Beyond attention: the causal effect of media on information production *

Bruce Yuan Li †

First version: May 30, 2018; Current version: December 14, 2018

Latest version can be found here

Abstract

This paper shows that media coverage causes investors to gather more information and analysts to produce more earning forecasts. I exploit random variation in the visual salience of corporate press releases to financial journalists as an instrument to media coverage. Doubling the amount of media coverage increases the number of EDGAR searches by 31% and the number of analysts issuing earnings forecasts by 78% in a two-day period. The evidence is most consistent with the theory of rational attention allocation: sophisticated investors acquire more information for media-covered events as media coverage signals higher variances of returns, and thus higher payoffs from having more precise information. Analysts cater to the increased information demand from institutional investors by responding to media-reported events. The results suggest that different information channels do interact, and financial media complements other channels.

Keywords: media, analyst, EDGAR search, institutional investor

JEL Codes: G12, G14

*I am indebted to my advisor, Gerard Hoberg, and my committee members, Kenneth Ahern, Oguz Ozbas, Yongxiang Wang, and TJ Wong, for their invaluable help and guidance. I also thank Heitor Almeida, William Cong, Itay Goldstein, Zhiguo He, Christopher Jones, E Han Kim, John Matsusaka, Sheridan Titman, Song Ma, Emily Nix, Christopher Parsons, Andrei Shleifer, Ben Zhang, two anonymous referees, and participants at the seminars at the University of Southern California for helpful comments and suggestions. All remaining errors are my own.

†PhD Student at Marshall School of Business, USC. Email: yli268@marshall.usc.edu

A well-functioning financial market requires information to flow efficiently between corporations and investors. Such information flows are often facilitated by a variety of information intermediaries, including financial media, analysts, and investors' information acquisition. While the literature has studied how each of these information intermediaries individually affects the information environment and stock trading, little is known about whether and how these intermediaries interact. Uncovering potential interactions helps us better understand the micro-process through which stock prices incorporate new information. In this paper, I study how media coverage affect the information acquisition of investors and earnings forecasts issuance of analysts.

Whether financial media complement, substitute, or do not affect other information intermediaries is ex-ante unclear. On the one hand, investors and analysts might focus on less salient events because these events are often associated with delayed market responses and possible abnormal returns (Cohen and Frazzini, 2008; Menzly and Ozbas, 2010). On the other hand, media attention could attract more retail traders and increase return volatility, thus increasing the reward of having more precise signals. Or, media coverage is simply a side show for more sophisticated market participants like analysts and institutional investors, who often have alternative and better access to information. Ultimately, the relationship between media and other information intermediaries is an open empirical question.

To answer the question, this paper studies that for a corporate announcement, does the media coverage of it affects amount of investors' information acquisition and analysts' earning forecasts about the firm. To construct a comprehensive sample of corporate announcements, this paper uses a dataset that includes almost all the press releases issued by the US public firms during 2004 to 2017. Using an event-study approach based on these press releases, I find that when media articles from Dow Jones Newswire cover an announcement, investors acquire more information about the firm through the SEC EDGAR system, and more analysts issue earning forecasts for that firm.

One challenge for convincingly answering the above question, however, is that endogenous event- or firm-characteristics would affect the coverage decisions of both media and other information intermediaries. For example, an earning surprise would attract media coverage, but at the same time also induce analysts to update their earning forecasts, thus creating a positive correlation between the responses from different information intermediaries.

This paper proposes a novel identification strategy based on the unique way that wire journalists process information and produce news. For a newswire journalist, the typical workflow is to monitor a real-time press release feed, select newsworthy events, and quickly replay the main points to their subscribers. Because different newswires compete over speed, newswire journalists need to produce news articles almost in real-time, and in many cases these news articles contain only a headline summarizing the event.

A particular challenge for these journalists is that press releases cluster at specific times within a day. For investors or analysts who only follow a handful of firms, such clustering is innocuous. However, newswire journalists typically don't specialize in any industry and thus cover the events for the *entire* market. As a result, when many unrelated firms issue press releases at the same time, the amount of information faced by newswire journalists could easily add up and cause cognitive burden.

Key to my identification is that within a busy cluster of press releases, some press releases will be more visually salient than similar others and thus receive more media coverage. Real-time systems, which the journalists monitor, follow a "first-in-first-out" rule: new content always shows up at the top of the user interface, pushing current content down and onto later pages. Thus, the length of time for which a press release stays in a prominent position (e.g. first page or the top of the screen) is determined by the speed of new releases replacing it. Now let's think of a cluster of press releases. For a press release queued near the beginning of the cluster, its "on-screen" time is short because the whole cluster follows and quickly replace it. In comparison, a press release queued near the end of the cluster is followed by much fewer releases, thus it gets a longer on-screen time than the releases queued near the beginning. In this paper, I define each busy cluster as the first 10 seconds of an hour, and within such a tight time frame, firms neither have incentives nor, in many cases, the ability to precisely control the order within the queue[1]. On top of that, it is the number of press releases from *other* firms that shifts the size of the cluster, thus the ultimate on-screen time is unlikely determined by the actions of any single firm. Indeed, using variable balancing tests, I find my measure of on-screen time is uncorrelated with observable firm- or event-characteristics, while the same characteristics significantly predict the amount of media coverage.

I find evidence for the on-screen time to be a valid instrument for media coverage. On the relevance condition, I find that a shorter on-screen time leads to less media coverage. The effect remains significant and large even with a rigid set of controls like firm-year interacted, date-hour, and detailed topic fixed effects. On the exclusion restriction condition, Moreover, I find that when the shorter on-screen time is caused by press releases from private firms, which likely only impact wire journalists but not analysts and investors, the same results obtain.

Using an instrumental variable approach, I find that media coverage of an event significantly increases the number of requests on the SEC EDGAR system. Doubling the amount of media coverage on the press release day increases the number of abnormal EDGAR requests by 31% in a two-day period. Perhaps more surprisingly, I also find more analysts issue earning forecasts for media-covered events. Doubling the amount of media coverage leads to a 78% increase in the number of analysts issuing earnings forecasts. The effects are short-lived: the effect on analysts disappears after two days, and the effect on EDGAR searches disappears after five days. As a placebo test, I find no effect of media coverage on the day prior to the press release.

The effect of media coverage can be through two possible channels. First, media cov-

[1]For example, the user interface in PR Newswire only allows users to schedule press releases to the minute.

erage improves the information production technology of investors and analysts. One can think of such improvement as an additional signal that complements the research by investors or analysts (Goldstein and Yang, 2015), or extra attention make investors and analysts aware of the events, reducing the search problem. Second, media coverage might increase the reward of having a precise signal. Sophisticated investors would rationally allocate more resource to ”learn” about firms that have higher ex-ante payoff variance (Kacperczyk, Van Nieuwerburgh, and Veldkamp, 2016). The increased information demand from also impacts the payoff to analysts, and clients' demand is one of the most important determinants for issuing earning forecasts or not. As Brown et al. (2015) concludes from a survey of 182 analysts, "Demand from their clients is analysts' most important motivation for making profitable stock recommendations and their second most important motivation for issuing accurate earnings forecasts".

Empirical evidence is more consistent with the reward-side rather than the production technology-side explanations. First, if media coverage takes effect by improving information production technology, then we would expect lower-skilled investors or analysts should benefit more and exhibit larger effects. However, I find institutional investors actually react even stronger to media coverage. Analysts of higher skills, measured by accuracy or experience, also show either indistinguishable or larger responses to media coverage. Second, the attention channel is also unlikely. I find investors who have searched the same firm in the previous month, thus already know about the firm and has less of a ”search problem” (Barber and Odean, 2008), also show larger responses to media coverage. Furthermore, it is hard to imagine that analysts would not know about an event if not for the media coverage of it. After all, it is analysts' job to maintain an information advantage.

In comparison, I find consistent evidence with the reward-side explanation. First, I find in the effect of media on analysts are stronger in the subsample of firms with higher institutional ownerships. Second, the market outcomes are consistent with a tug of war between two types of investors, one attracted by media due to attention, and the other that consciously trades more media-covered firms, profiting by trading against these potentially uninformed traders. On the event day, media coverage significantly widens the effective spread, suggesting a relative increase in informed traders. In the next two days, while media coverage does not affect absolute returns, it significantly increases the intraday price ranges. Peress (2014) attributes similar effects to "less price-sensitive traders who transact at less favorable prices". I also find that media neither improves nor deteriorates price efficiency, measured by the delayed response ratio from Dellavigna and Pollet (2009). The result echoes Blankespoor, deHaan, and Zhu (2018) who also find no effect of media on price efficiency, and is consistent with sophisticated and less sophisticated investors having opposite effects on the price efficiency.

The most important contribution of the paper is to cleanly identify that media coverage has a positive effect on the information produced by sophisticated market participants. Existing empirical evidence shows that media coverage increases the volatility of expected returns (Peress, 2014; Blankespoor, deHaan, and Zhu, 2018), changes the investor base (Barber and Odean, 2008), and possibly increases mispricing (Hillert, Jacobs, and Müller, 2014; Ahern and Sosyura, 2015). As a result, sophisticated investors may interpret media coverage as a reliable signal which suggests a higher reward to their information production. The increased information demand from institutional investors also impacts analysts, who would cater to such demand by also shifting their information production to media-covered events.

This paper first contributes to the literature that studies the interaction between different information channels in the financial market. Existing literature mostly focuses on the interaction between sell-side analysts and sophisticated investors. Kacperczyk and Seru (2007) find that fund managers with higher skills rely less on the public information from stock analysts. Chen et al. (2017) find that when the information production from analysts exogenously decreases due to the closures and mergers of brokerage firms, sophisticated investors scale up their information acquisition. This paper provides novel evidence that media could also impact the information production of investors and analysts. The results challenge the conventional wisdom that media does not matter for sophisticated market participants who have better information access and higher information processing skills.

This paper also contributes the literature that studies the role of media in the financial market. I refer readers to Tetlock (2014) for an excellent review of the literature. In particular, this paper adds to the work that studies the causal effect of media coverage. Engelberg and Parsons (2011) use extreme local weather events as exogenous shocks to news delivery. Peress (2014) uses a set of newspaper strike events as exogenous shocks to news production. Blankespoor, deHaan, and Zhu (2018) use the staggered implementation of robo-journalism to study the causal impacts of synthesizing information from analysts and other sources. Fedyk (2018) uses the random positioning of news on Bloomberg terminals to study the effects of being on the front page. This paper introduces a novel identification strategy that stems from the inefficiency in media production. Compared with previous work, the strategy also applies to a more representative and larger sample.

1 Empirical setup

1.1 Institution background

The empirical setup of the paper centers around the events of corporate disclosures through press releases. Since Regulation Fair Disclosure (Reg FD), press releases become an increasingly popular method for corporate disclosures, given their fast delivery and broad reach[2]. Typically, public firms will choose one of the top four wires, namely, PR Newswire, Business Wire, Market Wire, and GlobeNewswire, to publish their announcements (Solomon

[2]Reg FD implicitly encourages the use of press releases due to its fast dissemination speed and wide reach of investors. Reg FD states that "technological developments have made it much easier for issuers to disseminate information broadly. Whereas issuers once may have had to rely on analysts to serve as information intermediaries, issuers now can use a variety of methods to communicate directly with the market. In addition to press releases, these methods include, among others, Internet webcasting and teleconferencing". Similar argument can be found in Neuhierl, Scherbina, and Schlusche (2013). The full content of Reg FD can be found at https://www.sec.gov/rules/final/33-7881.htm

and Soltes, 2012). These press releases can cover a wide range of topics, and in my sample, the top three topics are "earnings", "products-services" (e.g., new product releases), and "labor-issues" (e.g., executive appointments).

These corporate disclosures, many of which contain new information that has large market impacts, also spurs much follow-on information production from journalists, analysts, and investors. To see this point, in Figure 1, I plot the percentages of all the news articles, analysts earning forecasts, and web requests on EDGAR that are produced on different days following corporate press releases[3]. Over 35% of all the news articles are published on days when firms issued new press releases. Most earning forecasts by analysts are issued immediately following corporate disclosures, with almost 50% of all the forecasts being published within two days after press release issuance. The evidence on analysts is consistent with the findings in Altınkılıç, Balashov, and Hansen (2013), who show that over 50% of analysts forecasts are issued following earning or guidance reports. Even EDGAR requests peak after corporate disclosures; about 20% of all the requests are made within two days of press releases. Therefore, even though this paper uses an event-study approach based on press releases, it still captures a significant portion of all information production activities.

[Figure 1 here]

While media coverage shows similarly large responses to corporate disclosures as earning forecasts or EDGAR requests do, financial journalists, especially newswire journalists, have drastically different workflow than stock analysts and investors (for simplicity, I call analysts and investors as finance professionals hereafter) for monitoring and processing press releases.

There are three major differences. First, logistically, journalists and finance professionals use different software. In the newsroom, a common software is a press release

[3]The sample covers all the news articles and analysts forecasts covering the firms in my press release sample during 2004-2017. The EDGAR data is from Jan. 2004 to Jun. 2017. The press releases data includes all the press releases on the top 4 press release wires.

feed aggregator that shows press releases from a variety of press release wires in real-time. While in the trading room, traders and analysts often use professional services like Bloomberg or Thomson Reuters to gather real-time information. Typically, the information stream in press release wires is quite ”noisy” in the sense that anyone can pay to publish a press release, thus many releases are promotional or advertisement-like. It is exactly the job of wire journalists to filter out these ”noisy” events and pick material ones. As result, the second different is that the ”information inputs” are very different for journalists and finance professionals. Wire journalists monitor a much larger set of firms and events than financial professionals. The job of wire journalists is often not to conduct in-depth analysis. In many cases, journalists produce media articles that only contain a headline summarizing the event. Because these tasks require mostly general skills, wire journalists typically do not specialize in certain industries thus need to monitor the events for the whole market. In comparison, investors and analysts tend to focus on a much smaller set of firms. For example, Chen et al. (2017) find that fund managers tend to consistently acquire insider trading forms from a small set of firms. Similarly, analysts tend to be industry-focused and only produce research for a handful of firms. The third difference is that wire journalists monitor a passively determined and continuously updating event stream, while finance professionals often rely on other active active search, notification push, or custom event filters to focus on specific information. As a summary, we can see the unique workflow of wire journalists from a job description of a Dow Jones Newswire journalist[4]

> Ian currently manages the U.K. companies desk, overseeing corporate news flashes and quick fire fills for both the Dow Jones Newswire and The Wall Street Journal’s website. The desk covers all stocks from the largest FTSE100s to the smallest AIM companies across the whole range of industries and subject matters.

[4]See the journalist’s bio at http://www.wsj.com/news/author/8056

1.2 Data sources

This section introduces the sources of different datasets used in this paper.

Press release data

I start by compiling a comprehensive set of press releases as the main sample of this paper. The press release data come from the RavenPack PR (press release) Edition. The data contains press releases published in over 10 press release wires from 2004 to 2017. Importantly, the data includes all the top 4 press release wires (PR Newswire, Business Wire, Globe Newswire, and Marketwired). RavenPack also links these press releases to public firms and provides the CUSIP number(s) of associated firms, allowing me to link the press release data to other datasets of the paper.

I impose a set of standard data filters to generate the final sample. First, I require the `RELEVANCE` score of each press release-firm pair to be 100. This filter ensures that the firm is the main subject of the press release, and such filter is recommended by RavenPack user manual. Second, I only include the top 4 press release wires because other major wires in the data are mostly in the non-US regions, such as Canadian Newswire or LSE Regulatory News Service. Third, I remove duplicated releases by requiring the `ENS` score to be 100. Fourth, I require the firm to be in the Compustat/CRSP universe. Fifth, I require that the press release is issued on a trading day. Finally, to better associate later responses with specific press releases, I keep only observations where a firm only issues one press release on that day. Table 12 in the Appendix shows the detailed data cleaning process as well as how the number of observations changes after each step.

A key measure of this paper is the "on-screen" time of each press release. More formally, the "on-screen" time is the length of time of a press release staying at a prominent position of the interface of the software. It is impossible to precisely measure the on-screen time without introducing many assumptions. In this paper, I construct the proxy for on-screen time based on the following observation: the on-screen time of a press release is determined by the speed at which new press releases replacing it. Therefore for each press release, I use the number of New Press Releases After it in the next 30 seconds as an (inverse) proxy for its on-screen time. For parsimony, I refer to the measure as NPRA hereafter. As the NPRA of a press release increases, its on-screen time shortens. Although RavenPack PR Edition starts from 2004, the timestamp variable does not include the precise second information until April 1, 2006. Therefore to correctly measure NPRA, the final sample of the paper starts from April 1, 2006.

Media coverage data

The media coverage data comes from RavenPack DJ Edition, which includes all the news articles published on the Dow Jones Newswire from 2000 to 2017. I apply several data filters to generate the media coverage measures. First, Dow Jones Newswire also include press releases. There press releases are redistribution of the original content by automated algorithms. So I drop these observations from the news measure by excluding news articles whose `NEWS-TYPE` is `PRESS-RELEASE`. Second, I require that the `RELEVANCE` score of each news-firm pair to be 100. Third, I also drop news articles covering analysts analysis and stock market reactions to avoid reverse causality. To do that, I drop news articles whose topic variable `GROUP` is in 'analyst-ratings', 'credit-ratings', 'order-imbalances', 'technical-analysis', 'stock-prices', or 'price-targets'.

For firm i on day t, I define the abnormal number of media coverage, *AbnNews*$_{it}$, as the log of 1 plus the number of media articles from Dow Jones Newswire covering firm i on day t, minus the log of 1 plus the average number of media articles covering firm i per day in the days $[t-70, t-11]$. The definitions of other media coverage variables can be found in Table 11.

EDGAR Log data

To measure information acquisition, I use the EDGAR server log files from SEC[5]. Whenever someone accesses a web page or a filing on the EDGAR system, the log data will create an observation with (1) the IP address of the visitor, (2) the CIK number, a unique firm identifier used by EDGAR, (3) the accession number, a unique filing document ID, and (4) the timestamp of the web request. The data contains all the web traffic records from Feb. 2003 to Jun. 2017 and contains rich information about how investors access information through the EDGAR system. The data is also massive in size: 2016 alone contains over 6.6 billions observations.

Following Loughran and McDonald (2017), I apply several filters to clean the data. First, I omit index page requests (`idx = 1`) to avoid double counting. Second, I drop requests whose server code is of 300 or higher, as these requests are either redirection requests or error requests. The third task is to drop web requests from web crawlers, which access and download SEC filings through automated algorithms. I start by dropping web requests from IP addresses that explicitly reveal themselves as web crawlers (`crawler = 1`). However, not all web crawlers explicitly reveal itself in the user-agent. Existing literature proposes two methods in identifying web crawlers. First, proposed by Drake, Roulstone, and Thornock (2015), an IP address is a web crawler if it made more than five web requests in a minute more than a thousand web requests in a day. Second, proposed by Lee, Ma, and Wang (2015), an IP address is a crawler if it requested information for more than 50 firms in a day. The results of this paper are robust using either of the method. For simplicity, in the main text I only show results using the first method to identify web crawlers.[6]

[5]Details on the data can be found in https://www.sec.gov/dera/data/edgar-log-file-data-set.html. Also see Loughran and McDonald (2017) for a general discussion about the data. Papers using the same dataset include Lee, Ma, and Wang (2015) and Drake, Roulstone, and Thornock (2015) among others

[6]These methods will lead to conservative measures of EDGAR requests from human users, because these methods possibly have large type I errors in identifying web crawlers. The reason is that many institutions have more users than the IP addresses they have. Thus they adopt network address translation (NAT) method to route the web traffics, and as a result, many users might share a single outbound IP address, and the above

Similarly, for firm i on day t, I define abnormal number of EDGAR requests, $AbnEdgar_{it}$, as the log of 1 plus the number of EDGAR requests about the filings from firm i on day t, minus the log of 1 plus the average daily number of EDGAR requests about firm i in the days $[t-70, t-11]$. I augment the EDGAR database by matching the first three octets of IP addresses to the IP ranges of different institutions, which is provided by MaxMind[7].

Analyst data

The analysts' earning forecast data comes from the unadjusted detail file in I/B/E/S. I keep observations that have non-missing EPS, announcement date, analyst ID, and broker firm ID. On a single day, an analyst may issue multiple forecasts that have different period ends for the same firm, and these forecasts would be separate records in the I/B/E/S data. To avoid double counting, for each firm-day, I count the number of unique analysts who issue any EPS forecast as the relevant measure. Thus multiple forecasts from the same analyst would only be count once for a firm-day. I define the abnormal number of analyst forecasts, $AbnAnalyst_{it}$, as the log of 1 plus the number of unique analysts issuing earning forecasts for firm i on day t, minus the average daily number of analysts issuing earning forecasts in days $[t-70, t-11]$.

Other data

The rest of the data comes from a variety of sources. First, the stock return, trading volume, and price data from CRSP. Second, firm characteristics data comes from Compustat. Third, the institutional ownership and institutional names come from Thomson Reuters. Fourth, the effective spread is calculated using the trade and quote data from TAQ.

---

methods might falsely tag these IPs as web crawlers.

[7]One can download the data from the ASN (autonomous system number) data file from MaxMind. To generate a time-varying mapping file between IP blocks and institutions, I use WayBack Machine to extract the historical versions of the ASN file. Details of the data construction process can be found in Online Appendix B.

Summary statistics

Table 1 shows the summary statistics of the variables. To control for the effect of outliers, I winsorize all the variables at the 1% and 99% level[8]. My main regression sample contains 131,683 press releases from 7,503 unique firms. On average, a firm issues about 12.50 press releases in a year, and 3.81 of them will be in the first 10 seconds of an hour. On average 48% of the press releases will be covered by Dow Jones Newswire, and each press release receives an average of 1.75 media articles.

2 Identification strategy

This section introduces the identification strategy of the paper. More specifically, I first show that press releases cluster and over-crowd at specific times within a day. Next, most importantly, I show that within these busy clusters, some press releases will randomly stay on the prominent position of the computer screen ("on-screen" time) that journalists monitor for a longer time than similar releases. I then show that such a longer on-screen time leads to significantly more media coverage from Dow Jones Newswire. Finally I test possible violations of the exclusion restriction requirement that the "on-screen" time is exogenous to the actions of analysts or investors.

2.1 Press release clustering

Using 738,196 press releases from 8,756 firms over April 1, 2006 to December 31, 2017, I show that firms issue most of their press releases in non-trading hours. Moreover, within each hour, press releases heavily cluster at the exact-hour and half-hour points.

[Figure 2 here]

Figure 2(a) shows that the press releases cluster in non-trading hours. To plot the figure, I first split a 24-hour day into 288 5-minute intervals. Then for each firm, I calculate

[8] All of the results in this paper are robust to removing the winsorizations

the percentage of press releases that are published in each 5-minute bin. Finally for each 5-minute interval, I calculate the average percentage of 8,756 firms. Each bar represents the average percentage of press releases published in that 5-minute interval, and the dashed lines represent the 95% confidence intervals of the group means. The first observation from Figure 2(a) is the contrast between trading and non-trading hours. The pre-market period (7-9AM) and the post-market period (4PM) contain the majority of the press releases, while the trading hours contain a much smaller number of press releases.

Moreover, we see that the number of press releases also varies a lot within each hour. For example, while 8AM is a busy hour, the majority of the press releases are issued in the first five minutes (8:00 - 8:04) and the five minutes after the half-hour (8:30 - 8:34). To better see the pattern, I use a darker shade to denote the two five-minute bins after the exact-hour or half-hour points, and we can see that the darker bars stand out in almost every hour. The pattern that press releases cluster at the exact-hour or half-hour points is more obvious in Figure 2(b), which plots the percentage of press releases by the minute of its publication time. On average, over 25% of a firm's press releases are issued in the first minute of an hour, and almost 15% issued in the thirty-first minute. These two minutes collectively consist over 40% of all press releases.

The preference to make announcements at these "integral" points is not hard to understand. When we make appointments, we naturally tend to schedule events at these "integral" points like the exact hour or the half-hour. The same social convention applies when managers determine disclosure times. Think about a management team discussing when to disclose the new earning results. The plan very likely will be "to disclose at 8 o'clock sharp" rather than "let's do 8:03". The benefit of such social convention is that knowing the important releases will happen at the exact- or half-hour points, typical investors or analysts only need to pay close attention during these specific times, thus the convention greatly reduces the "idle" time of waiting for the announcement and frees investors or analysts to perform other tasks in other times.

However, such social conventions could negatively impact wire journalists. Investors and analysts typically monitor the disclosures from only a few firms, thus the possibility of these firms issuing press releases at the same time is low. Even if they do, the amount of releases is easy to manage. However, wire journalists cover the events for the *entire* market, at busy times like 4PM sharp, the number of firms issuing press releases could add up to a large number. As can be seen from Table 1, for press releases issued in the first 10 seconds of 7-9AM or 4PM, they are on average followed by 51 new press releases in the next 30 seconds. In the next section, I show that a shorter on-screen time significantly reduces the amount of media coverage even after controlling for a rigid set of firm and event characteristics. More importantly, the variation of the on-screen time is likely exogenous.

2.2 Exogenous variation of on-screen time

The on-screen time of a press release affects its media coverage because of the unique way that wire journalists produce news from press releases. A typical workflow of wire journalists is to monitor a real-time press release feed, select newsworthy events, and summarize the content and distribute the article to their subscribers. The real-time press release feed is constantly updating, and the amount of new press releases determine how quickly the user interface updates. Therefore, during busy times like 8:00 or 16:00, the computer screen in front of journalists would quickly update as new releases flood in, causing cognitive challenges.

The most important building block of the identification strategy is that during these busy times, some press releases will stay on a prominent position of the user interface longer than similar press releases. Real-time systems, which journalists monitor, follows a ”first-in-first-out” rule: new content always shows up at the top of the user interface, pushing current content down and then onto later pages. The length of time for which a press release stays in a prominent position (e.g. first page or the top of the screen) is determined by the speed of new releases replacing it. Now think of a cluster of press releases all issued at the 4 o'clock sharp. For a press release queued near the beginning of the cluster, its on-screen time is short because the whole cluster following it will quickly replace it. In comparison, a press release queued near the end of the cluster is followed by much fewer releases, thus it gets a longer on-screen time than the releases queued near the beginning. In this paper, I define each cluster as the first 10 seconds of each hour. The queuing of the press releases within such a tight time frame is likely exogenous, as firms neither have incentives nor, in many cases, the ability to precisely control the order within the queue[9]. On top of that, it is the number of press releases from *other* firms that shifts the size of the cluster, thus the ultimate on-screen time is unlikely determined by the actions of any single firm.

As a support for its randomness, I show that on-screen time is uncorrelated with many firm- or event-characteristics covariate balancing tests. More importantly, the same characteristics significantly predict the amount of media coverage from Dow Jones Newswire. To conduct the covariate balancing test, I first construct a sample of busy clusters. Figure 2(b) shows that the first minute of an hour is the busiest, thus I define the busy clusters as the first 10 seconds of each hour. Using a tight time frame to define busy clusters has two additional benefits. First, the press releases in these clusters are similar in nature. Second, the tight time frame greatly reduces the concern for strategic press release timing as discussed in Dellavigna and Pollet (2009) and Michaely, Rubin, and Vedrashko (2016). As a proxy for the on-screen time, I construct NPRA, or the number of new press releases issued in the next 30 seconds. As NPRA of a press release increases, its on-screen time decreases. I then test the following regression

$$\log(\text{NPRA}_{ij} + 1) = \beta X_{ij} + \boldsymbol{\alpha} + \varepsilon_{ij} \tag{1}$$

In the regression, the dependent variable is the log of 1 plus NPRA. The key inde-

[9]For example, the user interface in PR Newswire only allows users to schedule press releases to the minute.

pendent variable, X, represents firm-characteristics including market to book ratio (Q), total asset, and firm age. Solomon and Soltes (2012) show that these characteristics strongly correlated with the amount of media coverage. Further more, I also use the event-characteristics like event-sentiment or title length [10] as the independent variable X. RavenPack adopts sophisticated natural language processing techniques and expert reviews to generate a set of sentiment scores. The goal of these sentiment scores is to create sufficient event summaries and to assist in trading. For example, the press release in which Sanofi announces positive results for a trial study on June 6th, 2015[11], receives an event sentiment score (ESS) of 87. In comparison, the press release in which Micron disclosed decreases in demand receives an ESS score of 17[12]. The regression also controls for a moderate set of fixed effects, including firm-, date-, hour-, and broad topic fixed effects. The regression sample contains 131,683 press releases that are published in the first 10 seconds of each hour over the period of April 2006 to December 2017.

[Table 2 here]

Table 2 shows that none of these characteristics significantly predict NPRA, as can be seen from the insignificant coefficient estimates in the table. Furthermore, in Column (7), I use all the firm- and event-characteristics as independent variables, and the joint F-statistics is only 0.39 with a p-value of 0.886. However, I find that the same set of characteristics significantly predict the amount of media coverage. In Column (8) of Table 2, I use the log of 1 plus the number of media articles on the event day as the dependent variable. In sharp contrast to the previous columns, all the coefficient estimates are significant, and their jointly F-statistics is 147.7. The significant coefficient estimates in Column (8) shows that these variables are highly relevant in terms of the newsworthiness. Yet, none of them is significantly correlated with the proxy of on-screen time, consistent with

[10] While a better measure would be the length of the full press release, RavenPack does not provide any this or any similar measure.

[11] See http://mediaroom.sanofi.com/sanofi-announces-positive-results-for-toujeo-in-phase-iii-study-extension-in-japanese-people-with-uncontrolled-diabetes-2/ for the press release

[12] See http://investors.micron.com/releasedetail.cfm?ReleaseID=440412 for the press release

the conjecture that in these tightly defined busy clusters, the variation of on-screen time is largely random.

In the next section, I show that a shorter on-screen time leads to less media coverage, satisfying the relevance condition to use the on-screen time as an instrument to media coverage. In the final section, I discuss possible violations of the exclusion restriction condition, and find no evidence that reject that the exclusion condition holds.

2.3 On-screen time and media coverage

I test whether on-screen time of a press release affects its media coverage by estimating the following regression.

$$AbnNews_{ijt} = \beta \log(NPRA_{ijt} + 1) + \boldsymbol{\alpha} + \varepsilon_{ijt} \tag{2}$$

In the regression, $AbnNews_{ijt}$ is the abnormal media coverage measure defined in Section 1. The key independent variable is NPRA, which measures the number of new press releases issued immediately after the press release j in the next 30 seconds, and is an inverse proxy for the on-screen time. $\boldsymbol{\alpha}$ represents a set of fixed effects that I am going to include in the regression. As in Table 2, the regression sample contains all the press releases issued in the first 10 seconds of an hour.

[Table 3 here]

I find that a larger NRPA, thus a shorter on-screen time, significantly decreases the amount of media coverage even after controlling for a rigid set of firm- and event-characteristics. Table 3 shows the results. As a benchmark, in Column (1) the regression only includes the date-hour fixed effects to control for the differences across clusters. The coefficient estimate is highly significant and economically large. Doubling NPRA[13], thus the on-screen

[13]The standard deviation of log(NPRA+1) is 1.11, thus doubling NPRA is close to an increase of one standard deviation.

time of a press release decreases, will decrease the amount of abnormal news by 12.8%. I incrementally introduce more fixed effects to control for possible omitted variables at the firm and the press release level. Column (2) further includes firm fixed effects to control for firm-invariant characteristics. Compared with the coefficient estimate in Column (1), the new estimate slightly increases in its magnitude (from -0.128 to -0.132). Column (3) introduces firm-year interacted fixed effects, and the analysis is essentially to compare two press releases issued by the same firm in the same year. The more rigid firm-characteristics controls actually increase the magnitude of the coefficient estimate. The coefficient estimate in Column (3) is -0.163. These results reveal that the endogenous factors at the firm-level likely work against me finding the effect.

To control for the characteristics of different press releases, I utilize the topic measures developed by RavenPack. RavenPack applies sophisticated textual analysis to categorize press releases into different topic groups. These topic measures are essentially the summaries of the underlying events. In Column (4) of Table 3, I further include the fixed effects that control for broad topic classifications (28 unique groups). With the new fixed effects, Column (4) still shows a significant estimate of -0.100. In Column (5), I further control for a more detailed topic classification and include the new topic fixed effects (143 topic groups) in the regressions. The richer fixed effects only slightly decrease the magnitude of the coefficient estimate, which now becomes -0.093 with a t-statistics of -7.99. Overall, these results show that as the on-screen time decreases, or NPRA increases, the firm receives less media coverage.

Such results are robust to news measures, sample selection, and the functional form of the dependent variable. First, I find consistent and significant results when I use different news measures, including using non-duplicated news, using only "flash" news that only includes a title, or using only "full" articles that have at least a paragraph. Second, I find almost identical results when I use different samples to address the potential issues. Section 5 provides a more detailed discussion of these issues. Finally, I find robust results when I change the log measure of media coverage to raw counts or dummy variables as the dependent variables. Table 13 in the Appendix shows the results.

In the next section, I discuss possible violations of the exclusion condition.

2.4 Exclusion restriction requirement

To use the on-screen time as an instrument to media coverage, the exclusion restriction condition requires that the on-screen time does not correlate with the information production of investors and analysts. There are two possible cases that this exclusion restriction condition does not hold. First, some omitted variable might affect both the on-screen time and the information production of investors and analysts. Such a case is unlikely. Previous tests have shown that the on-screen time does not correlate with common observables. Furthermore, the regressions control for very rigid fixed effects that absorb many firm-level and event-level variations. Instead of adding more controls and showing the results hold, in this section, I first perform falsification tests to show that it is indeed the limited cognitive resources of journalists, rather than unobserved factors, that caused the effects.

The second concern is that investors and analysts might be similarly and directly impacted by the on-screen time. Such a concern is also unlikely due to the drastically different event stream that wire journalists and other finance professionals monitor. In the second part of this section, I provide evidence to show that the on-screen time does not correlate with the information exposure at the industry level, which typically applies to investors and particularly analysts. In Section 5, I provide more empirical support by exploiting the fact that journalists and finance professionals have different responses to the press releases from private firms.

Falsification test

In this section, I conduct a series of falsification tests to show that it is indeed the limited cognitive capacity of journalists, instead of unobserved omitted variables, that drives the effects documented in Table 3.

[Table 4 here]

First, in parallel to human journalists, automated algorithms in Dow Jones Newswire will also redistribute some press releases. These algorithms are typically based on editorial judgments about which firms or event types are relevant to the market. The exclusion condition is that NPRA should not correlate with these economic factors, thus we would expect NPRA to have little impacts on such automated coverage. Indeed, Column (1) of Table 4 shows an insignificant estimate of -0.005 when the dependent variable, DJPR, is a dummy that equals to 1 if a press release is covered by the automated coverage. Another natural placebo test is to see whether the media coverage prior to the press release day, which could represent the existing interest about the coming disclosures, changes with NPRA. Column (2) of Table 4 also shows an insignificant estimate of -0.003.

If the cognitive capacity is the cause, then the effect of on-screen time should be stronger in busier times, and disappear during times that are not busy. Consistent with the hypothesis, I find similar effects in the press releases published in the 31st minute of each hour. As shown in Figure 2(b), the 31st minute, which is the half-hour point, also holds clusters of press releases. Column (3) of Table 4 shows a significant estimate of -0.085, indicating a similar effect in the 31st minute as well.

The result shows a drastic change when I estimate the same regression using press releases from all other minutes (excluding the 1st and 31st minutes). Column (4) of Table 4 in fact shows a positive coefficient estimate of 0.051. The result shows that the endogenous factors very likely will work against me finding the effect. When issuing important press releases, firm would spend more efforts to ensure that the press releases are issued on time, thus the press release is more likely to be at the early part of a cluster and has higher NPRA. Such endogenous forces will work against me finding a negative effect of NPRA on media coverage.

I also separate the sample into the busy hours (7-9AM and 4PM) and all non-busy hours (all other hours), based on the pattern in Figure 2(a). Similarly, we would expect the effect to be stronger in the busy hours and insignificant in the non-busy hours. Column (5) and (6) of Table 4 show consistent results. The coefficient estimate is -0.112 in the busy hours, and becomes insignificant in all other hours. Finally, I also sort the sample into quintiles by NPRA, and test how the effect changes with quintiles. Column (7) of Table 4 shows that the effects are insignificant in the first two quintiles (lower NPRA) and become significant and stronger in quintiles that have higher NPRA.

Effect of on-screen time on investors and analyst?

First of all, Table 4 shows that the effect only happens in situations where the amount of new information is so large that the cognitive capacity is stretched to the limit. Compared with wire journalists, investors and analysts follow a smaller set of firms. Almost all the real-time systems allow for information filters, thus investors and analysts can monitor an event stream that is individually customized. As a result, when the amount of information for the whole market is too much to handle, the information exposure to investors and analysts is still easily manageable at the industry level. As supporting evidence, Column (6) of Table 2 shows that the on-screen time is uncorrelated with the total number of press releases issued by firms in the same industry (2-digit SIC[14]).

The second necessary condition for the effect is that the user needs to monitor the event stream in real-time. Wire journalists do not directly control the information inflow because their job is to screen the information as it comes. Investors and analysts, on the other hand, may adopt other information acquisition methods. They could actively search (Da, Engelberg, and Gao, 2011; Ben-Rephael, Da, and Israelsen, 2017) or set up alarms only for the firms or events that they intend to follow. In these modes, the on-screen time of press release becomes an irrelevant factor.

[14]The result is robust to using 1-digit SIC industry, or the text-based industry classification by Hoberg and Phillips (2016)

In Section 5, I provide additional empirical evidence to support the hypothesis that the on-screen time is not directly impacting investors or analysts. The idea for the tests comes from the assumption that wire journalists follow the press releases from private firms while investor and analysts don't. When the on-screen time decreases due to the new press release issued by private firms, the effect should be equally strong if all the effects are indeed through media. I find consistent evidence with the hypothesis, and Section 5 introduces the detailed results.

3 Main effect on information production

In this section, I use the on-screen time, which is inversely measured by NPRA, as an instrument and study the effect of media on investors and analysts. For the majority of the analyses, I use two-stage least square (2SLS) regressions. In the first stage, I estimate Equation 2, and in the second stage, I estimate the following regression

$$\mathrm{Y}_{ijt} = \beta \widehat{\mathrm{AbnNews}}_{ij0} + \boldsymbol{\alpha} + \varepsilon_{ijt} \tag{3}$$

In the regression, the dependent variable will be the information production of investors or analysts on day t, where day 0 is the event day. The key independent variable, $\widehat{AbnNews}_{ij0}$, is the predicted abnormal media coverage on day 0 from the first stage regression. The regression inclues the same set of fixed effects, including firm-year, date-hour, and detailed topic fixed effects. Because Table 4 reveals that the blink effect is only significant in busy hours (7-9AM and 4PM), the following analyses will only use the press releases published in the first 10 seconds of these busy hours to show sharper results, though all my results are robust to using all the hours.

Information acquisition through EDGAR searches

I first show that media coverage increases the amount of information acquisition from investors, which is measured by the abnormal number of of EDGAR requests made by human users, AbnEdgar. In Column (1) of Table 5, I first regress AbnEdgar on AbnNew, and both variables are from the press release publication day (day 0). The coefficient estimate is highly significant. Doubling the amount of media coverage increases the abnormal number of EDGAR requests by 19%. Note the regression controls for the same set of rigid fixed effects that aborb all annual firm measures. In Column (2), I then regress AbnEdgar of day 0 directly on the instrument variable, log(NPRA + 1). The coefficient estimate is -0.034, showing that when NPRA increases, thus the on-screen time of the press release decreases, the amount of EDGAR searches also decreases. In Column (3), I show the second stage regression result. The first stage result is the same as the result reported in Column (5) of Table 4[15], and the F-stat of the instrument variable is 25.5, well above the common threshold of 10. Consistent with the result in Column (2), I find a positive effect of media coverage on the amount of EDGAR searches. The 2SLS estimation shows that doubling the amount of media coverage on the event day leads to 29% more EDGAR searches on the same day. In Column (4), I replace the dependent variable by the cumulative abnormal number of EDGAR searches on days 0 and 1, and find a slightly larger estimate of 0.313. Overall, the results show that media coverage significantly increases the amount of EDGAR searches.

Figure 3(b) plots the coefficient estimate of β in Equation 3 using other event days. I re-estimate the Equation 3 and replace the dependent variable by the abnormal number of EDGAR searches on other event days. The figure provides three interesting observations. First, the coefficient estimate on the day prior to the press release is insignificant and close to 0. Such a result is consistent with the falsification test in Table 4, and further validates

[15]Note the regression in Table 5 Column (1) - (4) has slightly lower number of observations. This is becasue the EDGAR log data stops at June 2017. Thus the exact first stage result slightly differ from the results reported in Column (5) of Table 4.

the exclusion restriction condition. Second, the coefficient estimates on days 1 to 3 are also significant at the 95% level. Because these EDGAR searches happen after the event day, the direction of causation is more definite. Third, by day 5, the coefficient estimate almost completely converted back to 0, consistent with the effect driven by a transient and temporary shock.

Note there are two caveats of using EDGAR requests as measures for investors' information acquisitions. First, EDGAR requests represent a lower bound of the overall information acquisitions. Investors may acquire information from other alternative sources include professional services like Bloomberg or online services like Yahoo! Finance. However, since I am documenting a positive effect of media on EDGAR searches, this means I am also estimating a lower bound of the positive effect. Even so, the magnitudes of the coefficient estimates are still quite large (31% in a two-day period). The second caveat is that not all the EDGAR requests are from investors. To partly address the issue, in Section 4, I try to identify a set of IP addresses that belong to institutional investors, and importantly, their EDGAR searches also exhibit positive responses to media coverage.

Earning forecasts of analysts

I next test how media coverage changes the information production of analysts. In the Columns (5) - (8) of Table 5, the dependent variable is the abnormal number of analysts issuing earning forecasts for the firm. Following similar analyses, I first estimate an OLS regression by regressing AbnAnalyst on AbnNews. The coefficient estimate is both large and highly significant even controlling for the strict set of fixed effects. Doubling the amount of media coverage increases the number of analysts issuing earning forecasts for the firm by 28.8%. In Column (5), I regress AbnAnalyst directly on the instrument variable, and in Column (6), I estiamte the effect using 2SLS regressions. Both columns show consistent results: doubling the amount of media coverage increases the amount of analyst forecasts by 46.9%, and the t-stat for the second-stage estimate is over 6. In Column (7), I find a even larger effect for the cumulative abnormal number of analyst forecasts.

Figure 3(c) plots the coefficient estimate of β in Equation 3 for other event days. First, similar to the result in EDGAR searches, the coefficient estimate is insignificantly different from 0 on the day prior to the press release. Second, we note that the coefficient estimate is even larger on day 1 than the event day. Such pattern is consistent with the pattern in Figure 1. Third, the coefficient estimate quick drops to insignificant since day 2. Indeed, analysts often produce reports that are highly time sensitive, and they constantly make tradeoffs between speed and precision (Beyer et al., 2010).

The large effect of media coverage on stock analysts is surprising, given our conventional wisdom that analysts are among the most informed set of participants in the financial market. It is the job of analysts to closely follow corporate disclosures and provide timely analysis. Moreover, analysts typically follow a small number of firms and have superior information access (e.g. direct communication with management). Thus it is hard to think how media coverage might affect the production function of analysts. However, it is possible that media coverage might affect the demand of earning forecasts. As noted in a survey of 182 analysts from Brown et al. (2015), "Demand from their clients is analysts' most important motivation for making profitable stock recommendations and their second most important motivation for issuing accurate earnings forecasts". To better understand the complementary relationship between media and investors and analysts, I explore possible mechanisms and test their relative strengths in explaining the empirical evidence.

4 Mechanism

This section discusses the possible mechanisms that explain the complementary relationship between media coverage and investors and analysts. I first introduce the possible mechanisms. Next, I show how the effects change with analysts and investors of different information production skills. Then, I show how the effects change in different subsample of firms. Finally, I present the effect of media coverage on the market outcomes to shed light on the underlying mechanism.

4.1 Possible mechanisms

To guide the later analyses, I start by discussing a general framework to categorize the possible mechanisms. Suppose an information producer, being him an investor or analyst, could choose to produce information for one firm from a pool of candidate firms. The information producer optimally chooses the target firm for information production to maximize his expected payoff. Similar to the payoff to the producers of physical goods, the payoff function to this information producer contains two part: the reward part, which is determined by the demand for the information of the chosen firm, and the cost part, which is determined by the producer's own production technology. The empirical result shows that when media covers a specific firm, the information producer is more likely to switch his optimal decision to the media-covered firm. The goal of this section is to understand why.

The effect might be through the cost part by changing the production technology. First, there could exist a *learning mechanism*, where the media coverage provides additional information that lower the cost for the information producer to generate more precision signals. Goldstein and Yang (2015) show a model where there are two fundamentals that affect the security payoff. They show that as the signal of one fundamental become more precise, investors will have incentives to acquire more information about the other fundamental. One possible information that media articles can provide is sentiment. For example, Tetlock (2007) documents the sentiment expressed in a Wall Street Journal column can predict even the aggregate trading next day. If media coverage helps investors to better gauge investor sentiment, then investors may also extend their research about the fundamental value of the firm. However, news articles produced by wire journalists are supposed to be factual and precise. In many cases, the article is simply a headline and contains little additional information to the original announcement, thus the learning mechanism is unlikely in the context of this paper.

In addition to the learning mechanism, media coverage could also impact the production part through an *attention mechanism*. Investors face a fundamental searching problem to pick stocks from thousands of candidate stocks. Thus it is possible that the investors would not have known about the underlying event if not for the media coverage of it. The existing literature in media has shown that the media coverage has a significant impact on investors' attention, particularly the retail investors. Yet, it is again hard to imagine that analysts rely on media to know about the existence of corporate disclosures. After all, analysts are compensated for having an information advantage. It is possible that investors might be subject to inattention, and if it is indeed the case, we would expect the effect to be weaker for less inattention investors.

The effect might alternatively be through the reward part to the information producer's payoff. For the same signal generated by the information producer, media coverage could make the signal more valuable by shifting the demand of it. I call this the *demand-side mechanism*. Existing empirical evidence shows that media coverage increases the volatility of expected returns (Peress, 2014; Blankespoor, deHaan, and Zhu, 2018), changes the investor base (Barber and Odean, 2008), and possibly increases mispricing (Hillert, Jacobs, and Müller, 2014; Ahern and Sosyura, 2015). As a result, sophisticated investors may interpret media coverage as a reliable signal which suggests a higher reward to their information production. Indeed, in the framework from Kacperczyk, Van Nieuwerburgh, and Veldkamp (2016), investors first choose a subset of firms that they will later "learn". One prediction of the model is that more resource is allocated to assets that have high prior payoff variance, and that is exactly the effects of media documented in prior papers. As investors increase their information demand for media-covered events, such demand increase might also shift the payoff to stock analysts, whose clients are mostly institutional investors.

To shed light on the above mechanisms, I conduct three sets of analyses. First, I test how the the results change with the skills of the information producer. Second, I test how the results change with the characteristics of the underlying firm. Finally, I test the effect of media coverage on market outcomes.

It is important to note that these different mechanisms need not to be mutually exclusive. In fact, these mechanisms likely co-exist as the market contains information producers of different level of sophistication. It is also beyond the scope of this paper to reject a mechanism or full attribute the observed the effects to a single mechanism. What I try to achieve here is to evaluate the relative prevalence of the mechanisms and test whether the data supports the prediction of these mechanisms. This line of inquiry would benefit from future research that could better distinguish among these different channels.

4.2 Information producers with different skills

In this section, I test whether the effects of media are different on information producers of different skills. If the effect of media is mainly through changing the production technology, we would expect investors and analysts of different skills react differently. Moreover, because the media articles in the context of this paper are often simple headlines and contain little additional information, we would expect the results to be stronger for lower-skilled or less-resourceful information producers who lack alternative channels and might rely more on the media article to gather information.

However, I find contradicting evidence for the EDGAR searches, as higher-skilled institutional investors exhibit even larger responses to the media coverage. Institutional investors are among the most sophisticated participants in the financial market, and relative to retail investors, they should have suffer less from inattention. To identify searches from institutional investors, I first match IP addresses to known institutions which have an autonomous system number (ASN). The IP-block and ASN link file comes from MaxMind[16]. The link file specifies different ranges of IP addresses that are assigned to different insti-

[16] https://dev.maxmind.com/geoip/geoip2/geolite2/

tutions. Because these institutions are not limited to the finance industry, so the second step is to identify which ones are actually financial institutions. I use two methods to do that. First, I directly search for finance-related words like "bank" or "fund" in the institution's names. Second, I compile a list of names from all 13F institutions, and use a name matching algorithm[17] to identify institutions that are in the universe of 13F institutions. Online Appendix B further describes these two methods in details. Columns (1) and (2) of Table 6 show the regression result for the EDGAR searches from financial institutions. The dependent variables are the cumulative abnormal number of EDGAR requests on days 0 and 1. In Column (1), I identify institutional investors by searching for finance-related words, while in Column (2), I identify institutional investors by matching to the names of 13F institutions. Both columns show significant and larger magnitude (0.374 and 0.385) than the baseline estimates (0.313).

I also find little empirical support for the attention mechanism. I find that investors who already follow the firm on EDGAR systems also significantly increases their EDGAR searches after media coverage. As shown by Chen et al. (2017), investors (mutual fund managers) exhibit very persistent searching activities on the EDGAR system. Moreover, the investors who already searched for a firm's filings in the past suffer less from the search problem associated with inattention (Barber and Odean, 2008) because they already know about the firm. If inattention is the only cause for the effect of media on EDGAR searches, we would expect a weaker effect from these existing EDGAR users. However, I actually find a slightly stronger effect for these investors. In Column (3) of Table 6, I only include human EDGAR searches where the same IP address has accessed any document from the same firm in the previous month. The coefficient is 0.402, larger than the baseline effect of 0.313.

The evidence from analysts also exhibit little support for the production technology channels, as higher-skilled analysts react equally or perhaps stronger to media coverage.

[17] The algorithm is based the code written by Jim Bessen, available at http://goo.gl/m4AdZ.

I first measure analysts' skills by their relative forecast accuracy in the previous year. I calculate the relative accuracy following the procedures in Hong and Kubik (2003) and Ljungqvist et al. (2007)[18]. In Column (4) of Table 6, the dependent variable is the abnormal number of analyst forecasts from analysts whose relative accuracy is above the median accuracy in the previous year. The dependent variables in Columns (4) to (9) are the cumulative abnormal number of analyst forecasts on days 0 and 1. Column (4) shows an coefficient estimate of . In Column (5), I test for the analysts who have below-median accuracy in the previous year, and the estimated effect, , is close to the estimate in Column (4). In Column (6), the dependent variable is the difference of the two dependent variables in Columns (4) and (5), and the coefficient estimate is insignificant. In Columns (7) to (9), I alternatively measure the skills of analysts by their experience with the firm, which is defined as the number of years that the analyst has been covering the same firm in the past five years. I find that analysts who have a longer experience with the firm, as shown in Column (7), reacts even stronger to the media coverage than the less experienced analysts, as shown in Column (8).

4.3 Results by different firm characteristics

In this section, I test whether the effects of media coverage change in different subsample of firms.

I first split the firms into two groups based on their average institutional ownership in the previous year. Then for each subsample, I separately estimate Equation 3 and Table 7 reports the second stage regression results. The dependent variables are all cumulative abnormal measures on days 0 and 1. In Columns (1) and (2), EDGAR searches show a

[18]The procedure exactly follows the footnote 6 in Ljungqvist et al. (2007). For analyst i covering firm k in year t, I first calculate the absolute forecast error using the following steps. (1) get the analysts most recent forecast of year-end EPS issued between Jan. 1 and Jun. 30, (2) calculate the difference with the subsequent realized earnings, (3) scale the difference by previous year-end price. Then for all the analysts covering firm k in year t, I re-scale the absolute forecast errors so that the most and least accurate analysts scores one and zero, respectively. Finally, analyst i's relative forecast accuracy in year t is his/her average score across the the stocks he/she covers over years t-2 to t.

slightly higher and much more significant responses in the subsample with above-median institution ownership. Such a result again lends little support to the attention mechanism. If it is indeed the searching problem that media coverage helps to solve, firms that are already owned by many institutional holders should be ”easily searched”, yet the subsample shows even larger responses to media coverage.

Columns (5) and (6) shows strong support for the demand-side mechanism for the effects of media on analysts. In Column (5), I find that the effect of media (1.01) is much stronger in the subsample where the institutional ownership is higher, compared with Column (6) (0.51) or the baseline model (0.78). Indeed, a higher institutional ownership suggest that the increased information demand is more likely to extend to analysts. In untabulated results, I find consistent results when I use the number of institutional investors, rather than the level of ownership, to split the sample.

I also split the sample into two groups based on the idiosyncratic volatility. In Columns (3) and (4) of Table 7, I find that the effect of media coverage is larger and more significant in the low IVOL subsample (0.395) and the high IVOL (0.275) subsample. The high IVOL firms typically have higher cost of arbitrage, thus investors, especially arbitrageurs, might choose to not participate in trading the stock in the first place. In Comparison, I find the analyst forecasts show similar responses in the two subsamples. Considering high-IVOL and low-IVOL firms might have different difficulty in predicting earnings, the results further suggest the effect of media is unlikely from the production side.

Overall, the results are consistent with the demand-side explanation, particularly for the results with stock analysts. Analysts show a stronger responses to media coverage in firms with larger institutional holdings, consistent with the increased likelihood that institutional clients request information from analysts. Note that while the results for EDGAR requests are in general consistent with the demand-side mechanism, they could not definitely reject the mechanisms in the production side. In the next section, I provide additional evidence from the market outcomes to further shed light on the underlying channels.

4.4 Market outcomes

The rational attention allocation hypothesis in this paper suggests that sophisticated investors focus on media-covered events to obtain higher returns. This section shows that media coverage attracts trade from both possible informed and noise traders. In addition, while informed investors enter the market right on the event day, I find more possibly noise traders on the following two days. Consistent with a tug of war between these two type of investors, I find that while the trading volume significantly increases, the overall price efficiency is not affected. The evidence in this section lends strong support to the demand-side mechanism.

The market outcome variables come from a variety of sources. The trading volume, intraday price range, and daily stock return data comes from CRSP. I calculate the abnormal daily turnover following Tetlock (2010) as the log of 1 plus the turnover minus the log of 1 plus the average daily turnover in the past 60 trading days. The effected spread variables, both equally-weighted and value-weighted, is generated from TAQ data following Goyenko, Holden, and Trzcinka (2009). More specifically, it is calculated as $2|\log(P_k) - \log(M_k)|$, where P_k is the price of the trade, and M_k is the mid-point of the consolidated BBO at the time of the trade. I define abnormal spread following Blankespoor, deHaan, and Zhu (2018) as the effective spread over the average daily effective spread in the past 60 trading days[19]. Following Peress (2014), I define daily range as the log of the daily high price minus the log of the daily low. Finally, I calculate abnormal returns by subtracting the CRSP value-weighted index return from the daily raw returns. Because the sample contains press releases from both pre- and post-market hours, I now define the event day, or day 0, as the first trading day after the press release is issued.

[19]This definition slightly differs from the definition in Blankespoor, deHaan, and Zhu (2018), who use the window [-40, -11] to calculate the average effective spread. I use the window of past 60 trading days for the consistency with other tests. My results are robust to using the window of [-40, -11] to calculate the average daily effective spread.

[Table 8 here]

I start by showing that media coverage significantly increases trading volumes. I estimate Equation 3 with the abnormal trading volume as the dependent variables. Column (1) of Table 8 shows the regression results where the dependent variable is the abnormal turnover on the event day. Doubling the amout of media coverage, the abnormal turnover will increase by 49.2%. The large effect on the trading volume is consistent with a growing literature that find similar effects. Barber and Odean (2008) find that news coverage significantly increases the buying activities of retail investors. Engelberg and Parsons (2011) find that investors trade a stock more when local newspapers cover the firm's earnings. Peress (2014) documents a 12% decrease in trading volume during strike days when newspapers could not be produced or delivered. Blankespoor, deHaan, and Zhu (2018) find that trading volume increases by approximately 11% when the firm is covered by machine-generated news articles. Fedyk (2018) shows that news articles at the front-page induces 280% higher trading volumes than other similar news at less prominent positions in the next 10 minutes. The large effect documented in this paper confirms that my unique setting is highly relevant to the financial market.

As evidence for the participation of informed traders, I find that the effective spread widens on the event day with more media coverage. In Columns (2) and (3), I use the equal-weighted and value-weighted effective spread on the event day as the dependent variable. The coefficient estimates in both columns are significant. Doubling the amount of media coverage will increase the equal-weighted (value-weighted) abnormal effective spread by 23.9% (26.6%), suggesting an increase in the relative proportion of informed trades. The result is consistent with the increased information production of institutional traders documented before. The effect on the effective spread disappears in the next two trading days, as shown in Columns (7) and (8) of Table 8. The coefficient estimate for the average daily abnormal effective spread is insignificant and quantitatively tiny (-0.036 for equal-weighted and -0.014 for value-weighted). Indeed, as analysts earning forecasts increase the amount of public information and more retail traders participate in the trading, the information asymmetry for market makers should decrease.

The evidence also shows that media coverage attracts more price-insensitive traders on days after the press release day. More specifically, I find on days [1, 2], while media coverage does not change the overall level of return, it increases the intra-day trading price range. Column (9) of Table 8 shows regression result where the absolute cumulative abnormal return on days [1, 2] is the dependent variable. The coefficient estimate is only 0.003 and insignificant. In contrast, Column (10), where the dependent variable is the average daily price range on days [1, 2], shows significant coefficient estimate. Such results are similar to the findings in Peress (2014), who finds that the absence of newspaper decreases the intraday price range while the aggregate level of return is unchanged. Peress (2014) attributes similar effects to "less price-sensitive traders who transact at less favorable prices".

Collectively, the results suggest a tug of war between two types of investors, one attracted by media due to attention, and the other that consciously trades more media-covered firms, profiting by trading against these potentially uninformed traders. The lead-lag responses, where the informed trade happens on the event day while less price-sensitive trade happens on the next two days, are possibly caused by the slow information diffusion to retail traders, who might rely on mass media for information (Blankespoor, deHaan, and Zhu, 2018). The pattern also echoes the findings in Ben-Rephael, Da, and Israelsen (2017), who document a similar lead-lag relationship in information searches. It is likely that these two types of traders would have the opposite effects on price efficiency. Consistently, I find the overall price efficiency is unchanged by media coverage. To measure price efficiency, I use the delayed response ratio measure from Dellavigna and Pollet (2009). While the original work of Dellavigna and Pollet (2009) uses a period of 75 days to measure the total cumulative returns, I construct the delayed response ratio measure using a range of periods, as my sample is not restricted to earning announcements, which is what Dellavigna and Pollet (2009) focus on. Across all the measures, as shown in Table 9, none of them show significant results. The results are in line with the findings in Blankespoor, deHaan, and Zhu (2018) who find that news does not improve nor impede the price efficiency.

[Table 9 here]

5 Robustness

This section addresses several concerns that might bias the results of the paper.

5.1 Are the effects through media coverage?

This paper argues that all the changes in EDGAR searches and analyst forecasts are causally impacted by media coverage. Such interpretation might not be valid if variation in the on-screen time can directly impact investors and analysts. As a result, the observed effects are not necessarily through journalists but from a common shock to all. However, this paper argues that the measured on-screen time is unique to wire journalists due to the way they monitor and process information. The measure is uncorrelated with the amount of information exposure at the industry level. In this section, I provide additional empirical support to show that the effects are mostly likely through journalists.

An ideal experiment to test whether the effects are through journalists is to find another shock that only impacts journalists. If the effects are not through journalists, such exogenous shock should cause no changes in analysts and investors. One possible candidate for the shock is the number of press release from private firms. While some investors might also care about news from private firms, it is unlikely that they will monitor the press releases from these private firms in real-time, because such information is generally not tradable. In comparison, wire journalists also cover news events from private firms, and the press release from private firms affect wire journalists in the same way as the releases from public firms do. Therefore, the variation in the number of press releases from private firms is likely a shock that only affects wire journalists.

Given the above assumption, I first test whether the composition of NPRA, that whether the following press releases are from private or public firms, changes the effect on analysts and investors. Suppose the effects are not through journalists, then for a press release, as the proportion of its following press releases that are from private firms increases, the effects on investors and analysts should decrease (as the absolute number of press release from public firms decreases), while the effect on journalists should stay unchanged (as the total number of press releases does not change). I test this hypothesis by including an interaction term, log(NPRA + 1) x % Priv, where % Priv is the percentage of press releases in NPRA that are issued by private firms. The dependent variables come from the main tests before: the abnormal news coverage, the two-day cumulative abnormal number of EDGAR searches, the two-day cumulative abnormal number of analyst forecasts, the abnormal turnover on the event day, the absolute abnormal return on the event day, and the price range on the event day. Panel (A) of Table 10 shows that the coefficient estimate for the interaction term is not significant in all the tests, while the coefficient estimates for log(NPRA+1) are almost identical as in my baseline results.

[Table 10 here]

The insignificant effect could also be due to the power of the test. Therefore I next test whether the variation in the number of press releases from private firms can generate similar effects as documented in my main tests. In Panel (B) of Table 10, I directly use log(NPRPriv + 1), the number of press releases from private firms in the next 30 seconds following the press release, as the instrument variable. The results show that all coefficient estimates for log(NPRPriv+1) are significant, and the magnitude are also similar to the effects in previous tables.

Altogether, the results show that the investors and analysts do not react differently to press releases issued by private firms. The evidence is most consistent with the hypothesis that the effects in this paper are through journalists, meaning interpreting the results as the causal effects of media is valid.

5.2 Pre-scheduled press releases

One limitation of the paper is that I could not distinguish press release that are pre-scheduled from the ones that are manually submitted at the spot. A valid concern is that most pre-scheduled press releases will show up in the first few seconds, thus typically have higher NPRA and lower on-screen time. Should the pre-scheduled press releases be very different than manually submitted releases, then our estimates might be biased. To evaluate such a concern, I rerun the main tests using a sample that exclude all the press releases that are published in the first second of an hour. The logic is that the first second would contain a much larger proportion of pre-scheduled press releases, and if there indeed exists any biases, we would find very different estimates with this new sample. Panel (C) of Table 10 shows the regression results. All the coefficient estimates are very close to the results in the baseline cases, suggesting that the concern about pre-scheduled press releases is unlikely impacting the results.

5.3 Endogenous disclosure

Another limitation of the paper is that I could not address the endogenous incentives for issuing press releases or not, thus there could exist some sample selection bias. However, since this paper uses a very tight time frame to construct the sample and includes a rigid set of fixed effects, it is not clear whether and how the endogenous press release issuance might bias the results. To partially speak to the issue, I focus on the subset of press releases that are likely required by regulation, thus the concern for endogenous disclosure is less. More specifically, I include press releases that are about earnings or are accompanied by a new EDGAR filing on the same day. Panel (D) of Table 10 shows similar results using this subsample.

6 Conclusion

This paper documents that wire journalists may inefficiently select and report corporate events when they reach their cognitive limits. The problem arises because corporate press releases overcrowd at specific times, forcing journalists to process a lot of information quickly. During these busy times, the exogenous variation of how visually salient a press release stays on the journalists' computer screen affects the amount of media coverage.

This paper finds strong support that the salience of press releases is affected by exogenous shocks that most directly impact wire journalists. This unique setting thus allows me to study the causal effect of media coverage on other market participants. I find that media coverage significantly increases the information acquisition of investors, and the effect also exists in investors who are less likely inattentive. The increased information demand from institutional investors also impacts stock analysts, who also issue more earning forecasts for media-covered events. The effect is most prominent for resource-constrained analysts and becomes stronger in firms with higher institutional ownership.

The results of the paper are consistent with the rational attention framework in Kacperczyk, Van Nieuwerburgh, and Veldkamp (2016). Sophisticated investors may rationally allocate their learning to media-covered events, which are empirically associated with higher price volatility and possibly more mispricing. Their increased information demand also affects analysts, especially the resource-constrained analysts who would crowdsource their coverage decision based on the clients' needs. The market outcomes are consistent with the tug of war between these two types of investors. The effective spread increases with media coverage only on the event day. On the next two days, the intra-day price range increases with more media coverage while the overall absolute return is unchanged.

As new data and technology brings more information to the financial market, a division of labor between different information intermediaries, including media, analysts, and investors themselves, could greatly improve the efficiency to process new information. However, as different information intermediaries become more intertwined, the inefficiency from one member could now affect others as well. This paper provides novel evidence that inefficiency in the wire media could also impact other information channels and cause large market reaction. Moreover, I find that the equilibrium forces do not necessarily reverse such exogenous shift in attention. Even those more sophisticated market participants, who suffer less from inattention, might find it more rewarding to chase the media-covered events.

References

Ahern, K. R., and D. Sosyura. 2015. Rumor has it: Sensationalism in financial media. *Rev. Financ. Stud.* 28:2050–93.

Altınkılıç, O., V. S. Balashov, and R. S. Hansen. 2013. Are analysts' forecasts informative to the general public? *Management Science* 59:2550–65.

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

Ben-Rephael, A., Z. Da, and R. D. Israelsen. 2017. It depends on where you search: Institutional investor attention and underreaction to news. *Review of Financial Studies* 30:3009–47.

Beyer, A., D. A. Cohen, T. Z. Lys, and B. R. Walther. 2010. The financial reporting environment: Review of the recent literature. *Journal of Accounting and Economics* 50:296–343.

Blankespoor, E., E. deHaan, and C. Zhu. 2018. Capital market effects of media synthesis and dissemination: evidence from robo-journalism. *Review of Accounting Studies* 23:1–36.

Brown, L. D., A. C. Call, M. B. Clement, and N. Y. Sharp. 2015. Inside the "black box" of Sell-Side financial analysts. *Journal of Accounting Research* 53:1–47.

Chen, H., L. Cohen, U. Gurun, D. Lou, and C. Malloy. 2017. IQ from IP: Simplifying search in portfolio choice. *Working paper* .

Cohen, L., and A. Frazzini. 2008. Economic links and predictable returns. *Journal of Finance* 63:1977–2011.

Da, Z., J. Engelberg, and P. Gao. 2011. In search of attention. *Journal of Finance* 66:1461–99.

Dellavigna, S., and J. M. Pollet. 2009. Investor inattention and friday earnings announcements. *Journal of Finance* 64:709–49.

Drake, M. S., D. T. Roulstone, and J. R. Thornock. 2015. The determinants and consequences of information acquisition via EDGAR. *Contemp Account Res* 32:1128–61.

Engelberg, J. E., and C. A. Parsons. 2011. The causal impact of media in financial markets. *Journal of Finance* 66:67–97.

Fedyk, A. 2018. Front page news: The effect of news positioning on financial markets. *Working paper* .

Goldstein, I., and L. Yang. 2015. Information diversity and complementarities in trading and information acquisition. *Journal of Finance* 70:1723–65.

Goyenko, R. Y., C. W. Holden, and C. A. Trzcinka. 2009. Do liquidity measures measure liquidity? *J. financ. econ.* 92:153–81.

Hillert, A., H. Jacobs, and S. Müller. 2014. Media makes momentum. *Review of Financial Studies* 27:3467–501.

Hoberg, G., and G. Phillips. 2016. Text-Based network industries and endogenous product differentiation. *Journal of Political Economics* 124:1423–65.

Hong, H., and J. D. Kubik. 2003. Analyzing the analysts: Career concerns and biased earnings forecasts. *Journal of Finance* 58:313–51.

Kacperczyk, M., and A. Seru. 2007. Fund manager use of public information: New evidence on managerial skills. *Journal of Finance* 62:485–528.

Kacperczyk, M., S. Van Nieuwerburgh, and L. Veldkamp. 2016. A rational theory of mutual funds' attention allocation. *Econometrica* 84:571–626.

Lee, C. M. C., P. Ma, and C. C. Y. Wang. 2015. Search-based peer firms: Aggregating investor perceptions through internet co-searches. *Journal of Financial Economics* 116:410–31.

Ljungqvist, A., F. Marston, L. T. Starks, K. D. Wei, and H. Yan. 2007. Conflicts of interest in sell-side research and the moderating role of institutional investors. *Journal of Financial Economics* 85:420–56.

Loughran, T., and B. McDonald. 2017. The use of EDGAR filings by investors. *Journal of Behavioral Finance* 18:231–48.

Menzly, L., and O. Ozbas. 2010. Market segmentation and cross-predictability of returns. *Journal of Finance* 65:1555–80.

Michaely, R., A. Rubin, and A. Vedrashko. 2016. Further evidence on the strategic timing of earnings news: Joint analysis of weekdays and times of day. *Journal of Accounting and Economics* 62:24–45.

Neuhierl, A., A. Scherbina, and B. Schlusche. 2013. Market reaction to corporate press releases. *Journal of Financial and Quantitative Analysis* 48:1207–40.

Peress, J. 2014. The media and the diffusion of information in financial markets: Evidence from newspaper strikes. *Journal of Finance* 69:2007–43.

Solomon, D. H., and E. F. Soltes. 2012. Managerial control of business press coverage. *Working paper* .

Tetlock, P. C. 2007. Giving content to investor sentiment: The role ogf media in the stock market. *Journal of Finance* 62:1139–68.

———. 2010. Does public financial news resolve asymmetric information? *Review of Financial Studies* 23:3520–57.

———. 2014. Information transmission in finance. *Annual Review of Financial Economics* 6:365–84.

Figure 1: Information Production after Press Releases

The figure plots the percentage of media articles, analyst estimates, and EDGAR searches that are produced on different days after the most recent press release from the related firm. The sample includes all the news articles from Dow Jones Newswire, analyst estimates from I/B/E/S, and EDGAR searches from 2004 to 2017 (EDGAR data ends at June 2017). The sample only includes firms that are in both of the RavenPack database and CRSP/Compustat Universe. The x-axis in the figure represents the number of days after the most recent press release, and the y-axis represents the percentage of observations.

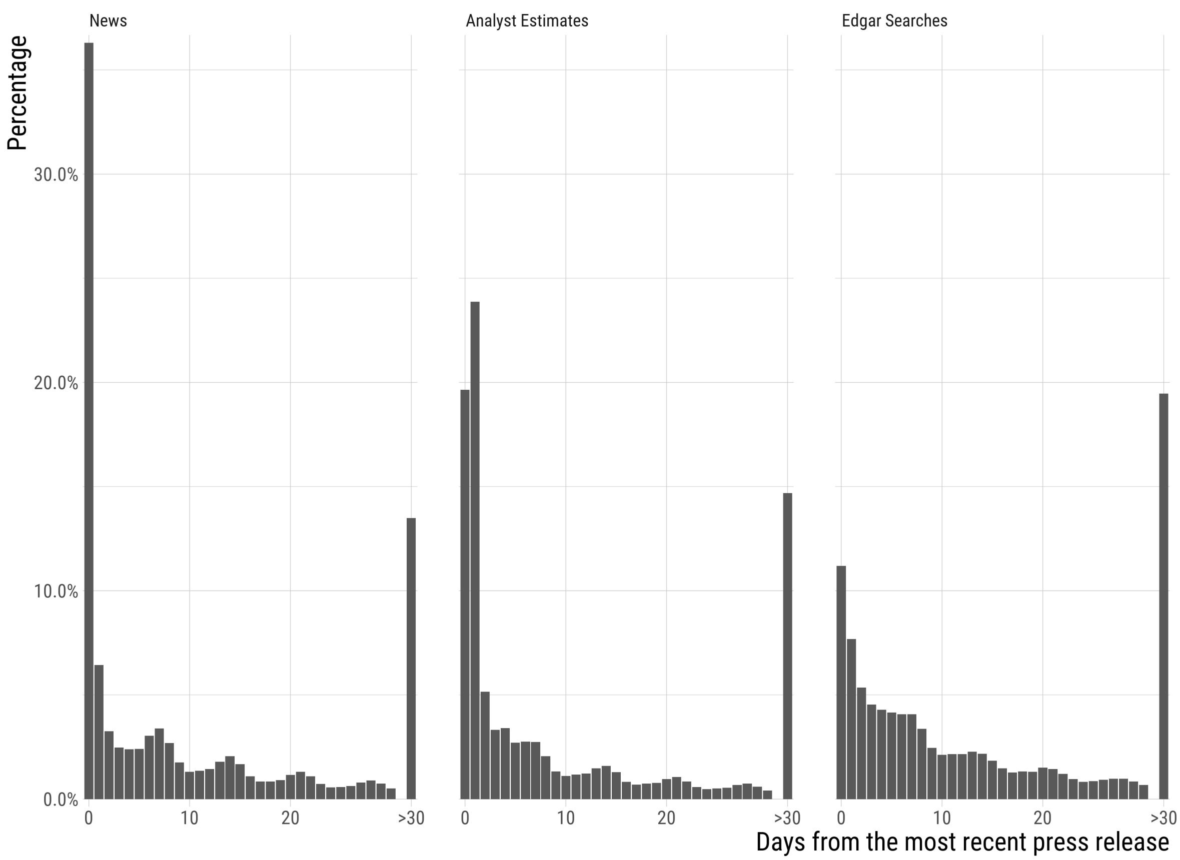

Figure 2: Publication Time of Press Releases

The figures below plot the distribution of press releases publication time within a day or an hour. Panel (a) plots the average percentage of press releases that are published in different 5-minute intervals within a day. I first split a 24-hour day into 288 5-minute intervals. Then for each firm, I calculate the percentage of press releases that are published in each 5-minute bin. Finally for each 5-minute interval, I calculate the average percentage of 8,756 firms. The x-axis denotes the publication time, and the y-axis denotes the percentage. Each bar represents the average percentage of press releases published in that 5-minute interval, and the dashed lines represent the 95% confidence intervals of the group means. For purpose of the exhibition, I only plot the confidence intervals for hours between 6AM to 8PM. The darker bars denote the minutes [0, 5) or [30, 34) of each hour, or the first five minutes after the exact hour or half hour point. To show the distribution of publication time within an hour, Panel (b) plots the average percentage of press releases that are published in each minute across firms. The calculation method is similar as in Panel (a). I first calculate the percentage of press release published in each minute for each firm, and then calculate the average of 8,756 firms. The bars denote the average percentage of press releases of all firms, and the error bars represent the 95% confidence intervals of the group means.

(a) distribution within a day

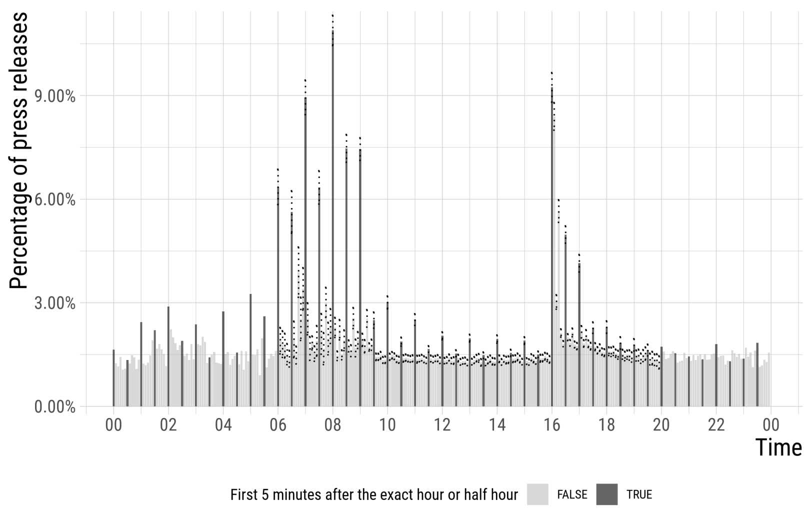

(b) distribution within an hour

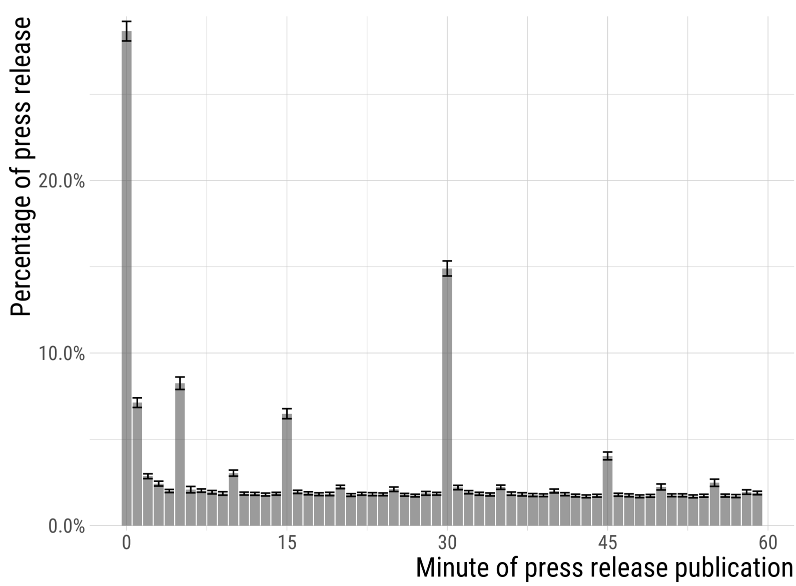

Figure 3: Coefficient Estimates for Different Event Days

This figure plots the coefficient estimates for β in the following two regressions using different event days:

$$AbnNews_{ijt} = \beta \log(NPRA_j + 1) + \boldsymbol{\alpha} + \varepsilon_{ijt}$$
$$AbnEdgar_{ijt} | AbnAnalyst_{ijt} = \beta \widehat{AbnNews}_{ij0} + \boldsymbol{\alpha} + \varepsilon_{ijt}$$

The dependent variables, AbnNews, AbnEdgar, and AbnAnalyst, measure the abnormal number of media coverage, EDGAR requests, and analysts issuing earning forecasts, respectively. Detailed variable definitions can be found in Table 11. The subscripts i denote the firm, j denote the press release, and t denote the event time. For the press release j, NPRA measures the number of following press releases that are published in the next 30 seconds. $\widehat{AbnNews}_{ij0}$ is the predicted value of the abnormal media coverage on the event day (day 0) using the first equation. $\boldsymbol{\alpha}$ includes the firm-year fixed effects, hour-date fixed effects, and detailed press release topic fixed effects. The regression sample contains all the press releases that are published in the first 10 seconds in 7-9AM and 4PM. Figure (a) plots the β estimate for the first equation, where AbnNews is the dependent variable. Figure (b) and (c) plot the β estimates from the second equation, where the dependent variables are AbnEdgar and AbnAnalyst, respectively. The triangle points denote the coefficient estimate, and the error bars represent the 95% confidence intervals. The x-axis denotes the event time t, where $t = 0$ is the press release publication day. All the standard errors are clustered by firm and date.

(a) AbnNews (b) AbnEdgar (c) AbnAnalyst

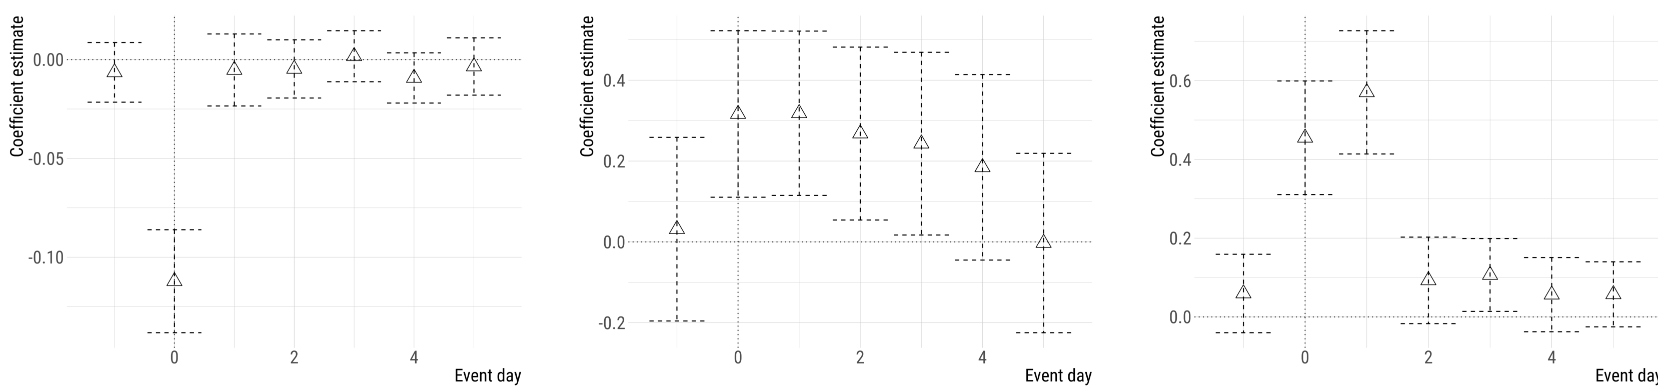

Figure 4: Coefficient Estimates for Cumulative Effects

This figure plots the coefficient estimates for β in the following two regression using different cumulative period

$$AbnNews_{ij}^{-1\to t} = \beta \log(NPRA_j + 1) + \boldsymbol{\alpha} + \varepsilon_{ijt}$$
$$AbnEdgar_{ij}^{-1\to t} | AbnAnalyst_{ij}^{-1\to t} = \beta \widehat{AbnNews}_{ij0} + \boldsymbol{\alpha} + \varepsilon_{ijt}$$

The dependent variables, AbnNews, AbnEdgar, and AbnAnalyst, measure the cumulative abnormal number of media coverage, EDGAR requests, and analysts issuing earning forecasts from day -1 to day t, respectively. Detailed variable definitions can be found in Table 11. The subscripts i denote the firm, j denote the press release, and t denote the event time. For the press release j, NPRA measures the number of following press releases that are published in the next 30 seconds. $\widehat{AbnNews}_{ij0}$ is the predicted value of the abnormal media coverage on the event day (day 0). $\boldsymbol{\alpha}$ includes the firm-year fixed effects, hour-date fixed effects, and detailed press release topic fixed effects. The regression sample contains all the press releases that are published in the first 10 seconds in 7-9AM and 4PM. Figure (a) plots the β estimate for the first equation, where cumulative AbnNews is the dependent variable. Figure (b) and (c) plot the β estimates from the second equation, where the dependent variables are cumulative AbnEdgar and cumulative AbnAnalyst, respectively. The triangle points denote the coefficient estimate, and the error bars represent the 95% confidence intervals. The x-axis denotes the event time t, where $t = 0$ is the press release publication day. All the standard errors are clustered by firm and date.

(a) AbnNews

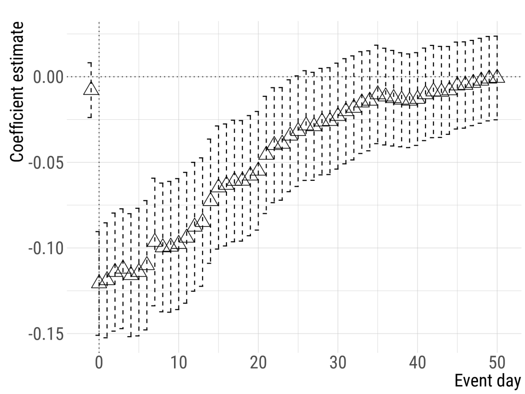

(b) AbnEdgar

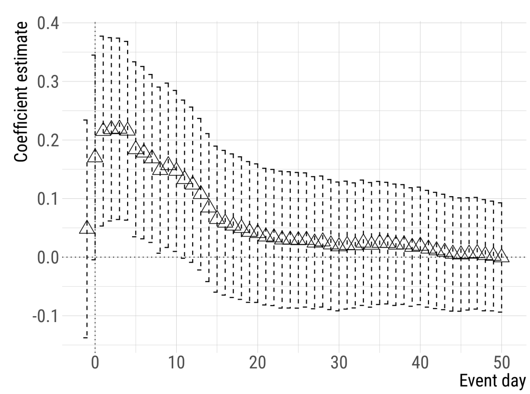

(c) AbnAnalyst

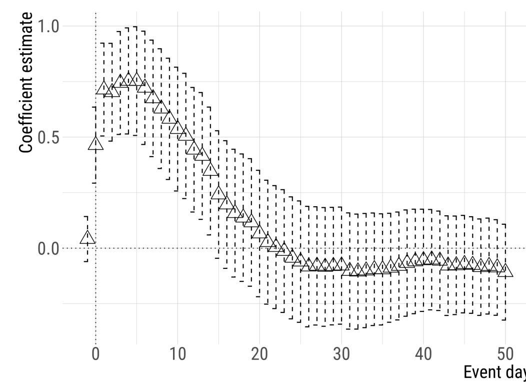

Table 1: Summary statistics

	N	Mean	Std Dev	10th	50th	90th
Panel A: Press release issuance						
Full sample						
# of PRs	738,196					
# of unique firms	8,756					
# of PRs per firm-year	59,035	12.50	9.08	3.00	11.00	23.00
# of PRs per day	2,958	242.00	113.16	113.70	226.00	400.00
Issued in the first 10 seconds of each hour						
# of PRs	131,683					
# of unique firms	7,503					
# of PRs per firm-year	34,532	3.81	3.97	1.00	2.00	8.00
# of PRs per day	2,958	43.04	21.97	19.00	39.00	72.00
Issued in the first 10 seconds of 7-9AM and 4PM						
# of PRs	80,246					
# of unique firms	6,560					
# of PRs per firm-year	24,766	3.24	3.43	1.00	2.00	7.00
# of PRs per day	2,949	26.35	15.93	10.00	23.00	48.00
Panel B: Sample summary statistics						
Issued in the first 10 seconds of each hour						
Age	131376	20.54	16.06	5.00	15.00	48.00
Q	125557	2.05	1.52	0.97	1.52	3.81
AT	130417	21006.78	86222.38	56.95	1093.07	32476.00
ESS	80246	55.35	7.90	50.00	50.00	69.00
NWord	131683	11.56	4.30	7.00	11.00	17.00
NPRA	131683	36.49	31.95	4.00	27.00	84.00
log(NPRA + 1)	131683	3.17	1.11	1.61	3.33	4.44
News	131683	1.75	2.96	0.00	0.00	5.00
AbnNews	131683	0.34	0.82	-0.46	-0.03	1.59
News_Dummy	131683	0.48	0.50	0.00	0.00	1.00
DJPR	131683	0.85	0.36	0.00	1.00	1.00
Issued in the first 10 seconds of 7-9AM and 4PM						
NPRA	80246	51.07	32.24	13.00	46.00	96.00
log(NPRA + 1)	80246	3.71	0.79	2.64	3.85	4.57
AbnNews	80246	0.41	0.83	-0.43	0.00	1.67
AbnEdgar_Human	75673	0.35	0.75	-0.34	0.03	1.41
AbnEdgar_Exist	75673	-0.48	0.99	-1.83	-0.16	0.51
AbnEdgar_Ins	75673	-0.85	1.22	-2.53	-0.64	0.20
AbnEdgar_13f	75673	-0.69	1.20	-2.35	-0.28	0.48
AbnAnalyst	80246	0.18	0.62	-0.30	-0.03	1.20

Continued on next page

Table 1 – *Continued from previous page*

	N	Mean	Std Dev	10th	50th	90th
AbnAnalyst_MoreAccu	80246	0.04	0.59	-0.53	-0.06	0.94
AbnAnalyst_LessAccu	80246	0.01	0.56	-0.57	-0.06	0.89
AbnAnalyst_MoreExp	80246	-				
AbnAnalyst_LessExp	80246	-				
AbnTurnover	80117	0.02	0.80	-0.79	0.00	0.96
AbnSpread_EW	41832	1.11	11.48	0.64	0.97	1.54
AbnSpread_VW	41832	1.04	0.51	0.58	0.92	1.63
|CAR|	80142	2.84	5.35	0.21	1.36	6.65
Range	80142	5.09	5.62	1.33	3.53	10.35
% Priv	80246	0.04	0.04	0.00	0.03	0.09
log(NPRPriv + 1)	80246	0.83	0.71	0.00	0.69	1.79

Table 2: Covariate Balancing Test of The On-screen Time

This table shows that the observable firm and event characteristics do not correlate with NPRA, which measures the number of new press releases published in the top 4 press release wires after the press release j in the next 30 seconds. The dependent variable for Columns (1)-(7) is log of 1 plus NPRA, and for Column (8) is the log of 1 plus the number of media articles on the Dow Jones Newswire covering the firm on the event day. The regression sample includes all the press releases published in the first 10 seconds of each hour. Standard errors are clustered by firm and day, and t-statistics are reported in the parentheses. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level, respectively. Detailed variable definitions can be found in Table 11.

	log(NPRA + 1)							News
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Q	0.001						0.001	0.009***
	(0.409)						(0.299)	(3.503)
log(AT)		−0.000					−0.001	0.014*
		(−0.065)					(−0.151)	(1.899)
log(Age + 1)			0.002				−0.013	0.039*
			(0.091)				(−0.524)	(1.774)
log(ESS)				−0.025			−0.022	0.496***
				(−1.403)			(−1.201)	(17.274)
NWords					−0.000		−0.000	−0.019***
					(−0.191)		(−0.276)	(−22.174)
log(NPRInd + 1)						−0.005	−0.005	0.027***
						(−0.658)	(−0.620)	(3.295)
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Date FE	Y	Y	Y	Y	Y	Y	Y	Y
Hour FE	Y	Y	Y	Y	Y	Y	Y	Y
Broad Topic FE	Y	Y	Y	Y	Y	Y	Y	Y
F-stat							0.3906	147.7
p-value							0.886	<0.001
Observations	127,844	130,407	130,407	131,683	131,683	131,683	127,844	127,844
Adjusted R^2	0.778	0.778	0.778	0.777	0.777	0.777	0.778	0.464

Table 3: Press Releases' On-screen Time and Media Coverage

This table shows that if a press release stays on the screen for shorter time due to new press releases replacing it, the issuing firm will receive less media coverage. I estimate the following regression

$$AbnNews_{ijt} = \beta \log(NPRA_j + 1) + \boldsymbol{\alpha} + \varepsilon_{ijt}$$

The dependent variable is the abnormal media coverage, defined as log of 1 plus the number of news articles that cover the press release issuing-firm on the issuing day, minus log of 1 plus the average number of news articles per day for the firm in the past 60 days, skipping 10 days. The independent variable, NPRA, measures the number of new press releases published in the top 4 press release wires after the press release j in the next 30 seconds. This is the proxy for the on-screen time of each press release. The regression controls for the fixed effects including: date-hour, firm, firm-year, broad topic, and/or detailed topic. The regression sample contains all the press releases published in the first 10 seconds of an hour. Standard errors are clustered by firm and day, and t-statistics are reported in the parentheses. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level, respectively. Detailed variable definitions can be found in Table 11 in the Appendix.

	AbnNews				
	(1)	(2)	(3)	(4)	(5)
log(NPRA + 1)	-0.128^{***}	-0.132^{***}	-0.163^{***}	-0.100^{***}	-0.093^{***}
	(-12.252)	(-12.877)	(-11.263)	(-8.332)	(-7.990)
Date-Hour FE	Y	Y	Y	Y	Y
Firm FE		Y			
Firm-Year FE			Y	Y	Y
Broad Topic FE				Y	
Detailed Topic FE					Y
Observations	131,683	131,683	131,683	131,683	131,683
Adjusted R^2	0.174	0.303	0.327	0.473	0.491

Table 4: Falsification tests

This table shows the falsification tests of the effect of on-screen time on media coverage. The dependent variable in Column (1) is a dummy variable that equals to 1 if the automated algorithm from Dow Jones Newswire republishes the press release. The dependent variable in Columns (2) is the abnormal media measure, AbnNews, on the day *before* the event (press release publication) day, and the dependent variables in Columns (3) - (7) are the same abnormal media measure on the event day. Columns (3) and (4) show regression results using press releases that are published in the first 10 seconds of minutes other than the first minute. Column (3) includes all the press releases published in the first 10 seconds of the 31st minute in each hour. Column (4) includes all the press published in the first 10 seconds of all the minutes except for the 1st and the 31st minutes. Column (5) and (6) show regression results using press releases that are published in different hours. Column (5) includes all the press releases published in the first 10 seconds in 7-9AM and 4PM, and Column (6) includes all the press releases published in the first 10 seconds of all other hours. In Column (7), I first sort the press releases by NPRA into quintiles, and Q2 - Q5 in the independent variables are quintile dummies, where Q5 represents the highest-NPRA quintile. Standard errors are clustered by firm and day, and t-statistics are reported in the parentheses. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level, respectively. Detailed variable definitions can be found in Table 11.

	DJPR	AbnNews					
		t-1	31st minute	other minutes	7-9AM & 4PM	Other Hours	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log(NPRA+1)	−0.006	−0.006	−0.081***	0.047***	−0.104***	−0.047	−0.021
	(−0.774)	(−0.910)	(−3.867)	(17.207)	(−8.523)	(−0.841)	(−0.889)
log(NPRA+1) x Q2							−0.034
							(−1.168)
log(NPRA+1) x Q3							−0.062**
							(−2.200)
log(NPRA+1) x Q4							−0.097***
							(−3.340)
log(NPRA+1) x Q5							−0.101***
							(−3.113)
Date-Hour FE	Y	Y	Y	Y	Y	Y	Y
Firm-Year FE	Y	Y	Y	Y	Y	Y	Y
Detailed Topic FE	Y	Y	Y	Y	Y	Y	Y
Second FE	Y	Y	Y	Y	Y	Y	Y
Observations	131,683	131,683	69,185	252,532	80,246	51,437	131,683
Adjusted R^2	0.284	0.084	0.455	0.488	0.539	0.105	0.496

Table 5: Media Coverage and Information Production

This table shows that media coverage increases the number of EDGAR web requests and the number of analysts issuing earning forecasts. The dependent variables in Columns (1)-(4) are the abnormal number of EDGAR requests, defined as the log of 1 plus the number of requests made on a day minus the log of 1 plus the average number of EDGAR requests per day in the past 60 days, skipping 10 days. I exclude EDGAR searches which are on the index page, have a server code above 300, or from IP addresses that have made more than 5 requests in a minute or 1000 requests in a day. The dependent variables in Columns (5)-(8) are the abnormal number of earning forecasts, defined as log of 1 plus the number of analysts who issued any forecasts for the firm on a day minus the log of 1 plus the average of daily number of analysts issuing forecasts for the firm in the past 60 days, skipping 10 days. The dependent variables in Columns (1) - (3) and (5) - (7) are measures on day 0, while the dependent variables in Column (4) and (8) are cumulative measures on days 0 to 1, where the press release day is day 0. Columns (1) and (5) show OLS regressions where the abnormal media coverage is the independent variable. Columns (2) and (6) show OLS regressions where the instrument, log(NPRA+1), is the independent variable. Columns (3), (4), (7), and (8) show the second-stage results of two-stage least square regressions, where the first-stage is shown in Column (5) of Table 4. The regression sample includes all the press releases published in the first 10 seconds of 7-9AM and 4PM. Standard errors are clustered by firm and day, and t-statistics are reported in the parentheses. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level, respectively. Detailed variable definitions can be found in Table 11.

	AbnEdgar				AbnAnalyst			
	OLS	IV	2SLS		OLS	IV	2SLS	
			day 0	days 0-1			day 0	days 0-1
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AbnNews	0.195***				0.288***			
	(22.370)				(29.663)			
log(NPRA+1)		−0.034***				−0.053***		
		(−2.821)				(−5.486)		
$\widehat{\text{AbnNews}}$			0.290***	0.313***			0.469***	0.784***
			(2.870)	(3.374)			(6.136)	(7.609)
Date-hour FE	Y	Y	Y	Y	Y	Y	Y	Y
Firm-Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Detailed Topic FE	Y	Y	Y	Y	Y	Y	Y	Y
Observations	75,667	75,667	75,667	75,667	80,246	80,246	80,246	80,246
Adjusted R^2	0.300	0.272	0.294	0.328	0.479	0.398	0.446	0.495

Table 6: The Effect of Media Coverage and Information Producer Characteristics

This table shows the effect of media coverage on the information producers of different characteristics. The table reports the second-stage regression coefficients from Equation 3, where the first-stage regression result is shown in the Column (5) of Table 4. The dependent variables in Columns (1) and (3) are the cumulative abnormal number of EDGAR requests on days 0 and 1, and the dependent variables in Columns (4) and (9) are the cumulative abnormal number of analysts issuing earning forecasts on days 0 and 1. Columns (1) and (2) include EDGAR requests that are from financial institutions. I identify financial institutions by matching the IP addresses that belong to known institutions registered with the autonomous system (ASN). Then from the names of the institutions, provided by the ASN registration list, I search for finance-related words in their names (Column 1) or match the name to 13F institutions (Column 2) to identify financial institutions. Column (3) only includes human web requests are from IP addresses that have searched for the same firm in the previous month. For more details about the EDGAR log data and the institution matching process, please refer to the Online Appendix B. Columns (4)-(9) show the effect of media on analysts of different characteristics. In Columns (4) - (6), I group analysts by their accuracy, defined as the average relative accuracy in the previous three years. Column (4) only includes analysts who have an above-median accuracy, and Column (5) includes analysts who have below-median accuracy. Column (6) uses the difference of the dependent variables in Columns (4) and (5) as the dependent variable. In Columns (7) - (9), I sort analyst-firm pairs into two groups based on the experience, which is defined as the number of years that the analyst has been covering the same firm in the past five years. Column (7) includes analysts who have an above-median experience in a year, and Column (8) includes analysts who have an below-median experience. Column (9) uses the difference of the dependent variables in Columns (7) and (8) as the dependent variable. The regression sample includes all the press releases published in the first 10 seconds of 7-9AM and 4PM. Standard errors are clustered by firm and day, and t-statistics are reported in the parentheses. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level, respectively. Detailed variable definitions can be found in Table 11.

	AbnEdgar			AbnAnalyst					
	Ins	13F	Existing follower	Accuracy			Experience		
				>50th	<50th	Diff	>50th	<50th	Diff
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\widehat{\text{AbnNews}}$	0.374**	0.385**	0.402***	0.398***	0.344***	0.060	0.653***	0.439***	0.222***
	(2.496)	(2.503)	(2.920)	(5.179)	(5.215)	(0.915)	(7.161)	(6.532)	(2.837)
Date-hour FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Firm-Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Detailed Topic FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	75,667	75,667	75,667	80,246	80,246	80,246	80,246	80,246	80,246
Adjusted R^2	0.242	0.608	0.516	0.320	0.369	0.097	0.452	0.383	0.224

Table 7: The Effect of Media Coverage and Firm Characteristics

This table shows how the effect of media coverage change with firm characteristics. The table reports the second-stage regression coefficients from Equation 3. The dependent variables in Columns (1)-(4) are the cumulative abnormal number of EDGAR requests on days 0 and 1, and the dependent variables in Columns (5)-(8) are the cumulative abnormal number of analysts issuing earning forecasts on days 0 and 1. In Columns (1), (2), (5), and (6), I first sort firms into two groups based on their average institutional ownership in the previous year. Columns (1) and (5) include observations where the institutional ownership is above median in the previous year, while Columns (2) and (6) use the below-median sample. Similarly, in Columns (3), (4), (7), and (8), I first sort firms into two groups based on their average monthly idiosyncratic volatility in the last three month. Columns (3) and (7) include observations where the IVOL is above the median measure using the same last three months, while Columns (4) and (8) use the below-median sample. The regression sample includes all the press releases published in the first 10 seconds of 7-9AM and 4PM. Standard errors are clustered by firm and day, and t-statistics are reported in the parentheses. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level, respectively. Detailed variable definitions can be found in Table 11.

	AbnEdgar				AbnAnalyst			
	Institution Holding		IVOL		Institution Holding		IVOL	
	$>$ 50th	$<$ 50th	$>$ 50th	$<$ 50th	$>$ 50th	$<$ 50th	$>$ 50th	$<$ 50th
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\widehat{\text{AbnNews}}$	0.395***	0.325	0.275	0.395**	1.009***	0.506***	0.920***	0.927***
	(2.707)	(1.405)	(1.263)	(2.091)	(4.971)	(2.993)	(4.269)	(3.793)
Date-hour FE	Y	Y	Y	Y	Y	Y	Y	Y
Firm-Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Detailed Topic FE	Y	Y	Y	Y	Y	Y	Y	Y
Observations	36,220	36,223	30,420	30,429	38,258	38,304	30,420	30,429
Adjusted R^2	0.340	0.286	0.300	0.337	0.486	0.476	0.378	0.450

Table 8: Media coverage and market reaction

This table shows that media coverage significantly affects trading volume, announcement returns, effective spread, and price ranges. The table reports the second-stage regression coefficients from Equation 3. The dependent variables in Columns (1) - (5) are from the first trading day after the press releases, and the dependent variables in Columns (6) - (10) are average measures (except for CAR) or cumulative measures (CAR) in days 1-2, the second and third trading days after the press release. The dependent variables for Columns (1) and (6) are abnormal turnover, defined as log of 1 plus the turnover minus the log of 1 plus the average daily turnover in the previous 60 trading days. The dependent variables in Columns (2), (3), (7) and (8) are abnormal effective spread measures, with equal-weighted measures in Columns (2) and (7) and value-weighted measures by dollar amount in Columns (3) and (8). The dependent variables in Columns (4) and (9) are the absolute value of cumulative abnormal returns, where abnormal returns are calculated by subtracting the CRSP value-weighted index return on the same day. The dependent variables in Columns (5) and (10) are the price ranges, defined as the log of the daily highest price minus the log of the daily lowest prices. The regression sample includes all the press releases published in the first 10 seconds of 7-9AM and 4PM. Standard errors are clustered by firm and day, and t-statistics are reported in the parentheses. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level, respectively. Detailed variable definitions can be found in Table 11.

	Day 0					Days 1-2				
	AbnTurnover	AbnSpread		|CAR|	Range	AbnTurnover	AbnSpread		|CAR|	Range
		EW	VW				EW	VW		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\widehat{\text{AbnNews}}$	0.492***	0.239***	0.266***	0.021***	0.023***	0.222*	0.036	−0.014	0.003	0.007**
	(3.519)	(3.006)	(2.636)	(4.349)	(4.353)	(1.932)	(0.544)	(−0.171)	(0.699)	(2.063)
Firm-Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Date-hour FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Second FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Detailed Topic FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	80,142	41,798	41,798	80,142	80,142	80,173	41,924	41,924	80,160	80,173
Adjusted R^2	0.372	0.197	0.153	0.310	0.517	0.352	0.245	0.203	0.269	0.609

Table 9: Media coverage and delayed response ratio

This table shows that media coverage has insignificant effects on the price efficiency. The table reports the second-stage regression coefficients from Equation 3. The dependent variables are delayed response ratios, as used in Dellavigna and Pollet (2009). The delayed response ratio over days [0, X] is calculated as $R^{(2,X)}/R^{(0,X)}$, where $R^{(2,X)}$ is the cumulative abnormal returns over days [2, X], and $R^{(0,X)}$ is the cumulative abnormal returns over the period [0, X]. Columns (1) - (6) show results with different period length. The regression sample includes all the press releases published in the first 10 seconds of 7-9AM and 4PM. Standard errors are clustered by firm and day, and t-statistics are reported in the parentheses. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level, respectively. Detailed variable definitions can be found in Table 11.

	Delayed response ratio					
	[2,5]	[2,15]	[2,30]	[2,45]	[2,60]	[2,75]
	(1)	(2)	(3)	(4)	(5)	(6)
$\widehat{\text{AbnNews}}$	−0.097	0.032	1.958	−0.346	−0.329	−0.452
	(−0.158)	(0.061)	(0.429)	(−0.979)	(−1.140)	(−1.614)
Firm-Year FE	Y	Y	Y	Y	Y	Y
Date-hour FE	Y	Y	Y	Y	Y	Y
Detailed Topic FE	Y	Y	Y	Y	Y	Y
Observations	80,127	80,127	80,127	80,127	80,127	80,127
Adjusted R^2	0.019	−0.004	0.135	0.009	0.005	−0.023

Table 10: Robustness tests

This table shows that investors and analysts unlikely experience similar blink effect. I estimate similar regressions as in previous tables. Both Panel (A) and (B) use all the press releases published in the first 10 seconds of 7-9AM and 4PM. Column (1) uses the abnormal media coverage on day 0 as the dependent variable. Column (2) uses the cumulative abnormal number of EDGAR searches on days 0 and 1 as the dependent variable. Column (3) uses the cumulative abnormal number of analyst forecasts on days 0 and 1 as the dependent variable. Column (4) uses the abnormal turnover on the event day as the dependent variable. Column (5) uses the abnormal spread on the event day as the dependent variable. Column (6) uses the absolute abnormal return on the event day as the dependent variable. Column (7) uses the intraday price range on the event day as the dependent variable. In Panel (A), I include an interaction term, log(NPRA + 1) x % Priv, where % Priv is the percentage of press releases in NPRA that are issued by private firms. In Panel (B), I directly regress the dependent variables on log(NPRPriv + 1), where NPRPriv is the number of press releases issued by private firms in the next 30 seconds. I also control for log(NPRInd + 1), where NPRInd is the number of press releases issued by firms in the same 2-digit SIC industry on the same day. In Panel (C), the regression sample includes all the press releases issued in the first 10 seconds, excluding the first second, of 7-9AM and 4PM. In Panel (D), the regression sample only includes press releases that are on earnings or are accompanied by an EDGAR filing on the same day. Standard errors are clustered by firm and day, and t-statistics are reported in the parentheses. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level, respectively. Detailed variable definitions can be found in Table 11.

	AbnNews	AbnEdgar	AbnAnalyst	AbnTurnover	AbnSpread
	(1)	(2)	(3)	(4)	(5)
Panel A: Interaction with the ratio of press releases from private firms					
log(NPRA+1)	−0.127***	−0.045***	−0.100***	−0.065***	−0.038***
	(−8.482)	(−3.792)	(−7.138)	(−3.537)	(−3.175)
log(NPRA+1) x % Priv	−0.007	0.019	−0.016	0.004	0.008
	(−0.286)	(0.998)	(−0.767)	(0.138)	(0.305)
% Priv	−0.012	−0.036	0.004	−0.032	−0.041
	(−0.284)	(−1.287)	(0.133)	(−0.666)	(−0.863)
FE	Y	Y	Y	Y	Y
Observations	80,108	75,536	80,108	80,004	41,760
Adjusted R^2	0.540	0.316	0.465	0.356	0.241
Panel B: direct test using press releases from private firms					
log(NPRA+1)	−0.126***				
	(−9.170)				
$\widehat{\text{AbnNews}}$		0.269***	0.867***	0.569***	0.249**
		(2.707)	(7.658)	(3.673)	(2.456)
FE	Y	Y	Y	Y	Y
Observations	80,246	75,667	80,246	80,142	41,798
Adjusted R^2	0.539	0.335	0.462	0.365	0.191
Panel C: Drop the 1st second					
log(NPRA+1)	−0.120***				
	(−6.944)				
$\widehat{\text{AbnNews}}$		0.355***	0.787***	0.552***	0.242**
		(3.095)	(5.764)	(2.895)	(2.281)
FE	Y	Y	Y	Y	Y
Observations	62,809	59,420	62,809	62,725	33,879
Adjusted R^2	0.548	0.301	0.504	0.376	0.187
Panel D: Required disclosures					
log(NPRA+1)	−0.159***				
	(−3.992)				
$\widehat{\text{AbnNews}}$		0.445**	0.808***	0.716**	0.329
		(2.205)	(3.821)	(2.403)	(1.167)
FE	Y	Y	Y	Y	Y
Observations	33,780	31,949	33,780	33,758	17,812
Adjusted R^2	0.527	0.347	0.661	0.314	−0.058

Appendix A. Variable Definition

Table 11: Variable definitions

Variable	Definition	Source
NPRA_j	The number of new press releases published immediately after the press release j in the next 30 seconds. I only include press releases published on the top 4 press release wires, namely, PRNewswire, BusinessWire, GlobeNewswire, Marketwired.	RavenPack PR Edition
News_{it}	The total number of news articles covering firm i on day t. The relevance score needs to be 100, meaning that firm i is the main subject of the news article.	RavenPack DJ Edition
AbnNews_{it}	Abnormal news coverage. The number is calculated by subtracting the average daily number of media articles covering firm i in the previous 60 days from News_{it}	RavenPack DJ Edition
$\text{AbnNews\_Novel}_{it}$	From AbnNews, I only include news articles whose novelty score from RavenPack (`ESS`) is 100	RavenPack DJ Edition
$\text{AbnNews\_Flash}_{it}$	From AbnNews, I only include news articles whose news type in RavenPack is `NEWS-FLASH`	RavenPack DJ Edition
AbnNews_Full_{it}	From AbnNews, I only include news articles whose news type in RavenPack is `FULL-ARTICLE`	RavenPack DJ Edition
News_Dummy_{it}	A dummy variable that equals to 1 if there is any media coverage on that day	RavenPack DJ Edition
AbnEdgar_{it}	Abnormal number of EDGAR requests about firm i on day t. I exclude requests where `idx` equals 1 (search on the index page) or the server code `code` is above 300. I also exclude web requests from possible web crawlers. To qualify as a web crawler, an IP address makes more than 5 requests in a minute or 1000 requests in a day. The abnormal measure is calculated as the log of 1 plus the number of searches on day t minus the average daily number of searches in the past 60 days.	SEC EDGAR log
$\text{AbnEdgar\_Exist}_{it}$	Abnormal number of human EDGAR searches from IP addresses which have accessed the filings from the same firm in the previous month.	SEC EDGAR log

Continued on next page

Table 11 – *Continued from previous page*

Variable	Definition	Source
AbnEdgar_Ins$_{it}$	Abnormal number of EDGAR searches from institutional investors. I identify institutional investors first by matching IP addresses to known institutions which have autonomous system numbers. The IP-ASN organization link file comes from MaxMind. I then search for finance-related words in the names of the institutions. Details can be found in the Online Appendix B	SEC EDGAR log
AbnEdgar_13F$_{it}$	Abnormal number of EDGAR searches from institutional investors. I identify institutional investors first by matching IP addresses to known institutions which have autonomous system numbers. The IP-ASN organization link file comes from MaxMind. I then match the names of these institutions to the names of all 13F institutions. Details can be found in the Online Appendix B	SEC EDGAR log
AbnAnalyst$_{it}$	Abnormal number of analyst forecasts issued on day t. For each firm-day, I count the unique number of analysts who issue any earning forecasts for firm i. Then the abnormal measure is calculated as the log of 1 plus the number of analysts issuing any earning forecast for firm i on day t, minus the log of 1 plus the average number of analysts issuing forecasts per day in the past 60 calendar days.	IBES
AbnAnalyst$_{it}$	Abnormal number of analyst forecasts issued on day t. For each firm-day, I count the unique number of analysts who issue any earning forecasts for firm i. Then the abnormal measure is calculated as the log of 1 plus the number of analysts issuing any earning forecast for firm i on day t, minus the log of 1 plus the average number of analysts issuing forecasts per day in the past 60 calendar days.	IBES

Continued on next page

Table 11 – *Continued from previous page*

Variable	Definition	Source
AbnAnalyst_MoreAccu$_{it}$	The abnormal number of analyst forecasts from analysts whose relative forecast accuracy is above median in the previous year. The relative forecast accuracy is constructed following Ljungqvist et al. (2007). For analyst i covering firm k in year t, I first calculate the absolute forecast error using the following steps. (1) get the analysts most recent forecast of year-end EPS issued between Jan. 1 and Jun. 30, (2) calculate the difference with the subsequent realized earnings, (3) scale the difference by previous year-end price. The for all the analysts covering firm k in year t, I re-scale the absolute forecast errors so that the most and least accurate analysts scores one and zero, respectively. Finally, analyst i's relative forecast accuracy in year t is his/her average score across the the stocks he/she covers over years t-2 to t.	IBES
AbnAnalyst_LessAccu$_{it}$	The abnormal number of analyst forecasts from analysts whose relative forecast accuracy is below median in the previous year.	IBES
AbnAnalyst_MoreExp$_{it}$	The abnormal number of analyst forecasts from analysts who are experienced. Experienced analysts are defined as the analysts who cover the firm for an above-median number of years in the past five years.	IBES
AbnAnalyst_LessExp$_{it}$	The abnormal number of analyst forecasts from analysts who are not experienced. Experienced analysts are defined as the analysts who cover the firm for an above-median number of years in the past five years.	IBES
AbnTurnover	Abnormal turnover, calculated as the log of 1 plus the turnover minus the log of 1 plus the average daily turnover in the past 60 trading days.	CRSP
AbnSpread_EW	Equal-weighted abnormal effective spread. For each trade, I calculate the effect spread by $2\|\log(P_k)-\log(M_k)\|$, where P_k is the price of the trade, and M_k is the mid-point of the consolidated BBO at the time of the trade. I exclude corrected orders, trades with zero price or zero size, trades with condition codes B, G, J, K, L, O, T, W, or Z, and quotes in which the bid-ask spread is negative or greater than 50% of the quote midpoint. Trades on all exchanges are included. Then for each day, I calculate the equal-weighted average effective spread over all the trades. I define abnormal spread as the effective spread on day t over the average daily effective spread in the past 60 trading days.	TAQ

Continued on next page

Table 11 – *Continued from previous page*

Variable	**Definition**	**Source**
AbnSpread_VW	value-weighted abnormal effective spread, calculated as the dollar-value weighted average of all the effective spread within the day.	TAQ
|CAR|	Absolute value of the abnormal return. I calculate abnormal returns by subtracting the CRSP value-weighted index return from the daily raw returns.	CRSP
Range	Daily price range, defined as the log of the daily high price minus the log of the daily low	CRSP
% Priv	The percentage of press releases in NPRA that are issued by private firms.	RavenPack PR Edition
NPRPriv$_{it}$	The number of press releases that are from private firms and published in the next 30 seconds after the press relese j.	RavenPack PR Edition
NPRInd$_{it}$	The total number of press releases that are issued by firms from the same 2-digit SIC industry as firm i, minus 1 (the press release from firm i itself).	RavenPack PR Edition
ESS	Event sentiment score. The score is generated by RavenPack. A group of experts first read and score a sample of stories to determine the direction of impacts (positive or negative) and the degree of different event types. In total they have over 2000 types of events. New articles are then compared to these tagged events to calculate the score. Stories with scores higher than 50 are positive, and lower than 50 are negative.	RavenPack PR Edition
AT	Total asset	Compustat
Q	Tobin's Q, defined as `(at + csho x prcc_f - ceq) / at`	Compustat
Age	The number of years since publication	Compustat, CRSP
NWord	The number of words in the title of a press release	RavenPack PR Edition
DJPR	A dummy variable that equals to one if the automated algorithm from Dow Jones Newswire republishes the press release	RavenPack PR/DJ Edition

Appendix B. Data Cleaning Summary

Table 12: Data cleaning process to generate the main sample

The table below shows the steps I take to compile the sample of press releases used in this paper.

Filtering criteria	# of press releases	# of unique firms
Keep if RELEVANCE = 100 and in CRSP/Compustat	1,620,046	9,406
Keep Top 4 press release wires	1,502,021	9,386
Remove duplicated releases (ENS = 100)	1,068,148	9,373
Keep only one press release per firm-day	909,874	9,368
Keep only trading days	901,774	9,366
Keep if after April 1, 2006	738,196	8,756
Keep if issued in the first 30 seconds of an hour	188,981	7,911
Keep if issued in the first 10 seconds of an hour	131,683	7,503
Keep if issued in 7AM-9AM or 4PM	80,246	6,560

Appendix C. Robustness

Table 13: Robustness of Table 3

This table shows that the results in Table 3 are robust to news measures, sample selection, and functional form of the dependent variable. Columns (1) - (3) re-estimate Equation 2 using different news measures. Column (1) uses non-duplicated news only, whose ENS score is 100 in RavenPack; Column (2) uses flash news, which only contains a headline; Column (3) uses full news, which contains a headline and at least one paragraph. Columns (4) - (8) re-estimate Equation 2 using different samples. Dependent variables in Columns (4) - (8) are the abnormal number of media articles covering the firm on the press release day. From all the press releases published in the first 10 seconds in each hour: Column (4) excludes the releases published in the first second, Column (5) excludes firm-years where the total number of press releases is above the 75th-percentile, Column (6) excludes firm-years where there is only one release, Column (7) excludes date-hours where there is only one release. Column (8) includes all the press releases published in the first 30 seconds of each hour. Columns (9)-(11) change the functional form of the dependent variable. Column (9) uses the log(News + 1) as the dependent variable, rather than the abnormal news measure. Column (10) uses the raw level of the abnormal news, rather than the log change measure, as the dependent variable. Column (11) uses a dummy variable that equals to 1 if there are any news coverage on the event day, as the dependent variable. The independent variable, NPRA, measures the number of new press releases published in the top 4 press release wires after the press release j in the next 30 seconds. The standard errors in all regressions are clustered by firm and date, and t-statistics are reported in the parentheses. Coefficients marked with ***, **, and * are significant at the 1%, 5%, and 10% level, respectively. Detailed variable definitions can be found in Table 11.

Robustness to	news measures			sample					functional form		
	novel	flash	full	excl. 1st second	excl. top 25% of PR issuers	excl. if only 1 PR per firm-year	excl. if only 1 PR per date-hour	first 30 seconds	no abnormal	no log	dummy
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
log(NPRA + 1)	−0.085***	−0.039***	−0.040***	−0.094***	−0.088***	−0.094***	−0.094***	−0.046***	−0.093***	−0.365***	−0.045***
	(−8.688)	(−4.802)	(−4.978)	(−6.023)	(−6.043)	(−8.620)	(−8.751)	(−6.986)	(−8.278)	(−7.897)	(−6.218)
Date-Hour FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Firm-Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Detailed Topic FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	131,683	131,683	131,683	101,115	103,673	120,424	116,450	188,981	131,683	131,683	131,683
Adjusted R^2	0.501	0.157	0.155	0.492	0.492	0.510	0.517	0.501	0.490	0.412	0.416