**	-0.301**	-0.303**	-0.310**
Standard error	*(0.041)*	*(0.041)*	*(0.042)*	*(0.023)*	*(0.023)*	*(0.026)*
# FP SI articles	435	435	435	435	435	435
# Non-FP SI articles	2,494	2,494	2,494	2,494	2,494	2,494

**, *, and † denote significance at the 1%, 5%, and 10% levels, respectively.

Table 6: Trading volumes and absolute price changes following PI and SI front page news articles.
Panel 1 compares trading volumes over five-, ten-, and sixty-minute windows following PI front page news articles against trading volumes following SI front page news articles.
Panel 2 compares absolute price changes over five-, ten-, and sixty-minute windows following PI front page news articles against absolute price changes following SI front page news articles.

Panel 1: Trading Volume

	Front Page SI News	PI News	Difference
First 5 min	0.10%	0.18%	0.09%
Standard Error	*(0.0002)*	*(0.0006)*	*(0.0007)*
# Observations	408	438	–
First 10 min	0.19%	0.30%	0.11%
Standard Error	*(0.0010)*	*(0.0004)*	*(0.0011)*
# Observations	415	446	–
First 60 min	0.79%	1.02%	0.22%
Standard Error	*(0.0018)*	*(0.0032)*	*(0.0036)*
# Observations	432	492	–

Panel 2: Absolute Price Changes

	Front Page SI News	PI News	Difference
First 5 min	0.45%	0.80%	0.35%**
Standard Error	*(0.0006)*	*(0.0007)*	*(0.0009)*
# Observations	408	438	–
First 10 min	0.60%	0.97%	0.37%**
Standard Error	*(0.0008)*	*(0.0008)*	*(0.0011)*
# Observations	415	446	–
First 60 min	1.08%	1.39%	0.31%†
Standard Error	*(0.0014)*	*(0.0011)*	*(0.0019)*
# Observations	432	492	–

** and * denote significance at the 1% and 5% levels, respectively.

Table 7: Serial correlation in returns following PI and SI front page news articles. Each column estimates the following specification:
$Ret_{s,i,[t+t_1,t+t_2]} = \alpha + \beta_1 Ret_{s,i,[t,t+t_1]} + \beta_2 PI_s + \beta_3 Ret_{s,i,[t,t+t_1]} \times PI_s + \gamma X_{i,t} + \epsilon_{s,i,[t+t_1,t+t_2]}$,
where PI_s is a dummy variable equal to 1 if the news article s is marked as "primary important"; $Ret_{s,i,[t,t+t_1]}$ denotes the return on security i during the immediate period $[t, t+t_1]$ after publication of the news article s, and $Ret_{s,i,[t+1,t+t_2]}$ is the return during the delayed period $[t+t_1, t+t_2]$. The main coefficient of interest is β_3 on the interaction term $Ret_{s,i,[t,t+t_1]} \times PI_s$ (highlighted in blue). The vector of controls is empty in columns marked with (1), consists of month and hour fixed effects in columns marked with (2), and also includes log firm size and industry fixed effects in columns marked with (3). The considered time intervals are $(t_1, t_2) \in \{(5\,min, 10\,min), (5\,min, 15\,min)\}$.

	$t_1 = 5\,min, t_2 = 10\,min$			$t_1 = 5\,min, t_2 = 15\,min$		
	(1)	(2)	(3)	(1)	(2)	(3)
$Ret_{s,i,[t,t+5\,min]}$	0.056	0.059	0.061	0.140**	0.131**	0.134**
Standard error	*(0.038)*	*(0.038)*	*(0.039)*	*(0.047)*	*(0.048)*	*(0.049)*
$Ret_{s,i,[t+5\,min]} \times PI_s$	0.0	0.015	0.029	0.037	0.044	0.058
Standard error	*(0.047)*	*(0.048)*	*(0.049)*	*(0.059)*	*(0.059)*	*(0.060)*
# PI articles	446	446	446	458	458	458
# FP SI articles	415	415	415	422	422	422

**, *, and † denote significance at the 1%, 5%, and 10% levels, respectively.

Table 8: Comparison of trading volumes and absolute price changes within ten minutes of non-front page SI articles during quiet and busy times. **Panel 1** considered trading volumes, while **Panel 2** looks at absolute price changes.

Panel 1: 10-Minute Trading Volume

	News in Quiet Times	News in Busy Times	Difference
Window: day	0.04%	0.06%	-0.02%**
Standard Error	*(0.00003)*	*(0.00006)*	*(0.00007)*
# Observations	1,289	1,056	–
Window: 5 hours	0.04%	0.06%	-0.02%*
Standard Error	*(0.00002)*	*(0.00006)*	*(0.00008)*
# Observations	1,414	931	–
Window: 2 hours	0.04%	0.05%	-0.01%*
Standard Error	*(0.00003)*	*(0.00005)*	*(0.00006)*
# Observations	1,247	1,098	–

Panel 2: 10-Minute Absolute Price Changes

	News in Quiet Times	News in Busy Times	Difference
Window: day	0.19%	0.23%	-0.04%**
Standard Error	*(0.0001)*	*(0.0001)*	*(0.0002)*
# Observations	1,289	1,056	–
Window: 5 hours	0.20%	0.22%	-0.02%
Standard Error	*(0.0001)*	*(0.0001)*	*(0.0001)*
# Observations	1,414	931	–
Window: 2 hours	0.20%	0.21%	-0.01%
Standard Error	*(0.0001)*	*(0.0001)*	*(0.0001)*
# Observations	1,247	1,098	–

** and * denote significance at the 1% and 5% levels, respectively.

Table 9: This table lists and labels the topics identified from a large and representative dataset of financial news from Reuters. The topics are listed in order from most to least common in the corpus. For each topic code, I provide the fifteen terms in the vocabulary that are most likely to appear conditional on that topic, a well the as the frequency of that topic in the corpus.

	Topic Label	Most Common Terms	Frequency
#1	Technology	data, technology, companies, security, information, comment, including, according, government, card, software, credit, declined, store, did	23.83%
#2	Earnings & Performance	percent, year, sales, quarter, million, analysts, share, rose, revenue, estimates, profit, earnings, fell, cents, average	9.84%
#3	Financial Services	million, year, bank, today, officer, financial, chief, according, statement, executive, firm, largest, new york, investment, unit	8.29%
#4	Automobile	vehicles, cars, downturn, sales, automaker, deliveries, turnover, air, current, safety, in-house, auto, backlog, switches, parts	6.48%
#5	Air transport	internet, service, search, aircraft, today, flight, plane, contract, engine, carrier, air, airline, satellite, web, traffic	6.22%
#6	Litigation	court, case, judge, federal, workers, law, claims, million, filed, trial, state, lawsuit, ruling, lawyers, attorney	5.96%
#7	Management	ceo, president, job, board, women, chairman, director, vice, named, executive, world, role, according, chief, leave	5.70%
#8	Healthcare	drug, patients, care, percent, flu, health, treatment, disease immunize, today, study, research, treatments, medical, medicines	4.92%
#9	Operational	according, years, got, little, long, later, industry, great, left, good, costs, international, commercial, saying, end	4.69%
#10	Business & Strategy	year, percent, executive, chief, market, officer, brand, today, products, global, world, plans, month, second, sales	4.63%
#11	Mergers & Acquisitions	deal, offer, price, people, bid, shares, comment, buy, companies, analyst, takeover, shareholders, matter, investors, call	4.40%
#12	Advertising	tv, like, food, video, according, subscribers, products, review, pay, media, content, digital, cable, website, advertising	4.15%
#13	Regulations	offer, regulator, today, agency, government, information, review, adjudicate, statement, public, rules, letter, asked, questions, mailed,	3.66%
#14	Retail	stores, chain, retailer, sales, retail, years, home, online, customers, holiday, shoppers, foods, black, season, target	3.62%
#15	Employees	companies, time, people, make, week, including, work, interview, want, just, need, way, making, does	3.61%

Table 10: This table presents the results from pairwise comparisons between sets of news articles from different positions and different levels of importance. Each cell gives the p-value of a χ-square test of independence between the topic distributions in the two specified sets of news articles. Topics are estimated on the large training dataset from Reuters, and the results display robustness to the number of topics varying from 10 to 25.

Panel 1 compares front page SI news articles against non-front page ones.

Panel 2 compares PI news articles against front page SI news articles.

Panel 1: Front Page SI versus Non-Front Page SI

# Topics in Model	p-value
10 topics	0.8670
15 topics	0.8776
20 topics	0.8731
25 topics	0.7801

Panel 2: PI versus Front Page SI

# Topics in Model	p-value
10 topics	0.1236
15 topics	0.0836†
20 topics	0.0526†
25 topics	0.0417*

* and † denote significance at the 5% and 10% levels, respectively.

Table 11: Summary statistics of the financial experts surveyed regarding the news. The table presents the breakdown of the survey respondents across MBA students and active finance professionals, as well as the breakdown of precise affiliations in each group across specific schools and institutions. [Note: The survey is currently ongoing, and the results are updated on an ongoing basis.]

Affiliation Type	Institution	Percentage
MBA Students		**5% of Total**
Breakdown:		
	Harvard Business School	20%
	Columbia Graduate School of Business	20%
	Chicago Booth	20%
	UVA Darden	20%
	Georgetown McDonough	20%
Professionals		**95% of total**
Breakdown:		
	Investment Managers	19%
	BlackRock	
	The Vanguard Group	
	State Street	
	Pacific Investment Management Co	
	Wellington Management Company	
	Northern Trust Company	
	T. Rowe Price	
	Dodge & Cox Funds	
	Acadian Asset Management	
	Eachwin Capital	
	Crane Asset Management	
	Hedge Funds	8%
	Bridgewater Associates	
	AQR Capital Management	
	Tudor Investment Corp	
	BlueMountain Capital Management	
	Blue Ridge Capital	
	QTrade Capital	
	Bluegrass Capital	
	One East Partners	
	Private Equity & Venture Capital	13%
	Blackstone Group	
	Warburg Pincus	
	Motive Partners	
	Garrison Investment Group	
	ATL Partners	
	Tamarisc	
	Private Investors	1%
	Pension Funds	1%
	North Carolina Retirement System	

Table 11 (Continued): Summary statistics of the financial experts surveyed regarding the news. The table presents the breakdown of the survey respondents across MBA students and active finance professionals, as well as the breakdown of precise affiliations across specific schools and institutions. [Note: The survey is currently running, and the results are updated on an ongoing basis.]

Affiliation Type	Institution	Percentage
Professionals		
Breakdown:		
	Banks and Broker Dealers	39%
	JP Morgan	
	Morgan Stanley	
	Goldman Sachs	
	Bank of America Merrill Lynch	
	BNP Paribas	
	Credit Suisse	
	Deutsche Bank	
	Wells Fargo	
	Royal Bank of Canada	
	UBS	
	Standard Chartered Bank	
	Citizens Bank	
	The NEX Group	
	Investment Banks	3%
	Barclays Capital	
	Lazard	
	Consulting	1%
	Boston Consulting Group	
	Insurance	2%
	Massachusetts Mutual Life Insurance	
	Voya Financial	
	Real Estate	1%
	Condor Partners	
	Government Agencies	1%
	Federal Reserve Board	
	Corporations	4%
	Nike, Inc.	
	Shaw's Supermarkets	
	The Walt Disney Company	
	Tiffany & Co.	
	Financial and Tax Advisory	2%
	Princeton Tax Services	
	DK Partners	
	Media	1%
	American Outlook, Inc.	
	Non-Profit	1%
	Ford Foundation	
	Other	3%

Table 12: This table presents the aggregated responses of financial experts to the news survey. The results are displayed for the full sample, and separately for the sub-samples of MBA students and active finance professionals. For all samples, the results are presented with and without the respondents who do not complete the survey in full.
Panel 1 reports the frequency with which the financial professionals and MBA students identify front page SI news articles as more impactful than the non-front page ones.
Panel 2 reports the incidence of the financial experts choosing PI news articles as more impactful than the front page SI news articles.
[Note: The survey is currently running, and the results are updated on an ongoing basis.]

Panel 1: Front Page SI versus Non-Front Page SI

Respondent Type	Choosing Front Page	Standard Error	# Respondents
Finance Professionals	48.46%	*(1.27%)*	104
MBA students	47.83%	*(5.21%)*	5
All Respondents	48.42%	*(1.23%)*	109
Finance Professionals (excl. attritors)	48.31%	*(1.28%)*	93
MBA Students (excl. attritors)	47.83%	*(5.21%)*	5
All Respondents (excl. attritors)	48.27%	*(1.25%)*	98

Panel 2: PI versus Front Page SI

Respondent Type	Choosing PI	Standard Error	# Respondents
Finance Professionals	63.53%**	*(1.68%)*	104
MBA Students	55.17%	*(6.57%)*	5
All Respondents	63.02%**	*(1.63%)*	109
Finance Professionals (excl. attritors)	64.23%**	*(1.72%)*	93
MBA Students (excl. attritors)	55.17%	*(6.57%)*	5
All Respondents (excl. attritors)	63.65%**	*(1.66%)*	98

** denotes a proportion differing from 50% with significance at the 1% level.

Appendix A Technical Details

A.1 Latent Dirichlet Allocation

I briefly present the Latent Dirichlet Allocation methodology for identifying representative topics covered by the financial news in the training corpus. For additional details on the methodology, please refer to Blei et al (2003).

Let D denote the set of financial news documents in the training corpus, with $d \in D$ representing an individual document. The document d is a sequence of N words: $d = (w_1, ..., w_N)$, where w_n is the nth term from the vocabulary W to appear in the document d. The vocabulary W is constructed as described in Appendix A.2.

The latent set of topics is denoted by T, where each element $t \in T$ is a unit vector in the k-dimensional space. The parameter k is the desired number of topics, specified by the researcher.

The Latent Dirichlet Allocation algorithm conceptualizes each document as a sequence of words drawn from a latent distribution $\mathcal{D}_d$ over topics. The distribution $\mathcal{D}_d$ is itself randomly determined for each document: in particular, for each document d, $\mathcal{D}_d$ is a multinomial distribution whose parameters are a random variable drawn from a pre-specified Dirichlet prior.

Note that unlike other methods such as the probabilistic Latent Semantic Indexing approach (see Hofmann (1999)), the Latent Dirichlet Allocation method does not require the parameters to be estimated individually for each document. This offers two advantages highlighted by Blei et al (2003). First, by reducing the number of estimated parameters, the Latent Dirichlet Allocation approach reduces the computational complexity of the estimation problem. Second, and more importantly, the Latent Dirichlet Allocation method allows for the generation of any arbitrary document and facilitates the evaluation of the likelihood of out-of-sample documents. This ability to represent out-of-sample documents in terms of the identified topics is essential to the application in this paper.

The generative process assumed by the Latent Dirichlet Allocation algorithm is as follows.

- Pre-specify model parameters: ξ, α, β.
- To construct each new document d:
  1. Choose the document length $N_d \sim Poisson(\xi)$.
  2. Choose a distribution over topics $\theta_d \in Dir(\alpha)$.
  3. Fill the N words in the document d by sequentially choosing each word w_n as follows:

(a) Choose a topic $t_n \sim Multinomial(\theta_d)$.

(b) Choose a word w_n from $\mathbb{P}\{w_n|t_n, \beta\}$, the conditional probability distribution over words in the vocabulary given the chosen topic t_n.

The model relies on three parameters, ξ, α, and β. The parameter ξ is chosen to best match the set of document lengths in the corpus, assuming that the lengths are drawn from a Poisson distribution. This parameter is independent of the rest of the process, and therefore I forego it in the remainder of the discussion.

The key model parameters of interest are α and β: α is a k-dimensional vector that governs the relative frequencies of the k topics, and β is a k-by-$2,000$ matrix that specifies the likelihood of each word in the vocabulary conditional on each of the k topics. Hence, the element in the ith row and jth column of β is $\beta_{i,j} = \mathbb{P}\{w_j = 1|z_i = 1\}$.

In theory, the parameters α and β are estimated to maximize the likelihood of observing the actual corpus of documents D. The conditional probability of observing a document d given the model parameters α and β is given by:

$$\mathbb{P}\{d|\alpha, \beta\} = \frac{\Gamma(\Sigma_{i=1}^{k}\alpha_i)}{\Pi_{i=1}^{k}\Gamma(\alpha_i)} \int \left(\Pi_{i=1}^{k}\theta_i^{\alpha_i - 1}\right) \times \left(\Pi_{n=1}^{N}\Sigma_{i=1}^{k}\Pi_{j=1}^{V}(\theta_i\beta_{i,j})^{w_n^j}\right) d\theta, \tag{A.1}$$

where w_n^j denotes the jth component of the nth word vector w_n, and $\Gamma(\cdot)$ is the Gamma function.

A.2 Text Preprocessing

The Latent Dirichlet Allocation algorithm takes as its input a set of documents, each represented by a sequence of terms from a pre-specified vocabulary. Before applying the topic modeling methodology, I need to identify a relevant vocabulary to represent the financial news documents. I proceed in three steps.

First, in order to focus on the set of relevant terms, I begin by stripping out all "stop words." To identify "stop words," I use the list provided by the University of Glasgow Information Retrieval Group.[9]

Second, I construct the vocabulary using not only single words appearing in the TRC2 news corpus, but also common pairs of words. The Latent Dirichlet Allocation method is a bag-of-words method, meaning that the algorithm ignores the ordering of terms within a document and treats each term as an independently drawn random variable. In theory, this may be a problematic assumption, particularly for financial news, where some concepts are captured by phrases, for example "traditional enterprise" or "stock exchange." In order to

[9]The full list of stop words can be accessed at http://ir.dcs.gla.ac.uk/resources/linguistic_utils/stop_words.

account for this feature of the data, I augment the vocabulary of unigrams (single words) appearing in the corpus with bigrams (pairs of words).

Lastly, I limit my attention to the most common and representative terms. In particular, I focus on the terms that appear in at least two distinct documents and that appear in no more than 70% of the documents in the training corpus. Furthermore, the terms are ranked according to frequency in order to capture relative importance; the final vocabulary is comprised of the top $2,000$ terms.

A.3 Topic Model Estimation

I estimate the model varying the number of iterations and the number of identified topics. The model's fit flattens out at around fifteen topics.

In practice, the expression in (A.1) is intractable, and hence parameter estimation relies on approximate inference methods. Following Griffiths and Steyvers (2004), I estimate the parameters from the training corpus of documents using a collapsed Gibbs sampling algorithm.

I vary the number of iterations of the sampling algorithm from 30 to 1,000, and find that the marginal improvement in the model's fit is largest up to approximately 250 iterations, and mostly flattens out after 500 iterations (see Panel 1 of Figure 7, which plots the log likelihood as a function of the number of iterations for a model with $k = 15$ topics). The results in the paper come from the estimation algorithm with 500 iterations for all considered specifications.

[FIGURE 7 AROUND HERE]

The results from estimating the parameters of the Latent Dirithlet Allocation topic model for $k \in \{10, ..., 40\}$ indicate that the model's fit is best at around fifteen topics. Panel 2 of Figure 7 plots the model fit for $k \in \{10, ..., 40\}$. For each number of topics, the model is estimated using the collapsed Gibbs sampler with 500 iterations. The figure shows the final log likelihood for each specification. The model's fit improves somewhat as the number of topics increases from ten to fifteen, with an increase in log likelihood from -6.01×10^5 to -6.00×10^5. Increasing the number of topics to 20, 25, or 30 does not offer marginal improvements over the $k = 15$ specification. Increasing the number of topics further to 35 or 40 markedly decreases the estimated log likelihood. Overall, the $k = 15$ specification achieves the best fit after 500 iterations.

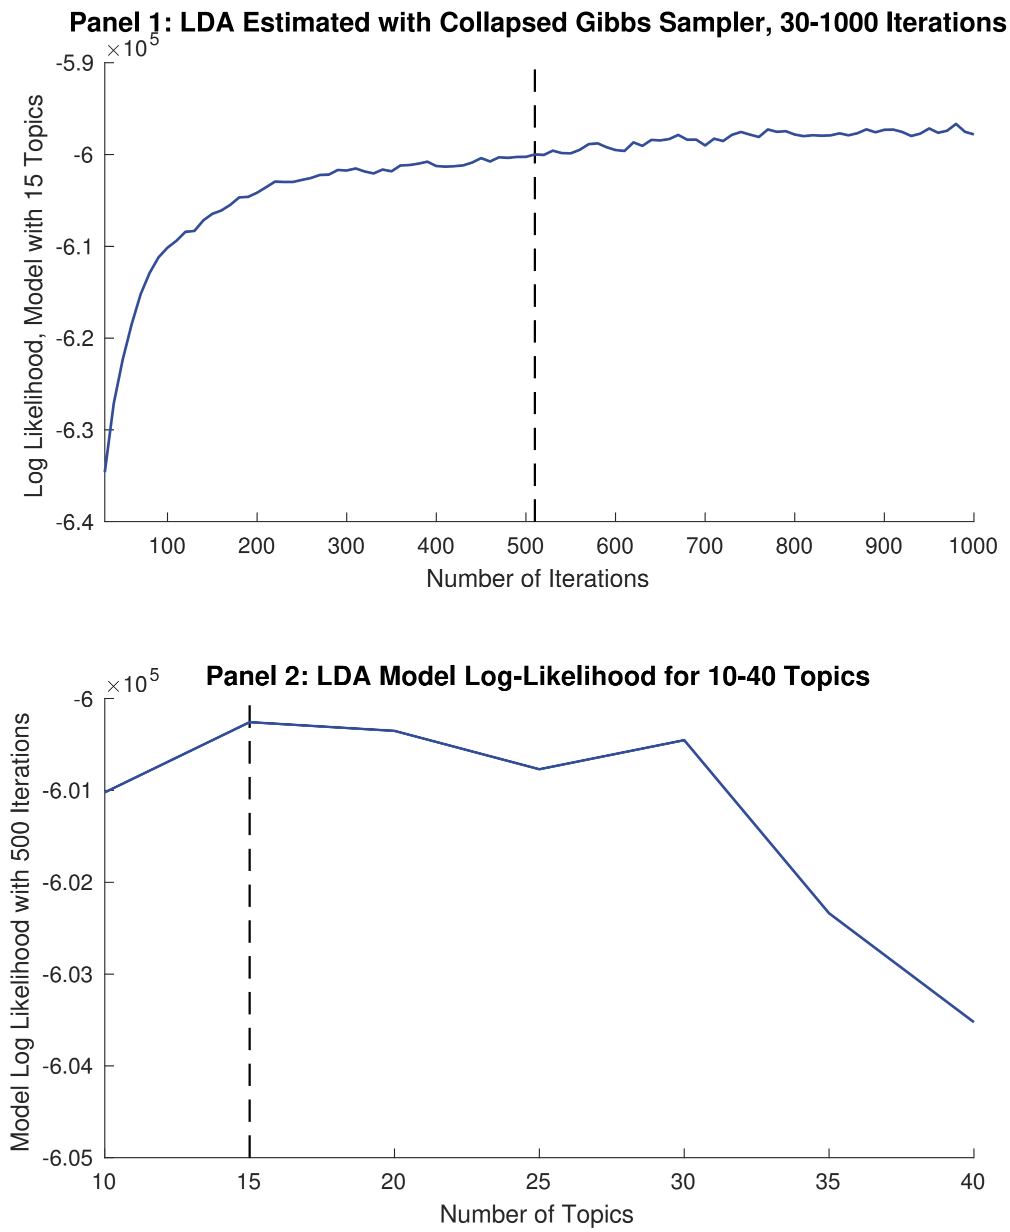

Figure 9: Log likelihood for different estimations of the Latent Dirithlet Allocation model. **Panel 1** plots model fit as a function of the number of iterations, using collapsed Gibbs sampling estimation and fixing the number of topics at $k = 15$. **Panel 2** displays the model log likelihood for the number of topics varying from $k = 10$ to $k = 40$, estimating each specification with 500 iterations using the collapsed Gibbs sampler.