Journal of Financial Economics 145 (2022) 239–258

Contents lists available at ScienceDirect

Journal of Financial Economics

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

A theory of financial media[☆]

Eitan Goldman[*], Jordan Martel, Jan Schneemeier

Indiana University, 1309 E 10th St, Bloomington, IN 47401, USA

ARTICLE INFO

Article history:
Received 2 October 2020
Revised 13 May 2021
Accepted 10 June 2021
Available online 5 July 2021

JEL classification:
D82
G14
M40

Keywords:
Financial journalism
Disclosure
Manipulation
Price quality

ABSTRACT

We present a model of media coverage of corporate announcements. Firms strategically use the media to communicate corporate announcements to a group of traders who observe announcements not directly but through media reports. Journalists strategically select which announcements to report to readers. Media coverage inadvertently incentivizes firms to manipulate the underlying announcements. In equilibrium, media coverage is tilted towards less manipulated negative news. The presence of financial journalists leads to more manipulation but makes stock prices more informative on average. We provide additional predictions regarding the media's impact on the quality of firm announcements and stock prices.

1. Introduction

Financial media plays an important economic role. A growing body of empirical research shows that financial journalists reach a broad swath of market participants, affect trading in financial markets, and help form stock prices (Fang and Peress, 2009; Engelberg and Parsons, 2011; Tetlock, 2011; Peress, 2014; Kaniel and Parham, 2017). Theory, however, provides little insight into their economic role. As a result, our understanding of the equilibrium interactions between the financial media, traders, and firms is somewhat limited.

In this paper, we aim to take a first step in filling this gap by explicitly modeling a *financial journalist* whose strategic actions affect her readers, the firms on which she reports, and the asset prices that result. We start with the basic premise that *some* traders (henceforth *readers*) are only made aware of firm announcements if the announcements get reported by financial journalists. Thousands of US firms file 10-K statements with the SEC, free for the world to see, yet few individual traders have the time to read each statement. For this reason, a financial journalist sifts through the many announcements made by firms and reports on those that she finds to be of greatest value to her readers.

In our model, there is a firm manager, a journalist, and a stock market populated by three kinds of traders. The first are *informed* traders, who observe the universe of all firm announcements and private signals about payoffs. The second are *liquidity* traders, who trade for reasons unrelated to information. The third are the *readers* of financial media, who cannot observe firm announcements directly.

[☆] We thank an anonymous referee, Bill Schwert (the editor), Adelina Barbalau, Sivan Frenkel, Diego Garcia, Christian Heyerdahl-Larsen, Craig Holden, Bo Hu, Dmitry Orlov, Lubos Pastor, Uday Rajan, Jacob Sagi, Chaojung Wang, Basil Williams, and Wenjie Xue. We would also like to thank seminar and conference attendees at Indiana University, University of Notre Dame, the 2019 Annual Conference on Financial Economics and Accounting, the 2019 Junior Accounting Theory Conference, the 2020 Jackson Hole Finance Group Conference, the 2020 IDC Herzliya Conference in Financial Economics Research, the 2020 Midwest Finance Association Annual Conference, the 2021 Financial Intermediation Research Society Conference, and the 2021 MIT Asia Conference in Accounting.

[*] Corresponding author.
E-mail address: eigoldma@indiana.edu (E. Goldman).

https://doi.org/10.1016/j.jfineco.2021.06.038
0304-405X/

The readers rely exclusively on the journalist to become aware of the reported financial information.

The firm manager receives some information and prepares a public announcement (a press release, for example). He attempts to inflate the stock price by *manipulating* the announcement. If the journalist decides to report on the announcement, the readers observe the report and trade on its potentially biased information. The existing empirical literature has highlighted several channels through which firms can manipulate their announcements and mislead traders about firm fundamentals. For instance, Huang et al. (2014) emphasize the tone of words in earnings press releases, while Li (2008) and Bushee et al. (2018) highlight the role of complex language.[1]

The main role of the financial journalist in our framework is to consider each firm announcement and report the announcements that yield the greatest informational benefit—and thus trading profits—to her readers. The journalist's optimal decision balances the informational benefit from reporting an announcement against the cost to readers of reading a manipulated report. As a consequence, more informative and less manipulated announcements are more likely to get reported.[2]

Readers are fully rational and form expectations about the potential degree of manipulation in each firm announcement that the journalist reports. If they suspect that a report is heavily manipulated, they rationally trade less aggressively on it. The firm manager's incentive to manipulate, in turn, depends on the journalist's strategic reporting decision. The optimal level of manipulation in the announcement balances the positive impact of manipulation on the stock price against the negative impact that a more heavily manipulated announcement is less likely to be reported by the journalist. In addition, we allow for an exogenous cost of manipulation, which can represent either a direct cost of being caught at manipulating information or an indirect reputational cost.

We embed this strategic interaction between the firm manager and the journalist in a standard trading model. Informed and liquidity traders always participate in the financial market, while readers only trade if the journalist reports. The equilibrium stock price clears the market and sets the aggregate demand equal to the fixed asset supply. The informed traders' information is partially revealed by the price and allows readers to correct the firm's manipulated signal in some cases. Based on the financial market equilibrium, we solve for the unique reporting and manipulation equilibrium. This equilibrium generates several key results. Some of these results confirm existing empirical findings, while others give rise to novel empirical predictions.

First, the model generates an equilibrium probability with which the journalist reports news. The journalist considers the actual content of the firm's announcement as well as the extent to which the firm tries to manipulate it. As a result, she only reports if the benefit to the readers from the information provided in the report outweighs the cost from the expected level of manipulation in that report. We find that this *reporting probability* depends crucially on the firm's signal. Positive news gets reported with a positive probability, which decreases with the expected degree of manipulation. Negative news gets reported with a higher probability than positive news because negative news does not depend on the degree of manipulation. This result stems from the strategic actions of the journalist and the firm manager, and occurs because a negative firm announcement is less likely to be manipulated than a positive one. Our model predicts that across all firm announcements at a given date, those that are more negative are more likely to be reported on because they are expected to contain a lower level of manipulation and hence are more useful to readers. This result is consistent with evidence in Tetlock (2007); Garcia (2013), and Niessner and So (2018) that financial media reports tend to be negative.[3]

Second, we find that the presence of a journalist is associated with an increased incentive for firms to manipulate their announcements. In our setting, a report by the journalist and a larger degree of manipulation appear jointly. Intuitively, readers of the newspaper trade only based on the information provided by the journalist. The journalist's report encourages the firm to try to manipulate its announcement because reporting increases the number of traders that are susceptible to potentially manipulated public information.[4]

It is important to note, however, that readers form rational expectations about the degree of manipulation; hence, stock prices are not systematically biased on average. Instead, we find that the presence of a journalist results in a reduction in the deviation of the stock price from the realized cash flow or an increase in "price quality." This means that stocks that are more likely to be covered by a journalist are more efficiently priced relative to stocks with less media coverage. The reason for this finding is that the journalist strategically reports only if the benefit to her readers outweighs the cost associated with trading on a manipulated announcement.

Finally, our paper provides additional results concerning the relation between stock returns and media reports.

[1] The importance of strategic bias or manipulation has led to a debate about different ways to measure it. For example, Li (2008) uses the Fog Index to measure the information content of various firm disclosures, while Loughran and McDonald (2014) construct a readability index to measure the extent to which a firm's disclosure is informative. These papers demonstrate that firms use language to hide or highlight financial information in their disclosed statements.

[2] In a model extension, we consider an additional role for the journalist and allow her to clarify announcements to minimize her readers' exposure to manipulated information. Having such an "investigative" reporter reduces the firm's incentive to manipulate but does not eliminate manipulation entirely.

[3] A large literature argues that readers suffer from loss aversion and that this explains why general media reports focus on bad news ("if it bleeds it leads"). For example, Garz (2014) shows this in the reporting of unemployment news, and Soroka et al. (2019) demonstrate that this phenomenon is true across many countries. Our model offers an alternative, rational explanation for this phenomenon.

[4] For ease of exposition, readers are the only traders that are affected by a manipulated signal. However, we could add an additional group of traders that trades on the firm's signal even in the absence of a media report. In this case, the presence of a journalist would *amplify* the manager's incentive to manipulate.

For example, our model suggests that following a journalist's report, prices might overreact in the short-term, reflecting a slight bias, then revert back to the true unbiased value. This will be true when the news is positive, as the stock price following a positive news report is biased. This time series price response is consistent with the findings in Tetlock (2011) and others. We provide an economic explanation for this finding and generate more granular implications. As another example, our paper relates to the work of Huberman and Regev (2001) and Tetlock (2011), who show that traders respond to stale news that is reported by the media. While their empirical work implicitly takes reporting as given and then argues that traders are irrational, we offer an alternative interpretation. Our model suggests that a journalist optimally decides to report stale news because she believes that her readers have not incorporated this (public) information into their past trading.

In Section 5, we provide a detailed description of the empirical predictions of our model about the probability of reporting, the magnitude of manipulation, and the overall impact on stock prices. In sum, our paper helps to answer questions such as *What kind of news should be reported by the financial media? How does the media's presence alter the firm's incentive to release accurate information? Are individual traders better off with media reporting?* And *What are the implications for stock prices when journalists are present?*

The model makes two important assumptions. First, we consider a journalist who makes a reporting decision based on its impact on her readers' ability to trade. This is a benchmark under which the journalist's ability to attract readers depends on whether or not they view her information as useful in the long term. In the context of financial news, this would mean that the information she provides helps readers make better financial decisions, which we model as better trading outcomes. Because the usefulness of *financial* news differs fundamentally from that of *political* news, we do not consider a journalist whose objective is to maximize reputation or who caters to her readers' beliefs (see, e.g., Gentzkow and Shapiro, 2006; Mullainathan and Shleifer, 2005). We believe that our specific objective function uniquely characterizes financial news reporting.[5]

Second, our baseline assumption is that readers are unaware of a firm's announcement unless the announcement is reported by the journalist. Thus, readers do not trade absent news from the journalist.[6] Hence, the main role of the journalist is to *disseminate* existing information by highlighting to her readers a small subset of available information that is of higher importance. Our focus is on the day-to-day reporting that happens in financial newspapers such as the *Wall Street Journal*, rather than on investigative reporting, which happens less frequently but usually receives more public attention.[7]

Note that in the baseline model we assume that readers are rational about the firm's incentive to manipulate announcements and that the journalist does not directly address the manipulation in her report. In a model extension, we consider an additional investigative role for the journalist. In particular, we allow for the possibility of a journalist who is also able to *clarify* firm announcements by identifying, with some probability, cases in which the firm announcements are manipulated. As discussed earlier, there is a growing body of empirical work suggesting that firms are strategic in writing firm announcements (see, e.g., Huang et al., 2014; Bushee et al., 2018). The role of the financial journalist is to detect these distortions and to provide a clearer picture to readers. The main finding in this extension is that the firm endogenously manipulates less in anticipation of the journalist's corrective action. In addition, the probability of coverage increases with the journalist's skill. Furthermore, journalists have a stronger incentive to investigate when the firm is expected to manipulate more.

Our paper takes a first step towards a more complete understanding of the role of *financial* news. The theoretical work of Mullainathan and Shleifer (2005) explores the incentive of the media to bias news more generally in order to cater to the beliefs of its readers. Gentzkow and Shapiro (2006) focus on the media's political bias. In both of these papers, the journalist chooses to engage in biased reporting. In our equilibrium, we also find the existence of a media bias, but in contrast to these papers, we argue that bias in financial reporting occurs *despite* the efforts of the journalist to eliminate it. Furthermore, our model generates two distinct types of media bias.

First, the journalist is more likely to report negative news than positive news (an ex post bias). Second, the firm manipulates its announcements to make them rosier than the truth (an ex ante bias). Given the unique features of reporting on financial news, our paper also highlights a novel interaction between the journalist's reporting decision and the firm manager's incentive to manipulate information, which is absent in the work above. Therefore, the specific financial market environment creates novel endogenous forces with non-trivial implications for the media's reporting incentives.

More broadly, our paper contributes to the theoretical literature studying the role of *public information* on stock market trading, price formation, and price quality. Building on early contributions such as Diamond (1985); Admati and Pfleiderer (1986), and Fishman and Hagerty (1989), several recent papers study the impact of corporate disclosure in a market with sophisticated traders and liquidity traders [see Goldstein and Yang, 2017 for a recent survey of this literature]. For instance, Gao and Liang (2013); Han et al. (2016), and

[5] We also abstract from quid pro quo incentives but acknowledge empirical evidence that journalists sometimes pander to the firms on which they report (Dyck and Zingales, 2003; Call et al., 2018; Baloria and Heese, 2018). However, we think that the incentive to report news that is useful to her readers is of first-order importance for the journalist.

[6] An alternative assumption could be that, absent a news report, readers do not trade stocks they are unaware of, but do trade stocks they have heard of before. In this case, one would assume that, absent a news report, readers trade based on their estimate of expected news. We find this alternative assumption less appealing, but have verified that our main findings continue to hold in this setting.

[7] There is some empirical evidence suggesting that retail traders buy stocks that are covered in the media (e.g., Barber and Odean, 2008), as well as that stock prices respond to the media's reporting of stale news (e.g., Tetlock, 2011; Drake et al., 2014). Both are consistent with the media's role as an information "pass-through."

Goldstein and Yang (2019) study the impact of corporate disclosure on private information acquisition and real efficiency. These papers emphasize the delicate interaction between public information provision and private information acquisition. Moreover, Kurlat and Veldkamp (2015) analyze an alternative cost of public information and show that it can lead to a reduction in trading opportunities. In our framework public information is also endogenous. Unlike the aforementioned papers, however, we consider a setting where information must be disclosed and could be manipulated by the firm manager to inflate the firm's stock price (see, e.g., Goldman and Slezak, 2006; Gao and Zhang, 2019).

Cohn et al. (2018) study a setting in which firms can manipulate information that is subsequently observed by a strategic credit rating agency, and analyze the implications for rating accuracy and manipulation.[8] Similar to us, they consider strategic interactions between a firm that manipulates information and an intermediary that screens the information. The key difference is that in our paper the journalist decides whether or not to report only after considering the trading profits of her readers. This choice leads to an endogenous reporting probability, which is the main focus of our paper. In contrast, their focus is on the precision of screening.

Our paper also relates to models of financial analysts who can be viewed as another type of information intermediary (e.g., Langberg and Sivaramakrishnan, 2010; Einhorn, 2018; Frenkel et al., 2020). In contrast to these papers, our key modeling assumption is that it is the journalist, not the firm, who decides which corporate announcements should be made public. This assumption results in a very different set of predictions, which better match the economic role of an information intermediary who disseminates existing information (the journalist), instead of creating new information (the analyst).

2. Economic framework

The model features a journalist, a firm manager, and three types of traders. In this section, we first discuss the strategic actions of each of these players. We conclude the section with a characterization of the financial market equilibrium.

2.1. Model setup

There are four dates $t \in \{0, 1, 2, 3\}$ and two assets, one risk-free and the other risky. The risk-free asset serves as the numeraire and is in unlimited supply. The risky asset is in zero net supply and represents a claim to a firm's liquidating dividend d_θ, which is paid at $t = 3$. A fundamental shock θ takes on values L and H with equal probability. Without loss of generality, we assume that $d_H > d_L$. In the following, we will often refer to the standard deviation of d_θ as *payoff uncertainty*, which is given by $\sigma_d = \frac{1}{2}(d_H - d_L)$. The risky asset is traded in a secondary financial market at $t = 2$, and we denote its equilibrium price by p.

The model features three types of traders: (i) a unit mass of informed traders (I), (ii) a mass $\chi > 0$ of readers (R), and (iii) liquidity traders. All traders are risk-neutral and trade competitively. In addition to these three types of traders, there is also a firm manager (F, "he") and a journalist (J, "she"). Fig. 1 summarizes the key model elements, and Fig. 2 provides a timeline for the main model.

2.1.1. Firm manager

At $t = 0$, the firm's manager observes the fundamental shock θ and issues a public signal $s_F \in \{L, H\}$.[9] The tractable binary structure for the asset payoff and signals builds on the frameworks of Chen et al. (2007); Strobl (2013); Cohn et al. (2018), and Gao and Zhang (2019). As in these papers, the manager only has an incentive to issue $s_F \neq \theta$ if the fundamental is low ($\theta = L$). If the fundamental shock is high ($\theta = H$), then the firm's public signal s_F is always accurate. Hence, the probability that the manager sends a high signal given that the fundamental shock is high is $\mathbb{P}(s_F = H|\theta = H) = 1$. If the fundamental shock is low, the realization of s_F depends on the firm's choice of the *intensity of manipulation* $m \in [0, 1]$. We assume that with probability $\mathbb{P}(s_F = H|\theta = L) = m$, the manager successfully manipulates the signal and reports H instead of L. With probability $\mathbb{P}(s_F = L|\theta = L) = 1 - m$, manipulation is unsuccessful and the firm reports L.

The manager chooses the intensity of manipulation to maximize the firm's expected stock price p net of a private manipulation cost $\mathcal{C}(m)$:[10]

$$\max_{m\in[0,1]} \mathbb{E}[p|\theta = L] - \mathcal{C}(m). \tag{1}$$

As shown below, the impact of manipulation on the expected stock price is affected by the strategic decision of the journalist of whether or not to report the firm's announcement.

The assumption that the manager maximizes the expected stock price rather than the long-term liquidation value of the firm can reflect concerns for managerial reputation, as in Narayanan (1985) and Scharfstein and Stein (1990), or managerial myopia, as in Stein (1989). For ease of exposition, we choose a simple linear cost $\mathcal{C}(m) = c_m m$ with $c_m > 0$. As in the existing literature, we assume that the manager's manipulation choice m is privately observed. Other market participants, such as the journalist and readers, base their actions on a conjecture $\widehat{m}$ about m, which is equal to m in equilibrium.

The signal structure implies that a low managerial reported signal ($s_F = L$) is perfectly informative about the firm's fundamental shock. Conversely, a high managerial

[8] Other papers studying the role of credit rating agencies (CRAs) include Bolton et al. (2012); Fulghieri et al. (2014); Frenkel (2015), and Piccolo and Shapiro (2017). In this literature, the focus is usually on the attempt of the CRA to manage its reputation as an information provider with its ability to maintain a positive interaction with the firm it is rating.

[9] The results are robust to the alternative assumption that the payoff contains an additional, unpredictable component. Moreover, given that the manager always receives the signal d_θ, he does not have an incentive to withhold negative news, as doing so would perfectly reveal θ to be L.

[10] We consider a *private* manipulation cost for ease of exposition, but we could also incorporate $\mathcal{C}(m)$ into the firm's payoff.

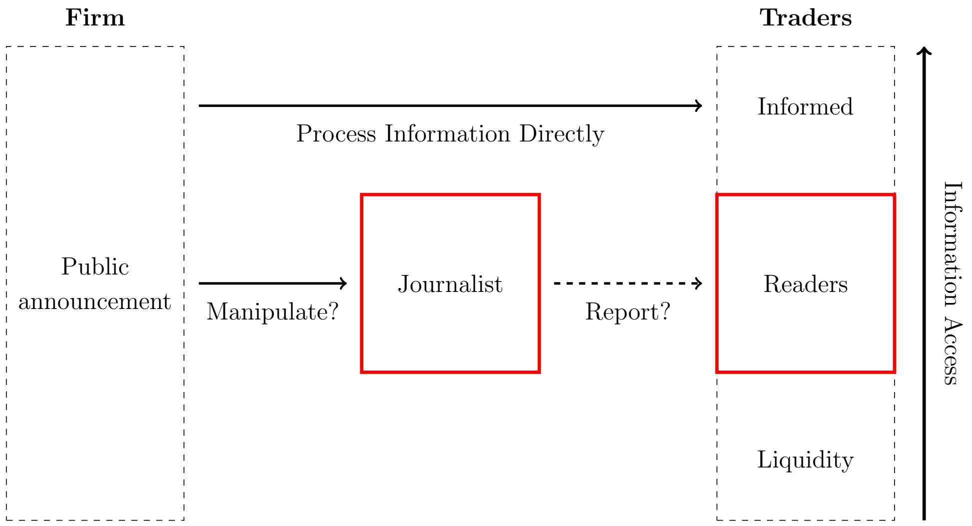

Fig. 1. Three types of traders. We distinguish informed from uninformed traders ("readers") by their ability to observe the firm's public announcement.

reported signal ($s_F = H$) could result from either the fundamental being high ($\theta = H$) or the manager's successful manipulation of a low fundamental signal ($\theta = L$). In the context of our model, s_F can be interpreted as a public announcement such as an earnings report or a press release. The manager can inflate the content of this signal through his choice of m. We think of the manager's manipulation efforts quite generally as any activity he can use to hide bad information or to emphasize good information. As mentioned in the introduction, the existing empirical literature has, for instance, highlighted the use of tone management and complex language in corporate announcements. We interpret m not necessarily as *illegal* manipulation or fraud but rather as a tool to mislead some traders in the market.

2.1.2. Journalist

The journalist observes the firm's announcement s_F at $t = 1$ and decides whether to report it to her readers. She observes neither θ nor m, and her readers do not observe s_F unless she reports it. If $s_F = L$, the journalist knows that $\theta = L$, but if $s_F = H$, she is uncertain about θ because the manager might have successfully manipulated the announcement. Given the firm's announcement and her conjecture $\widehat{m}$ of the manager's manipulation choice m, the journalist decides whether to report the announcement ($\mathcal{D}_r = 1$) or not ($\mathcal{D}_r = 0$). If she decides to report, she issues a report $s_J = s_F$. Otherwise, she does not issue a report, and $s_J = \emptyset$.

The journalist's report is observed by all agents, but only readers rely on s_J in their trading decision. The other two types of traders do not rely on the journalist's report. Informed traders are endowed with superior information about the firm's payoff and cannot learn any additional information from the journalist's report. Liquidity traders trade for exogenous reasons that are assumed to be independent of the firm's payoff.

It should be noted that in contrast to some of the existing literature, such as Mullainathan and Shleifer (2005) and Gentzkow and Shapiro (2006), the journalist does not sensationalize the firm's report (for example, by adding a "media bias") in our setting. Rather, we view the journalist as a benevolent transmitter of information who tries to report as accurately as possible on the firm. In Section 4, we consider a journalist who can also choose to conduct an independent investigation of the firm's public announcement. The investigative role implies that the signal reported by the journalist can be more accurate than the one announced by the firm.

Assumption 1 (Journalist's objective). The journalist's reporting decision is made to maximize expected reader utility net of a private reporting cost c_r.

Two factors determine the journalist's decision to report. The first factor is the anticipated utility gain for her readers, and the second is her opportunity cost. The utility benefit to readers who observe a report comes from their ability to trade on the new information. Consequently, their utility gain from observing a report is equal to their expected trading profits given $s_J = s_F$ minus that given $s_J = \emptyset$. The journalist thus acts in the best interest of her readers and only reports on the firm if reporting generates a gain in expected trading profits for her readers.[11] Our interpretation of the journalist is thus similar to that used in the literature on information sales such as Admati and Pfleiderer (1986) and Admati and Pfleiderer (1988).

The second factor that influences the journalist's reporting decision is an independent stochastic opportunity cost c_r that is uniformly distributed on $[0, \bar{c}_r]$. This cost can be interpreted as the journalist's utility from reporting on a different topic, such as another firm. The parameter $\bar{c}_r$ governs the average appeal of these alternative topics.[12] The introduction of an opportunity cost allows us to capture

[11] Our main results are robust to the alternative specification in which the journalist maximizes the sum of all readers' trading profits, not just those of a representative reader.

[12] One way to endogenize c_r would be to consider a multi-firm setup. A capacity constraint on the journalist's time or attention would then force her to report on the firm that creates the greater benefit for her readers.

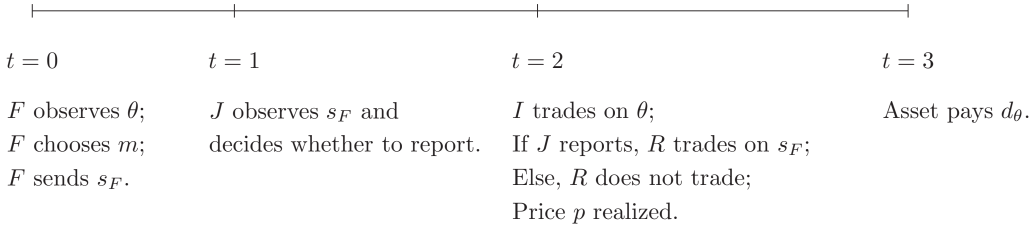

Fig. 2. Timeline for the main model.

the fact that not all corporate announcements can be reported on the front page. If a certain announcement lacks credibility or simply confirms a widely held view, it should be in the best interest of the reader to shift the focus to a different story.[13]

According to Assumption 1, the journalist compares the increase in the expected utility of a representative reader with her opportunity cost.[14] We can summarize the journalist's reporting strategy as follows:

$$\mathcal{D}_r = \begin{cases} 1 & \text{if } \mathbb{E}[U_R|s_J = s_F] - \mathbb{E}[U_R|s_J = \emptyset] > c_r \\ 0 & \text{if } \mathbb{E}[U_R|s_J = s_F] - \mathbb{E}[U_R|s_J = \emptyset] \le c_r. \end{cases} \tag{2}$$

2.1.3. Readers and informed traders

At $t = 2$, informed traders and readers submit asset demand schedules x_i, where $i \in \{I, R\}$, conditional on the stock price p, to maximize their expected trading profits. To keep their demands finite, we introduce a quadratic trading cost $\frac{\kappa}{2}x_i^2$ with $\kappa > 0$ as in Pouget et al. (2017) and Banerjee et al. (2018).[15] We can thus write the trading profits for informed traders and readers as:

$$U_i = x_i(d_\theta - p) - \frac{\kappa}{2}x_i^2 \tag{3}$$

for $i \in \{I, R\}$. It follows that the optimal demand for type i is

$$x_i = \kappa^{-1}(\mathbb{E}[d_\theta|\Omega_i] - p) \tag{4}$$

where Ω_i denotes type i's information set. Informed traders condition on the fundamental shock and the stock price: $\Omega_I = \{\theta, p\}$. Readers rely solely on the journalist's report and the stock price: $\Omega_R = \{s_J, p\}$. Informed traders are therefore perfectly informed in our model. They observe the true payoff d_θ, and their optimal demand is given by:

$$x_I = \kappa^{-1}(d_\theta - p) \tag{5}$$

so that each informed trader observes the mispricing of the firm's stock $(d_\theta - p)$ and trades against it. The convex trading cost prevents traders from taking extremely large positions and generates limits to arbitrage, which is captured by the constant factor κ^{-1} in their optimal demand. The lower the trading cost, the higher the traders' aggressiveness in exploiting mispricing.

Assumption 2 (Readers' observed signals). Readers observe the firm's signal only if it is reported by the journalist. If the signal is reported, they supplement the firm's signal with stock price information.

Readers have inferior information compared to informed traders. They do not observe the firm's announcement directly and depend on the journalist to write a report in order to receive information about d_θ. If the journalist does not report, readers are unaware of the firm's announcement. If the journalist reports, they trade on the reported signal and behave fully rationally. This means they also understand that the equilibrium stock price partially reflects the informed traders' private information. Readers are, therefore, fully-rational traders that require the journalist's reporting to become aware of a specific firm's signal.

Assumption 3 (No-reporting benchmark). Readers do not trade in the absence of a report.

Assumption 3 states that absent any news, readers prefer not to trade in the firm's stock, and their expected utility in that case is equal to zero: $\mathbb{E}[U_R|s_J = \emptyset] = 0$. This assumption seems intuitive and can be justified in a more general setting where trading has a small fixed cost and where there is a positive probability that firms do not have new information they need to announce. In unreported analysis, we show that allowing readers to trade, absent a news report, based on their understanding of the expected value of d_θ results in a slightly more complex model setting, in which the main results of the paper remain.

We model the journalist as an information intermediary who transmits information from the firm to a group of uninformed traders. In actual markets, these types of traders have a limited attention budget and might be overwhelmed by the amount of information provided by firms. They rely on a journalist to determine the relevance and substance of these signals. If a journalist does not cover a specific firm, the firm is not within the readers' investment opportunity set. Empirically, there is ample evidence that corporate announcements require media coverage if they are to reach parts of the market, and that media reporting matters for traders (e.g., Huberman and Regev, 2001; Engelberg and Parsons, 2011; Tetlock, 2011).

Readers' equilibrium demand, if the journalist reports $(\mathcal{D}_r = 1)$, is given by:

$$x_R = \kappa^{-1}(\mathbb{E}[d_\theta|\Omega_R] - p) \tag{6}$$

[13] In line with this intuition, Fang and Peress (2009) document that less than 75% of NYSE stocks are covered (by four major newspapers) in a typical year.

[14] As mentioned earlier, we could model the journalist's cost c_r as a function of the underlying news. To avoid any "baked-in" asymmetries, we keep the distribution of c_r constant.

[15] We could alternatively use a mean-variance objective function for these two types of traders at the cost of less tractable equilibrium expressions. Our qualitative results are robust to this alternative objective.

and $x_R = 0$ if $\mathcal{D}_r = 0$.

In addition to informed traders and readers, there is also a unit continuum of liquidity traders with exogenous net demand u. To obtain tractable solutions and to allow readers to learn additional information from the stock price, we assume that for some $\Delta > 0$, u is uniformly distributed on $[-\Delta, \Delta]$. Hence, average liquidity demand is equal to zero, and its variance is given by $\sigma_u^2 = \frac{\Delta^2}{3}$. Throughout the paper, we assume that

$$\Delta > \kappa^{-1}\sigma_d \tag{7}$$

where κ and σ_d denote the trading cost coefficient and payoff uncertainty, respectively. This condition ensures that the stock price is never fully revealing for all realizations of liquidity demand. If it were violated, readers would *always* be able to perfectly learn the realization of θ from the equilibrium stock price, and there would be no difference between informed traders and readers, conditional on reporting.

The role of liquidity demand u is twofold. First, u adds non-fundamental variation to the stock price and prevents it from perfectly revealing the informed traders' information to readers. Second, u also allows readers to make positive trading profits in equilibrium, which is necessary to incentivize the journalist to report.

The market clearing condition sets the asset demands of the three types equal to the fixed zero supply:[16]

$$x_I + \chi x_R + u = 0. \tag{8}$$

Definition 1. An equilibrium consists of (i) a trading policy by informed traders and readers, (ii) a reporting policy by the journalist, and (iii) a manipulation policy by the firm manager such that:

1. The informed traders' demand x_I maximizes $\mathbb{E}[U_I|\Omega_I]$;
2. The readers' demand x_R maximizes $\mathbb{E}[U_R|\Omega_R]$;
3. The journalist's reporting policy $\mathcal{D}_r \in \{0, 1\}$ maximizes $\mathcal{D}_r\mathbb{E}[U_R|s_F] + (1 - \mathcal{D}_r)c_r$;
4. The manager's manipulation policy $m^* \in [0, 1]$ maximizes $\mathbb{E}[p|\theta = L] - \mathcal{C}(m)$;
5. The stock price $p(\theta, s_J, u)$ clears the stock market as in Eq. (8);
6. The conjecture $\widehat{m}$ is correct in equilibrium, i.e., $\widehat{m} = m^*$.

2.2. Financial market equilibrium

As a first step, we solve for the financial market equilibrium at $t = 2$ and take the journalist's reporting decision ($t = 1$) and the manager's manipulation decision ($t = 0$) as given. We solve for these two equilibrium choices afterwards in Section 3.

We plug in the optimal demands for informed traders and readers into the market clearing condition in Eq. (8) to solve for the equilibrium stock price p as a function of the true fundamental $\theta \in \{L, H\}$, the journalist's report $s_J \in \{s_F, \emptyset\}$, and liquidity trading u. As a result, the stock price is given by:

$$p(\theta, s_J, u) = \begin{cases} (1+\chi)^{-1}(d_\theta + \chi\mathbb{E}[d_\theta|\Omega_R] + \kappa u) & \text{if } s_J = s_F \\ d_\theta + \kappa u & \text{if } s_J = \emptyset. \end{cases} \tag{9}$$

If the journalist does not report, only informed and liquidity traders trade. Therefore the stock price is equal to the firm's fundamental value plus noise: $p(\theta, \emptyset, u) = d_\theta + \kappa u$.

If the journalist reports, the stock price also depends on the readers' demand, which might be influenced by the firm's manipulation efforts. More specifically, readers' demand for the stock depends on the reported signal s_J, the stock price p, and the readers' estimate of manipulation $\widehat{m}$.

If the journalist reports a low signal ($s_J = L$), the readers rationally infer that the fundamental is low ($\theta = L$); hence the equilibrium stock price in this case is equal to $p(L, L, u) = d_L + (1+\chi)^{-1}\kappa u$. Regardless of whether $s_J = L$ or $s_J = \emptyset$, p is unbiased and $\mathbb{E}_u\left[p(\theta, s_J, u)\right] = d_\theta$. It is important to note, however, that when $s_F = L$, a mass $1 + \chi$ of traders trade knowing that $\theta = L$, whereas when $s_F = \emptyset$, only a unit measure do so. As a result, the price becomes more efficient, in the sense that it deviates less from the realized payoff on average. We analyze price quality, defined as the squared deviation of p from d_θ, in detail in Section 3.

If the firm issues a high signal ($s_F = H$) and the journalist reports, then readers rationally infer that the underlying fundamental shock could be either high or low. To gauge the likelihood of both scenarios, readers conjecture an equilibrium manipulation intensity $\widehat{m} \in [0, 1]$. Moreover, they understand that informed traders have superior information and that information about the fundamental shock is contained in the equilibrium stock price. Conditional on $s_J = H$, readers invert the equilibrium stock price and back out the following price signal:

$$s_p \equiv (1+\chi)p - \chi\mathbb{E}[d_\theta|\Omega_R] = d_\theta + \kappa u. \tag{10}$$

Note that conditional on $\theta = L$, the price signal is uniformly distributed between $d_L - \kappa\Delta$ and $d_L + \kappa\Delta$. Conditional on $\theta = H$, it is uniformly distributed between $d_H - \kappa\Delta$ and $d_H + \kappa\Delta$. Condition (7) implies that the intersection of these two sets is non-empty, which ensures that readers do not always learn θ from the price signal. It follows that the readers' conditional expectation of d_θ is given by:

$$\mathbb{E}[d_\theta|s_J = H, s_p] = \begin{cases} d_H & \text{if } s_p > \bar{s} \\ \mathbb{E}[d_\theta|s_J = H] & \text{if } \underline{s} < s_p < \bar{s} \\ d_L & \text{if } s_p < \underline{s} \end{cases} \tag{11}$$

where $\bar{s} \equiv d_L + \kappa\Delta$ and $\underline{s} \equiv d_H - \kappa\Delta$. A very positive price signal perfectly reveals that $d_\theta = d_H$, while a very negative price signal perfectly reveals that $d_\theta = d_L$. In the intermediate range $(\underline{s}, \bar{s})$, the price signal is uninformative.

It will be useful to define the probability with which the price reveals d_θ as π_p. This measure of price informativeness is explicitly given by

$$\pi_p \equiv \mathbb{P}(s_p > \bar{s}|\theta = H) = \mathbb{P}(s_p < \underline{s}|\theta = L) = (\kappa\Delta)^{-1}\sigma_d. \tag{12}$$

From condition (7), we have that $\pi_p < 1$. With probability $1 - \pi_p$, the price signal is uninformative, and readers up-

[16] The assumption that the asset is in zero net supply is without loss of generality in our setting due to the traders' risk neutrality.

date their belief about θ solely based on the journalist's report and their conjecture $\widehat{m}$ of m:

$$\mathbb{E}[d_\theta | s_J = H] = \frac{1}{1+\widehat{m}} d_H + \frac{\widehat{m}}{1+\widehat{m}} d_L \tag{13}$$

where $1/(1+\widehat{m})$ and $\widehat{m}/(1+\widehat{m})$ are the Bayesian weights.[17]

Intuitively, readers rationally discount a high signal because they understand that it might have been manipulated by the firm manager. If the manager is expected *not* to manipulate in equilibrium, readers take the report at face value, and their expectation is equal to d_H. The higher the conjectured manipulation intensity, the closer the conditional expectation gets to the unconditional expectation $\mathbb{E}[d_\theta] = (d_H + d_L)/2$.

The expected price conditional on θ and $s_J = H$ is given by:

$$\mathbb{E}_u[p(\theta, H, u)] = \pi_p d_\theta + (1-\pi_p) \times \left[\frac{1}{1+\chi} d_\theta + \frac{\chi}{1+\chi}\left(\frac{1}{1+\widehat{m}} d_H + \frac{\widehat{m}}{1+\widehat{m}} d_L\right)\right]. \tag{14}$$

With probability π_p the price perfectly reveals θ and is consequently unbiased, i.e., equal to d_θ in expectation. Otherwise, only informed traders trade based on the true fundamental θ, while readers trade on $s_J = H$. In the latter case, the expected stock price accounts for the fact that the journalist's signal might have been manipulated by the firm. Therefore, $\mathbb{E}_u[p(\theta, H, u)]$ is decreasing in the conjectured intensity of manipulation $\widehat{m}$.

Next, we compute expected trading profits for readers at $t = 1$. We take an expectation of U_R conditional on s_J by integrating over liquidity demand u and the fundamental shock θ:

$$\mathbb{E}[U_R | s_J] = \mathbb{E}\left[x_R(d_\theta - p) - \frac{\kappa}{2} x_R^2 | s_J\right]. \tag{15}$$

We plug in the readers' optimal demands given s_J and the equilibrium stock price p. The following lemma formalizes the resulting expressions for readers' expected trading profits.

Lemma 1 (Readers' expected trading profits). *Conditional on $s_J \in \{s_F, \emptyset\}$, readers' expected trading profits are given by:*

$$\mathbb{E}[U_R | s_F] = \begin{cases} \frac{\kappa\Delta^2}{6(1+\chi)^2} - \frac{2\widehat{m}\sigma_d^2(\Delta\kappa - \sigma_d)}{\Delta\kappa^2(1+\widehat{m})^2(1+\chi)^2} & \text{if } s_F = H \\ \frac{\kappa\Delta^2}{6(1+\chi)^2} & \text{if } s_F = L \end{cases} \tag{16}$$

and $\mathbb{E}[U_R | s_J] = 0$ if $s_J = \emptyset$.

Proof. See Appendix A.1. □

Lemma 1 provides closed-form solutions for the readers' expected profits. Trivially, if the journalist does not report, then readers do not trade, and their profits are equal to zero. If the journalist reports, the expected profits depend on the reported signal $s_J = s_F$. If the journalist reports $s_F = L$, readers always trade on accurate information and their expected profits are equal to $\frac{\kappa\Delta^2}{6(1+\chi)^2}$. The positive dependence on Δ highlights the fact that readers benefit from the presence of liquidity traders. An increase in readership χ reduces trading profits per reader because more traders are trading on the same information.

If the journalist reports $s_F = H$, readers might trade on a manipulated signal (the report about the firm's announcement). Hence, their expected trading profits are smaller than those obtained when the signal is $s_F = L$. Furthermore, the profits are strictly decreasing in the conjectured level of manipulation $\widehat{m}$ because more manipulation reduces the informational content of the journalist's report.

As we demonstrate below, the adverse impact of $\widehat{m}$ on readers' profits plays an important economic role in the strategic choice of the journalist. If the journalist conjectures a higher manipulation intensity, she anticipates lower trading profits for readers and therefore becomes less likely to report the firm's announcement. This response in turn changes the firm's manipulation efforts.

3. Equilibrium manipulation and reporting

In this section, we endogenize the journalist's reporting decision and the firm's manipulation decision. We start with the journalist's decision, which depends on two factors: the readers' utility gain from reporting $\mathbb{E}[U_R|s_F]$ and the stochastic opportunity cost c_r. More specifically, the journalist reports on the firm if and only if the utility gain exceeds the reporting cost: $\mathbb{E}[U_R|s_F] > c_r$. Since the opportunity cost is privately observed by the journalist at $t = 1$, the firm views the reporting decision as uncertain, ex ante. The firm manager can only compute a *reporting probability*:

$$\pi_r(s_F, \widehat{m}) \equiv \mathbb{P}(\mathcal{D}_r = 1 | s_F, \widehat{m}) = \mathbb{P}(\mathbb{E}[U_R|s_F] > c_r | s_F, \widehat{m}). \tag{17}$$

To compute π_r, we use the fact that c_r is uniformly distributed between 0 and $\bar{c}_r$. Hence, the journalist never reports if $\mathbb{E}[U_R|s_F] \le 0$ and always reports if $\mathbb{E}[U_R|s_F] \ge \bar{c}_r$. If $0 < \mathbb{E}[U_R|s_F] < \bar{c}_r$, then the reporting probability is given by:

$$\pi_r(s_F, \widehat{m}) = \bar{c}_r^{-1}\mathbb{E}[U_R|s_F]. \tag{18}$$

We restrict our attention to the case in which the highest opportunity cost for the journalist is sufficiently high:

$$\bar{c}_r > \max_{\widehat{m}} \mathbb{E}[U_R|s_F]. \tag{19}$$

In other words, the highest possible opportunity cost for the journalist always exceeds readers' maximum expected trading profits. This assumption ensures that the reporting probability is strictly less than one. Lemma 1 implies that readers' trading profits, conditional on reporting, are strictly positive. It follows that $\pi_r \in (0, 1)$, so the reporting probability is given by Eq. (18). Intuitively, there is always a chance that the journalist might not cover a given announcement because her opportunity cost is larger than her readers' utility gain from reporting. This assumption simplifies some of our derivations because it makes π_r differentiable (with respect to $\widehat{m}$).

[17] Specifically, H and L occur with equal probability; conditional on reporting, the probability that $s_F = s_J = H$ given that $\theta = H$ is 1; the probability that $s_F = s_J = H$ given that $\theta = L$ is $\widehat{m}$.

Next, we combine the expression for π_r above with the results in Lemma 1 to obtain the reporting probability as a function of model parameters and the conjectured manipulation intensity.

Proposition 1 (Reporting strategy). *Given the firm's signal $s_F \in \{L, H\}$ and a conjectured manipulation choice $\widehat{m} \in [0, 1]$, the journalist's equilibrium reporting probability is given by:*

$$\pi_r(s_F, \widehat{m}) = \begin{cases} \frac{1}{\bar{c}_r}\left(\frac{\kappa\Delta^2}{6(1+\chi)^2} - \frac{2\widehat{m}\sigma_d^2(\Delta\kappa-\sigma_d)}{\Delta\kappa^2(1+\widehat{m})^2(1+\chi)^2}\right) & \text{if } s_J = H \\ \frac{1}{\bar{c}_r}\frac{\kappa\Delta^2}{6(1+\chi)^2} & \text{if } s_J = L. \end{cases} \tag{20}$$

The reporting probability has the following properties:

1. *for any $\widehat{m} > 0$, the journalist is more likely to report bad news: $\pi_r(H, \widehat{m}) < \pi_r(L, \widehat{m})$;*
2. *for $s_J \in \{L, H\}$, it is decreasing in the mass of readers χ and the reporting cost coefficient $\bar{c}_r$;*
3. *if $s_J = H$, it is decreasing in $\widehat{m}$; if π_p is sufficiently large, it is increasing in liquidity uncertainty Δ and the trading cost κ and decreasing in payoff uncertainty σ_d; otherwise it is decreasing in Δ and κ and increasing in σ_d;*
4. *if $s_J = L$, it is increasing in κ and Δ.*

Proof. See Appendix A.2. □

Proposition 1 provides closed-form solutions for the journalist's reporting probability given the firm's equilibrium manipulation choice $\widehat{m}$, which is chosen at $t = 0$. The reporting probability is proportional to the readers' trading gain because the journalist acts in the readers' best interest, net of her reporting cost. She is more likely to report negative news because $s_F = L$ is inherently more informative than $s_F = H$. A negative announcement is necessarily truthful, while a positive announcement could either correspond to a high or a low fundamental. Furthermore, the informational content of $s_F = H$ decreases in the manager's manipulation efforts. As a consequence, the reporting probability is strictly decreasing in $\widehat{m}$.

For a given manipulation choice $\widehat{m}$, the journalist's reporting probability is strictly decreasing in the mass of readers χ. An increase in readership crowds out trading profits for an individual reader and makes it less profitable for the journalist to report. The impact of the remaining three key parameters σ_d, Δ, and κ is more nuanced. For negative news ($s_J = L$), the reporting probability is strictly increasing in κ and Δ because both parameters increase the liquidity component in the equilibrium stock price. For positive news ($s_J = H$), there is an opposing force. Higher liquidity leads to a less informative stock price, which makes it harder for readers to learn about θ from the price signal s_p. The net effect is ambiguous. As a consequence, κ and Δ might have a negative impact on the reporting probability.

Since a low report $s_J = L$ always reveals $\theta = L$, the reporting probability does not depend on payoff uncertainty, σ_d. However, if the journalist reports $s_J = H$, then there are two opposing effects associated with σ_d. On the one hand, an increase in σ_d makes it riskier for readers to trade on a high signal, because their loss to informed traders, who observe θ, is large if the firm has successfully manipulated the signal. On the other hand, an increase in σ_d allows traders to learn more efficiently from the stock price, because it becomes more likely that they can learn θ from the price signal s_p. It follows that the impact of σ_d on the reporting probability is ambiguous. Fig. 3 plots π_r as a function of $\widehat{m}$ and Δ for a fixed set of model parameters. This figure emphasizes that the reporting probability is always higher for $s_F = L$ (dotted blue line) than for $s_F = H$ (solid black line). It also highlights the negative dependence of π_r on $\widehat{m}$ and the U−shaped dependence on Δ (if $s_F = H$).

Next, we move back to $t = 0$ and analyze the manager's manipulation choice. The manager chooses m to maximize the firm's expected stock price, conditional on θ, net of the linear manipulation cost $c_m m$.[18] If $\theta = L$, the manager can only bias the equilibrium stock price if the journalist covers the firm. Informed traders observe θ and are thus unaffected by s_F. We can write the expected stock price given $\theta = L$, by taking an expectation over liquidity demand u, the firm's signal s_F, and the journalist's reporting decision, which depends on the reporting cost c_r:

$$\begin{aligned}\mathbb{E}[p|\theta = L] = {} & (1 - m)[\pi_r(L, \widehat{m})\mathbb{E}_u[p(L, L, u)] \\ & + (1 - \pi_r(L, \widehat{m}))\mathbb{E}_u[p(L, \emptyset, u)]] \\ & + m[\pi_r(H, \widehat{m})\mathbb{E}_u[p(L, H, u)] \\ & + (1 - \pi_r(H, \widehat{m}))\mathbb{E}_u[p(L, \emptyset, u)]] \\ = {} & d_L + m\pi_r(H, \widehat{m})(\mathbb{E}_u[p(L, H, u)] - d_L)\end{aligned} \tag{21}$$

where we have used the previous result that the price $p\big(\theta, s_J, u\big)$ is unbiased if $s_J \in \{L, \emptyset\}$. If the firm manipulates successfully and the journalist reports $s_J = H$, the expected price is equal to:

$$\mathbb{E}_u[p(L, H, u)] = (1 + \chi)^{-1}(d_L + \chi\mathbb{E}_u[\mathbb{E}[d_\theta|\Omega_R]]). \tag{22}$$

The expected price is thus a weighted average of the expected payoffs from the perspective of informed traders and readers. It follows from Eqs. (21) and (22) that the manager's marginal benefit of manipulation is given by:

$$\frac{\partial\mathbb{E}[p|\theta = L]}{\partial m} = \chi(1 + \chi)^{-1}\pi_r(H, \widehat{m})(\mathbb{E}_u[\mathbb{E}[d_\theta|\Omega_R]] - d_L). \tag{23}$$

The marginal benefit can be decomposed into three components. First, it is more profitable for the firm to manipulate if there is a greater mass of readers in the market. As a consequence, the marginal benefit increases in the proportion of readers (relative to the sum of informed traders and readers) $\chi/(1 + \chi)$. Second, the marginal benefit is higher if the journalist is more likely to report s_F (that is, if π_r is high). Third, the marginal benefit also depends positively on $\mathbb{E}_u[\mathbb{E}[d_\theta|\Omega_R]] - d_L$, which captures the price wedge that results from successful manipulation.

The manager's marginal cost of manipulation is equal to the positive constant c_m. It follows that the optimal manipulation choice is equal to $m^* = 1$ if the marginal benefit exceeds c_m at $\widehat{m} = 1$. Similarly, $m^* = 0$ if c_m is greater than the marginal benefit at $\widehat{m} = 0$. At an interior optimum, we

[18] Note that the manager does not manipulate if $\theta = H$, so we can focus on the case $\theta = L$.

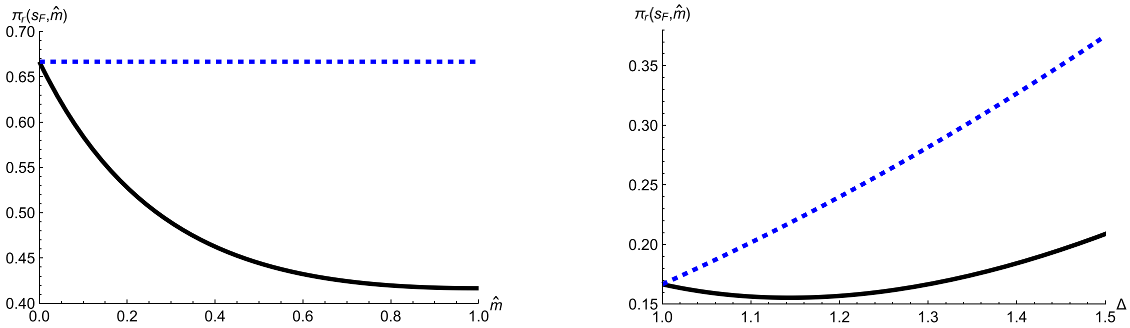

Fig. 3. Reporting probability. This figure plots the reporting probability in Proposition 1 against the conjectured manipulation intensity $\widehat{m}$ (left panel) and liquidity uncertainty Δ (right panel). Solid black line: $s_F = H$; dotted blue line: $s_F = L$. Parameters: $\sigma_d = \kappa = \chi = 1$, $\bar{c}_r = 0.25$. Left panel: $\Delta = 2$; right panel: $\widehat{m} = 0.9$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

obtain the optimal degree of manipulation by setting the marginal benefit in Eq. (23) equal to the marginal cost. In this case, m^* is implicitly given by the following equation:

$$\chi(1+\chi)^{-1}(1-\pi_p)\pi_r(H,m^*)\big(\mathbb{E}[d_\theta|s_J=H]-d_L\big)=c_m \tag{24}$$

where we have set the journalist's (and the readers') conjecture $\widehat{m}$ equal to m^*. Moreover, we have used the fact that the price signal reveals $\theta = L$ with probability π_p. In this case, readers' expected payoff is equal to d_L, and the price wedge $\mathbb{E}_u[\mathbb{E}[d_\theta|\Omega_R]] - d_L$ vanishes.

We formally show in the Appendix, that the left-hand side of Eq. (24) is decreasing in m^*. As a consequence, there is a unique manipulation equilibrium, which is fully characterized next.

Proposition 2 (Equilibrium manipulation). *There exists a unique equilibrium manipulation intensity $m^* \in [0,1]$ with the following properties:*

1. *it increases in liquidity uncertainty Δ and the transaction cost κ; it increases in readership χ if and only if $\chi < \frac{1}{2}$; it increases in payoff uncertainty σ_d if and only if π_p is sufficiently large;*
2. *it decreases in the manipulation and reporting cost coefficients c_m and $\bar{c}_r$.*

Proof. See Appendix A.3. □

The equilibrium extent of manipulation m^* maximizes the firm's expected stock price, given $\theta = L$, net of the manipulation cost. Naturally, an increase in the manipulation cost parameter c_m lowers m^*. One of our main insights is that the journalist's threat not to report the firm's signal is an additional cost of manipulation. Therefore, m^* is also decreasing in the journalist's reporting cost parameter $\bar{c}_r$. An increase in $\bar{c}_r$ makes it less likely that the journalist covers a manipulated signal. As a consequence, the stock price is less likely to respond to the signal and less profitable for the firm to manipulate.

The manager is more likely to manipulate if there is a lot of uncertainty about liquidity demand or if trading costs are high. Both parameters increase the non-fundamental component in the equilibrium stock price and render it less likely that readers infer that $\theta = L$ from the price signal. The firm is therefore more likely to manipulate because it is less likely that their manipulation efforts will be revealed by the price.

Readership χ has an ambiguous impact on m^*. On the one hand, an increase in readership lowers readers' trading profits and reduces the reporting probability. On the other hand, it also increases the effective bias in the equilibrium stock price. For small values of $\chi < \frac{1}{2}$, the second effect dominates and χ increases m^*; otherwise, the first effect dominates. Similarly, payoff uncertainty σ_d also has an ambiguous impact on the manager's manipulation efforts. On the one hand, σ_d increases the spread in payoff realizations $(d_H - d_L)$ and thus the benefit from successful manipulation. On the other hand, an increase in σ_d also makes the stock price signal more informative and makes it more likely that readers will detect a manipulated signal. Proposition 2 shows that the first effect dominates if the price signal is sufficiently informative and $\pi_p = \frac{\sigma_d}{\kappa\Delta}$ is large enough. The hump-shaped relationship of m^* with respect to χ and σ_d is depicted in Fig. 4.

3.1. Equilibrium reporting

Proposition 1 characterizes the journalist's optimal reporting policy for a given conjectured manipulation choice $\widehat{m}$. Next, we analyze this policy in equilibrium, i.e., we set the conjecture equal to m^* in Proposition 2.

Corollary 1 (Equilibrium reporting). *In the equilibrium in Proposition 2, the journalist's reporting choice has the following properties:*

1. *conditional on reporting, the firm's fundamental is more likely low than high:*
$$\mathbb{P}(\theta = L|\mathcal{D}_r = 1) \geq \mathbb{P}(\theta = H|\mathcal{D}_r = 1) \tag{25}$$
2. *unconditionally, the journalist's report is more likely high than low:*
$$\mathbb{P}(s_J = H) \geq \mathbb{P}(s_J = L) \tag{26}$$
3. *the fraction of firm announcements that are positive is more than the fraction of journalist reports that are positive. Equivalently, the fraction of firm announcements that are negative is less than the fraction of journalist reports that are negative.*
$$\mathbb{P}(s_F = H|\mathcal{D}_r = 1) \leq \mathbb{P}(s_F = H) \tag{27}$$

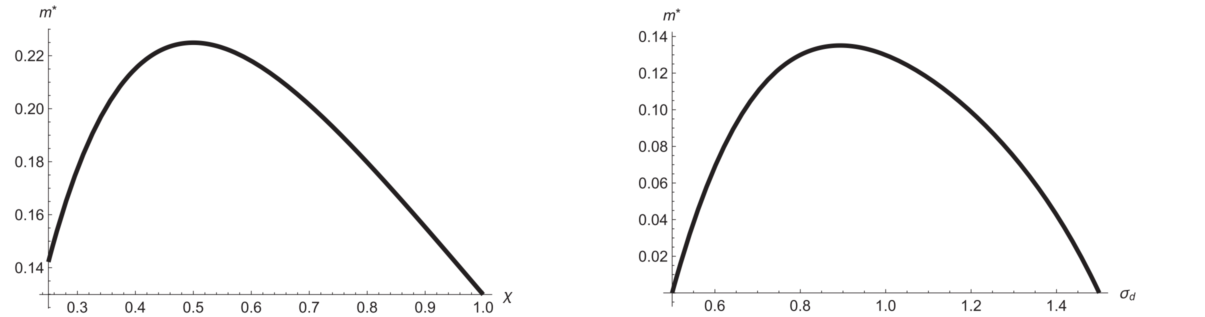

Fig. 4. Manipulation intensity. This figure plots the equilibrium manipulation intensity in Proposition 2 against the mass of readers χ (left panel) and payoff uncertainty σ_d (right panel). Parameters: $\kappa = 1$, $\Delta = 2$, $\bar{c}_r = c_m = 0.25$. Left panel: $\sigma_d = 1$; Right panel: $\chi = 1$.

$$\mathbb{P}(s_F = L|\mathcal{D}_r = 1) \geq \mathbb{P}(s_F = L) \tag{28}$$

4. *the reporting probability $\pi_r(s_F, m^*)$ has the following properties:*
   (a) *it is increasing in the manipulation cost c_m if $s_F = H$;*
   (b) *the comparative statics with respect to the remaining parameters are the same as in Proposition 1.*

Proof. See Appendix A.4. □

We discuss Corollary 1 by way of a numerical example. Consider 100 realizations of θ (the fundamental), 50 of which are positive ($\theta = H$) and 50 of which are negative ($\theta = L$). Suppose that 20 of the 50 negative news events are successfully manipulated ($m^* = 40\%$). The journalist therefore observes 70 positive announcements ($s_F = H$) and 30 negative announcements ($s_F = L$). Suppose further that the journalist reports on 40% of positive announcements and 60% of negative announcements. We should therefore expect the journalist to report $40\% \times 70 = 28$ positive announcements (of which $40\% \times 50 = 20$ are actually positive and 8 are actually negative) and $60\% \times 30 = 18$ negative announcements.

The first result in Corollary 1 reflects the previous finding in Proposition 1 that the journalist is more likely to report if $s_F = L$. As a result, conditional on reporting ($\mathcal{D}_r = 1$) it is more likely that the underlying fundamental is low ($\theta = L$). In our example, the journalist ultimately reported $28 + 18 = 46$ announcements, of which 20 are *actually* positive and $18 + 8 = 26$ are *actually* negative.

The second result says that the journalist reports more positive announcements *on average*. In our example, she reports 28 positive announcements and 18 negative ones. It is important to note that there are two opposing forces. On the one hand, there are more positive firm announcements than negative ones due to manipulation. On the other hand, the journalist is less likely to report positive firm announcements in order to reduce the number of manipulated reports. In our model, the first effect dominates because the journalist understands that allowing some manipulated announcements to be reported does not fully harm readers, as readers also rationally expect the correct level of manipulation when they trade.[19]

[19] We formally show that the first effect dominates in Eq. (A.24) in the Appendix.

The third result is more applied in nature. An econometrician could observe firm announcements in general (s_F) and firm announcements that are reported on by journalists in particular (s_J). In our example, 70% of firm announcements are positive and 30% are negative, but among those that are reported, $28/46 \approx 61\%$ are positive and $18/46 \approx 39\%$ are negative.

In addition to the comparative statics in Proposition 1, we find that the firm's manipulation cost c_m positively affects the journalist's equilibrium reporting policy. Intuitively, an increase in this cost lowers the incentive to manipulate. As a consequence, it increases readers' trading profits and makes it more attractive for the journalist to report the firm's announcement.

3.2. Equilibrium stock prices

Next, we revisit the firm's equilibrium stock price $p = p\big(\theta, s_J, u\big)$, which is formally stated in Eq. (9), and characterize its properties given the equilibrium manipulation choice m^*. As we have shown above, for $s_J \in \{L, \emptyset\}$ the equilibrium stock price is unbiased and equal to d_θ, on average. If, however, the journalist reports $s_J = H$, then the price systematically deviates from d_θ because readers might trade on a manipulated signal. More specifically, readers trade less aggressively on this signal if they do not learn θ from the equilibrium stock price. In this case, their asset demands are lower than those of informed traders who observe $\theta = H$, but higher than those of informed traders who observe $\theta = L$. This rational response to potentially manipulated information is then reflected in the equilibrium stock price and leads to mispricing. Below, we introduce two measures of mispricing.

The asset payoff d_θ is realized at $t = 3$ and can therefore be interpreted as the firm's long-run stock price. We formally define the difference between d_θ and the expected short-run stock price p as the expected *price drift* τ.

Definition 2 (Price drift). Price drift is defined as the difference between the asset's payoff and the expected price:

$$\tau\big(s_J, \theta\big) \equiv d_\theta - \mathbb{E}_u\big[p|\theta, s_J\big]. \tag{29}$$

This measure of mispricing depends on the realized value of the journalist's report and the firm's fundamental. It captures the extent to which the asset is overpriced or underpriced, at $t=2$. We integrate over liquidity demand u to obtain the average, or expected, drift in the stock price.

Corollary 2 (Equilibrium price drift). *If the journalist reports $s_J=L$ or $s_J=\emptyset$, then the expected price drift is equal to zero. If the journalist reports $s_J=H$, then the expected drift is given by:*

$$\tau(H,\theta)=\frac{\chi}{1+\chi}(1-\pi_p)\times\left[d_\theta-\left(\frac{1}{1+m^*}d_H+\frac{m^*}{1+m^*}d_L\right)\right],\tag{30}$$

where $m^*\in[0,1]$ is the equilibrium manipulation choice given in Proposition 2.

Proof. See Appendix A.5. □

Conditional on a positive report $s_J=H$, the stock price systematically drifts up between $t=2$ and $t=3$ if $\theta=H$ and drifts down if $\theta=L$. The drift is smaller when the price is more informative (π_p is larger) and larger when there are more readers (χ is larger). The dependence on the firm's equilibrium manipulation effort m^* is always increasing as readers rationally lower their trading aggressiveness on $s_J=H$.

Next, we integrate over $\theta\in\{L,H\}$ and $s_J\in\{L,H,\emptyset\}$ and compute the unconditional expectation of p:

$$\mathbb{E}[p]=\sum_{\theta,s_J}\mathbb{P}(\theta,s_J)\mathbb{E}_u\left[p\left(\theta,s_J,u\right)\right]=\mathbb{E}[d_\theta].\tag{31}$$

Evidently, the price is an unbiased estimator of the asset payoff. If the journalist reports a high signal, the expected price drift is equal to zero. Intuitively, the deviations from the expected payoff "washout" on average and the firm's manipulation efforts do not influence the unconditional expectation of p, which is standard in the "signal-jamming" literature. Although the firm's manipulation efforts and the journalist's reporting do not lead to a systematic ex ante bias in the stock price, they do influence the price's accuracy. To formalize this measure of mispricing in the stock, we introduce the concept of price quality below.

Definition 3 (Price quality). Price quality is defined as the negative expected squared deviation of the price from the asset's payoff:

$$\Lambda(\theta,s_J)=-\mathbb{E}_u\left[(d_\theta-p)^2|\theta,s_J\right].\tag{32}$$

Our measure of price quality $\Lambda(\theta,s_J)$ corresponds to the mean-squared error of the equilibrium stock price and has been considered in the existing literature (see, e.g., Banerjee et al., 2018; Frenkel et al., 2020). It is maximized at $\Lambda=0$ if the equilibrium stock price is always equal to the asset payoff. In Eq. (32), we define price quality for a given realization of $\theta\in\{L,H\}$ and $s_J\in\{L,H,\emptyset\}$, but below we will also analyze

$$\lambda(s_J)\equiv\mathbb{E}_\theta\left[\Lambda(\theta,s_J)\right]\tag{33}$$

which integrates over θ and only depends on the journalist's reported signal.

If the journalist's report is equal to $s_J=L$ or if she does not report, price quality is given by

$$\Lambda\left(\theta,s_J\right)=\begin{cases}-\frac{\kappa^2\sigma_u^2}{(1+\chi)^2} & \text{if } s_J=L\\ -\kappa^2\sigma_u^2 & \text{if } s_J=\emptyset.\end{cases}\tag{34}$$

Eq. (34) follows by substituting Eq. (9) into Eq. (32). In this expression, $\sigma_u^2=\frac{\Delta^2}{3}$ captures the noise in the stock price caused by liquidity trading. In these two cases, the stock price is more accurate with reporting, because reporting leads to a greater mass of informed traders in the market and lowers the relative impact of nonfundamental trading.

If the journalist reports $s_J=H$, the signal might be manipulated, and readers rationally base their demands on the conjectured manipulation intensity $\widehat{m}$. As before, we use the expression for p in Eq. (9) and the definition of price quality to obtain:

$$\Lambda(\theta,H)=\begin{cases}-\frac{\kappa^2\sigma_u^2+(1-\pi_p)\left[4\chi^2\sigma_d^2\frac{\widehat{m}^2}{(1+\widehat{m})^2}+4\chi\sigma_d^2\frac{\widehat{m}}{1+\widehat{m}}\right]}{(1+\chi)^2} & \text{if } \theta=H\\ -\frac{\kappa^2\sigma_u^2+(1-\pi_p)\left[4\chi^2\sigma_d^2\frac{1}{(1+\widehat{m})^2}+4\chi\sigma_d^2\frac{1}{1+\widehat{m}}\right]}{(1+\chi)^2} & \text{if } \theta=L\end{cases}\tag{35}$$

where π_p measures the probability with which the price reveals θ to readers. If the price is always perfectly revealing ($\pi_p\to 1$) or if there are no readers ($\chi=0$), price quality given $s_J=H$ is equal to price quality if $s_J=L$.

Corollary 3 (Equilibrium price quality). *In the equilibrium in Proposition 2, price quality $\lambda(s_J)$ has the following properties:*

1. *it is highest when $s_J=L$, followed by the case in which $s_J=H$, and lowest if $s_J=\emptyset$:*

$$\lambda(L)\geq\lambda(H)>\lambda(\emptyset)\tag{36}$$

2. *it increases in readership χ and decreases in liquidity uncertainty Δ and the transaction cost κ;*
3. *if $s_J=H$, it decreases in payoff uncertainty σ_d if and only if π_p is sufficiently large; it increases in the manipulation and reporting cost coefficients c_m and $\bar{c}_r$.*

Proof. See Appendix A.6. □

Corollary 3 shows that stock prices are always more efficient when the journalist reports, even though $s_J=H$ could be a manipulated signal. Intuitively, readers take the possibility into account that the signal is manipulated and trade less aggressively. In equilibrium, the stock price is closer to d_θ than in the scenario in which the journalist does not report. Prices are most efficient if the journalist reports $s_J=L$, because this signal is always accurate. For the same reason, price quality is increasing in the mass of readers χ.

Liquidity uncertainty Δ and the transaction cost κ reduce price efficiency because they increase the degree of nonfundamental variation in the equilibrium stock price. If the journalist reports $s_J=H$, price quality is also affected by payoff uncertainty σ_d. In particular, higher payoff uncertainty reduces price quality if and only the price is sufficiently likely to reveal d_θ to readers. An increase in the manipulation cost c_m and the reporting cost coefficient $\bar{c}_r$

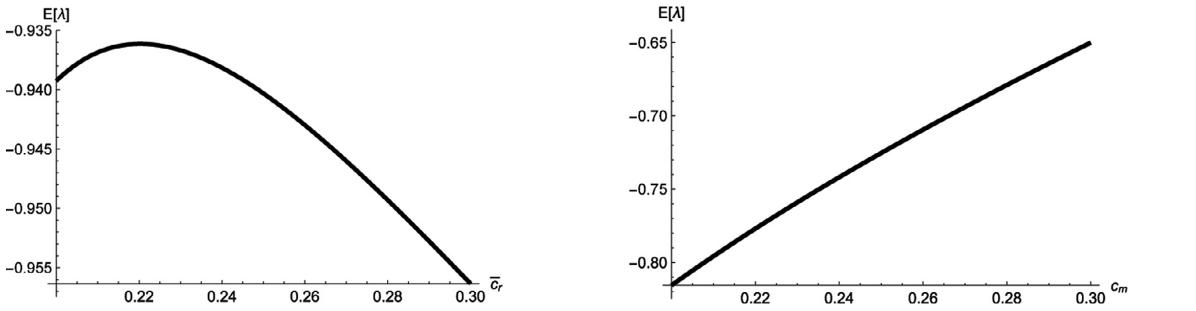

Fig. 5. Unconditional expectation of price quality. This figure plots the unconditional expectation of price quality $\lambda(s_J)$ against the cost coefficients $\bar{c}_r$ (left panel) and c_m (right panel). Parameters: $\kappa = \sigma_d = \chi = 1$, and $\Delta = 2$. Left panel: $c_m = 0.15$; Right panel: $\bar{c}_r = 0.2$.

lowers the firm's manipulation intensity and leads to a more precise public signal, which in turn improves price quality.

Fig. 5 plots the unconditional expectation of λ against the cost coefficients $\bar{c}_r$ and c_m. The left panel shows that an increase in the journalist's reporting cost has an ambiguous impact on expected price quality $\mathbb{E}[\lambda]$. On the one hand, an increase in $\bar{c}_r$ lowers the reporting probability because it becomes less likely that a given announcement is sufficiently interesting relative to the journalist's opportunity cost. According to the results in Corollary 3, this channel reduces price quality. On the other hand, a decrease in the reporting probability also lowers the firm manager's incentive to manipulate (see Proposition 2). Therefore, the underlying firm announcement is less likely to be manipulated, and the journalist's report becomes more accurate. The second effect, which is positive, dominates for small values of $\bar{c}_r$. The right panel confirms that an increase in the manager's direct manipulation cost c_m always leads to a more efficient stock price. An increase in this cost coefficient lowers the level of manipulation *and* makes it more likely that the journalist reports. Hence, the impact of c_m on $\mathbb{E}[\lambda]$ is unambiguously positive.

4. Investigative reporting

In this section, we consider an extension of the main model. We assume that, at $t = 1$, the journalist receives a private signal $y \in \{\theta, \emptyset\}$, which reveals the true fundamental θ with probability $\alpha \in [0, 1]$. Except for this additional signal, the economic framework is the same as before. This extension collapses to the main model if $\alpha = 0$. The journalist's private signal can be interpreted as the outcome of her investigative efforts, and α governs the precision of her signal or the intensity of her efforts. At first, we take α as exogenous, but at the end of the section we analyze the journalist's optimal choice of it.

4.1. Exogenous investigative effort

We rewrite the journalist's report as $s_J \in \{(s_F, y), \emptyset\}$, so that she adds her private signal to the firm's signal if she chooses to report. If the firm reports $s_F = L$, the journalist's signal is redundant. But if $s_F = H$, the journalist's signal helps readers to gauge the veracity of the firm's signal.

We can follow the same steps as in the main model and write the equilibrium stock price as follows:

$$p(\theta, s_J, u) = \begin{cases} \frac{1}{1+\chi}(d_\theta + \chi \mathbb{E}[d_\theta|\Omega_R] + \kappa u) & \text{if } s_J \neq \emptyset \\ d_\theta + \kappa u & \text{if } s_J = \emptyset \end{cases} \tag{37}$$

where $\Omega_R = \{s_J, p\}$ and s_J incorporates both s_F and y, if the journalist reports. More specifically, if the journalist's investigative efforts are successful and she reports, then her readers learn θ and their payoff expectation equals d_θ. If the journalist's efforts are unsuccessful, reader expectations are identical to those in the main model.

The firm's expected stock price, conditional on $\theta = L$, is given by:

$$\mathbb{E}[p|\theta = L] = d_L + m(1-\alpha)\pi_r((H, \emptyset), \widehat{m}) \times (\mathbb{E}_u[p(L, (H, \emptyset), u] - d_L), \tag{38}$$

where the reporting probability $\pi_r((H, \emptyset), \widehat{m})$ is conditional on $s_F = H$ and $y = \emptyset$. The accuracy of the journalist's private, investigative signal therefore diminishes the marginal benefit of manipulation, because the firm can only manipulate the stock price if the journalist's investigation efforts are unsuccessful, which happens with probability $1 - \alpha$.

To find the equilibrium manipulation intensity m^*, we compare the marginal benefit $\frac{\partial \mathbb{E}[p|\theta=L]}{\partial m}$ to the marginal cost c_m, as before. It follows from the expression in Eq. (38) that the manager's problem is equivalent to that of the baseline model, if we rewrite the manipulation cost as $c_m/(1-\alpha)$. Hence, all of our main results continue to hold in this extension. Moreover, an increase in the journalist's investigative efforts makes it costlier for the firm to manipulate. As a consequence, m^* decreases and π_r increases in α.

Corollary 4 (Exogenous investigative effort). *An increase in the journalist's detection probability* $\alpha \in [0, 1]$ *leads to a decrease in equilibrium manipulation* m^* *and an increase in the reporting probability* π_r.

Proof. See Appendix A.7. □

4.2. Optimal investigative effort

In the preceding analysis, we assumed that the journalist costlessly observed a signal y that perfectly revealed

θ with some exogenous probability α, which we interpret as investigative effort on the part of the journalist. Evidently, effort is not costless in reality. The journalist must expend time and resources to verify claims made by the firm. In an alternative setup, the journalist pays a cost $C(\alpha)$, with $C \geq 0$ and $C', C'' > 0$, to increase the precision of y. Now in equilibrium, if $s_F = L$, then the journalist correctly infers that $\theta = L$; hence, as before, the journalist chooses $\alpha = 0$. If, however, $s_F = H$, then y is valuable and the journalist might find it profitable to expend investigative effort so that she can observe θ with some probability. In this sense, the journalist investigates potentially biased announcements ($s_F = H$) but not announcements that are known to be unbiased ($s_F = L$).

To solve for the optimal α, we first define the expected utility of the journalist, given $s_F = H$, as the expected gain in reader utility net of the private reporting cost:

$$U_J(H) = \mathbb{E}[(\mathbb{E}[U_R|s_J] - c_r)\mathbf{1}_{c_r \leq \mathbb{E}[U_R|s_J]}] \tag{39}$$

where the first expectation is taken over the uniformly distributed reporting cost c_r and the journalist's private signal y. For a given realization of y, we can write her expected utility as $\frac{1}{2\bar{c}_r}\mathbb{E}[U_R|s_J]^2$, which is proportional to the (squared) expected trading profits of her readers, given $s_J = (H, y)$.

For a given $\alpha \in [0, 1]$ and $s_F = H$, the journalist observes θ with probability α and observes no additional information otherwise. Therefore, she observes $\theta = H$ with probability $\alpha\frac{1}{1+\widehat{m}}$ and $\theta = L$ with probability $\alpha\frac{\widehat{m}}{1+\widehat{m}}$. Otherwise, she observes $y = \emptyset$.

The journalist's expected utility is thus given by:

$$U_J(s_F = H) = \frac{1}{2\bar{c}_r}\Big(\alpha\mathbb{E}[U_R|s_J = (H, \theta)]^2 + (1-\alpha)\mathbb{E}[U_R|s_J = (H, \emptyset)]^2\Big) \tag{40}$$

where expected trading profits for readers follow from the results in Lemma 1. If the journalist's investigation efforts are successful, readers learn θ perfectly, and their trading profits are given by

$$\mathbb{E}[U_R|s_J = (H, \theta)] = \frac{\kappa\Delta^2}{6(1+\chi)^2}. \tag{41}$$

Trading profits in this case are thus identical to those in the baseline model if the journalist reports $s_F = L$. If the journalist's efforts are unsuccessful, this extension collapses to the baseline model, and we obtain:

$$\mathbb{E}[U_R|s_J = (H, \emptyset)] = \frac{\kappa\Delta^2}{6(1+\chi)^2} - \frac{2\widehat{m}\sigma_d^2(\Delta\kappa - \sigma_d)}{\Delta\kappa^2(1+\widehat{m})^2(1+\chi)^2}. \tag{42}$$

Below we formally summarize our findings for this extension.

Corollary 5 (Optimal investigative effort). *In the extension with endogenous investigative effort, we find that:*

1. *If $s_F = L$, the journalist never exerts investigative effort.*
2. *If $s_F = H$, her effort choice is increasing in the conjectured manipulation intensity $\widehat{m}$.*
3. *A decrease in the marginal cost of investigation leads to a higher reporting probability.*

Proof. See Appendix A.8. □

As a result, the journalist is more likely to investigate potentially manipulated signals. Moreover, she is more likely to investigate the firm's signal if she conjectures a higher manipulation intensity. A decrease in the journalist's marginal cost of investigation leads to an increase in the optimal effort choice α. As a consequence, the journalist is more likely to learn the firm's true fundamental through her private investigation. Expected profits for readers are higher in this case, and the journalist's reporting probability increases.

5. Empirical implications

We discuss empirical implications that relate to (1) firm behavior, (2) journalist behavior, and (3) stock prices. Below, we provide the details of each implication and discuss how it could be tested empirically.

5.1. Firm manipulation

The first set of predictions relates to firm managers and their incentives to manipulate corporate announcements. As mentioned before, we interpret manipulation not necessarily as outright fraud. Instead, we have in mind that managers present facts in such a way that the journalist and market walk away with the interpretation of facts that the manager wants, not the interpretation they would have come to had they discovered the facts themselves.

Manipulation is inherently difficult to measure in practice. However, several recent empirical studies have developed indices that measure the extent to which the truth of a financial disclosure has been obfuscated. For instance, Li (2008) develops a *Fog Index*, and Loughran and McDonald (2014) develop a *Readability Index*.

The novel mechanism in our model is that the presence of a financial journalist changes firms' incentives to manipulate their announcements. To highlight this effect, we assume that readers are the only traders that are exposed to the firm's manipulated signal. However, the mechanism is robust to an alternative assumption that an additional group of traders are affected by the firm's announcement even in the absence of reporting. Under this assumption, there would be a positive degree of manipulation with and without media reporting, and the presence of a journalist would *amplify* the firm manager's incentive to manipulate. Our novel mechanism is thus purely driven by the journalist's core function to screen firm announcements on the basis of their informational content. This result is presented in our first empirical prediction:

Prediction 1. Manipulation increases following an exogenous increase in media coverage. The effect is stronger when (i) a greater fraction of traders are readers and (ii) stock prices are less informative.

Eq. (23) represents the firm's marginal benefit of manipulation. The firm manager is more likely to manipulate an announcement when the journalist is more likely to report it. How much *more* likely depends on the ratio of institutional traders to newspaper readers. Suppose there are many more institutional traders than newspaper readers. Then a small increase in the likelihood of reporting makes the manager slightly more likely to manipulate. Much of the market is already informed, so whether or not the journalist reports is of less concern to the manager. If, however, there are many more newspaper readers than institutional traders, a small increase in the likelihood of reporting makes the manager much more likely to manipulate.

Similarly, the impact of an exogenous increase in media coverage is mitigated by the informational content of stock prices. If stock prices are particularly informative about the firm's fundamental, it is more likely that a manipulated announcement is detected. As a result, the firm manager is less likely to manipulate.

Prediction 2. Firms that have yet to receive media coverage manipulate less than those that have already received media coverage.

Firms that are either household names or have operated for decades have almost certainly received media coverage in the past. Such firms expect that their announcements will be more likely to be reported, relative to firms that have never been reported on. We therefore posit that (according to Proposition 2) the firm that has yet to receive media coverage will manipulate *less*. Because reporting on the firm is already so costly, the firm's announcements must contain very little manipulation to entice the journalist to report on them.

Prediction 3. Manipulation is lower when the journalist covering the firm is more skilled or her opportunity cost is higher.

Our third prediction follows from Corollary 4 and Proposition 2, Item 2. A more highly skilled journalist is more likely to detect the firm's manipulation efforts. As a consequence, the firm manager's anticipated marginal benefit of manipulation is lower. One way to test this prediction is to construct a novel measure of *media skill* and combine it with existing proxies for manipulation. The second part of the prediction follows from the fact that manipulated corporate announcements reach a greater mass of traders if there is media reporting. An increase in the journalist's opportunity cost makes it less likely that a certain announcement is covered and therefore lowers the firm manager's incentives to manipulate. Empirically, the journalist's opportunity cost might be affected by the total number of corporate announcements on a given day or by attention-grabbing events (Peress and Schmidt, 2020).

5.2. Media reporting

Our second unique set of predictions relate to the equilibrium choice of the journalist on what firm announcements to cover. Our model provides several predictions for this relationship.

Prediction 4. Negative announcements are more likely to be reported than positive ones.

This prediction follows from Corollary 1, Item 3. In our model, the journalist's decision to cover a given corporate announcement depends on the consequences for readers' trading profits. Positive announcements might have been manipulated by the firm manager and are therefore inherently less accurate than negative announcements. As a result, the journalist is more likely to cover a given announcement if it is negative. An empirical test of this prediction requires (i) a comprehensive data set of firm announcements and corresponding media reports, and (ii) a measure of the "tone" of each announcement.

Prediction 5. Announcements that are more manipulated are less likely to be reported.

If the journalist expects a higher extent of manipulation in a certain firm announcement, she anticipates lower trading profits for her readers and is therefore less likely to cover the announcement. Formally, this prediction is stated in Proposition 1, Item 3. There are several ways to test this prediction empirically. First, one can analyze the cross-section of corporate announcements and the associated degree of media coverage. Second, one could also compare media coverage for corporate financial news, such as earnings announcements, to coverage for other financial news, such as unemployment statistics. Because earnings announcements can be more easily manipulated or obfuscated, our model predicts that their coverage should be (1) lower and (2) more heavily tilted towards negative news.

Prediction 6. Announcements are more likely to be reported if the journalist is more skilled.

A more highly skilled journalist is better able to protect her readers from manipulated corporate announcements. This prediction can be seen in Corollary 5, Item 3 in Section 4, where we extend the baseline model and allow for investigative reporting. Skilled reporters therefore reduce the firm's incentive to manipulate. In the data this type of skill could be proxied for by a journalist's tenure or past accuracy.

5.3. Stock prices

The third set of predictions relates to the relationship between media coverage and asset prices. This has been the main focus of previous empirical work on journalism in finance. A well-established finding in this literature is that media coverage oftentimes causes a temporary overreaction in stock prices (see, e.g., Vega, 2006; Tetlock, 2007; Ahern and Sosyura, 2014).

Our model also demonstrates the existence of short-term mispricing associated with media coverage. Readers trade based on an informative, but also manipulated, report.

Prediction 7. An exogenous increase in media coverage leads to an increase in the magnitude of the post-announcement price drift. This effect is stronger when (1) the news is positive and (2) a greater share of traders are readers.

This effect can be seen from Eq. (30) in Corollary 2. Without media coverage, mis-pricing is purely due to liquidity trading and thus unrelated to the firm's announcement. If the journalist covers a given firm announcement, however, readers potentially trade on a manipulated signal. As a result, the magnitude of the price drift increases for positive firm announcements. Because our readers are assumed to be fully rational about the extent of manipulation in the report, the price is unbiased on average. Sometimes the price is underpriced relative to the truth, and sometimes it is overpriced.

Another interesting result in the empirical literature is that media coverage of stale news also affects traders (see, e.g., Huberman and Regev, 2001; Tetlock, 2007). These papers interpret their findings as indicative of an investor behavioral bias, such as limited attention. Our model suggests that the reporting of stale news is an *optimal* decision of a journalist. Hence, the journalist will report on news only if she believes that her readers would still benefit from the information. This implies, for example, that reporting of stale information will be higher if the past stock response to the announcement was low.

Finally, we show that the media has a causal impact on the trading behavior of its readers.

Prediction 8. Media coverage leads to a positive correlation in the orders by retail traders and institutional traders. This correlation is stronger for negative news.

In our model, readers trade more aggressively if the journalist reports negative news, because this signal is more accurate. In this case, readers and informed investors trade on the same signal. Empirically, this effect could be measured by investigating the correlation between trades of retail traders (who proxy for readers) and institutional investors (who proxy for informed traders). For positive news, our model predicts that readers sometimes face manipulated signals, which in turn lowers their trading aggressiveness. As a result, we predict that there is more disagreement between retail investors and institutional investors following good news.

6. Conclusion

Financial journalists are part of the ecosystem of agents who take the vast amount of publicly available financial information and process this information to their readers. We consider a model in which the role of the financial journalist is to identify, to her readers, the most important financial information put out by the firm. The resulting equilibrium demonstrates the type of news that a strategic journalist will choose to report as well as how her presence affects her readers' ability to trade, the incentive of firms to manipulate their announcements, and equilibrium stock prices. We have enumerated a plethora of predictions that should be readily testable by empirical researchers.

Appendix A. Proofs

For ease of exposition, we will use the following notation for some of the proofs:

1. For the asset payoff we set

$$d_H = \mu + \sigma_d \tag{A.1}$$

$$d_L = \mu - \sigma_d \tag{A.2}$$

for constants μ and $\sigma_d > 0$. Note that $d_H - d_L = 2\sigma_d > 0$.

2. We define the error in the readers' expectation by

$$\varepsilon_\eta = \mathbb{E}[d_\theta | \Omega_R] - d_\eta \tag{A.3}$$

where $\Omega_R = \{s_J, p\}$ and $\eta \in \{H, L\}$.

3. If the price signal is uninformative and $s_J = H$, we can compute the readers' conditional payoff expectation using Bayes' rule

$$\mathbb{E}[d_\theta | s_J = H] = \frac{1}{1+\widehat{m}}(\mu + \sigma_d) + \frac{\widehat{m}}{1+\widehat{m}}(\mu - \sigma_d) = \mu + \widetilde{\sigma_d} \tag{A.4}$$

where $\widetilde{\sigma_d} = \sigma_d \frac{1-\widehat{m}}{1+\widehat{m}} \in [0, \sigma_d]$.

As shown in the main text, the readers' expected payoff given $s_J = L$ is equal to $\mathbb{E}[d_\theta | s_J = L] = \mu - \sigma_d$.

A1. Proof of Lemma 1

We compute the readers' expected profits given $s_J = s_F$. First, we use the definition of U_R in Eq. (3) and apply the law of iterated expectations to get:

$$\mathbb{E}[U_R | s_F] = \mathbb{E}\left[x_R(d_\theta - p) - \frac{\kappa}{2} x_R^2 \middle| s_F \right] = \frac{\kappa}{2}\mathbb{E}\left[x_R^2 \middle| s_F \right] \tag{A.5}$$

where $x_R = \kappa^{-1}\mathbb{E}[d_\theta - p | s_F, p]$. Next, we plug in the expression for p given in Eq. (9) to rewrite squared reader demand as:

$$x_R^2 = \frac{\left(\varepsilon_\eta - \kappa u\right)^2}{\kappa^2(1+\chi)^2} = \frac{\varepsilon_\eta^2 - 2\varepsilon_\eta \kappa u + \kappa^2 u^2}{\kappa^2(1+\chi)^2} \tag{A.6}$$

where ε_η is defined in Eq. (A.3). Next, we take an expectation over u and ε_η for a given $s_F \in \{L, H\}$:

1. For $s_F = L$, we have $\varepsilon_\eta = 0$ and therefore

$$\mathbb{E}[U_R | s_F = L] = \frac{\kappa \mathbb{E}[u^2]}{2(1+\chi)^2} = \frac{\kappa \Delta^2}{6(1+\chi)^2} = \frac{\kappa \sigma_u^2}{2(1+\chi)^2}. \tag{A.7}$$

2. For $s_F = H$, ε_η is either equal to 0, if p is perfectly revealing, or equal to $\mu + \widetilde{\sigma_d} - d_\theta$, if it is not. Next,

$$\mathbb{E}[U_R | s_F = H] = \frac{1}{1+\widehat{m}}\mathbb{E}[U_R | s_F = H, \theta = H] + \frac{\widehat{m}}{1+\widehat{m}}\mathbb{E}[U_R | s_F = H, \theta = L]. \tag{A.8}$$

If $\theta = H$, then $\varepsilon_H = 0$ with probability π_p and $\varepsilon_H = -2\sigma_d m/(1+m)$ with probability $1 - \pi_p$, and hence

$$\mathbb{E}[U_R | s_F = H, \theta = H] = \frac{\kappa^2\sigma_u^2 + (1-\pi_p)\left[(\widetilde{\sigma}_d - \sigma_d)^2 + 2\sigma_d(\widetilde{\sigma}_d - \sigma_d)\right]}{2\kappa(1+\chi)^2}. \tag{A.9}$$

If $\theta = L$, then $\varepsilon_L = 0$ with probability π_p and $\varepsilon_L = 2\sigma_d/(1+m)$ with probability $1 - \pi_p$, and hence

$$\mathbb{E}[U_R|s_F = H, \theta = L] = \frac{\kappa^2\sigma_u^2 + (1-\pi_p)\left[(\widetilde{\sigma}_d + \sigma_d)^2 - 2\sigma_d(\widetilde{\sigma}_d + \sigma_d)\right]}{2\kappa(1+\chi)^2}. \quad (A.10)$$

Combining these two expressions and using $\widetilde{\sigma}_d = \sigma_d \frac{1-\widehat{m}}{1+\widehat{m}}$ leads to:

$$\mathbb{E}[U_R|s_F = H] = \frac{\kappa\Delta^2}{6(1+\chi)^2} - \frac{(1-\pi_p)2\widehat{m}\sigma_d^2}{\kappa(\widehat{m}+1)^2(\chi+1)^2}, \quad (A.11)$$

with $\pi_p = \frac{\sigma_d}{\kappa\Delta}$.

A2. Proof of Proposition 1

The expression for $\pi_r(s_F, \widehat{m})$ follows directly from Lemma 1 and the definition of the reporting probability in Eq. (18). Note that $\pi_r(L, \widehat{m}) = \pi_r(H, 0)$ and that $\frac{\partial\pi_r(H,\widehat{m})}{\partial\widehat{m}} < 0$. It immediately follows that $\pi_r(H, \widehat{m}) < \pi_r(L, \widehat{m})$ for $\widehat{m} > 0$. It also follows that the reporting probability is maximized at $\widehat{m} = 0$ and minimized at $\widehat{m} = 1$. The parametric assumption on $\bar{c}_r$ in Eq. (19), ensures that $\pi_r < 1$. Furthermore, $\pi_r(H, 0) > 0$ because the parametric assumption in Eq. (7) ensures that $\Delta > \frac{\sigma_d}{\kappa}$.

If $s_J = L$, it immediately follows, from the expression for $\pi_r(L, \widehat{m})$, that the reporting probability is decreasing in χ and $\bar{c}_r$ and that it is increasing in κ and Δ.

If $s_J = H$, we can factor out $\frac{1}{(1+\chi)^2}$ from the expression for $\pi_r(H, \widehat{m})$ to see that the reporting probability is decreasing in χ and $\bar{c}_r$. The comparative statics for σ_d, Δ, and κ are as follows:

1. Comparative statics with respect to σ_d:
$$\frac{\partial\pi_r(H,\widehat{m})}{\partial\sigma_d} = \frac{\sigma_d(3\sigma_d - 2\Delta\kappa)}{2\Delta\kappa^2(\chi+1)^2\bar{c}_r} \quad (A.12)$$
which is positive if and only if $\pi_p = \frac{\sigma_d}{\kappa\Delta} > \frac{2}{3}$.
2. Comparative statics with respect to Δ:
$$\frac{\partial\pi_r(H,\widehat{m})}{\partial\Delta} = \frac{2\Delta^3\kappa^3 - 3\sigma_d^3}{6\Delta^2\kappa^2(\chi+1)^2\bar{c}_r} \quad (A.13)$$
which is positive if and only if $\pi_p = \frac{\sigma_d}{\kappa\Delta} < \sqrt[3]{\frac{2}{3}}$.
3. Comparative statics with respect to κ:
$$\frac{\partial\pi_r(H,\widehat{m})}{\partial\kappa} = \frac{\Delta^3\kappa^3 + 3\Delta\kappa\sigma_d^2 - 6\sigma_d^3}{6\Delta\kappa^3(\chi+1)^2\bar{c}_r} \quad (A.14)$$
which is positive if and only if $\pi_p = \frac{\sigma_d}{\kappa\Delta} < \frac{1}{x_0}$ where $x_0 \approx 1.29$ is the real solution to the cubic equation $-6 + 3x_0 + x_0^3 = 0$.

A3. Proof of Proposition 2

We start from Eq. (23) in the main text. With probability π_p, the stock price is perfectly revealing and $\mathbb{E}[d_\theta|\Omega_R] = d_L$. With probability $1-\pi_p$, the stock price is uninformative and the readers' expected payoff is given by $\mathbb{E}[d_\theta|s_J = H]$. It follows that we can rewrite the marginal benefit as:

$$\frac{\chi}{1+\chi}\pi_r(H, m^*)(1-\pi_p)\left(\mathbb{E}[d_\theta|s_J = H] - d_L\right) \quad (A.15)$$

where π_r is characterized in Proposition 1 and $\mathbb{E}[d_\theta|s_J = H] = \frac{1}{1+m^*}d_H + \frac{m^*}{1+m^*}d_L$. It follows that the marginal benefit is decreasing in m^*. As a consequence, it is maximized at $m^* = 0$:

$$\bar{c}_m \equiv \frac{2\chi(1-\pi_p)\sigma_d}{1+\chi}\pi_r(H, 0) \quad (A.16)$$

and minimized at $m^* = 1$:

$$\underline{c}_m \equiv \frac{\chi(1-\pi_p)\sigma_d}{1+\chi}\pi_r(H, 1). \quad (A.17)$$

It immediately follows that $m^* = 0$ if $c_m > \bar{c}_m$ and $m^* = 1$ if $c_m < \underline{c}_m$. For $c_m \in [\underline{c}_m, \bar{c}_m]$, we can find the optimal m by setting the marginal cost equal to the marginal benefit:

$$c_m = \frac{2\chi(1-\pi_p)\sigma_d}{1+\chi}\frac{\pi_r(H, m^*)}{1+m^*} \equiv G(m^*), \quad (A.18)$$

which yields a unique equilibrium (due to the intermediate value theorem) because (i) $G(m^*)$ is decreasing in m^*, (ii) $G(0) > c_m$, and (iii) $G(1) < c_m$.

We obtain the following comparative statics:

1. With respect to c_m: c_m increases only the left-hand side above and thus decreases m^*;
2. With respect to $\bar{c}_r$: $\bar{c}_r$ decreases only $G(m^*)$ and thus decreases m^*;
3. With respect to χ: $G(m^*)$ increases in χ if and only if $\chi < \frac{1}{2}$, in which case it increases m^*; and it decreases m^* otherwise;
4. With respect to Δ: $G(m^*)$ increases in Δ and hence m^* increases in Δ;
5. With respect to κ: $G(m^*)$ increases in κ and hence m^* increases in κ;
6. With respect to σ_d: it follows that $G(m^*)$ increases in σ_d if and only if:
$$m^*(-2m^*\pi_p + m^* - 4\pi_p(3(\pi_p - 1)\pi_p(5\pi_p - 3) + 1) + 2) - 2\pi_p + 1 > 0 \quad (A.19)$$
where $\pi_p \in (0, 1)$. It follows from $m^* \in (0, 1)$ that the inequality above is satisfied if π_p is sufficiently small.

A4. Proof of Corollary 1

1. Conditional on reporting, we obtain the following probabilities:
$$\mathbb{P}(\theta = L|\mathcal{D}_r = 1) = \frac{\mathbb{P}(\mathcal{D}_r = 1|\theta = L)}{\mathbb{P}(\mathcal{D}_r = 1|\theta = L) + \mathbb{P}(\mathcal{D}_r = 1|\theta = H)} = \frac{m^*\pi_r(H, m^*) + (1-m^*)\pi_r(L, m^*)}{(1+m^*)\pi_r(H, m^*) + (1-m^*)\pi_r(L, m^*)} \quad (A.20)$$
which is greater than
$$\mathbb{P}(\theta = H|\mathcal{D}_r = 1) = \frac{\pi_r(H, m^*)}{(1+m^*)\pi_r(H, m^*) + (1-m^*)\pi_r(L, m^*)} \quad (A.21)$$
because $\pi_r(H, m^*) \le \pi_r(L, m^*)$ as shown in Proposition 1.

2. Note that the probability of $s_J = L$ is equal to:

$$\mathbb{P}\big(s_J = L\big) = \frac{1}{2}(1 - m^*)\pi_r(L, m^*) \tag{A.22}$$

and the probability of $s_J = H$ is equal to:

$$\mathbb{P}\big(s_J = H\big) = \frac{1}{2}(1 + m^*)\pi_r(H, m^*). \tag{A.23}$$

We can plug in the expressions for π_r in Proposition 1 to obtain:

$$\begin{aligned}\mathbb{P}\big(s_J = H\big) - \mathbb{P}\big(s_J = L\big) &= \frac{m^*\big(\Delta^3\kappa^3(m^*+1) + 6\sigma_d^3 - 6\Delta\kappa\sigma_d^2\big)}{3\Delta\kappa^2(m^*+1)(\chi+1)^2\bar{c}_r} \\ &= \frac{m^*\Delta^3\kappa^3\big(m^*+1-6\pi_p^2(1-\pi_p)\big)}{3\Delta\kappa^2(m^*+1)(\chi+1)^2\bar{c}_r} \geq 0\end{aligned} \tag{A.24}$$

where we have used $\pi_p = \frac{\sigma_d}{\kappa\Delta} \in [0,1]$, which implies that $6\pi_p^2\big(1-\pi_p\big) < 1$. Therefore, we obtain that $\mathbb{P}\big(s_J = H\big) - \mathbb{P}\big(s_J = L\big)$ is positive *regardless* of the choice of the firm manager.

3. Conditional on reporting, we obtain the following probabilities:

$$\begin{aligned}\mathbb{P}(s_F = H|\mathcal{D}_r = 1) &= \frac{\pi_r(H,m^*)\mathbb{P}(s_F = H)}{\pi_r(H,m^*)\mathbb{P}(s_F = H) + \pi_r(L,m^*)\mathbb{P}(s_F = L)} \\ &\leq \mathbb{P}(s_F = H)\end{aligned} \tag{A.25}$$

and

$$\begin{aligned}\mathbb{P}(s_F = L|\mathcal{D}_r = 1) &= \frac{\pi_r(L,m^*)\mathbb{P}(s_F = L)}{\pi_r(H,m^*)\mathbb{P}(s_F = H) + \pi_r(L,m^*)\mathbb{P}(s_F = L)} \\ &\geq \mathbb{P}(s_F = L)\end{aligned} \tag{A.26}$$

because $\pi_r(H,m^*) \leq \pi_r(L,m^*)$ as shown in Proposition 1.

4. Comparative statics for $\pi_r(H, m^*)$:
   (a) The cost coefficient c_m affects π_r only through m^*. We have shown before that m^* is decreasing in c_m and that π_r is decreasing in m^*. Hence, π_r is increasing in c_m.
   (b) If $s_J = L$, π_r does not depend on m^*, so the comparative statics are the same as in Proposition 1.
   (c) If $s_J = H$, we can use the implicit function theorem to confirm that the comparative static results in Proposition 1 continue to hold.

A5. Proof of Corollary 2

If $s_J = L$ or if $s_J = \emptyset$, we have shown in the main text that $\mathbb{E}_u[p] = d_\theta$ which implies that $\tau = 0$. If $s_J = H$, the expected stock price is equal to d_θ with probability π_p and equal to $(1+\chi)^{-1}\big(d_\theta + \chi\mathbb{E}[d_\theta|s_J = H]\big)$ otherwise. Substituting this expression into the definition of τ leads to the expression for $\tau(H,\theta)$ given in the Corollary.

A6. Proof of Corollary 3

Note that there are five possible combinations for (θ, s_J):

1. (H, H) with probability $\frac{1}{2}\pi_r(H, m^*)$;
2. $(H, \emptyset)$ with probability $\frac{1}{2}(1 - \pi_r(H, m^*))$;
3. (L, H) with probability $\frac{1}{2}m^*\pi_r(H, m^*)$;
4. (L, L) with probability $\frac{1}{2}(1 - m^*)\pi_r(L, m^*)$;
5. $(L, \emptyset)$ with probability $\frac{1}{2}[m^*(1 - \pi_r(H, m^*)) + (1 - m^*)(1 - \pi_r(L, m^*))]$.

Next, we compute the associated expressions for price quality $\Lambda(\theta, s_J)$:

1. If $\theta = H$ and $s_J = H$, we obtain:

$$\begin{aligned}d_H - p(H,H,u) &= \frac{d_H + \chi d_H - (d_H + \chi\mathbb{E}[d_\theta|\Omega_R] + \kappa u)}{1+\chi} \\ &= \frac{-(\chi\varepsilon_H + \kappa u)}{1+\chi}\end{aligned} \tag{A.27}$$

with $\varepsilon_H = \mathbb{E}[d_\theta|\Omega_R] - d_H$. It follows that:

$$\begin{aligned}\Lambda(H,H) &= -\frac{\mathbb{E}\big[(\chi\varepsilon_H + \kappa u)^2\big]}{(1+\chi)^2} \\ &= -\frac{\mathbb{E}\big[\chi^2\varepsilon_H^2 + 2\chi\varepsilon_H\kappa u + \kappa^2u^2\big]}{(1+\chi)^2}\end{aligned} \tag{A.28}$$

Note that ε_H is equal to zero with probability π_p. With probability $1 - \pi_p$ it is equal to:

$$\frac{1}{1+m^*}d_H + \frac{m^*}{1+m^*}d_L - d_H = -2\sigma_d\frac{m^*}{1+m^*}. \tag{A.29}$$

Also note that the expectation of κu if the price is nonrevealing and $\theta = H$ is equal to $-\sigma_d$. It thus follows that:

$$\Lambda(H,H) = -\frac{\kappa^2\sigma_u^2 + (1-\pi_p)\Big[4\chi^2\sigma_d^2\frac{(m^*)^2}{(1+m^*)^2} + 4\chi\sigma_d^2\frac{m^*}{1+m^*}\Big]}{(1+\chi)^2} \tag{A.30}$$

2. If $\theta = H$ and $s_J = \emptyset$, we obtain $\Lambda(H,\emptyset) = -\kappa^2\sigma_u^2$ as shown in the main text.
3. If $\theta = L$ and $s_J = H$, we obtain:

$$\begin{aligned}d_L - p(L,H,u) &= \frac{d_L + \chi d_L - (d_L + \chi\mathbb{E}[d_\theta|\Omega_R] + \kappa u)}{1+\chi} \\ &= \frac{-(\chi\varepsilon_L + \kappa u)}{1+\chi}\end{aligned} \tag{A.31}$$

with $\varepsilon_L = \mathbb{E}[d_\theta|\Omega_R] - d_L$. It follows that:

$$\begin{aligned}\Lambda(L,H) &= -\frac{\mathbb{E}\big[(\chi\varepsilon_L + \kappa u)^2\big]}{(1+\chi)^2} \\ &= -\frac{\mathbb{E}\big[\chi^2\varepsilon_L^2 + 2\chi\varepsilon_L\kappa u + \kappa^2u^2\big]}{(1+\chi)^2}.\end{aligned} \tag{A.32}$$

Note that ε_L is equal to zero with probability π_p. With probability $1 - \pi_p$ it is equal to:

$$\frac{1}{1+m^*}d_H + \frac{m^*}{1+m^*}d_L - d_L = 2\sigma_d\frac{1}{1+m^*}. \tag{A.33}$$

Also note that the expectation of κu if the price is non-revealing and $\theta = L$ is equal to σ_d. It thus follows that:

$$\Lambda(L,H) = -\frac{\kappa^2\sigma_u^2 + (1-\pi_p)\left[4\chi^2\sigma_d^2\frac{1}{(1+m^*)^2} + 4\chi\sigma_d^2\frac{1}{1+m^*}\right]}{(1+\chi)^2} \tag{A.34}$$

4. If $\theta = L$ and $s_J = L$, we obtain $\Lambda(H,\emptyset) = -\frac{\kappa^2}{(1+\chi)^2}\sigma_u^2$ as shown in the main text.
5. If $\theta = L$ and $s_J = \emptyset$, we obtain $\Lambda(L,\emptyset) = -\kappa^2\sigma_u^2$ as shown in the main text.

Next, we define $\lambda(s_J) = \mathbb{E}_\theta[\Lambda(\theta, s_J)]$. It follows that:

$$\lambda(\emptyset) = -\kappa^2\sigma_u^2, \tag{A.35}$$

and

$$\lambda(L) = -\frac{\kappa^2\sigma_u^2}{(1+\chi)^2}. \tag{A.36}$$

For $s_J = H$:

$$\begin{aligned}\lambda(H) &= \mathbb{P}\big(\theta = H|s_J = H\big)\Lambda(H,H) \\ &\quad + \mathbb{P}\big(\theta = L|s_J = H\big)\Lambda(L,H) \\ &= \frac{1}{1+m^*}\Lambda(H,H) + \frac{m^*}{1+m^*}\Lambda(L,H) \\ &= -\frac{\kappa^2\sigma_u^2 + (1-\pi_p)4\chi(\chi+2)\sigma_d^2\frac{m^*}{(1+m^*)^2}}{(1+\chi)^2}\end{aligned} \tag{A.37}$$

It follows that $\lambda(L) \geq \lambda(H) > \lambda(\emptyset)$.

Comparative statics:

1. The comparative statics for $\lambda(L)$ and $\lambda(\emptyset)$ follow directly from the expressions derived above.
2. For $\lambda(H)$, we apply the implicit function theorem together with the results in Proposition 2 to confirm the results presented in the corollary.

A7. Proof of Corollary 4

The expression for $\mathbb{E}[p|\theta = L]$ given in Eq. (38) implies that the manager's marginal benefit is given by:

$$\frac{\partial\mathbb{E}[p|\theta = L]}{\partial m} = (1-\alpha)\pi_r(\mathbb{E}_u[p] - d_L). \tag{A.38}$$

At an interior solution, m^* is determined by $\frac{\partial\mathbb{E}[p|\theta=L]}{\partial m} = c_m$. We have shown in the proof of Proposition 2 that the marginal benefit is decreasing in m^*. Hence, an increase in α reduces the marginal benefit and leads to a lower equilibrium level of manipulation. If the journalist's investigation efforts are successful, her reporting probability is equal to that in the main model, if $s_F = L$. If her efforts are unsuccessful, the reporting probability is equal to that in the main model, if $s_F = H$. As shown above, the firm's equilibrium level of manipulation is decreasing in α. Since, π_r is decreasing in $\widehat{m}$, it follows that an increase in α leads to an increase in the equilibrium reporting probability.

A8. Proof of Corollary 5

As shown in the text, the journalist's expected utility does not depend on α if $s_F = L$. Hence, the optimal α is equal to zero in this case. If, however, $s_J = H$, then the journalist's marginal benefit of investigation is equal to:

$$\frac{\partial U_J(s_F = H)}{\partial\alpha} = \frac{1}{2\bar{c}_r}\left(\mathbb{E}[U_R|s_J = (H,\theta)]^2 - \mathbb{E}[U_R|s_J = (H,\emptyset)]^2\right) \tag{A.39}$$

which is positive, based on the expressions given in the main text. Next, we compute the change in the marginal benefit if we increase $\widehat{m}$:

$$\frac{\partial^2 U_J(s_F = H)}{\partial\widehat{m}\partial\alpha} = -\frac{1}{\bar{c}_r}\mathbb{E}[U_R|s_J = (H,\emptyset)]\frac{\partial\mathbb{E}[U_R|s_J = (H,\emptyset)]}{\partial\widehat{m}} \tag{A.40}$$

Note that since $\mathbb{E}[U_R|s_J = (H,\emptyset)] > 0$ and $\mathbb{E}[U_R|s_J = (H,\emptyset)]$ is decreasing in $\widehat{m}$ (according to Proposition 1), the expression in Eq. (A.40) is positive, i.e., the marginal benefit of α is increasing in $\widehat{m}$. It follows from the (standard) assumptions on the cost function $C(\alpha)$ that the optimal α is increasing in $\widehat{m}$.

A decrease in the marginal cost $C'(\alpha)$ leads to a higher equilibrium choice of α because the marginal benefit of investigation in Eq. (A.39) does not depend on α. A higher α, in turn, increases the probability that the journalist learns θ from her private signal y. Reader profits given $y = \theta$ are higher than those if $y = \emptyset$. As a result π_r increases in response to a decrease in C'.

References

Admati, A.R., Pfleiderer, P., 1986. A monopolistic market for information. J. Econ. Theory 39 (2), 400–438.

Admati, A.R., Pfleiderer, P., 1988. Selling and trading on information in financial markets. Am. Econ. Rev. 78 (2), 96–103.

Ahern, K.R., Sosyura, D., 2014. Who writes the news? Corporate press releases during merger negotiations. J. Finance 69 (1), 241–291.

Baloria, V.P., Heese, J., 2018. The effects of media slant on firm behavior. J. Financ. Econ. 129 (1), 184–202.

Banerjee, S., Davis, J., Gondhi, N., 2018. When transparency improves, must prices reflect fundamentals better? Rev. Financ. Stud. 31 (6), 2377–2414.

Barber, B.M., 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 (2), 785–818.

Bolton, P., Freixas, X., Shapiro, J., 2012. The credit ratings game. J. Finance 67 (1), 85–111.

Bushee, B.J., Gow, I.D., Taylor, D.J., 2018. Linguistic complexity in firm disclosures: obfuscation or information? J. Account. Res. 56 (1), 85–121.

Call, A.C., Emett, S.A., Maksymov, E., Sharp, N.Y., 2018. Meet the Press: Survey Evidence on Financial Journalists as Information Intermediaries. Unpublished working paper. Arizona State University..

Chen, Q., Hemmer, T., Zhang, Y., 2007. On the relation between conservatism in accounting standards and incentives for earnings management. J. Account. Res. 45 (3), 541–565.

Cohn, J.B., Rajan, U., Strobl, G., 2018. Credit Ratings: Strategic Issuer Disclosure and Optimal Screening. Unpublished working paper. University of Michigan.

Diamond, D.W., 1985. Optimal release of information by firms. J. Finance 40 (4), 1071–1094.

Drake, M.S., Guest, N.M., Twedt, B.J., 2014. The media and mispricing: the role of the business press in the pricing of accounting information. Account. Rev. 89 (5), 1673–1701.

Dyck, A., Zingales, L., 2003. The bubble and the media, pp. 83–104.

Einhorn, E., 2018. Competing information sources. Account. Rev. 93 (4), 151–176.

Engelberg, J.E., Parsons, C.A., 2011. The causal impact of media in financial markets. J. Finance 66 (1), 67–97.

Fang, L., Peress, J., 2009. Media coverage and the cross-section of stock returns. J. Finance 64 (5), 2023–2052.

Fishman, M.J., Hagerty, K.M., 1989. Disclosure decisions by firms and the competition for price efficiency. J. Finance 44 (3), 633–646.
Frenkel, S., 2015. Repeated interaction and rating inflation: a model of double reputation. Am. Econ. J. 7 (1), 250–280.
Frenkel, S., Guttman, I., Kremer, I., 2020. The effect of exogenous information on voluntary disclosure and market quality. J. Financ. Econ. 138 (1), 176–192.
Fulghieri, P., Strobl, G., Xia, H., 2014. The economics of solicited and unsolicited credit ratings. Rev. Financ. Stud. 27 (2), 484–518.
Gao, P., Liang, P.J., 2013. Informational feedback, adverse selection, and optimal disclosure policy. J. Account. Res. 51 (5), 1133–1158.
Gao, P., Zhang, G., 2019. Accounting manipulation, peer pressure, and internal control. Account. Rev. 94 (1), 127–151.
Garcia, D., 2013. Sentiment during recessions. J. Finance 68 (3), 1267–1300.
Garz, M., 2014. Good news and bad news: evidence of media bias in unemployment reports. Public Choice 161 (3–4), 499–515.
Gentzkow, M., Shapiro, J.M., 2006. Media bias and reputation. J. Polit. Econ. 114 (2), 280–316.
Goldman, E., Slezak, S.L., 2006. An equilibrium model of incentive contracts in the presence of information manipulation. J. Financ. Econ. 80 (3), 603–626.
Goldstein, I., Yang, L., 2017. Information disclosure in financial markets. Annu. Rev. Financ. Econ. 9, 101–125.
Goldstein, I., Yang, L., 2019. Good disclosure, bad disclosure. J. Financ. Econ. 131 (1), 118–138.
Han, B., Tang, Y., Yang, L., 2016. Public information and uninformed trading: implications for market liquidity and price efficiency. J. Econ. Theory 163, 604–643.
Huang, X., Teoh, S.H., Zhang, Y., 2014. Tone management. Account. Rev. 89 (3), 1083–1113.
Huberman, G., Regev, T., 2001. Contagious speculation and a cure for cancer: anonevent that made stock prices soar. J. Finance 56 (1), 387–396.
Kaniel, R., Parham, R., 2017. WSJ category kings–the impact of media attention on consumer and mutual fund investment decisions. J. Financ. Econ. 123 (2), 337–356.
Kurlat, P., Veldkamp, L., 2015. Should we regulate financial information? J. Econ. Theory 158, 697–720.
Langberg, N., Sivaramakrishnan, K., 2010. Voluntary disclosures and analyst feedback. J. Account. Res. 48 (3), 603–646.
Li, F., 2008. Annual report readability, current earnings, and earnings persistence. J. Account. Econ. 45 (2–3), 221–247.
Loughran, T., McDonald, B., 2014. Measuring readability in financial disclosures. J. Finance 69 (4), 1643–1671.
Mullainathan, S., Shleifer, A., 2005. The market for news. Am. Econ. Rev. 95 (4), 1031–1053.
Narayanan, M., 1985. Managerial incentives for short-term results. J. Finance 40 (5), 1469–1484.
Niessner, M., So, E.C., 2018. Bad News Bearers: The Negative Tilt of the Financial Press. Unpublished working paper. Yale School of Management.
Peress, J., 2014. The media and the diffusion of information in financial markets: evidence from newspaper strikes. J. Finance 69 (5), 2007–2043.
Peress, J., Schmidt, D., 2020. Glued to the TV: distracted noise traders and stock market liquidity. J. Finance 75 (2), 1083–1133.
Piccolo, A., Shapiro, J.D., 2017. Credit Ratings and Market Information. Unpublished working paper. Oxford University.
Pouget, S., Sauvagnat, J., Villeneuve, S., 2017. A mind is a terrible thing to change: confirmatory bias in financial markets. Rev. Financ. Stud. 30 (6), 2066–2109.
Scharfstein, D.S., Stein, J.C., 1990. Herd behavior and investment. Am. Econ. Rev. 80 (3), 465–479.
Soroka, S., Fournier, P., Nir, L., 2019. Cross-national evidence of a negativity bias in psychophysiological reactions to news. Proc. Natl. Acad. Sci. 116 (38), 18888–18892.
Stein, J.C., 1989. Efficient capital markets, inefficient firms: a model of myopic corporate behavior. Q. J. Econ. 104 (4), 655–669.
Strobl, G., 2013. Earnings manipulation and the cost of capital. J. Account. Res. 51 (2), 449–473.
Tetlock, P.C., 2007. Giving content to investor sentiment: the role of media in the stock market. J. Finance 62 (3), 1139–1168.
Tetlock, P.C., 2011. All the news that's fit to reprint: do investors react to stale information? Rev. Financ. Stud. 24 (5), 1481–1512.
Vega, C., 2006. Stock price reaction to public and private information. J. Financ. Econ. 82 (1), 103–133.