Journal of Accounting and Economics 66 (2018) 415–418

Contents lists available at ScienceDirect

Journal of Accounting and Economics

journal homepage: www.elsevier.com/locate/jae

Discussion of "earnings announcement promotions: A Yahoo Finance field experiment"

Joseph Engelberg

Rady School of Management, University of California, San Diego, USA

A B S T R A C T

Lawrence et al. run a field experiment on Yahoo Finance which randomly promotes some earnings stories. They find that firms in promoted stories have higher contemporaneous returns. I begin by reviewing the media and finance literature and discuss the difficulty in solving the identification problem Lawrence et al. wish to solve. I then discuss how their results fit with existing theories of attention and suggest some changes that could help pin down which theory best explains the results. Finally, I discuss possible directions for future research.

I would like to thank John Core (Editor) and participants at the 2017 Journal of Accounting and Economics conference for their comments and suggestions. All errors here are my own.

1. Media and asset pricing

There is now a well-established relationship between media coverage and stock price movements. When a firm has media coverage, it typically has more trading volume and more volatility compared to days when it has no media coverage (Tetlock, 2010). Moreover, firms with high levels of media coverage have lower ex-post returns (Fang and Peress, 2009) and the textual sentiment of media articles predicts returns in the days ahead in the same direction as the sentiment (Tetlock et al., 2008).

The key question – and the question in Lawrence, Ryans, Sun and Laptev (LRSV) – is whether the media is *causing* any of this market response. After all, media coverage is the non-random choice of profit-maximizing firms (Gentzkow and Shapiro, 2010). The factors that influence these media coverage decisions could be the same ones that affect the market response.

Take, for example, a day when the *Wall Street Journal* (WSJ) writes an article about Walmart's earnings announcement. Given the previous literature, we know Walmart is more likely to have an extreme return and more trading volume on such a day. It is possible that the WSJ article *causes* the abnormal volume and return. It is also possible that some other factor – such as the earnings announcement - jointly causes both the article and the market response. In the former case, financial media play an important role and should be the subject of future research into the incentives, biases and consequences of their reporting. In the latter case, financial media play only an incidental role and become far less interesting. Solving this identification problem has been an important objective of the media-related literature in finance and accounting.

2. Solving the identification problem

To date, the literature has taken three approaches to solve the problem.

2.1. The time-series approach

Under the time series approach, authors compare media coverage and the corresponding market response at different points in

E-mail address: jengelberg@ucsd.edu.

https://doi.org/10.1016/j.jacceco.2018.08.005
Received 10 July 2018; Received in revised form 28 July 2018; Accepted 2 August 2018
Available online 13 August 2018
0165-4101/ © 2018 Elsevier B.V. All rights reserved.

time. In the context of the example above, Walmart's trading volume and return would be compared on days it received WSJ coverage with days it did not.

The challenge with this approach is clear: days with media coverage are likely to have more information arrival than days without media coverage, and it is this difference in information arrival which causes the differential market response. Thus, the time series approach will only be as successful as the authors' ability to show there are no differences in information arrival between the days with and without media coverage. For example, Huberman and Regev (2001) describe the case of Entremed stock which increased fourfold overnight following a *New York Times* (NYT) story about its cancer treatment process. The authors go through the NYT story carefully to be sure that no new information is being disclosed; in fact, all of the information in the NYT story had been released in the journal *Nature* and other media outlets months earlier.

The time series approach can also be successful if authors can show that the choice of media coverage is unrelated to the information arrival. For example, Dougal, Engelberg, Garcia and Parsons (2012) exploit the scheduling of columnists at the WSJ in order to vary media content in a way which is arguably unrelated to information arrival. Peress (2014) exploits newspaper strikes in France, Greece, Italy and Norway in order to achieve plausibly exogenous variation in media coverage.

2.2. The cross-sectional approach

Under the cross-sectional approach, authors compare the behavior of those who are treated with media coverage with those that are not following the same information event. Returning to our Walmart and WSJ example, the cross-sectional approach would compare the trading behavior of those who saw the WSJ article about the earnings announcement with those that did not. Critically, the event, an earnings announcement, is held fixed and the relevant source of variation comes from those who saw the WSJ coverage of the event and those that did not.

The cross-sectional approach has been used successfully in the economics literature which considers the effect of media bias on political outcomes. For example, DellaVigna and Kaplan (2007) employ this approach when comparing the voting patterns of the 2000 presidential election in areas whose cable companies carried Fox News with those that did not.

Engelberg and Parsons (2011) employ this approach in a study of local newspaper coverage of earnings announcements. For a given earnings announcement, they compare the trading behavior of investors whose local newspaper covered the announcement with the trading behavior of investors whose local newspaper did not.

The benefit of the cross-sectional approach is that we no longer have to worry about differences in media coverage coming from differences in informational arrival *because there is only one event in the cross-sectional approach*. However, the cross-sectional approach has its own problem: the cross-sectional treatment of coverage may not be random. For example, perhaps areas that choose to carry Fox News are correlated with factors related to voting outcomes, and perhaps the areas which choose to cover an earnings announcement are correlated with factors related to trading outcomes. For example, maybe the *Atlanta Journal-Constitution* is more likely to cover a Walmart earnings announcement than the *Seattle Post-Intelligencer* because there are more Walmart stores in Atlanta and investors in Atlanta are more interested in Walmart stock. Thus, the cross-sectional approach will only be as successful as the authors' ability to show that the cross-sectional treatment is random or their ability to exploit some exogeneity in treatment among the cross-section. DellaVigna and Kaplan (2007) show that the rollout of Fox News across the country had more to do with cable company mergers and acquisitions than local politics. Engelberg and Parsons (2011) exploit local weather events, like blizzards and thunderstorms, to create exogenous variation in the delivery of local media to local investors.

2.3. The experimental approach

Both the time-series approach and the cross-sectional approach have a weakness. Specifically, coverage choice in the time-series or cross-section may be non-random and related to the market response. The experimental approach attacks this weakness directly *by randomizing coverage explicitly*.

Under the experimental approach, media coverage is randomly assigned to agents and, by doing so, the econometrician can be sure that the choice of coverage is unrelated to the underlying events. Now if the econometrician finds a correlation between (randomized) media coverage and a market response she can conclude that the media coverage is *causing* the market response. This is the approach taken in LRSV and, to the best of my knowledge, the first paper to execute the experimental approach in an accounting or finance journal in order to solve the identification problem described above.

2.3.1. The experiment in LRSV

LRSV execute the experiment on Yahoo Finance between May 12, 2016, and July 28, 2016. As of April 2017, Yahoo Finance was the most popular financial news website in the U.S.[1] For firms with earnings announcements during this period, LRSV randomly select a group of firms to promote to 1% of Yahoo Finance users. Promotion means placing the firms' earnings stories on the homepage for these Yahoo Finance users. Control firms are size-matched firms which do not receive promotion.

Although the executed experiment has some shortcomings, which I discuss later, the authors deserve credit for introducing this approach to the media and finance literature. If future authors are able to convince media companies to execute experiments like these, LRSV will almost certainly be a springboard for this work.

[1] See https://www.similarweb.com/blog/us-media-publications-ranking-april-2017

3. Interpreting LRSV's results

LRSV's experiment randomly promotes some stocks' earnings stories to the front page of Yahoo Finance and looks for the effect on these stocks' trading volume and return. Intuitively, these earnings stories will receive more attention from investors. Therefore, the natural place to start when thinking about LRSV's results is the literature which concerns investor attention. That literature argues that investors have limited attention and can only allocate so much attention across a near-infinite amount of information events and sources (Simon, 1971; Hirshleifer and Teoh, 2003).

The attention literature makes an important distinction between two cases when making predictions for volume and return. In the first case, *news is value-relevant*. In this case, the empirical evidence suggests investors underreact. The academic literature finds stock-market underreaction to earnings news (Bernard and Thomas (1989)), equity issuance news (Spiess and Affleck-Graves, 1995), equity repurchase news (Ikenberry et al., 1995), insider trading news (Cohen et al., 2012) and customer earnings news (Cohen and Frazzini, 2008). The attention literature predicts that attention shocks towards value-relevant news will lead to *less underreaction* to that news. For example, Hirshleifer et al. (2009) find larger earnings response coefficients and less post-earnings announcement drift on earnings announcement days with fewer distracting events (i.e., more attention).

In the second case, when information is *not* value-relevant, Barber and Odean (2008) argue that attention shocks lead to net buying by retail investors because retail investors rarely short-sell. This uninformed buying creates price pressure which temporarily pushes up stock prices, leading to an eventual reversal.

Because earnings news seems value-relevant, I would expect the LRSV experiment to pertain to the first case. In this case, I would expect to find the media effect causes more page views, larger earnings response coefficients (ERCs) and less post-earnings announcement drift (PEAD). *Oddly enough, the authors find none of this*. They find no more page views among the treatment group (Figure 3.A and Table 4). They find larger ERCs for positive earnings but smaller ERCs for negative earnings (Figure 3.E) and find no evidence of drift in either sample within 10 days of the event (page 18).

Even if we view the Yahoo Finance earnings articles as not value-relevant (the second case), the results do not entirely line-up with the Barber and Odean (2008) hypothesis either. The authors do find price increases for the treated firms (Figure 3.C and Table 4) which is consistent with the Barber and Odean (2008) hypothesis. But they find no evidence of abnormal trading (Figure 3.B and Table 4) and no evidence of reversal within 10 days of the event (page 18).

In short, LRSV has its strengths and weaknesses. From a procedural point of view, it does an excellent job designing and executing the first experiment of its kind meant to solve the identification problem in the media and finance literature. Its weakness lies in connecting the results of the experiment to a coherent asset-pricing theory. The authors cite Barber and Odean (2008) as the theory that explains their evidence, but it is only consistent with some of that theory's predictions. Most notably, how can LRSV's results be consistent with net buying following an attention shock, as in Barber and Odean (2008), if there is no evidence of net buying?

4. Suggestions for future work

One reason why LRSV's results may not paint a coherent story is the experiment's lack of power. Specifically, the 169 treatment firms and the promotion to only 1% of Yahoo Finance users may not have been powerful enough to detect an effect on volume and page views. In fact, LRSV perform some calculations which suggest this is the case (page 18–19). Unfortunately for LRSV, the size and scope of the experiment were not under the authors' control: ``The sample size for the experiment in this study was restricted by Yahoo's internal prioritization of resources and was not under the control of the authors." (page 18)

4.1. More powerful experiments

The lesson for future experiments of this kind is to increase the ex-ante power of the experiment. Specifically, future authors should use the results in LRSV and in other empirical work which relates views/viewership and returns and volume (e.g., Engelberg et al. 2012) in order to make ex-ante calculations about the sample size necessary to detect effects. Researchers should be reluctant to enter into an agreement with a media company that cannot promise a sample size necessary to clear the thresholds determined in these calculations.

Moreover, researchers should consider *targeted* treatment in order to have the power necessary to detect effects. For example, less power will be necessary to detect effects among smaller, less liquid stocks. This is because retail traders are more likely to traffic in smaller stocks (Lee et al., 1991) and less trading will be necessary to have price-pressure effects in these stocks (Engelberg et al., 2012). Thus, if future researchers have constraints concerning sample size, they should allocate larger treatment samples among smaller stocks.

4.2. A deeper examination of financial media as a medium for disclosure

With LRSV, we now have evidence from three different approaches – time series, cross-sectional and experimental - that financial media have causal effects on important market outcomes. It does not appear that financial media merely reflect market conditions. Media can cause them to change. Sometimes these changes can be quite dramatic, as in the case of Entremed (Huberman and Regev, 2001).

Establishing casual effects for financial media is the point of departure for a larger research agenda. If media can cause investors to behave in a way they otherwise would not, then understanding the incentives and potential biases financial media have when

making their reporting choices is an important next step. While this work is beginning (e.g., Reuter and Zitzewitz, 2006; Gurun and Butler, 2012; Ahern and Sosyura, 2014), it is in its infancy. As a comparison, there is ample evidence that financial analysts have causal effects on market outcomes and many papers have been written which examine the incentives, biases and mistakes of these information intermediaries (e.g., Lin and McNichols, 1998; Michaely and Womack, 1999). As of the writing of this discussion, I find 538 papers in the top three finance journals and the top three accounting journals with the word ``analyst'' in the title. Among those same six journals, I find 51 papers with the word ``media'' in the title.

After establishing causality from financial media, there are many avenues to explore. First, journalists have several choices to make when they write an article: Which firms should be reported on?; Which events should be covered?; Should the article be positive or negative?; and so on. Do we observe any systematic biases among journalists when making these choices and, if so, which firms win or lose from these biases?

Second, the relationship between firms and journalists should be scrutinized in the same way the relationship between firms and analysts has been. Financial media want access to newsworthy firm information, and firms want positive media coverage (Dyck and Zingales, 2003). How do these joint incentives affect the articles we see journalists write in equilibrium? Do we find evidence that investors understand these incentives when consuming financial media?

Third, the media landscape has changed dramatically over the years. Fifty years ago, the power of financial media was concentrated in the hands of a handful of media companies. Today, the landscape is decentralized, with hundreds of media sources competing for the eyeballs of the investor universe. How has this changing landscape affected the kind of information being reported by financial media?

After establishing the causality of media effects, these questions, and others which concern financial media as a medium for disclosure, seem like the next logical step for the literature to take.

References

Ahern, K., Sosyura, D., 2014. Who writes the news? corporate press releases during merger negotiations. J. Finance 69, 241–291.

Barber, B., Odean, T., 2008. All that glitters: the effect of attention and news on the buying behavior of individual and institutional investors. Rev. Financ. Stud. 21, 785–818.

Cohen, L., Frazzini, A., 2008. Economic links and predictable returns. J. Finance 63, 1977–2011.

Cohen, L., Malloy, C., Pomorski, L., 2012. Decoding inside information. J. Finance 67, 1009–1043.

DellaVigna, S., Kaplan, E., 2007. The fox news effect: media bias and voting. Q. J. Econ. 122, 1187–1234.

Dyck, A., Zingales, L., 2003. The Media and Asset Prices. HBS and University of Chicago Working paper.

Engelberg, J., Parsons, C., 2011. The causal impact of media in financial markets. J. Finance 66, 67–97.

Engelberg, J., Sasseville, C., Williams, J., 2012. Market madness? the case of mad money. Manage. Sci. 58 (2), 351–364.

Fang, L., Peress, J., 2009. Media coverage and the cross-section of stock returns. J. Finance 64, 2023–2052.

Gentzkow, M., Shapiro, J., 2010. What drives media slant? Evidence from U.S. daily newspapers. Econometrica 78, 35–71.

Gurun, U., Butler, A., 2012. Don't believe the hype: local media slant, local advertising, and firm value. J. Finance 67 (2), 561–598.

Hirshleifer, D., Lim, S., Teoh, S.H., 2009. Driven to distraction: Extraneous events and underreaction to earnings news. J. Finance 64, 2289–2325.

Huberman, G., Regev, T., 2001. Contagious speculation and a cure for cancer: a nonevent that made stock prices soar. J. Finance 56, 387–396.

Ikenberry, D.L., Lakonishok, J., Vermaelen, T., 1995. Market underreaction to open market share repurchases. J. Financ. Econ. 39 (2), 181–208.

Lee, C., Shleifer, A., Thaler, R., 1991. Investor sentiment and the closed-end fund puzzle. J. Finance 46, 75–109.

Lin, H., McNichols, M., 1998. Underwriting relationships, analysts' earnings forecasts and investment recommendations. J. Account. Econ. 25, 101–127.

Michaely, R., Womack, K., 1999. Conflicts of interest and the credibility of underwriter analyst recommendations. Rev. Financ. Stud. 12, 653–686.

Peress, J., 2014. The media and the diffusion of information in financial markets: evidence from newspaper strikes. J. Finance 69, 2007–2043.

Reuter, J., Zitzewitz, E., 2006. Do ads influence editors? Advertising and bias in the financial media. Q. J. Econ. 121, 197–227.

Simon, H., 1971. “Designing organizations for an information-rich world. In: Greenberger, M. (Ed.), Computers, Communications, and the Public Interest. John Hopkins Press, Baltimore, pp. 38–52.

Spiess, D., Affleck-Graves, J., 1995. Underperformance in long-run stock returns following seasoned equity offerings. J. Financ. Econ. 38 (3), 243–267.

Tetlock, P., Saar-Tsechansky, M., Macskassy, S., 2008. More than words: quantifying language to measure firms' fundamentals. J. Finance 63, 1437–1467.

Tetlock, P., 2010. Does public financial news resolve asymmetric information? Rev. Financ. Stud. 23, 3520–3557.