Creating Controversy in Proxy Voting Advice

Finance Working Paper N° 772/2021
June 2021

Andrey Malenko
Boston College, CEPR and ECGI

Nadya Malenko
Boston College, NEBR, FTG, CEPR and ECGI

Chester S. Spatt
Carnegie Mellon University

ECGI Working Paper Series in Finance

Creating Controversy in Proxy Voting Advice

Working Paper N° 772/2021
June 2021

Andrey Malenko
Nadya Malenko
Chester S. Spatt

We thank Shashwat Agrawal and Filipp Prokopev for research assistance. We have read The Journal of Finance disclosure policy and have no con?icts of interest to disclose.

Abstract

We analyze how a profit-maximizing proxy advisor designs vote recommendations and research reports. The advisor benefits from producing informative, unbiased reports, but only partially informative recommendations, biased against the a priori likely alternative. Such recommendations induce close votes, increasing controversy and thereby the relevance and value of proxy advice. Our results suggest shifting from an exclusive emphasis on recommendations, highlighting the importance of both reports and recommendations in proxy advisors' information provision. They rationalize the one-size-fits-all approach and help reinterpret empirical patterns of voting behavior, suggesting that proxy advisors' recommendations may not be a suitable benchmark for evaluating shareholders' votes.

Andrey Malenko

Boston College, CEPR and ECGI
Professor of Finance
Boston College
e-mail: malenkoa@bc.edu

Nadya Malenko

Boston College, NEBR, FTG, CEPR and ECGI
Professor of Finance
Boston College
e-mail: malenko@bc.edu

Chester S. Spatt

Carnegie Mellon University
Pamela R. and Kenneth B. Dunn Professor of Finance
Carnegie Mellon University, Tepper School of Business
e-mail: cspatt@andrew.cmu.edu

Creating Controversy in Proxy Voting Advice

ANDREY MALENKO, NADYA MALENKO, and CHESTER SPATT*

Journal of Finance, forthcoming

ABSTRACT

We analyze how a profit-maximizing proxy advisor designs vote recommendations and research reports. The advisor benefits from producing informative, unbiased reports, but only partially informative recommendations, biased against the a priori likely alternative. Such recommendations induce close votes, increasing controversy and thereby the relevance and value of proxy advice. Our results suggest shifting from an exclusive emphasis on recommendations, highlighting the importance of both reports and recommendations in proxy advisors' information provision. They rationalize the one-size-fits-all approach and help reinterpret empirical patterns of voting behavior, suggesting that proxy advisors' recommendations may not be a suitable benchmark for evaluating shareholders' votes.

*Andrey Malenko is at Boston College, CEPR, ECGI, and FTG. Nadya Malenko is at Boston College, CEPR, ECGI, FTG, and NBER. Chester Spatt is at Carnegie Mellon University and NBER. For helpful comments, we are grateful to the Editor (Philip Bond), Associate Editor, two anonymous referees, Sugato Bhattacharyya, Emiliano Catan, Briana Chang, Jonathan Cohn, Vicente Cunat, Hülya Eraslan, Yonca Ertimur, Fabrizio Ferri, Scott Guernsey, Yingni Guo, Tomislav Ladika, Richmond Mathews, Alan Morrison, Volkan Muslu, David Oesch, Gaizka Ormazabal, Cecilia Parlatore, Paul Povel, Amiyatosh Purnanandam, Uday Rajan, Miriam Schwartz-Ziv, Marshall Vance, Paul Voss, Alexander Wagner, Lucy White, Yan Xiong, Chris Yung, Yishu Zeng, Junyuan Zou; conference participants at the 2021 Conference on Markets and Economies with Information Frictions, Midwest Economic Theory Conference, FSU SunTrust Beach Conference, FTG, SDU Finance Workshop, SFS Cavalcade, BFI Economic Theory Conference, FIRS, WFA, OxFIT, CICF, NBER Corporate Finance SI, UW Summer Finance Conference, EFA, ICGS, Yale Conference on Voting and Governance, Economic Theory Conference VIII, 2022 ECGI Annual Meeting; and seminar participants at the University of Michigan, Rochester, Warwick, FTWebinar, QMUL, CMU, Vienna GSF, Rice, Georgetown, Aarhus, LBS, UIUC, GSU, Ole Miss, Utah, Rutgers, CityU, Purdue, Chicago Fed, Baylor, Iowa, Stanford, Villanova, Temple, Peking University HSBC Business School, UNIL, EPFL, Houston, UT Dallas, NYU, TAU, Hebrew University, University of Western Australia, BC Micro Theory Seminar, UQ Theory Seminar, and Harvard Law, Economics and Organizations Workshop. We thank Shashwat Agrawal and Filipp Prokopev for research assistance. We have read *The Journal of Finance* disclosure policy and have no conflicts of interest to disclose.

Proxy advisory firms have emerged as major players in corporate governance. They make recommendations on how to cast votes and provide research reports for subscribing shareholders that contain the rationales for those recommendations, including information on various aspects of firms' governance practices. The clients of the largest proxy advisor, Institutional Shareholder Services (ISS), include more than 1,500 institutional investors, which cast votes in 50,000 shareholder meetings around the world.[1] While proxy advisors' research reports are only available to their subscribers, their vote recommendations are frequently made public, either by the media or by the party supported by the proxy advisor – the company or the activist investor.[2] These recommendations have been shown to have a large impact on voting outcomes (Alexander et al., 2010; Iliev and Lowry, 2015; Malenko and Shen, 2016).

Given the strong influence of proxy advisors, the quality and information content of their advice have become an important topic of discussion among market participants and policymakers. Over 2018 to 2020, the SEC adopted a number of regulatory changes to ensure "that investors who use proxy voting advice receive more transparent, accurate, and complete information on which to make their voting decisions."[3] Do proxy advisors have incentives to produce high-quality reports and recommendations? And, as the SEC's concept release on the U.S. proxy system questioned, "Does the lack of a direct pecuniary interest in the effects of their recommendations on shareholder value affect how they formulate recommendations?"[4]

Motivated by these questions, this paper studies the design of vote recommendations and research reports by a proxy advisor that maximizes its profits from information sale to shareholders. We show that the advisor has incentives to produce informative and unbiased research reports for its paying subscribers. At the same time, the advisor benefits from producing recommendations that are partially informative, but asymmetrically biased against the more likely alternative. Recommending against the a priori likely alternative "too often" increases the chance of close, contentious votes, and thereby the relevance and value of proxy advice. For example, if management is expected to win, the advisor benefits from recommending against management too often. Following the literature, we refer to such recommendations as "con-

[1]Source: ISS website. The second largest proxy advisor, Glass Lewis, has more than 1,300 clients, who collectively manage more than $40 trillion in assets (source: Glass Lewis website).

[2]For example, it is common for an activist running a proxy fight to announce proxy advisors' support for its directors through the campaign website or press release distribution services, and for management to announce proxy advisors' support for its proposals. See Section I.A for details.

[3]See Releases No. 34-86721 and No. 34-89372. In June 2021, the SEC suspended the enforcement of these new regulations.

[4]See p. 125 in Release No. 34-62495 (July 14, 2010).

troversial" and the proxy advisor's tendency to give them as "creating controversy."[5]

Our results suggest a more nuanced view of proxy advisors' information provision role. While their recommendations, which have been extensively criticized, are not fully informative in our model, we show that they are often at least partly informative, and banning them could lead to lower quality outcomes. Moreover, we emphasize the distinction between research reports and recommendations, suggesting that an exclusive focus on recommendations may limit understanding of proxy advisors' role. Our paper also advocates for a more careful interpretation of empirical evidence on shareholder voting, highlighting that proxy advisors' recommendations may not be the appropriate benchmark for evaluating investors' votes.

In our model, shareholders vote on a proposal whose value depends on an unknown state. The proxy advisor, which has information about the state, designs two signals: one, which we refer to as the "research report," is available only to its subscribers, and the other, which we refer to as the "vote recommendation," is publicly observed by all investors. One way to interpret the design of recommendations are voting guidelines, which proxy advisors announce each year (see Section I.A). We model the information design problem as Bayesian persuasion (Kamenica and Gentzkow, 2011; Rayo and Segal, 2010) and show that the chosen policies are time-consistent: because the advisor maximizes its ex-ante profits and has no vested interest in the vote outcome, it has no ex-post incentive to deviate from the designed signals. In addition to designing the two signals, the advisor sets a fee for subscribing to the report. Each shareholder decides whether to pay the fee and get the report or to rely on the recommendation. The state is then realized, the advisor's report and recommendation are produced, and shareholders vote based on the information they receive. Shareholders maximize the value of their shares.

The fact that the proxy advisor maximizes its profit from information sales highlights a fundamental tension. As former SEC Commissioner Daniel Gallagher put it, citing Choi, Fisch, and Kahan (2009), "proxy advisors do not have a financial stake in the companies about which they provide voting advice" (Gallagher, 2013). If, instead, the advisor were maximizing the value for the downstream entities – asset managers and operating companies, then its public recommendations would perfectly reveal all the advisor's information, but then no investors would need to purchase the reports and the advisor would earn zero revenue.

This raises the question of which design of recommendations and reports maximizes the

[5] The shareholder voting literature often refers to proposals for which proxy advisors recommend against management as "controversial" (e.g., Marquardt et al. 2018; Boone, Gillan, and Towner, 2020; Iliev and Vitanova, 2023) or "contentious" (e.g., Iliev and Lowry, 2015; Dimmock et al., 2018). Similarly, ex-post close votes or, more generally, votes that have generated substantial disagreement among shareholders are frequently labeled as "controversial" (e.g., Fos and Holderness, 2023; Fahlenbrach, Rudolf, and Wegerich, 2024).

fees that can be obtained from subscribers. One might naturally expect the advisor to produce completely uninformative public recommendations to avoid diluting the value of private reports, which are offered for a fee. Instead, we show that the advisor's profit is maximized if it produces (1) a fully informative and unbiased report, and (2) a public recommendation that is partially informative but favors the alternative that is a priori unlikely to be chosen. In particular, the advisor recommends against the more likely alternative too often, even in some cases when this alternative is value-increasing.

At the heart of these results is the proxy advisor's desire to increase the relevance of its advice, a point frequently mentioned in interviews with market participants.[6] When most votes in shareholder meetings are non-contentious and pass with overwhelming support, individual votes become less significant. As a result, shareholders have less motivation to be diligent stewards of their investments and spend resources on making informed voting decisions, which lowers the demand for proxy advisory services. To enhance the relevance of its advice, the advisor benefits from designing recommendations that increase the frequency of contentious, close votes. Since many shareholders rely solely on recommendations, the advisor can effectively coordinate those shareholders' votes and guide them to vote against the a priori likely alternative "too often." For shareholders who read the advisor's report and make their own voting decisions, recommendations may not offer much insight into the proposal itself. However, recommendations do influence these shareholders' perceptions of how *others* will vote, in a way that increases their attention to the vote and their value from proxy advisory services.

To see the intuition more precisely, consider a common type of proposal: the approval of directors nominated by the firm. Suppose there is a high prior probability that the nominees add value (in practice, directors endorsed by management are approved with high support rates in uncontested elections). If proxy advisors' recommendations did not exist or were uninformative, most shareholders would vote based on their priors and predominantly support the directors. The votes would always be non-contentious, and a shareholder deciding whether to subscribe to the proxy advisor's reports would have little incentive to do so. Suppose, instead, that the advisor issues partially informative recommendations that favor the less likely alternative of voting "against": it always recommends against bad directors, but sometimes recommends even against good directors. Then, shareholders who vote solely based on recommendations cannot infer from a negative recommendation whether the director is good or bad and are unsure how to vote, with some of them voting against and some voting in favor.

[6] Section III discusses the interview evidence from Hayne and Vance (2019) that supports our mechanism.

This leads to a high chance of a close vote, which, in turn, increases the incentives of other shareholders to invest in information and subscribe to the reports. We show that a recommendation appropriately designed this way is often optimal for the advisor and dominates any other information design, including issuing completely uninformative recommendations.[7] At the same time, the advisor produces fully informative and unbiased research reports because doing so maximizes revenue from fees charged to subscribers. In this sense, the interests of the advisor are aligned with those of its clients, to whom it sells subscriptions.

The combination of recommendations and reports is thus central to our mechanism: by frequently issuing recommendations against the unexpected alternative, the advisor raises shareholders' willingness to pay for the reports. Effectively, the advisor has two audiences: it provides full information in the reports and thus serves the needs of its clients who utilize the reports, while limiting the information released through recommendations for shareholders who do not read the reports. These results are consistent with the survey evidence by Bew and Fields (2012) on how institutional investors view proxy advisory services. According to the survey, institutions derive substantial value from proxy advisors' reports, highlighting their ability to synthesize and aggregate vast amounts of information, but perceive the recommendations as less valuable (see Section III for details).

Proxy advisors are often criticized for their "one-size-fits-all" approach, that is, giving recommendations that do not consider company-specific factors. This approach naturally emerges in our model as a way to produce recommendations that favor the less likely alternative, while avoiding apparent contradictions between the report and the recommendation. For example, one of ISS's guidelines concerns busy directors: "Generally vote against or withhold from individual directors who sit on more than five public company boards." While the factor of director busyness is indeed relevant, this guideline overlooks other director characteristics. Consider a nominee who sits on several boards but has extensive industry experience and whose other board seats are not too demanding. A typical research report provides a detailed summary of the key facts concerning the proposals. Hence, the report will describe all director qualifications, enabling shareholders who read it to independently assess and potentially form a positive opinion regarding the nominee's quality. At the same time, the guideline's one-size-fits-all approach permits a negative recommendation without the need to balance it against

[7] In particular, such recommendation design is optimal if the prior is sufficiently asymmetric. If the prior that the proposal is beneficial is close to 50%, a close vote is already likely, and the advisor's recommendations are then completely uninformative. See Section II.C.1 for details and Section III for a discussion of how the one-size-fits all approach can implement such uninformative recommendations.

other pertinent factors, thereby avoiding a contradiction with the report. A shareholder who only observes the recommendation and does not carefully read the report will find it challenging to discern whether the negative recommendation reflects any other concerns or is given purely because of the number of board seats, and will often vote against. This logic aligns with the results in Iliev and Lowry (2015), who show that relative to ISS, large shareholders, who are likely to be engaged voters, vote in a manner that is less one-size-fits-all (see Section III).

In our setting, there are two groups of investors: subscribers, who are assumed to read the entire report, and non-subscribers, who rely solely on recommendations. In reality, a third category exists: investors who subscribe to proxy advisors for non-informational reasons (e.g., vote execution or litigation protection), but lack the capacity to digest the entire report and base their votes solely on recommendations. Our model's logic applies in this scenario as well. By issuing recommendations that favor the less likely alternative, the advisor can effectively coordinate the votes of investors who rely solely on recommendations (both non-subscribers and those subscribing for other purposes), thereby increasing the relevance of its advice and the value of information to investors who learn from the entire report. Thus, our mechanism is more general and does not require vote recommendations to be made public.

The paper suggests a reinterpretation of the empirical evidence on funds' voting behavior. Voting for management when proxy advisors recommend "against" is often seen as a lack of monitoring and pro-management bias. In contrast, we highlight that proxy advisors' recommendations may not be the right benchmark, as they may be biased against management to increase the value of proxy advice. Shareholders who support management despite negative recommendations could simply be adjusting for this tendency, rather than voting in a conflicted manner. In fact, in our setting, the votes of large asset managers arise as a more suitable benchmark. Our model also predicts an asymmetry in shareholders' responsiveness to recommendations. Recommendations that align with the a priori likely alternative should be "rubberstamped," with few shareholders deviating from them. In contrast, recommendations against the priors should not be rubberstamped; instead, there should be substantial dispersion in shareholders' votes and closer outcomes. These and other predictions align well with the empirical evidence, as we discuss in Sections III and IV, and Section II of the Internet Appendix.

Our results are important for policy discussions of proxy advisors. First, we highlight a distinction between the information content of reports and recommendations, which we corroborate with empirical and anecdotal evidence (see the case studies of Tesla, Avalara, and HSBC described in Section III). Thus, policy discussions should focus not only on recommendations, but also on the reports and how shareholders use the information they contain: an exclusive focus on recommendations could neglect the high informational value of the reports. Second, our analysis suggests that the common policy concern, which concentrates on potential bias in proxy advisors' recommendations due to their corporate consulting services, may overlook another source of tension. The tendency towards "creating controversy" is inherent to the core business model of proxy advisors – selling information to voters. Despite this, proxy advisors' recommendations add value in our setting: even if they favor the less likely alternative, they contain valuable information and guide the votes of shareholders who do not read the detailed reports. Without these recommendations, small shareholders would follow their priors (e.g., predominantly vote with management). In Section IV, we formalize this argument and show that banning recommendations and limiting proxy advisors' reports to only factual content could backfire. While such a policy may increase the number of shareholders reading the report and making their own conclusions, it would make other shareholders' votes less informed, potentially yielding a net negative impact.

In our baseline model, shareholders are assumed to be solely focused on maximizing value and making informed voting decisions. This overlooks two important aspects of the proxy voting process: concerns about litigation and ideological preferences. First, it is often argued that many institutions subscribe to proxy advisors as a strategy to mitigate litigation risk associated with their voting practices. According to former SEC commissioner Daniel Gallagher, "relying on the advice from the proxy advisory firm became a cheap litigation insurance policy" (Gallagher, 2014). Our insight that proxy advisors benefit from a higher incidence of close votes applies in this scenario as well. When votes are non-contentious and pass with large margins, institutions likely face lower litigation risks, as potential litigation costs outweigh the benefits. For example, the 2019 SEC guidance specifically stresses the need for heightened analysis when "the matter is highly contested or controversial" (SEC Release No. IA-5325). By coordinating shareholders who rely on their recommendations to vote against the a priori likely alternative, proxy advisors increase the incidence of contentious votes, thereby enhancing the value of their services to shareholders. In Section V, we discuss the robustness of our results to an extension in which shareholders receive an additional, non-informational benefit from subscribing to proxy advice, which could be interpreted as the reduction in the risk of litigation.

In another extension, we study the role of ideological preferences, which could stem from shareholders' views on environmental and social issues (Bolton et al., 2022; Dikolli et al., 2022) or attitudes towards management (Matvos and Ostrovsky, 2010; Jackson and Zytnick, 2022). We assume that a portion of investors are ideological, voting in a particular direction regardless of information, while other investors are value-maximizers, who vote based on information. In this setting, the information content of recommendations is significantly influenced by whether investors' preferences are aligned or heterogeneous, and by whether they reinforce or contradict prior beliefs about the proposal's value. For example, if the proposal is a priori believed to be value-increasing but faces ideological opposition from many investors, close votes are likely despite strong prior beliefs. In such cases, the proxy advisor's motivation to issue controversial recommendations diminishes. Conversely, if investors' preferences align with prior beliefs or are too heterogeneous to counteract strong prior beliefs, partially informative recommendations against the likely alternative still benefit the advisor.

Related literature. Our paper contributes to the literature on shareholder voting,[8] including the growing literature on proxy advisors. Malenko and Malenko (2019) and Buechel, Mechtenberg, and Wagner (2024) analyze proxy advisors impact shareholders' independent research; both papers take the quality of recommendations as given and assume they are unbiased. Levit and Tsoy (2022) show how one-size-fits-all recommendations arise in a cheap talk setting where a biased expert (e.g., proxy advisor) wants to convince other agents (e.g., shareholders) to accept a proposal. Unlike these papers, we focus on the information design problem of an advisor who maximizes profits from selling information. Ma and Xiong (2021) also study information design and show that biased recommendations can arise, but do not distinguish between a public (recommendation) and private (research report) signal. As a consequence, in their model, recommendations are biased only if shareholders are biased, whereas if shareholders maximize firm value, recommendations are unbiased. In contrast, in our setting, biased recommendations arise even when all shareholders maximize firm value, as a way to increase the probability of a close vote through manipulation of public information. This also distinguishes our paper from Matsusaka and Shu (2021), who study how proxy advisors cater their recommendations to biased shareholders such as socially responsible investing funds, and analyze the resulting industry structure. Blonien et al. (2024) estimate a structural model to uncover the error rates in fund votes and ISS recommendations, finding evidence they interpret as consistent with our predictions: very low error rates in ISS recommendations favoring management but higher error rates in those opposing management.

[8] See Maug (1999), Brav and Mathews (2011), Levit and Malenko (2011), Esö, Hansen, and White (2014), Bar-Isaac and Shapiro (2020), Cvijanovic, Groen-Xu, and Zachariadis (2020), Bouton et al. (2021), and others.

In the literature on Bayesian persuasion, the closest papers to ours examine information design by a biased expert aiming to manipulate elections towards his preferred direction (Alonso and Camara, 2016; Bardhi and Guo, 2018; Chan et al., 2019; Kerman, Herings, and Karos, 2020). In contrast, in our paper, the designer is unbiased, with no vested interest in the vote's outcome. This implies, in particular, that its information design policies are time-consistent, which is different from most other papers on Bayesian persuasion. Another feature that distinguishes our paper is that the designer creates two signals for two different audiences – one public (for all shareholders) and one private (for subscribers); furthermore, the composition of the latter audience is endogenously chosen by the designer. Inostroza (2021) also considers a designer (regulator) creating two signals for multiple audiences, but both are public signals on two different dimensions of the bank's fundamentals. In Leitner and Yilmaz (2019), the designer (bank) designs two signals, one observed by the receiver (regulator), and the other possibly only by the designer. Michaeli (2017) studies a designer who, as in our paper, optimally chooses the fraction of users who see its information, but unlike in our paper, discloses the same signal to multiple users. Goldstein and Huang (2016), Inostroza and Pavan (2020), and Alonso and Zachariadis (2021) analyze a regulator designing information (stress test) for multiple receivers with private information, but differently from our paper, assume that the designer sends a single signal to all receivers.[9] Chang and Szydlowski (2020) study persuasion in a matching market with multiple heterogeneous senders (investment advisors) and receivers (their customers). Overall, the proxy advisory industry has unique features that distinguish our setting and mechanisms from those in the literature.

Our paper also relates to studies on the sale of information (e.g., Sarvary and Parker, 1997; Bergemann and Bonatti, 2019), including information sales to financial market traders (e.g., Admati and Pfleiderer, 1986, 1990; Fishman and Hagerty, 1995; Cespa, 2008; Garcia and Sangiorgi, 2011). An important conclusion in the latter literature is that the seller may benefit from adding noise to the information it sells, as a way to reduce the leakage of information through prices. In contrast, we show that the proxy advisor benefits from selling the most precise information to those subscribing to its reports, but to increase the value of this information, it adds a specific form of noise (against the more likely alternative) to the public information it reveals. These different results stem from the distinct interactions among information users: while traders compete, shareholders in our model are aligned on value maximization but prefer

[9]Other papers studying information design in the context of stress tests include Goldstein and Leitner (2018), Leitner and Williams (2023), and Orlov, Zryumov, and Skrzypacz (2023).

to free-ride on others' information acquisition efforts. This feature also distinguishes proxy advisors from credit rating agencies, which are another type of information provider to investors (see Sangiorgi and Spatt (2017) for an overview of the literature). Credit rating agencies and proxy advisors also differ in their pricing models: credit rating agencies are paid by issuers, while proxy advisors are paid by investors.[10] There are also certain similarities, such as the issue of multiple, albeit different from the proxy advisory setting, audiences (Frenkel, 2015; Bouvard and Levy, 2018); the provision of both paid and unpaid signals (Fulghieri, Strobl, and Xia, 2014); and the coarseness of ratings and recommendations (Goel and Thakor, 2015).

The paper proceeds as follows. Section I introduces the setup. Section II derives our main results. Section III presents evidence in line with our mechanism. Section IV discusses the implications. Section V considers several extensions of the basic model. Finally, Section VI concludes.

I. Model setup

Proposal to be voted on. The firm is owned by $N \geq 3$ shareholders, where N is odd. Each shareholder owns one share in the firm and has one vote. There is a proposal to be voted on, which is approved if at least $\frac{N+1}{2}$ shareholders vote in favor. Let $d \in \{0,1\}$ denote the decision on the proposal, where $d = 1$ ($d = 0$) corresponds to proposal approval (rejection).

The value of the proposal to shareholder i, $u_i(d,\theta)$, depends on the unknown state $\theta \in \{0,1\}$ and on the importance of the proposal to the shareholder, v_i, as follows:

$$u_i(d,\theta) = v_i \cdot u(d,\theta), \tag{1}$$

where

$$u(1,\theta) = \begin{cases} 1, & \text{if } \theta = 1, \\ -1, & \text{if } \theta = 0, \end{cases} \qquad u(0,\theta) = 0. \tag{2}$$

In other words, approving the proposal increases (decreases) shareholder value if $\theta = 1$ ($\theta = 0$), while rejecting the proposal and maintaining the status quo leaves firm value unchanged. The ex-ante probability that the proposal is value-increasing is $\Pr(\theta = 1) = \mu \in (0,1)$.

[10] In our model, the advisor strictly benefits from its recommendations being public (not making recommendations public is, in our setting, equivalent to making them uninformative, which is generally suboptimal). This presents an interesting comparison to credit rating agencies. It is often believed that credit rating agencies switched from the issuer-pay model to the investor-pay model because their ratings were easily reproducible and shareable among investors, hindering their ability to charge fees for these ratings. In the context of proxy advisors, the reproducibility of recommendations is advantageous: by designing the recommendations in a certain way and allowing them to be publicly accessible, the proxy advisor can increase its revenues.

Thus, all shareholders' interests are aligned, but the extent to which they care about the proposal, v_i, may differ across them. The role of heterogeneous v_i is to produce variation in shareholders' incentives to pay for advice. There are multiple reasons for this heterogeneity in practice. First, it can be attributed to varying levels of sensitivity in fund managers' compensation to the value of their portfolio firms, with hedge funds having the highest sensitivity, index funds the lowest, and active mutual funds in between. Second, heterogeneity in v_i can be due to degrees of scrutiny by regulators and market participants on shareholders' voting practices: for example, mutual funds are required to disclose their votes and may have higher v_i than other shareholders. Larger v_i can also capture investors with larger portfolios.[11]

We assume that v_i is an independent (across shareholders) draw from a distribution with a continuous and differentiable c.d.f. $H(\cdot)$ over $[\underline{v}, \overline{v}]$ with $0 \leq \underline{v} < \overline{v} \leq \infty$. When making his information acquisition decision, each shareholder knows his own v_i, but not v_j of other shareholders (in practice, shareholders do not perfectly know the ownership structure of their portfolio firms). This assumption is made for simplicity, as it allows us to focus on symmetric equilibria at the information acquisition stage.

Shareholders maximize $u_i(d, \theta)$ minus any costs of information acquisition.

Information structure. Each shareholder is initially uninformed and has prior μ that the proposal is value-increasing. There is a seller of information, the proxy advisor, which maximizes its profits from information sale to shareholders. The proxy advisor has an informative signal about the state. For simplicity, we assume that it knows the state with certainty (we discuss the extension to a noisy signal and costly information acquisition in Section V).

The advisor designs two signals, a private signal available only to the subscribers and a public signal available to everyone, and sets the fee it charges for the private signal.[12] The private signal, denoted $\mathcal{R} = (R, \{\phi(\cdot|\theta)\}_{\theta\in\{0,1\}})$, consists of a finite signal space R and two distributions $\{\phi(\cdot|\theta)\}_{\theta\in\{0,1\}}$ of signal realizations $r \in R$, which describe the probability of signal realization r in each state. The public signal, denoted $\mathcal{S} = \left(S, \{\gamma(\cdot|r)\}_{r\in R}\right)$, consists of a finite signal space S and a family of distributions $\{\gamma(\cdot|r)\}_{r\in R}$ of signal realizations $s \in S$, which describe the probability of signal realization s for each $r \in R$. We will refer to the

[11] More precisely, assuming that all firms have the same prior probability that a certain type of proposal is value-increasing, v_i would be proportional to the number of firms in the investor's portfolio that have this type of proposal on the agenda. In addition, v_i can indirectly capture the investor's position in the firm, although this interpretation needs to be used cautiously given our assumption of equal stakes across shareholders.

[12] The use of the term "private signal" is different from what is often called "private persuasion" in the information design literature, which is when the sender conditions the signal distribution on private information of the receiver (Kolotilin et al., 2017; Guo and Shmaya, 2019).

private signal as the *research report* and to the public signal as the *vote recommendation.* This formulation means that the research report is informative about the state θ, while policy $\mathcal{S}$ determines how the content of the report is mapped into the vote recommendation.

We first conjecture and later verify (see Section II.C.3) that it is optimal for the advisor to design a fully informative research report, that is, $R = \{0, 1\}$ and $r = \theta$. Thus, by subscribing to the advisor's services, a shareholder learns the state with certainty. Given this, the advisor's problem is how to design the public recommendation for each possible realization of state θ: $\mathcal{S} = (S, \{\gamma(\cdot|\theta)\}_{\theta\in\{0,1\}})$. For example, in the case of a binary recommendation space $S = \{0, 1\}$, which we will show to be optimal, the recommendation policy is characterized by two probabilities, $\Pr(s = 1|\theta = 1) = \gamma(1|1)$ and $\Pr(s = 1|\theta = 0) = \gamma(1|0)$.

The advisor chooses policies $\mathcal{R}$ and $\mathcal{S}$ to maximize its expected profits. Note that the optimal information policies are dynamically consistent, in contrast to Kamenica and Gentzkow (2011) and most other models of Bayesian persuasion. This is because the advisor maximizes ex-ante profits from selling information and does not obtain any ex-post benefit from a vote outcome in either direction. Thus, once the state is realized, the advisor does not gain from deviating and reporting a different signal for either the report or the recommendation. The time consistency of information policies could be undermined in a different setting, where the advisor exerts unobservable costly effort to produce information and interacts repeatedly with its customers. Then, the advisor's incentives to build a reputation for producing precise reports may lead to dynamic inconsistency in the information policy. This dynamic inconsistency can be resolved if recommendations are based on ex-post verifiable information, since the advisor's deviations from the announced policy can then be detected. Consistent with this, in practice, reports and recommendations are based on facts from proxy statements and other filings, activists' documents, and market data, and much of it is ex-post verifiable information.

Timeline. The timeline is shown in Figure 1. At stage 1, the advisor chooses the information policy $(\mathcal{R}, \mathcal{S})$ and fee f that it charges to shareholders for subscribing to its research report. At stage 2, having observed the information policy $(\mathcal{R}, \mathcal{S})$, fee f, and their realizations of proposal importance v_i, shareholders simultaneously and non-cooperatively decide on whether to pay fee f to subscribe to the report. At stage 3, the advisor observes θ and issues report $r \in R$ and recommendation $s \in S$ based on the information policy $(\mathcal{R}, \mathcal{S})$. At stage 4, all shareholders observe the realization of s, and shareholders that became the subscribers also observe r. Shareholders then simultaneously decide whether to vote "for" $(a_i = 1)$ or "against" $(a_i = 0)$ the proposal, and the proposal gets implemented if it is approved by the majority.

We assume that shareholders cannot abstain from voting. This assumption is consistent the idea that many institutional investors, such as mutual funds, do not abstain for reputational and regulatory reasons. See Section V for a detailed discussion.

The equilibrium concept is a symmetric Bayes-Nash equilibrium: at the information acquisition stage, shareholders with the same v_i follow the same subscription strategy, and at the voting stage, shareholders with the same information follow the same voting strategy. We focus on equilibria in weakly undominated strategies at the voting stage. In particular, when we write that there is a unique equilibrium, we mean a unique equilibrium in this class. We also assume that if, for a given fee f, there exist multiple equilibria at the information acquisition and voting stages, the advisor can induce his preferred equilibrium.[13]

[Figure 1 here]

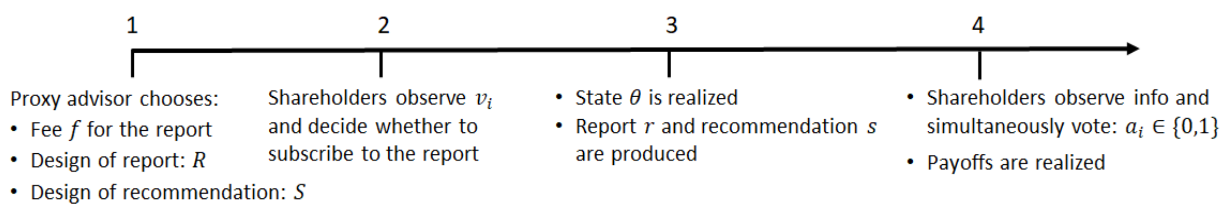

Figure 1. Timeline of the model.

[13]This equilibrium selection allows us to abstract from equilibria of the following form: the advisor charges any positive fee, no shareholder buys the report, and at the voting stage all shareholders vote the same way (for example, all vote in favor or all vote against). This is an equilibrium because a shareholder is never pivotal, so there is no profitable deviation at the voting stage and, since the value of information is zero, no shareholder wants to deviate and buy the report even if the fee is arbitrarily small.

A. Mapping between the model and proxy advisory practices

This information structure and timeline capture the observed voting practices. The interpretation of private signal r is the research report produced by the proxy advisor: these reports are only available to the subscribers. Figure A.4 in the Internet Appendix presents a screenshot of ISS's research report on a proxy contest organized by Trian Fund Management at Procter & Gamble. A typical report starts with the background information about the company, its performance relative to industry peers, and governance structure. Then, for each proposal on the agenda, the report presents detailed statistics and facts, as well as a discussion about the costs and benefits of the proposal. Importantly, the report also contains the proxy advisor's vote recommendation. In Figure A.4, this recommendation is prominently displayed at the top of the first page, as well as on page 24 following the extensive analysis on the previous pages: ISS recommends in favor of the activist's director nominee.

Unlike the report itself, which contains detailed analysis covering different aspects of the proposal, the vote recommendation in the report is binary. It is easy to communicate, and this is exactly what often happens in practice. First, proxy advisors' recommendations are routinely reported by the media. Figure A.5 in the Internet Appendix shows the article in Financial Times announcing that ISS backed the activist's nominee in the P&G proxy contest (panel A) and the article in Reuters announcing a negative ISS recommendation on say-on-pay (panel B). Second, the party supported by the proxy advisor – activist or management – has incentives to make this information public. Panel C of Figure A.5 presents a press release by the management of Avalara after ISS supported its proposed sale to another firm.

For these reasons, public signal s is best interpreted as the vote recommendation, which is part of every research report but, in reality, gets transmitted outside of the report.[14] However, as we discuss in the Introduction, our mechanism does not rely on recommendations being public: a tendency to favor the less likely alternative would also arise in a setting where some investors subscribe for non-informational reasons and recommendations are not public.

In our model, at the beginning of the game, the advisor announces its information design policy $\mathcal{S}$, which is a mapping from the report to the recommendation.[15] A good empirical analog of $\mathcal{S}$ are voting guidelines that proxy advisors announce every year. These guidelines describe, for various types of proposals, detailed rules and criteria that the advisor plans to

[14] Consistent with vote recommendations being publicly available, Brav et al. (2024) rely on public filings for their data collection of recommendations and write: "Because each party has an incentive to publicize a favorable recommendation from a proxy advisor, this process should reveal most of the advisors' recommendations."

[15] While in practice, proxy advisors' recommendations are binary, we do not impose this restriction. However, as we show below, the optimal recommendation is indeed binary, corresponding to the observed practices.

use when making its recommendations on such proposals during the upcoming proxy season. Hence, effectively, these are mappings from the content of the report to the recommendation. These guidelines are updated every year, and receive a lot of publicity. The timing of their release and public dissemination corresponds to that in Figure 1: proxy advisors typically announce their guidelines for the upcoming proxy season (with most shareholder meetings happening in March and April) at the end of the previous year. These guidelines are not firm-specific, are announced before the actual shareholder meetings, and hence are used by shareholders to decide whether to subscribe to proxy advisory services. Subscription decisions are made before the proxy season and usually apply to all firms in an investor's portfolio.

Figure A.6 in the Internet Appendix presents screenshots from ISS's guidelines. We elaborate on these guidelines in Section III, where we discuss how the predictions of our model can be mapped to proxy advisory practices. However, the information design problem of the advisor should be understood more generally than just the design of the guidelines.

B. Benchmark case: One shareholder

We start with the benchmark case of one shareholder, that is, a bilateral sender-receiver setting. This captures situations in which shareholders can perfectly coordinate with each other.

The shareholder estimates the probability of making the correct decision given the information in the report, $\Pr(d=\theta|\mathcal{R})$, and the information in the public recommendation, $\Pr(d=\theta|\mathcal{S})$, and subscribes to the report if and only if

$$v_i\left[\Pr(d=\theta|\mathcal{R})-\Pr(d=\theta|\mathcal{S})\right]\geq f. \tag{3}$$

Hence, the advisor wants to maximize $\Pr(d=\theta|\mathcal{R})$ and minimize $\Pr(d=\theta|\mathcal{S})$. The former is achieved by designing a fully informative report, and the latter by designing an uninformative recommendation. The optimal fee solves the standard problem of a monopolist, which faces a trade-off between quantity sold (i.e., probability of the shareholder subscribing to the report) and the price paid by the customers (i.e., the fee charged for the report). We conclude:

PROPOSITION 1 (Benchmark case): *If $N=1$, the proxy advisor always designs an uninformative recommendation and a fully informative research report.*

Thus, when there are no coordination frictions among shareholders, the advisor exclusively provides information through a paid report, without offering any information for free. In practice, however, coordination among shareholders is imperfect (see Section V). Then, as our further analysis reveals, the advisor's information design problem is no longer trivial. In particular, with multiple shareholders, it can be optimal to share some information for free.

II. Analysis

We solve the model by backward induction. We first assume that the report $\mathcal{R}$ is fully informative and solve for the equilibrium in the voting game, the equilibrium subscription decisions, pricing of information, and the public recommendation design. In Section II.C.3, we complete the solution by proving that making the report fully informative is optimal for the advisor.

A. Voting stage

Consider the voting stage following any given realization of the state θ and public recommendation s. First, consider a shareholder who subscribed to the report. As we conjecture above, the report conveys the state perfectly, so the shareholder knows the state with certainty. Hence, it is a weakly dominant strategy for him to vote based on the state: $a_i = \theta$.

Next, consider shareholder i who did not subscribe to the report. This shareholder observes the public recommendation s and forms his posterior belief $\mu_s = \Pr(\theta = 1|s)$ about the state knowing the information policy $\mathcal{S}$. In addition, as in the literature on strategic voting (Austen-Smith and Banks, 1996; Feddersen and Pesendorfer, 1998), the shareholder takes into account that his vote is only pivotal for the outcome if the number of "for" votes among other shareholders is exactly $\frac{N-1}{2}$, that is, if the votes of others are split (we denote this set of events by Piv). In particular, the shareholder rationally conditions his decision not only on the recommendation s but also on the information that must be true when he is pivotal.[16] The shareholder finds it optimal to vote "for" ("against") if his posterior belief $\Pr(\theta = 1|s, Piv)$ is above (below) $\frac{1}{2}$, and is indifferent if it is exactly equal to $\frac{1}{2}$, where by Bayes' rule,

$$\Pr(\theta = 1|s, Piv) = \frac{\Pr(Piv|\theta = 1, s)\,\mu_s}{\Pr(Piv|\theta = 1, s)\,\mu_s + \Pr(Piv|\theta = 0, s)\,(1 - \mu_s)}. \tag{4}$$

[16] Such strategic voting is related to the "winner's curse" in auctions: both hinge on the idea that an agent's action only matters in a particular situation – when his bid is the highest or his vote is pivotal. Since other agents (bidders and shareholders, respectively) have valuable information that the agent lacks, he rationally conditions his decision on the information that must be true when his action matters. In Section IV.D of the Internet Appendix, we show that our results continue to hold if shareholders do not condition on being pivotal.

Intuitively, the shareholder updates his beliefs based on two pieces of information: the recommendation s and the information he learns from the fact that he is pivotal. The former corresponds to terms μ_s and $1-\mu_s$ in (4), and the latter corresponds to terms $\Pr(Piv|\theta=1,s)$ and $\Pr(Piv|\theta=0,s)$. To find these two latter terms, let $\tilde{q}(v)$ denote the probability, as perceived by the shareholder, with which another shareholder subscribes to the report if his type is v, and let $q \equiv \int_{\underline{v}}^{\bar{v}} \tilde{q}(v)\, dH(v)$ denote the shareholder's unconditional perceived probability that each other shareholder subscribes to the report. In addition, let π denote the probability, as perceived by the shareholder, with which each non-subscribing shareholder votes for the proposal given recommendation s. In equilibrium, these perceived probabilities coincide with the actual probabilities of subscribing to the report and voting in favor given s.

Note that if $q=0$ or $q=1$, the advisor's profit is zero: if $q=0$, no shareholder buys the report, and $q=1$ is only possible if the fee is zero.[17] It follows that $q=0$ and $q=1$ never arise on equilibrium path, and we can focus on $q \in (0,1)$.[18]

Given q and π, the shareholder's perceived probability that each other shareholder votes "for" conditional on state θ and signal s is $q+(1-q)\pi$ if $\theta=1$ and $(1-q)\pi$ if $\theta=0$, so the probability that there are exactly $\frac{N-1}{2}$ "for" votes among other shareholders is

$$\Pr(Piv|\theta=1,s) = C_{N-1}^{\frac{N-1}{2}} \left[q+(1-q)\pi\right]^{\frac{N-1}{2}} \left[(1-q)(1-\pi)\right]^{\frac{N-1}{2}}, \quad (5)$$
$$\Pr(Piv|\theta=0,s) = C_{N-1}^{\frac{N-1}{2}} \left[(1-q)\pi\right]^{\frac{N-1}{2}} \left[1-\pi(1-q)\right]^{\frac{N-1}{2}}, \quad (6)$$

where $C_n^k = \frac{n!}{k!(n-k)!}$ is the binomial coefficient.

The next result characterizes the equilibrium in the voting game for any realization s:

PROPOSITION 2 (Shareholders' voting strategies): *Consider the voting game that follows a recommendation realization s, and suppose that each shareholder believes that other shareholders subscribe to the report with probability $q \in (0,1)$. This game has a unique equilibrium in undominated strategies, which takes the following form. If a shareholder is a subscriber, he votes according to the state, $a_i=\theta$. If a shareholder is not a subscriber and $\mu_s \in (0,1)$, he*

[17] Indeed, suppose $q=1$, so that all shareholders know the state with certainty. In any equilibrium in which a shareholder is pivotal with a strictly positive probability, he must vote according to the state, but then all shareholders always vote the same way, so such an equilibrium does not exist. Hence, the only possible equilibria are those where no shareholder is ever pivotal. In these equilibria, the value of information is zero, so shareholders have no incentive to pay for the report if $f>0$.

[18] Formally, this is because the advisor can set fee f for which there is an equilibrium with $q \in (0,1)$ and the advisor's profit is positive (we show this statement below). Thus, given our assumption that the advisor can induce his preferred equilibrium for a given fee, the advisor will not find it optimal to induce $q=0$ or $q=1$.

votes in favor with probability $\frac{1}{2}$ *if* $\mu_s = \frac{1}{2}$ *and with probability*

$$\pi(q, \mu_s) = \frac{z_s(1-2q) - 1 + \sqrt{(z_s - 1)^2 + 4q^2 z_s}}{2(z_s - 1)(1-q)} \tag{7}$$

if $\mu_s \neq \frac{1}{2}$, *where* $z_s \equiv \left(\frac{\mu_s}{1-\mu_s}\right)^{\frac{2}{N-1}}$, $\pi(q, \mu_s) \in (0,1)$ *and is increasing in* μ_s. *If a shareholder is not a subscriber and* $\mu_s = 1$ *(*$\mu_s = 0$*), he votes in favor with probability one (zero).*

Thus, all shareholders with precise information (subscribers) vote according to their information, whereas shareholders with imprecise information (non-subscribers) randomize between voting for and against, and the expected voting support by non-subscribers is increasing in the belief μ_s that the proposal is value-increasing.

Figure 2 illustrates Proposition 2. It shows that the sensitivity of non-subscribing shareholders' votes to posterior μ_s depends on the expected fraction of subscribers, q. Intuitively, a shareholder who observes s expects other non-subscribers to likely vote along the posterior belief μ_s. But then, the fact that the vote is split implies that the subscribers (who know the state) are relatively more likely to have voted against belief μ_s. Thus, by conditioning on being pivotal, the shareholder updates his beliefs in the direction opposite of μ_s, so his vote becomes less reliant on μ_s. The extent of this learning from being pivotal depends on q. If q is low, the shareholder learns little from the fact that the vote is split, since almost all other shareholders are non-subscribers and have the same information as him. As a result, his strategy relies heavily on whether posterior μ_s is above $\frac{1}{2}$, as illustrated by the solid (red and green) lines in Figure 2 for $q = 0.0001$ and 0.01, respectively. In particular, the red line shows that in the limit when $q \to 0$, $\pi(q, \mu_s)$ converges to a step integer function $1\left\{\mu_s > \frac{1}{2}\right\}$. In contrast, if q is high, a non-subscriber learns quite a bit from the fact that the votes of others are split, since the probability that each voter is perfectly informed is higher. As a result, the probability that he votes "for" is less sensitive to posterior μ_s around $\mu_s = \frac{1}{2}$. This is illustrated by the dashed and dash-dotted lines in Figure 2 for $q = 0.1$ and $q = 0.6$, respectively.

One may wonder whether shareholders are so rational that they incorporate the information learned from being pivotal. Maug and Rydqvist (2009) find evidence consistent with such strategic voting. However, our results do not rely on this: in Section IV.D of the Internet Appendix, we assume that investors (irrationally) do not learn from pivotal events and show that incentives to create controversy still arise. In the rest of the paper, we continue to treat shareholders as fully rational and incorporating the information from pivotal events.

[Figure 2 here]

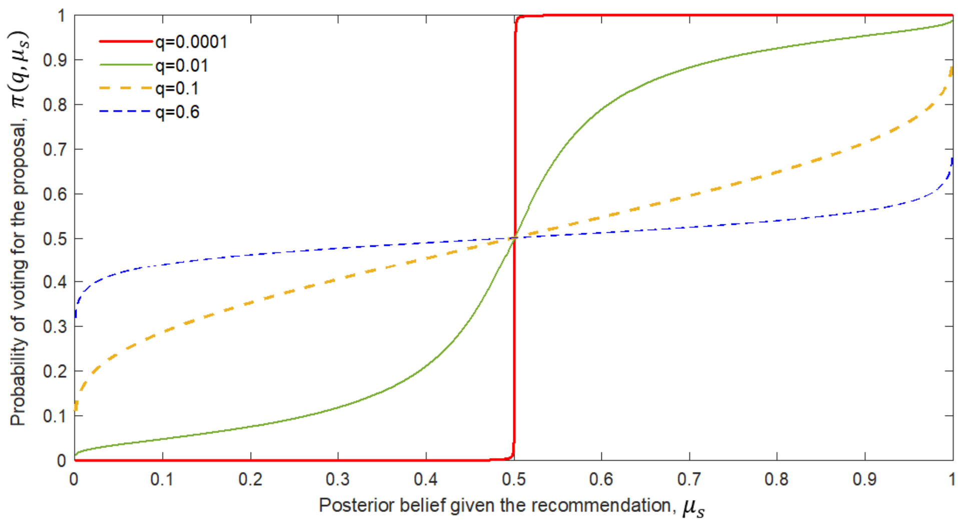

Figure 2. Shareholders' voting strategies. The figure illustrates Proposition 2 by plotting the probability $\pi(q, \mu_s)$ that a non-subscribing shareholder votes for the proposal as a function of his posterior belief μ_s for $N = 25$ and four different values of q: 0.0001, 0.01, 0.1, and 0.6.

B. *Information acquisition stage*

Since the vote recommendation is publicly observable, the shareholder's incentive to purchase the report reflects his incremental value from the additional information contained in the report, that is, from learning the state. In particular, it captures his incremental value from a more efficient voting outcome due to his own vote being more informed. This value is a function of probability q, with which he expects others to subscribe to the report, recommendation policy $\mathcal{S}$, and the shareholder's concern about the voting outcome v_i. Let $V_i(q, \mathcal{S})$ denote this value.

To derive $V_i(q, \mathcal{S})$, consider shareholder i's value from buying the report, and thus learning the state with certainty, for a particular realization of the public recommendation $s \in S$. If the posterior belief μ_s induced by this recommendation is one or zero, the shareholder knows the state with certainty from observing the recommendation, so his incremental value from the report is zero. If $\mu_s \in (0, 1)$, the shareholder's value from the report is positive. His vote changes the decision of the firm only if the votes of other $N-1$ shareholders are split. We denote the probability of this event by $\Pr(Piv|q,\mu_s)$ and note that

$$\Pr(Piv|q,\mu_s) = \Pr(Piv|\theta=1,s)\,\mu_s + \Pr(Piv|\theta=0,s)\,(1-\mu_s)\,, \tag{8}$$

where $\Pr(Piv|\theta,s)$ is given by (5)–(6) with $\pi = \pi(q,\mu_s)$ given by (7). Conditional on the shareholder being pivotal, learning the state with certainty changes his probability of voting correctly from $\frac{1}{2}$ to 1, and his expected value from doing so is $\frac{v_i}{2}$.[19] Thus, for any realization s, the shareholder's value from the report is $\frac{v_i}{2}\Pr(Piv|q,\mu_s)$. Aggregating over all possible realizations of $s \in S$, the value of the report to shareholder i is

$$V_i(q,\mathcal{S}) = v_i \cdot V(q,\mathcal{S})\,, \tag{9}$$

where

$$V(q,\mathcal{S}) = \frac{1}{2}\sum_{s\in S}\Pr(Piv|q,\mu_s)\,\tau_s \tag{10}$$

and $\tau_s \equiv \mu\gamma(s|1) + (1-\mu)\,\gamma(s|0)$ is the probability of recommendation s. Intuitively, $V(q,\mathcal{S})$ is the average (before the recommendation is realized) probability that a shareholder is pivotal multiplied by the benefit $\frac{1}{2}$ of learning the state conditional on being pivotal.

Hence, shareholder i buys the report if and only if $v_i \cdot V(q,\mathcal{S}) \geq f$, or equivalently, $v_i \geq \frac{f}{V(q,\mathcal{S})}$. It follows that, given fee f, the probability that each shareholder subscribes to the report is $1 - H\left(\frac{f}{V(q,\mathcal{S})}\right)$. We can rewrite this expression as an inverse demand function: to ensure that, on average, fraction q of shareholders subscribe, the fee must be:

$$f = V(q,\mathcal{S})\,H^{-1}(1-q)\,, \tag{11}$$

where $H^{-1}(\cdot)$ is an increasing function. We summarize these arguments as follows:

PROPOSITION 3 (Shareholders' information acquisition): *For a given fee f and public rec-*

[19] This is because, as explained above, the shareholder rationally conditions his decisions on the information that is true when he is pivotal. As Proposition 2 and its proof show, if $\mu_s \in (0,1)$, then the shareholder randomizes between voting for and against, and his posterior belief $\Pr(\theta=1|s,Piv)$ conditional on s and being pivotal is $\frac{1}{2}$. Thus, the shareholder believes that without the research report, he votes correctly ($a_i = \theta$) with probability $\frac{1}{2}$. With the research report, the shareholder votes correctly with probability 1. Since $u(\theta,\theta) - u(1-\theta,\theta) = 1$, the shareholder's expected value from learning the state (i.e., his willingness to pay for the report) conditional on s and being pivotal is $\frac{v_i}{2}$. In particular, this conditional willingness to pay does not depend on the prior μ. See Section III.D of the Internet Appendix for a detailed derivation.

ommendation policy S, shareholder i subscribes to the proxy advisor's report if and only if the proposal is sufficiently important to him, $v_i \geq f/V(q, \mathcal{S})$, where $V(q, \mathcal{S})$ is given by (10). If the equilibrium fraction of subscribers q is in $(0,1)$, it satisfies (11).

C. Proxy advisor's problem

The proxy advisor chooses fee f and information policy $\mathcal{S}$ to maximize its expected profit. Since q is the expected fraction of subscribers and each subscriber pays fee f, the advisor's expected profit is Nqf. Using (11), the advisor chooses q and $\mathcal{S}$ to solve

$$\max_{q,S} qH^{-1}(1-q)\left(\sum_{s\in S} \Pr(Piv|q, \mu_s)\,\tau_s\right). \tag{12}$$

We solve this problem by decomposing it into two steps. First, in Section II.C.1, we take the fraction of subscribers q as given and find the advisor's optimal public recommendation policy $\mathcal{S}^*(q)$ for any q. This optimal policy maximizes the average probability of a split vote, $\sum_{s\in S} \Pr(Piv|q, \mu_s)\,\tau_s$, which, according to (9), is exactly what determines shareholders' willingness to pay for the report. Second, in Section II.C.2, we find the advisor's optimal fraction of subscribers and the fee that induces it, taking into account $\mathcal{S}^*(q)$.

C.1 Public recommendation design

Note that any information policy $\mathcal{S}$ can be represented as a combination of the set of possible recommendations S, posterior beliefs $\{\mu_s, s \in S\}$ that these recommendations induce, and the frequencies with which each of these recommendations is produced $\{\tau_s, s \in S\}$. Since shareholders use Bayesian updating, the average posterior belief must be equal to the prior:

$$\sum_{s\in S} \mu_s \tau_s = \mu. \tag{13}$$

Condition (13) is referred to as the Bayes plausibility constraint (Kamenica and Gentzkow, 2011). Kamenica and Gentzkow (2011) show that the reverse is also true: any combination $\{S, \{\mu_s, \tau_s\}_{s\in S}\}$ that satisfies (13) can be induced by some information policy $\mathcal{S} = (S, \gamma(\cdot|\theta=0), \gamma(\cdot|\theta=1))$. Thus, for a given q, we can rewrite the advisor's problem as

$$\begin{aligned} &\max\nolimits_{S,\{\mu_s,\tau_s\}} \textstyle\sum_{s\in S} \Pr(Piv|q, \mu_s)\,\tau_s \\ &s.t.\ \textstyle\sum_{s\in S} \mu_s \tau_s = \mu. \end{aligned} \tag{14}$$

In other words, taking q as given, how should the advisor design its recommendations to maximize the average probability that the vote is split? The intuition behind the Bayes plausibility constraint is that the advisor cannot systematically deceive the shareholders, which, in turn, limits its ability to inflate the probability of a split vote. Suppose, for example, that one of the advisor's recommendations induces $\mu_s \approx 0.5$, which creates a lot of uncertainty and is thereby likely to result in a split vote. If the prior belief μ is sufficiently high (e.g., $\mu = 0.8$), then the Bayes plausibility constraint implies that this "controversial" recommendation cannot be given too frequently: other recommendations, which induce a more positive posterior belief (above 0.8) and are less likely to result in a split vote, must be sufficiently frequent as well.

We solve (14) using the standard approach in the information design literature, referred to as "concavification" (Kamenica and Gentzkow, 2011). This approach considers the objective of the designer as a function of any given posterior belief (in our setting the objective function is $\Pr(Piv|q,\cdot)$, the probability of a shareholder being pivotal), and then takes the concave closure of this function, which we denote $P(q,\cdot)$. Then, the largest expected probability of a split vote that the advisor can achieve given prior belief μ is $P(q,\mu)$, and the information policy that achieves this maximum can be found graphically. To apply this approach to our setting, we start by deriving the properties of $\Pr(Piv|q,\cdot)$.

LEMMA 1: *There exist $\underline{\mu} \in \left(0,\frac{1}{2}\right)$, $\bar{\mu} \in \left(\frac{1}{2},1\right)$, and $\varepsilon > 0$, such that: (i) if $0 < q < \frac{1}{2}$, $\Pr(Piv|q,\mu_s)$ is strictly convex in μ_s for $\mu_s \in \left(0,\underline{\mu}\right)$ and $\mu_s \in (\bar{\mu},1)$, and strictly concave in μ_s for $\mu_s \in \left(\frac{1}{2}-\varepsilon,\frac{1}{2}+\varepsilon\right)$; (ii) if $\frac{1}{2} < q < 1$, $\Pr(Piv|q,\mu_s)$ is strictly concave in μ_s for $\mu_s \in \left(0,\underline{\mu}\right)$, $\mu_s \in (\bar{\mu},1)$, and $\mu_s \in \left(\frac{1}{2}-\varepsilon,\frac{1}{2}+\varepsilon\right)$.*

These properties are illustrated in Figure 3: the solid (blue) line in each panel plots $\Pr(Piv|q,\mu_s)$ as a function of μ_s for $N = 25$ and different values of the fraction of subscribers q. To see the intuition, recall from Proposition 2 and Figure 2 that non-subscribers' voting strategy $\pi(q,\mu_s)$ is less sensitive to μ_s around $\mu_s = \frac{1}{2}$ when q is higher. If q is small, most shareholders do not subscribe to the report, and each non-subscriber bases his vote mostly on μ_s: if $\mu_s > 0.5$ ($\mu_s < 0.5$), he is likely to vote in favor (against). Hence, the probability of a split vote is small (zero in the limit of $q \to 0$) except in a narrow interval around $\mu_s = 0.5$. This case is illustrated in panels A and B for $q = 0.01$. As q increases, non-subscribers' votes become less sensitive to μ_s around $\frac{1}{2}$ and, in addition, there is a higher fraction of subscribers, who vote according to the true state. As a result, the probability of a split vote becomes less sensitive to μ_s around $\frac{1}{2}$ (see panels C and D for $q = 0.1$ and $q = 0.6$).

[Figure 3 here]

To see how we can use Lemma 1 to derive the optimal information policy using the concavification approach, consider several information policies, illustrated in panels A and B of Figure 3 for $q = 0.01$. First, consider a *fully informative* policy, which can be implemented by giving two recommendations: one negative (given if $\theta = 0$) and one positive (given if $\theta = 1$). This policy corresponds to two posterior beliefs, $\mu_0 = 0$ and $\mu_1 = 1$, and is represented by two orange circles in panel A. Since the probability of a split vote is zero for both posterior beliefs, the average probability of a split vote is zero as well, so such a policy is never optimal.

An alternative information policy is an *uninformative* recommendation. For example, it can be implemented by always recommending against, always recommending in favor, or randomizing between recommending in favor and against in a way uncorrelated with the state. This policy corresponds to a single posterior belief equal to the prior μ, and is illustrated by the blue circle in panel A. If μ is close to 0.5 (as in panel A), there is a lot of a priori uncertainty, so the probability of a split vote given an uninformative recommendation is relatively high, and moreover (as we show in Proposition 5), is higher than for any other information policy. However, if μ is sufficiently far from 0.5, for example, $\mu = 0.8$ (as in panel B), the proposal is a priori likely to be value-increasing, so the probability of a split vote upon an uninformative recommendation is small. In this case, the advisor can increase the average probability of a split vote by making its recommendations *partially informative*.

To see this, suppose the advisor gives two recommendations: positive (denoted $s = 1$) and negative (denoted $s = 0$), which induce posteriors $\mu_1 = 1$ and μ_0 close to 0.5, illustrated in panel B. The advisor can implement this policy by always giving a negative recommendation if the proposal is value-decreasing, but also sometimes giving a negative recommendation if it is value-increasing.[20] Bayes plausibility implies that the frequencies of the two recommendations, τ_1 and τ_0, are such that $\tau_1 + \tau_0\mu_0 = \mu$, that is, $\tau_0 = \frac{1-\mu}{1-\mu_0}$. The positive recommendation reveals the state with certainty and thus never results in a split vote, whereas the negative recommendation generates a lot of uncertainty and results in a split vote with a high probability. The average probability of a split vote is the weighted average between 0 and $\Pr(Piv|q, \mu_0)$

[20] More specifically, this information policy has $\gamma(0|0) = 1$ and $\gamma(0|1)$ such that $\mu_0 = \Pr(\theta = 1|s = 0) = \frac{\gamma(0|1)\mu}{\gamma(0|1)\mu + 1 - \mu}$, which implies $\gamma(0|1) = \frac{1-\mu}{\mu}\frac{\mu_0}{1-\mu_0}$.

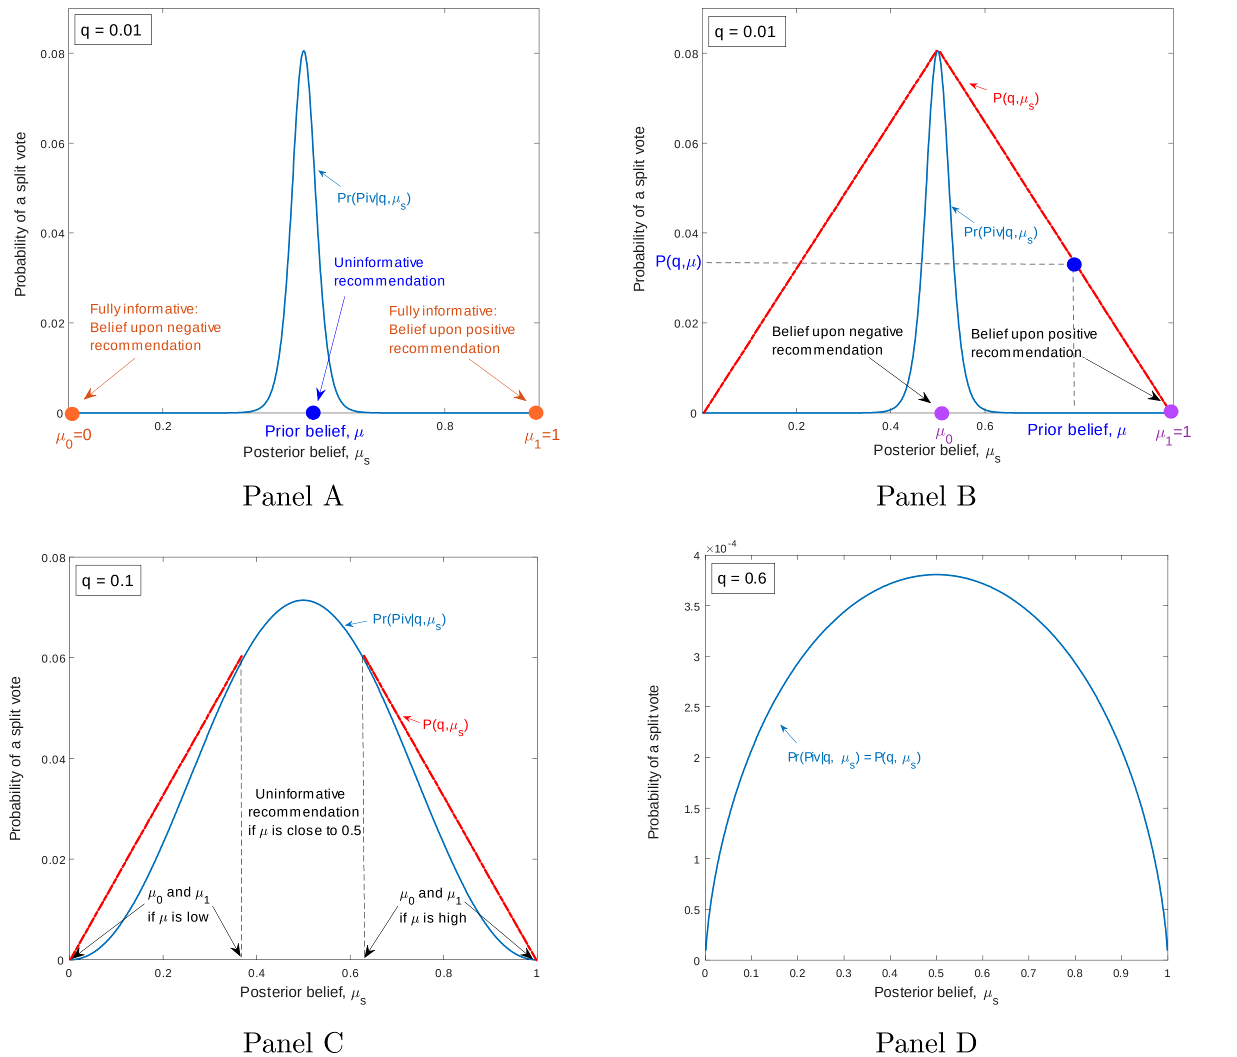

Figure 3. Design of public recommendations for different values of q. The solid (blue) lines in all panels plot the probability of a shareholder being pivotal, $\Pr(Piv|q, \mu_s)$, as a function of the posterior belief μ_s for N=25. The dashed (red) lines in panels B and C present the concave closure of this function, $P(q, \mu_s)$. Panels A and B illustrate the case of q=0.01. Panel A shows a fully informative recommendation and an uninformative recommendation for μ close to 0.5. Panel B shows the optimal (partially informative) recommendation for μ=0.8. Panel C illustrates the optimal recommendation for q=0.1 and different μ: if μ is low (below 0.38), it induces posteriors 0 and 0.38; if μ is high (above 0.62), it induces posteriors 0.62 and 1; if μ is between the two cut-offs, it is uninformative. Panel D illustrates the case of q=0.6, in which the optimal recommendation is uninformative for all μ.

with weights τ_1 and τ_0, and the Bayes plausibility constraint pinpoints this average probability to the one illustrated in panel B: it is the value of the linear function (depicted by the red dashed line connecting the two points) at prior belief μ. As is clear from panel B, this partially informative policy leads to a higher average probability of a split vote than an uninformative policy. Finally, note that for the average probability of a split vote to be maximized, the dashed red line should be tangent to the solid blue curve at point μ_0, that is, it should concavify function $\Pr(Piv|q,\cdot)$. Hence, the maximum average probability of a split vote is $P(q,\mu)$.

The following result uses Lemma 1 and formalizes this intuition.

PROPOSITION 4 (Optimality of partially informative recommendations): *The optimal public recommendation is binary, $S=\{0,1\}$. If $q\in\left(0,\frac{1}{2}\right)$ and priors are sufficiently asymmetric, the optimal recommendation takes the following form:*

1. *If $\mu\le\mu_1(q)\equiv\frac{\left(1-4q^2\right)^{\frac{N-1}{2}}}{1+\left(1-4q^2\right)^{\frac{N-1}{2}}}$, recommendation $s=0$ induces belief $\mu_0=0$, and recommendation $s=1$ induces belief $\mu_1(q)$ and is given with probability $\frac{\mu}{\mu_1(q)}$.*

2. *If $\mu\ge\mu_0(q)\equiv\frac{1}{1+\left(1-4q^2\right)^{\frac{N-1}{2}}}$, recommendation $s=0$ induces belief $\mu_0(q)$ and is given with probability $\frac{1-\mu}{1-\mu_0(q)}$, and recommendation $s=1$ induces belief $\mu_1=1$.*

Proposition 4 implies that if priors are sufficiently asymmetric and there are not many subscribers, the optimal recommendation is partially informative. For example, consider $\mu\ge\mu_0(q)$. A positive recommendation (which is in line with the prior) is only given if the proposal is value-increasing ($\theta=1$) and thus reveals that $\theta=1$ with certainty. Thus, recommendations in line with the prior are *rubberstamped*: all shareholders vote in favor, regardless of whether they are subscribers or not. In contrast, a negative recommendation (which is against the prior) is always given if $\theta=0$ but also sometimes given if $\theta=1$. Hence, it does not fully reveal θ and induces a disagreement in shareholders' votes; in this sense, it is controversial. The likelihood of a negative recommendation is $\Pr(s=0)=\frac{1-\mu}{1-\mu_0(q)}$, which exceeds the probability that the proposal is value-decreasing ($\Pr(\theta=0)=1-\mu$) by a factor of $\frac{1}{1-\mu_0(q)}$. In this sense, the advisor's recommendation policy favors the a priori unlikely alternative. The frequency of recommendations against the prior depends on the expected fraction of subscribers q:

COROLLARY 1: *For any μ, there exists $\bar{q}(\mu)$ such that partially informative recommendations of the form in Proposition 4 are optimal if and only if $q\le\bar{q}(\mu)$. For $q\le\bar{q}(\mu)$, if q increases:*

(i) there is a higher frequency of recommendations against the prior (measured by $\frac{\Pr(s=1)}{\Pr(\theta=1)}$ *or* $\Pr(s=1) - \Pr(\theta=1)$ *for small* μ*, and by* $\frac{\Pr(s=0)}{\Pr(\theta=0)}$ *or* $\Pr(s=0) - \Pr(\theta=0)$ *for large* μ*);*

(ii) recommendations against the prior become "less convincing," in the sense that the posterior belief upon them is closer to the prior.

This result is illustrated in Figure 3, and we explain it for the case of large μ. The comparison between panels B and C reveals that as q increases, the posterior belief upon the recommendation against the prior (i.e., belief μ_0 upon $s = 0$) moves farther from 0.5 and closer to μ. Intuitively, the advisor faces the following trade-off when picking μ_0. One option is to induce a very uncertain posterior (i.e., μ_0 close to 0.5), so that the probability of a split vote upon the controversial recommendation is very high. However, Bayes plausibility implies that such a recommendation cannot be given too frequently, which is costly for the advisor because the "rubberstamped" recommendation $s = 1$, which never results in a split vote, must be given more frequently. An alternative is to induce a less uncertain, that is, closer to the prior, posterior μ_0, which leads to a lower probability of a split vote but is given more frequently. Hence, the advisor's trade-off is between a higher probability of a split vote conditional on the controversial recommendation and a higher frequency of a split vote. When q increases, shareholders' votes become less sensitive to their posterior beliefs around $\frac{1}{2}$ (see Figure 2), so the probability of a split vote is relatively high even if the posterior is far from $\frac{1}{2}$. As a result, the advisor finds it optimal to pick a less uncertain posterior but with a higher frequency. Overall, as q increases, recommendations against the prior become more frequent, but "less convincing" (in the sense of inducing a posterior closer to the prior) and result in a lower likelihood of a split vote. Once q becomes high enough (such that $\mu_0(q) = \mu$), this partially informative recommendation policy becomes uninformative: $s = 0$ is given with probability one and induces belief μ.

While partially informative recommendations are optimal when priors are sufficiently asymmetric, they are not always optimal. If the prior belief is close to $\frac{1}{2}$, there is already a lot of uncertainty about the correct decision, so the probability of a split vote is high and shareholders have incentives to become informed. In this case, there is no benefit to manipulating public information, as formalized in the following proposition:

PROPOSITION 5 (Optimality of uninformative recommendations): *There exists* $\varepsilon > 0$ *such that if* $\mu \in \left(\frac{1}{2} - \varepsilon, \frac{1}{2} + \varepsilon\right)$*, the optimal public recommendation is uninformative.*

This result is illustrated in panel C of Figure 3 for $q = 0.1$ and $N = 25$. It shows that when μ is between the two cutoffs, $\mu_1(q)$ and $\mu_0(q)$, the optimal recommendation is uninformative. Comparing panels B–D also shows that as the expected fraction of subscribers increases, the range of priors for which the optimal recommendation is uninformative widens. Moreover, if q increases to over $\frac{1}{2}$ (Lemma 1), it becomes suboptimal to give information for free, regardless of the priors, as seen in panel D for $q = 0.6$. Intuitively, as q increases, shareholders' votes become less sensitive to the posterior beliefs μ_s around $\frac{1}{2}$, reducing the marginal benefit of inducing beliefs closer to 50% with a negative recommendation. Since there is also a cost of doing so—the split vote never occurs with a positive recommendation—the range of priors for which the optimal recommendation is informative shrinks. An uninformative recommendation policy can be implemented by always recommending the same action (e.g., always against) for certain proposals. We discuss such examples and supporting evidence in Section III.

C.2 Pricing of information

We now analyze the fee charged by the advisor for the research report. The advisor faces the standard trade-off between price and quantity: a lower fee attracts shareholders with lower valuations (i.e., lower v_i) and increases quantity sold, but leaves more rents to shareholders with high v_i. In addition, the unique feature of the proxy advisory setting is that a shareholder's valuation of the seller's product depends on the number of other shareholders purchasing the product (i.e., on q) and on the seller's design of public recommendations (see Eq. (9)–(10)). These are the two margins the advisor uses to further increase its revenues.

Suppose the advisor chooses a fee such that it sells, in expectation, to a fraction q of shareholders. As shown above, the highest probability of a split vote that the advisor can achieve (when the prior is μ and the expected fraction of subscribers is q) is $P(q, \mu)$, where $P(q, \cdot)$ is the concave closure of $\Pr(Piv|q, \cdot)$. Thus, using (12), the advisor's expected profit is $\frac{N}{2} qH^{-1}(1-q) P(q, \mu)$, so the optimal fee induces q that solves[21]

$$\max_q qH^{-1}(1-q) P(q, \mu) . \tag{15}$$

[21] In general, a given fee f can induce multiple equilibria at the information acquisition and voting stages, corresponding to different values of q. However, given our earlier assumption that the advisor can induce his preferred equilibrium, we simply optimize over q. The fee that induces this optimal q is then given by (11).

The next result provides sufficient conditions under which the advisor's optimal fee and information design result in recommendations that create controversy.

PROPOSITION 6: *Let q^* be the maximum of $qH^{-1}(1-q)$. If $q^* < \frac{1}{2}$ and priors are sufficiently asymmetric, the advisor sets the fee and designs public recommendations to create controversy:*

1. *If $\mu \leq \mu_1(q^*)$, recommendation $s = 0$ induces belief $\mu_0 = 0$, and $s = 1$ induces belief $\mu_1(q^*)$ and is given with probability $\frac{\mu}{\mu_1(q^*)}$. The price of the research report is $f = 2^{1-N}\mu C_{N-1}^{\frac{N-1}{2}} H^{-1}(1-q^*)$, and each shareholder subscribes to it with probability q^*.*

2. *If $\mu \geq \mu_0(q^*)$, recommendation $s = 0$ induces belief $\mu_0(q^*)$ and is given with probability $\frac{1-\mu}{1-\mu_0(q^*)}$, and $s = 1$ induces belief $\mu_1 = 1$. The price of the research report is $f = 2^{1-N}(1-\mu) C_{N-1}^{\frac{N-1}{2}} H^{-1}(1-q^*)$, and each shareholder subscribes to it with probability q^*.*

The condition $q^* < \frac{1}{2}$ in Proposition 6 implies that the distribution $H(\cdot)$ is such that the advisor optimally sells the report to a relatively small fraction of shareholders and demands a high price from them. To increase these shareholders' willingness to pay for the report, the advisor designs public recommendations that favor the less likely alternative. The condition $q^* < \frac{1}{2}$ is material: as shown in Lemma 1, the function $\Pr(Piv|q, \mu_s)$ is strictly concave around $\mu_s = 0$ and $\mu_s = 1$ when $q > \frac{1}{2}$, and hence a partially informative recommendation of the form described by Proposition 4 may no longer be optimal. The next result shows this.

PROPOSITION 7: *Let q^* be the maximum of $qH^{-1}(1-q)$, and suppose that $q^* \geq \frac{1}{2}$ and the distribution $H(\cdot)$ has an increasing hazard rate. Then, for any prior belief μ, any pair $(q, \mathcal{S}^*(q))$, where $S^*(q)$ is a partially informative recommendation policy of the form in Proposition 4, is dominated by some pair $(\hat{q}, S_{\text{uninf}})$, where S_{uninf} is an uninformative recommendation.*

To illustrate the intuition for Propositions 6 and 7, consider an example in which the intensity of shareholders' concerns about the proposal v_i follows the power distribution.

Example 1: $H(x) = x^{\alpha}$. Note that $qH^{-1}(1-q) = q(1-q)^{\frac{1}{\alpha}}$ is increasing in q if and only if $q < \frac{\alpha}{\alpha+1}$, so $q^* = \frac{\alpha}{\alpha+1}$. Thus, if $\alpha < 1$ and the prior belief is sufficiently asymmetric, the advisor finds it optimal to design a partially informative recommendation that favors the less likely alternative and sell the report to an expected fraction $\frac{\alpha}{\alpha+1} < \frac{1}{2}$ of shareholders. Intuitively, $\alpha < 1$ means that the distribution of v_i has a positive skew: most shareholders care little about the proposal, but some care quite a lot. In this case, it is optimal to sell the report to a relatively small fraction of shareholders who care significantly about the proposal, and to increase these shareholders' value from the report by a public recommendation that induces close votes. In contrast, if $\alpha > 1$, the distribution of v_i has a negative skew: there are many shareholders who care a lot about the proposal and some who care very little. In this case, it is optimal to sell the report to a large fraction of shareholders, and any recommendation of the form in Proposition 4 is dominated by an uninformative recommendation.

Properties of strategic recommendation design. We conclude this section by exploring the comparative statics of the public recommendation policy designed by the advisor.

First, consider a change in the distribution $H(\cdot)$. This change affects the trade-off between price and quantity sold and thus alters q^*, the optimal expected fraction of subscribers. Corollary 1 and Propositions 6–7 imply that as q^* increases, recommendations against the prior become more frequent but less convincing (i.e., induce posterior beliefs closer to the prior). Moreover, once q^* reaches a certain point, such recommendation design is no longer optimal. This is because, as explained earlier, a higher q^* leads to stronger learning from being pivotal, which in turn leads to a lower sensitivity of shareholders' votes to posterior μ_s around $\frac{1}{2}$.

In Section IV of the Internet Appendix, we derive other properties of recommendation design, which rely on a related intuition. First, we analyze the comparative statics in N, the number of shareholders. As N increases, recommendations against the prior become more frequent but less convincing, and once N becomes large enough, partially informative recommendations of the form in Proposition 4 are no longer optimal. This, together with Proposition 1, implies that controversial recommendations arise only with multiple shareholders, but that the number of shareholders cannot be too large. The intuition behind the comparative statics in N is similar to that for q^*: as N increases, shareholders learn more from being pivotal, so their voting strategies become less sensitive to their posterior belief μ_s around $\frac{1}{2}$. Second, we consider an extension where some shareholders obtain information from other sources (e.g., via independent research) and show that this increases the frequency of recommendations against the prior, but that the posterior beliefs upon such recommendations move closer to the prior.

C.3 Design of the private research report

So far, we have solved for the optimal public recommendation design, the fee charged for the report, and the equilibrium at the information acquisition and voting stages conjecturing that the advisor finds it optimal to design the private report that perfectly reveals the state. The next proposition verifies that this conjecture is correct:

PROPOSITION 8 (Optimality of a fully informative research report): *A fully informative private research report, that is, $R = \{0, 1\}$ and $r = \theta$, is optimal for the advisor.*

Loosely speaking, a fully informative report is optimal because adding noise dilutes its value and lowers the willingness to pay of subscribing shareholders. The more precise intuition, formalized in the proof, involves two steps. The first step points out that for any report with an arbitrary number of signals, there is a report with two signals for which each shareholder's willingness to pay is not lower. Intuitively, any signal induces one of three best responses: subscribers either vote for, against, or randomize between voting for and against. In the latter case, the subscriber's posterior belief must be $\frac{1}{2}$, so the value from buying the report is zero. Hence, the advisor is better off designing a report that avoids such signals. As for the former two types of signals, the advisor could equivalently combine all signals that induce subscribers to vote "for" into one, and all signals that induce them to vote "against" into the other. In the second step, we point out that since the advisor is the only source of information about the state, one can think of the game induced by the report as the baseline game (with a fully informative report) in which the signals in the report are referred to as "states." In this redefined game, an imperfectly informative report means that the value upon taking the "correct" decision (voting according to the report) is lower than one. The arguments in Section II.B then imply that a shareholder's value from the report is lower than if it were fully informative.

III. Evidence of the mechanism

This section discusses the evidence that aligns with the mechanism we are highlighting.

Our model predicts that shareholders may reach different conclusions when considering proxy advisors' reports alongside the recommendations, compared to voting based solely on those recommendations. There are several pieces of evidence corroborating this prediction.

Content of research reports and recommendations. Ertimur, Ferri, and Oesch (2013, 2018) study the content of ISS research reports on management proposals regarding executive compensation and director elections, respectively. They show substantial variability in the severity of concerns expressed in the report when ISS issues a negative recommendation. Moreover, shareholders are less likely to vote against management when the research report conveys less severe concerns. This aligns with the idea that many shareholders carefully read the reports, potentially arriving at a more favorable assessment of these management proposals when the report's content is less negative than the recommendation suggests.

Case studies. A compelling example of the contrast between the recommendation and the report is the 2024 high-profile vote on executive compensation at Tesla, when the company sought shareholder approval for Elon Musk's compensation package, previously endorsed by shareholders in 2018 but subsequently invalidated by court.[22] ISS issued a negative recommendation, but was more positive in the report underlying the negative recommendation:

- "Tesla's financial performance has skyrocketed along with the value of the earned grant... The company has generated significant value for shareholders and achieved many of the award's performance hurdles... Musk has not received additional compensation since the award, and investors have shared in mountainous market cap growth over the course of the performance award."
- "It is undeniable that the performance hurdles required under this grant structure represented extraordinary growth in market capitalization and earnings for Tesla and value for its shareholders. The grant structure arguably contributed to the company's significant market cap growth as well as a flip from negative to positive earnings."
- "Some investors may find the board's argument compelling, that it would be unfair for CEO Musk not to receive the award, which was previously approved by shareholders, and after achieving the high performance hurdles."

The proposal passed, although many small institutional investors voted against. Notably, the largest institutional owners of the firm, Vanguard and BlackRock, voted in favor, deviating from the negative recommendation. All these features are in line with our model's predictions.

Another example, involving the 2022 sale of Avalara, is described in Section I of the Internet Appendix. Avalara's board pushed for shareholder approval of the transaction, while a major shareholder strongly opposed it. ISS issued a "for" recommendation, leading Avalara's management to promptly disclose this information via a press release (Figure A.5 in the Internet Appendix), but its research report depicted a more pessimistic view.

[22] See "Tesla Shareholders Approve Big Stock Package for Musk," *The New York Times*, June 13, 2024.

Investors' perspectives. Our prediction that research reports are more informative and valuable to investors aligns with survey-based evidence on how investors perceive proxy advisory services. Bew and Fields (2012) found in their survey of large asset managers that "virtually unanimously, research participants highlighted the value they derive from ... [reports] ... digest[ing] and normaliz[ing] the vast quantities of data present in proxy statements in a short period of time," but concluded that the "value of ... voting recommendations is distinctly secondary." Similarly, Michelle Edkins, Global Head of BlackRock's Investment Stewardship team, has praised the information content of proxy advisors' research reports stating "we get to read a lot [of proxy statements], and it can be very hard to find the pertinent information ... so having that information synthesized and accessible is hugely important to us being able to take an informed decision," but emphasized that they rely less on specific recommendations.[23]

One interpretation of the information design policy $\mathcal{S}$ is proxy advisors' voting guidelines, which map the contents of the research report into vote recommendations. As we argue next, the way proxy advisors design their voting guidelines is consistent with our mechanism.

One-size-fits-all-approach. How can proxy advisors produce unbiased reports but recommendations that favor the less likely alternative, given that explicit disagreement between the report and recommendation can be problematic? We argue that one way to achieve this is through the "one-size-fits-all" approach in voting guidelines. Proxy advisors have been criticized for issuing recommendations according to pre-specified rules, neglecting the unique circumstances of each company.[24] One-size-fits-all (also referred to as "blanket") guidelines naturally arise in the context of our model, as a way to implement recommendations with controversy without creating a contradiction between the recommendations and the reports.

A good example of a one-size-fits-all approach is the "overboarded director" policy: ISS's guidelines state that it will recommend against directors who sit "on more than five public company boards." Other things equal, directors' busyness is likely to correlate negatively with their contribution to firm value, leading to partially informative recommendations. However, this guideline does not consider other director characteristics. Many directors who sit on multiple boards could have exceptional qualifications and be in high demand for this reason. Shareholders who merely observe a negative recommendation, without delving into the research

[23]See the transcript of the SEC's Proxy Advisory Services Roundtable in 2013. The valuable information role of research reports alings with the survey evidence by McCahery, Sautner, and Starks (2016): 55% of their respondents believe that proxy advisors help them make more informed voting decisions.

[24]See Hayne and Vance (2019) and the SEC's 2013 Proxy Advisory Services Roundtable discussion.

report, may find it challenging to discern whether the recommendation reflects the director's quality or solely his multiple board positions. Many such shareholders will follow the negative recommendation, increasing the likelihood of a close vote. In contrast, shareholders who read the report gain access to the underlying facts and thus can consider a broad range of factors: reports typically describe the director's qualifications, level of independence, potential conflicts of interest, attendance records at board meetings, and other characteristics, which may together lead shareholders to vote in favor. Importantly, the negative recommendation from the proxy advisor simply follows the pre-established guidelines, and thus does not contradict the content of the report. For example, typical language used in ISS reports on such cases is "ISS believes that directors who are overextended may be jeopardizing their ability to serve as effective representatives of shareholders," without a discussion of how this factor should be weighed against other director characteristics.

Glass Lewis has similar guidelines regarding overboarded directors. An illustrative example is its recommendation against two directors of HSBC Holdings in 2017, each of whom sat on six boards. The report noted: "we generally believe that the time commitment required by these combinations of board chairmanships/memberships ... may preclude these directors from dedicating the time necessary to fulfill the responsibilities required of a director of this Company." However, right after this statement, the report noted: "We recognise that shareholders may nonetheless choose to support these directors' election given their adequate attendance record during the past fiscal year, the depth of their relevant experience, and the absence of any other apparent conflicts of interest."

Proxy advisors use the one-size-fits-all approach for other issues, such as executive compensation.[25] Our results predict that large, active voters (those with high v_i in the model) are less likely to follow recommendations based on such blanket guidelines and will vote on a case-by-case basis. This prediction is consistent with Iliev and Lowry (2015), who find that larger funds and funds that are part of larger fund families are significantly less likely to take a blanket approach toward voting compared to ISS.

One-size-fits-all and uninformative recommendations. Our model predicts that it may be optimal for the proxy advisor to design completely uninformative recommendations

[25]For example, in addition to various qualitative and quantitative aspects of the compensation plan that ISS considers, there are several "overriding negative factors" that will lead to a negative recommendation regardless of how the company scores on other factors. One such overriding factor is concentration of pay at the top executive level: ISS will recommend against if a metric it refers to as "concentration ratio" exceeds 30% for the CEO (see "Equity Compensation Plans Frequently Asked Questions," ISS, January 30, 2023).

(Proposition 5). Such recommendations can be implemented by always recommending the same action (always against or always in favor) for certain proposals – an extreme form of the one-size-fits-all approach. Both ISS and Glass Lewis use such "uninformative" recommendations for certain types of proposals. A good example is board declassification: ISS guidelines in 2024 and earlier years include a "general recommendation" to "vote against proposals to classify (stagger) the board" and "vote for proposals to repeal classified boards." In line with this, our empirical analysis of shareholder proposals to declassify the board from 2010 to 2019 revealed zero instances of a negative ISS recommendation. Consequently, shareholders receiving a positive recommendation on board declassification can infer no information about the value of this proposal for the specific firm in question from the recommendation alone.

Figure A.3 in the Internet Appendix shows that, even with these uninformative recommendations, there is substantial disagreement among shareholders regarding the value of board declassification proposals, with average support levels of 73.7%. This disagreement aligns with the ongoing debate and inconclusive evidence on whether classified boards decrease value by entrenching management (Bebchuk and Cohen, 2005; Faleye, 2007) or increase value by enabling a long-term focus (Johnson, Karpoff, and Yi, 2015; Cremers, Litov, and Sepe, 2017): the literature concludes that whether the net benefits are positive crucially depends on firm characteristics. Thus, the fact that ISS provides effectively uninformative recommendations on board declassification supports Proposition 5, which predicts that uninformative recommendations emerge when prior beliefs about the proposal are sufficiently uncertain.

Consistency in the frequency of negative recommendations. Empirical evidence shows that firms gradually adjust their governance policies in line with proxy advisors' guidelines, to avoid receiving a negative recommendation (Ertimur, Ferri, and Oesch, 2013; Larcker, McCall, Ormazabal, 2015; Cabezon, 2022). One might then expect that over time, it would become harder for the advisor to issue negative recommendations without contradicting the report, as there would be fewer firms violating the guidelines. This would imply that the proportion of negative recommendations should decline over the years, but this is not what happens. Figure A.2 in the Internet Appendix shows that the frequency of negative ISS recommendations on say-on-pay proposals and director elections has remained stable over time.

One way to explain these patterns is through the fact that proxy advisors frequently evolve their guidelines, introducing new criteria and considerations. Our model can interpret this as the advisor's way to maintain recommendations that favor the less likely alternative without contradicting the reports. According to interviews by Hayne and Vance (2019), market participants indeed perceive this practice as consistent with our mechanism: interviewees suggest that proxy advisors aim to enhance their "relevance" and thereby profits by evolving their guidelines and consistently giving a significant proportion of negative ("controversial") recommendations. For example, a former proxy advisor employee noted, "When you devise new voting recommendations, it keeps you in business ... If you keep evolving your guidelines, ... you are keeping yourself relevant." Another referred to this practice as "[PA Firm] is a for-profit organization so, at the end of the day, they have their investors who are expecting that they will generate profit." According to one of the compensation consultants,

> "it sure seems that they have essentially a standard of "We need to vote against a certain percentage of companies every year to keep people on their toes about how we're evaluating things." ... performance has gotten better, but they're still recommending against the same 10% of companies every year."[26]

Rubberstamping. Our model predicts a distinct voting pattern: When the advisor recommends in favor of the expected alternative, shareholders will "rubberstamp" this recommendation. In contrast, if the recommendation goes against the expected alternative, there should be no rubberstamping and, instead, a disagreement among shareholders and a high chance of a close vote. Both say-on-pay proposals and management nominees in uncontested director elections receive, on average, very high support from shareholders, 91% and 95%, respectively (Table A.1 in the Internet Appendix), so it is reasonable to assume that "for management" is the expected alternative. Then, we expect shareholders to rubberstamp pro-management recommendations, but not recommendations against management, for both types of proposals.

The evidence is in line with this prediction. For example, in the sample of Malenko and Shen (2016), positive ISS recommendations on say-on-pay receive 93% average support with no failed proposals, consistent with "rubberstamping." In contrast, negative recommendations result in 69% average support and an 11% chance of the proposal failing. Figure A.1 in the Internet Appendix illustrates this distribution of voting outcomes, highlighting differences between positive and negative ISS recommendations on say-on-pay over 2011 to 2019.[27] Ertimur, Ferri, and Oesch (2018) document a similar pattern for director elections: directors endorsed

[26] These quotes are from Hayne and Vance (2019) and the Online Appendix to their paper.

[27] Many say-on-pay proposals with a negative ISS recommendation result in close votes, that is, votes with support ranging between 40% and 60%. Panel B of Table A.1 in the Internet Appendix extends this observation to all types of proposals in ISS Voting Analytics. It shows that there about 500 close votes per year, and for the vast majority of them, ISS recommendation opposes that of the management. In contrast, when ISS agrees with management, close votes are extremely rare.

by either ISS or Glass Lewis receive majority support, with most getting over 90%, whereas negative recommendations result in much greater vote dispersion.

The rubberstamping of recommendations that are consistent with the priors leads to another implication: lopsided votes (i.e., votes where almost all shareholders vote the same way) are relatively frequent. Lopsided votes are prevalent in practice (see the discussion above) but rarely arise in models of strategic voting, so their relatively high frequency is a distinguishing feature of our model. A more detailed discussion is in Section II of the Internet Appendix.

IV. Implications

Reinterpreting the evidence on votes and the right "benchmark". Empirical evidence shows that proxy advisors are often less supportive of management than major index funds (see Brav et al. (2024) for proxy contests and Bolton et al. (2020) and Bubb and Catan (2022) for other types of proposals). Large index funds are frequently criticized because their greater support for management relative to proxy advisors' recommendations is viewed as a sign of passivity or pro-management bias. In contrast, our results suggest a different interpretation of such voting behavior and emphasize that proxy advisors' recommendations may not be the most suitable benchmark (see also discussion in Spatt, 2021). Empirically, management proposals are typically approved with a high support rate, suggesting that "for management" is often the a priori more likely alternative (see Table A.1 in the Internet Appendix).[28] Assuming this is the case, our model predicts that proxy advisors may have an excessive tendency to recommend against management, so shareholders who deviate from negative recommendations and support management could be optimally correcting for this tendency.

If proxy advisors' recommendations are not always the suitable benchmark, what could be an alternative? Our results imply that the votes of large asset managers might reflect a more appropriate benchmark. Recall that one interpretation of v_i, the extent to which a shareholder cares about the proposal, is the shareholder's size. Under this interpretation, the model predicts that institutional investors that manage larger portfolios are more likely to vote based on their analysis of proxy advisors' reports, rather than purely based on recommendations. This aligns with the idea in Iliev and Lowry (2015), Iliev, Kalodimos, and Lowry (2021), and Gantchev and Giannetti (2021) that larger asset managers tend to be "active voters." Moreover, in our model, the votes of shareholders with large v_i are both informed and unbiased, and thus arise

[28] This is expected, as the proposals are endogenously put forward by management and thus are likely crafted to appeal to shareholders. Additionally, many of these votes concern matters that are usually non-contentious.

as a suitable benchmark. Consistent with this, the structural estimates in Blonien et al. (2024) suggest that the votes of Vanguard and BlackRock (post-2009) have a relatively high precision, while the quality of ISS recommendations features an asymmetry in line with our predictions: high (low) precision for recommendations supporting (opposing) management.

Our results also suggest several policy implications.

Focus on a different source of tension. Policy discussions often center on the concern that some proxy advisors receive consulting fees from the companies they evaluate, which may lead to a pro-management bias in recommendations (Li, 2018). Our paper underscores a very different source of tension, which is inherent to the core business model of proxy advisors – selling information to voters. The tendency to favor the less likely alternative is not related to the consulting business, even if it exists, and hence cannot be alleviated by separating the two businesses or disclosing the advisor's consulting relationships with companies, which have been the focus of policy proposals (SEC press release 2020-161). This is important because, according to Hayne and Vance (2019), proxy advisor representatives "indicated that selling research reports and voting recommendations to institutional investors is their primary source of revenue." Despite this tension and the tendency to favor the less likely alternative, our research highlights that recommendations contain valuable information, thereby improving the votes of small shareholders compared to the scenario without any recommendations at all.

Shifting attention towards research reports. A key focus of policy discussions has been on the quality and influence of proxy advisors' recommendations. In contrast, the quality of research reports has received less attention, possibly because of the perception that recommendations adequately represent the essence of the underlying reports. Our paper highlights an important distinction: the two are not necessarily the same, and an undue emphasis on recommendations may restrict a comprehensive understanding of proxy advisors' information role. As we show, proxy advisors have incentives to maintain the highest possible quality in their reports. Thus, even when recommendations fall short, research reports may still be a valuable source of information, and the evidence in Section III supports the idea that institutional investors find them very valuable. This implication also underscores the need for more empirical research on the content of reports and their use by investors.

Should recommendations be offered alongside research reports? In the case where reports are informative and unbiased while recommendations may favor the less likely alternative, a question arises: should proxy advisors continue to provide both? One policy that could address the issue highlighted in our paper would involve banning proxy advisors from including vote recommendations in their reports, since such recommendations can be easily followed. Instead, the reports would only contain relevant analysis and facts. However, as we argue, this is not a straightforward solution. Even if recommendations favor the less likely alternative, they are still partially informative. The alternative would entail many shareholders being left without any information, resulting in uninformed voting. We formally study the ban on recommendations in the Appendix. Our analysis confirms a downside to this approach: it would reduce the information available to small investors, leading to less informed voting by them and a potential decrease in firm value. Nevertheless, as we also show, there is a potential upside: such a ban would increase the proportion of shareholders reading the reports, meaning a larger share of shareholders would vote based on comprehensive information.

V. Discussion

In this section, we discuss several important factors that the basic model abstracts from, and study how they influence the design of proxy advisors' recommendations.

Litigation pressure. We have assumed that shareholders become subscribers solely to make more informed voting decisions. This assumption fits well in settings where investors face no regulatory pressure regarding their voting practices. In reality, investors also subscribe for non-informational reasons, such as to mitigate the risk of potential litigation. As Daniel Gallagher noted, "for the price of purchasing the proxy advisory firm's recommendations, an investment adviser could ward off potential litigation over its conflicts of interest."[29] In Section IV.E of the Internet Appendix, we extend the model to account for these non-informational motives. In this extension, shareholder i's value from subscribing is the sum of two components: the value of making an informed voting decision, which is the same as in the baseline model and given by (9), plus an additional component unrelated to information, capturing litigation protection. Intuitively, this implies that for small institutional investors (e.g., those with small v_i), litigation concerns are the primary motive for subscribing to proxy advisors. In contrast, for large investors (those with large v_i), informational motives are also important, consistent with the interview evidence in Bew and Fields (2012) and McCahery, Sautner, and Starks (2016).

[29] See Gallagher (2014). Additionally, investors may subscribe to access extra services such as vote execution and regulatory form filing (Shu, 2024).

In this setting, as in the baseline model, proxy advisor reports are fully informative and unbiased, while the information content of recommendations is smaller. If the litigation protection motive is moderate, recommendations are partially informative and favor the less likely alternative. With heightened litigation concerns, recommendations become completely uninformative if shareholders are sophisticated and learn from being pivotal. However, if shareholders are naive and do not learn from pivotal events, recommendations that favor the less likely alternative remain optimal, even in the presence of strong litigation motives.[30]

In recent years, the SEC has moved away from the notion that merely subscribing to proxy advisors provides adequate protection against litigation. For instance, in 2018, the SEC withdrew two 2004 no-action letters that had initially encouraged institutional investors to rely on proxy advisors,[31] and in 2019, the SEC stated that investment advisors relying on proxy advisors "must take reasonable steps to ensure the use of that advice is consistent with their fiduciary duties" (Press Release 2019-158). These changes could be interpreted as moving closer to the baseline model, where proxy advisors primarily serve an informational role for their clients. Our findings suggest that they should not have a strong effect if shareholders are naive but could lead to more controversial recommendations if shareholders are fully rational.

Ideological preferences. Our setting focuses on investors who care exclusively about value maximization. However, some shareholders' votes can be driven by ideological preferences, for example, due to their views on environmental and social issues (Bolton et al., 2022; Dikolli et al., 2022; Malenko and Malenko, 2024) or attitudes towards management (Matvos and Ostrovsky, 2010; Jackson and Zytnick, 2022). In Section IV.F of the Internet Appendix, we extend our model to include a fraction of shareholders who vote ideologically, favoring or opposing proposals regardless of the information. We show that the optimal design of recommendations crucially depends on whether investors' preferences are aligned or heterogeneous, and whether they support or contradict prior beliefs. For example, suppose the prior belief that the proposal is value-increasing is high, but many investors have strong ideological preferences against it. In this case, close votes are likely despite asymmetric prior beliefs. Essentially, ideological investors' votes counteract those of uninformed non-ideological investors, leading

[30] Intuitively, the stronger the benefit from litigation protection, the greater the equilibrium fraction of subscribers. With many subscribers, non-subscribers learn a lot from being pivotal, so their votes are not as sensitive to recommendations, making uninformative recommendations optimal. However, if shareholders do not learn from being pivotal, this effect does not arise, and recommendations with controversy remain optimal.

[31] These letters were the "Investment Advisers Act of 1940 - Rule 206(4)-6" 2004 letter to Egan Jones and the "Investment Advisors Act of 1940 - Rule 206(4)-6" 2004 letter to ISS.

to significant disagreement and a close outcome. As a result, the advisor's recommendations become uninformative, similar to Proposition 5. In contrast, if investors' preferences prompt them to vote in the same direction as prior beliefs, or if investors' preferences are heterogeneous and cannot counteract strong prior beliefs, close votes remain unlikely. In such scenarios, the advisor continues to benefit from recommendations that favor the less likely alternative.[32]

Coordination among shareholders. In our setting, shareholders would benefit from coordinating their information acquisition and voting decisions. At the information acquisition stage, they would benefit if only the shareholder with the highest concern for the vote outcome (v_i) acquired information. However, such coordination would require shareholders to truthfully reveal their v_i to each other, which would not be incentive compatible, as each of them benefits from free-riding on information acquisition of others. At the voting stage, shareholders would benefit if the vote outcome were only based on the information of those who buy the report. This can be achieved in two ways, either by informed shareholders communicating their information to others, or by uninformed shareholders abstaining from the vote.

First, consider communication. Even though sharing the advisor's report itself is likely not possible given contractual restrictions, subscribers could disclose how they are going to vote. We assume that this does not happen because, in practice, such communication is very rare. A key reason is "acting in concert" rules, which Puchniak and Varottil (2023) describe as a "significant legal hurdle" to shareholder coordination. According to these rules, communication among shareholders could be considered as forming a "group," potentially resulting in substantial expenses. This may include imposing significant disclosure costs under Schedule 13D, triggering a poison pill, or even, in some circumstances, requiring shareholders to make a mandatory takeover offer to the firm. According to Delaware Court of Chancery, acting in concert rules can block "potentially benign stockholder communications," giving the company "discretion to determine whether 'plus' factors as innocuous as 'exchanging information, attending meetings, [or] conducting discussions' can trigger" the poison pill.[33] Another hurdle to communication is that based on anecdotal evidence, publicly disclosing one's vote against

[32] When ideological considerations are present, proxy advisors sell not only benchmark recommendations, which are the focus of this paper, but also custom recommendations, which are tailored to shareholders' individual ideologies. For example, the votes of ideological investors in the extension could be driven by such custom recommendations. Hu, Malenko, and Zytnick (2024) analyze custom recommendations and their implications for shareholders' information acquisition decisions. Our paper, in contrast, focuses on how the design of *benchmark* recommendations is affected by shareholders' ideological considerations.

[33] See 2021 WL 754593 (Del. Ch.). Similarly, the 2011 report by Dechert LLP states that "shareholder concern about unintentionally forming a group has chilled communications among large holders of shares in U.S. public companies." Black and Coffee (1994) and Puchniak and Varottil (2023) offer extensive discussions.

management is viewed by management much more negatively than a negative (but private) vote per se, so institutional investors are reluctant to disclose their votes to avoid managerial retaliation. Finally, in reality, shareholders may have heterogeneous preferences about the proposal (e.g., due to differences in investment horizon or environmental and social attitudes), which can make credible communication challenging (Crawford and Sobel, 1982).

Next, consider abstention. Note that even if we allowed for abstention, the equilibrium analyzed in the paper, in which no shareholder abstains, would still exist.[34] At the same time, there can also be other equilibria, in which informed shareholders (who observe the report) vote, whereas the shareholders who only observe the recommendation abstain. The equilibrium without abstention is likely to be empirically more relevant, because institutional investors rarely abstain, largely because of perceived regulatory pressure.[35] Furthermore, while abstention can result in more efficient voting outcomes keeping shareholders' information endowments fixed, the option to abstain might actually decrease shareholders' information endowments due to endogenous information acquisition decisions and strategic pricing by proxy advisors.

Costly information acquisition. We have assumed that the advisor knows the state with certainty and at no cost. In reality, information production is costly and reveals the state with noise. Section IV.C of the Internet Appendix considers an extension with costly information acquisition by the advisor. As we show, the advisor underinvests in information because part of the surplus from a more precise signal is captured by the shareholders. Moreover, the advisor invests in information less when the prior belief about the state is strong (μ is close to 0 or 1) compared to the case when there is considerable uncertainty. Nevertheless, our results continue to hold: the advisor benefits from designing partially informative recommendations that favor the less likely alternative, but reveals all the information it acquires through the report.

Multiple proposals and firms. Our model features one proposal and one firm. In practice, proxy advisors sell their research as a bundle: an investor receives reports for all

[34] In particular, the equilibrium in Proposition 2 would be an equilibrium of the model in which shareholders can abstain, and in the event of a tie, the proposal is implemented randomly. Indeed, consider a non-subscribing shareholder. Conditional on being pivotal and the public recommendation, his posterior belief about the state is $\frac{1}{2}$. If the shareholder does not abstain, he randomizes between voting in favor and against. If he deviates and abstains, there is a tie, and the proposal is implemented randomly. Since the shareholder's posterior belief is $\frac{1}{2}$, his assessed probability of the correct decision being made is $\frac{1}{2}$ regardless of whether he abstains or not. Thus, he is indifferent between abstaining and not abstaining, so the equilibrium continues to exist. The same argument applies if, in the event of a tie, the proposal is always accepted or is always rejected.

[35] See https://www.sec.gov/rules/interp/2019/ia-5325.pdf.

firms in his portfolio, and the report for each firm covers all proposals on the agenda. Such bundling does not change our conclusions. To see this, suppose that there is still one firm, but with K proposals on the agenda, and the advisor combines its research on these proposals in one report. For any of K proposals, as in the baseline model, the value of the report to a shareholder increases in the probability that the vote on this proposal is close. Hence, for a given fee, the advisor's problem can be considered as K separate problems, and the incentives to create controversy arise for each of K proposals. (Of course, the fee charged for the report will depend on the combined value of K proposals.) A similar argument applies to a subscription that covers multiple firms in an investor's portfolio.

VI. Conclusion

Proxy advisors play a significant role as information providers to shareholders. Their advice has two key elements: detailed research reports and vote recommendations. Our paper highlights an important distinction between these two elements. We show that proxy advisors benefit from producing informative and unbiased reports, but only partially informative, asymmetrically biased recommendations, which favor the less likely alternative. Such recommendations increase the incidence of close votes, thereby enhancing the relevance of proxy advice and the value of proxy advisory services. Shareholders' concerns about litigation and the extent to which their preferences are ideological have important implications for the information content of proxy advisors' recommendations. Our results offer a new perspective on the empirical evidence on shareholder voting, suggesting that proxy advisors' recommendations may not be the right benchmark for evaluating shareholders' votes, and that the votes of large asset managers may potentially be a more suitable benchmark. They also suggest that an exclusive focus on recommendations can limit a comprehensive understanding of proxy advisors' information role: as we show, proxy advisors have incentives to produce reports of high quality. Thus, it is important to pay attention to the content of both recommendations and reports, and to gain insight into how shareholders employ each of these elements in their voting practices.

Appendix A. Ban on recommendations

One policy that could address the issue highlighted in our paper is to prohibit the advisor from issuing recommendations, as they are easy to follow. A potential way to implement this policy is to ban proxy advisors from including vote recommendations in their reports, and only allowing the reports to contain relevant analysis and facts. Shareholders would then have to read the entire report to understand the pros and cons of the proposal, and would make decisions based on more comprehensive information.

Formally, suppose the advisor is not allowed to generate any public signal $\mathcal{S}$. Note that the absence of a public signal is equivalent to an uninformative recommendation policy $\mathcal{S}_{uninf}$, that is, the one satisfying $\Pr(\mu_s = \mu) = 1$. Hence, this constrained model can be analyzed in the same way as the general model analyzed above. In particular, a fully informative private research report is still optimal if there is a ban on recommendations (the proof of Proposition 8 applies to any policy $\mathcal{S}$, including an uninformative one). In this sense, the ban on recommendations removes the advisor's tendency to favor the a priori unlikely alternative in the information provided to the shareholders.

Let q_{uninf} denote the expected fraction of subscribers that the advisor targets if the recommendation is uninformative, and let $(q^*, \mathcal{S}^*)$ denote the optimal expected fraction of subscribers and recommendation policy in the basic model. We are interested in comparing $\Pr(d = \theta | q^*, \mathcal{S}^*)$ and $\Pr(d = \theta | q_{uninf}, \mathcal{S}_{uninf})$, where $\Pr(d = \theta | q, \mathcal{S})$ is the probability of a correct decision being made when the recommendation policy is $\mathcal{S}$ and the probability of becoming a subscriber is q. Since the ban on recommendations only has bite if $\mathcal{S}^*$ is partially informative, we assume that the conditions of Proposition 6 are satisfied, so that $\mathcal{S}^*$ generates controversy.

Our model suggests that a ban on recommendations has both an upside and a downside, and the effectiveness of the ban depends on the trade-off between the two. The downside is that the ban on recommendations reduces information of non-subscribers (or small investors – those with a small v_i) and leads to less informative voting by them (partially informative recommendations that favor the less likely alternative are better than no information at all). If the fraction of subscribers stayed the same, the reduction of information of non-subscribers would lead to lower-quality voting outcomes. The upside of the ban is that it induces the proxy advisor to target a higher fraction of subscribers. Intuitively, when priors are asymmetric and there are no public recommendations, the probability of a split vote is relatively small because non-subscribing shareholders are all likely to vote according to the priors. By marginally increasing q, the advisor makes non-subscribing shareholders' votes less sensitive to the priors (see Figure 2) and, in addition, increases the probability that the votes of subscribers will counteract the votes of non-subscribers. Both effects increase the probability of a split vote and thereby shareholders' willingness to pay for the report. The next proposition formalizes this trade-off (the proof is in Section III.E of the Internet Appendix):

PROPOSITION A1 (Ban on public recommendations): *Suppose that S^* is a partially informative recommendation from Proposition 6, and suppose distribution $H(\cdot)$ has an increasing hazard rate. Then a ban on public recommendations:*

(i) decreases the quality of decision-making if the fraction of subscribers is kept constant: $\Pr(d = \theta | q^*, \mathcal{S}^*) > \Pr(d = \theta | q^*, S_{uninf})$*;*

(ii) increases the probability with which a shareholder becomes a subscriber, that is, $q_{\text{uninf}} > q^*$*.*

To illustrate this trade-off most starkly and show how either of the two effects can dominate, we consider an example in which there are only two types.

Example 2: Suppose $N = 25$, $\mu = 0.9$, and $v_i \in \{v_L, v_H\}$, where $v_L = 1 < v_H$ and $\Pr(v_i = v_H) = 0.1$. We can think of this distribution as the limit of a continuous distribution with two modes, and all the above arguments apply (the derivations for this example are in Section III.E.2 of the Internet Appendix). If v_H is sufficiently high (above 5), then without a ban on recommendations, the advisor finds it optimal to set a high fee, so that only shareholders of type v_H subscribe to the report (i.e., $q = 0.1$), and to produce partially informative recommendations. In particular, the positive recommendation induces a posterior of 1, while the negative recommendation induces a posterior of 0.62, as in Panel C of Figure 3. Then, voting results in the correct decision with probability 92.5%. If a ban on public recommendations is introduced, there are two cases. If v_H is not too high, for example, $v_H = 5$, the advisor finds it optimal to sell the report to some of the low types as well. Specifically, we show that the optimal fee induces $q = 0.257$ (0.1 of type v_H and 0.157 of type v_L), in which case voting results in the correct decision with probability 95.65% > 92.5%. Intuitively, since all low types have the same valuation, it is easy for the advisor to attract many additional subscribers by decreasing the fee by a small amount. As a result, the positive effect from additional subscribers, corresponding to part (ii) of Proposition A1, is very high and dominates the negative effect from the loss of a partially informative public signal. In contrast, if v_H is high, for example, $v_H = 7$, the advisor finds it optimal to sell the report to high types only, as without a ban. As a result, the positive effect from additional subscribers is non-existent, and the probability of the correct decision falls to 90.3% < 92.5%.

Appendix B. Proofs of main results

The appendix contains the proofs of the main results. The proofs of supplementary results and extensions are in the Internet Appendix.

Proof of Proposition 1: Proven in the main text.

Proof of Proposition 2: Since we focus on weakly undominated strategies, subscribing shareholders vote according to the state: $a_i = \theta$. Next, plugging (5)-(6) into (4) and simplifying, we obtain

$$\Pr(\theta = 1|s, Piv) = \frac{\left((q + (1-q)\pi)(1-\pi)\right)^{\frac{N-1}{2}} \mu_s}{\left((q + (1-q)\pi)(1-\pi)\right)^{\frac{N-1}{2}} \mu_s + \left(\pi(1 - \pi(1-q))\right)^{\frac{N-1}{2}} (1-\mu_s)}. \quad \text{(A1)}$$

First, consider a candidate equilibrium in which $\pi = 1$, that is, each non-subscribing shareholder votes in favor. Then $\Pr(Piv|\theta = 1, s) = 0$ (since all subscribers then also vote in favor) and $\Pr(Piv|\theta = 0, s) > 0$ (since $q \in (0,1)$ and all subscribers then vote against). These imply that if $\mu_s < 1$, then (1) shareholder i is pivotal with a strictly positive probability and (2) $\Pr(\theta = 1|s, Piv) = 0$, which together imply that the shareholder finds it strictly optimal to deviate and vote against, so this equilibrium cannot exist. If $\mu_s = 1$, then $\theta = 1$ with certainty, so both subscribing and non-subscribing shareholders vote in favor. Since the shareholder is then never pivotal, he is indifferent between voting for and against, and in particular, voting in favor is indeed optimal. Hence, equilibrium with $\pi = 1$ exists if and only if $\mu_s = 1$.

Second, consider a candidate equilibrium in which $\pi = 0$, that is, each non-subscribing shareholder votes against. Then $\Pr(Piv|\theta = 0, s) = 0$ (since all subscribers then also vote against) and $\Pr(Piv|\theta = 1, s) > 0$ (since $q \in (0,1)$ and all subscribers then vote in favor). These imply that if $\mu_s > 0$, then (1) shareholder i is pivotal with a strictly positive probability and (2) $\Pr(\theta = 1|s, Piv) = 1$, which together imply that the shareholder finds it strictly optimal to deviate and vote in favor, so this equilibrium cannot exist. If $\mu_s = 0$, then $\theta = 0$ with certainty, so both subscribing and non-subscribing shareholders vote against. Since the shareholder is then never pivotal, he is indifferent between voting for and against, and in particular, voting against is indeed optimal. Hence, equilibrium with $\pi = 0$ exists if and only if $\mu_s = 0$.

Third, consider a candidate equilibrium in which each non-subscribing shareholder votes $a_i = 1$ with probability $\pi \in (0,1)$. Since $q \in (0,1)$ and non-subscribing shareholders randomize, then for any $\mu_s \in [0,1]$, the shareholder is pivotal with a strictly positive probability. Hence, to find it optimal to randomize between voting for and against, he must be indifferent between proposal acceptance and rejection, which is the case if and only if $\Pr(\theta = 1|s, Piv) = \frac{1}{2}$, or equivalently, using (A1),

$$((q + (1-q)\pi)(1-\pi))^{\frac{N-1}{2}} \mu_s = (\pi(1-\pi(1-q)))^{\frac{N-1}{2}} (1-\mu_s). \tag{A2}$$

In particular, this equilibrium can only exist if $\mu_s \in (0,1)$: if the state is known with certainty ($\mu_s = 0$ or 1), and since the shareholder is pivotal with a strictly positive probability, he is never indifferent between voting for and against.

Denoting $z_s \equiv \left(\frac{\mu_s}{1-\mu_s}\right)^{\frac{2}{N-1}}$, we can rewrite (A2) as $w(\pi, z_s) = 0$, where

$$w(\pi, z_s) = (z_s - 1)(1-q)\pi^2 + (1 + (2q-1)z_s)\pi - z_s q. \tag{A3}$$

Note that

$$\begin{aligned} w(0, z_s) &= -z_s q < 0 & \text{(A4)} \\ w(1, z_s) &= (z_s - 1)(1-q) + 1 + (2q-1)z_s - z_s q = q > 0. & \text{(A5)} \end{aligned}$$

First, if $z_s = 1$, it is equivalent to $q(2\pi - 1) = 0 \Leftrightarrow \pi = \frac{1}{2}$. Second, if $z_s > 1$, then (since $-z_s q < 0$) this quadratic equation has a unique positive root $\pi(q, \mu_s)$, which is the larger of the two roots and is given by (7). Given (A4)-(A5), this root lies between 0 and 1. Third, if $z_s < 1$, $w(\pi, z_s)$ is an inverted parabola, and given (A4)-(A5), it has a unique root in $(0,1)$, which is the smaller of the two roots, and hence is also given by (7). Note that by l'Hopital's rule, $\lim_{\mu_s \to \frac{1}{2}} \pi(q, \mu_s) = \frac{1}{2}$, so the probability of voting in favor is continuous in μ_s.

Finally, we prove that $\pi(q, \mu_s)$ increases in μ_s. To see this, note that z_s increases in μ_s, and the derivative of $w(\pi, z_s)$ with respect to z_s is

$$\delta(\pi, q) \equiv (1-q)\pi^2 + (2q-1)\pi - q. \tag{A6}$$

Since $\frac{\partial \delta(\pi,q)}{\partial q} = -\pi^2 + 2\pi - 1 < 0$ and $\delta(\pi, 0) = \pi^2 - \pi < 0$ for $\pi \in (0,1)$, then $\delta(\pi, q) < 0$ for all $q \in (0,1)$. Hence, $w(\pi, z_s)$ is decreasing in z_s in the neighborhood of its root $\pi(q, \mu_s)$. When $z_s > 1$, $w(\cdot, z_s)$ is a parabola and $\pi(q, \mu_s)$ is the larger of the two roots, so it increases in μ_s. When $z_s < 1$, $w(\cdot, z_s)$ is an inverted parabola and $\pi(q, \mu_s)$ is the smaller of the two roots, so it also increases in μ_s.

Proof of Proposition 3: The statement of the proposition is proven in the main text. Here, we provide an equivalent representation of $\Pr(Piv|q,\mu_s)$, which will be useful in subsequent proofs. We can rewrite the indifference condition $\Pr(\theta=1|s,Piv)=\frac{1}{2}$ as

$$\Pr(Piv|\theta=1,s)\,\mu_s=\Pr(Piv|\theta=0,s)\,(1-\mu_s)\,. \tag{A7}$$

Using (A7) and (5)-(6), we get three equivalent formulas for $\Pr(Piv|q,\mu_s)$:

$$\begin{aligned}\Pr(Piv|q,\mu_s) &= \Pr(Piv|\theta=1,s)\,\mu_s+\Pr(Piv|\theta=0,s)\,(1-\mu_s) && \text{(A8)}\\ &= 2\,(1-\mu_s)\,C_{N-1}^{\frac{N-1}{2}}\left(\varrho_0(q,\mu_s)\,(1-\varrho_0(q,\mu_s))\right)^{\frac{N-1}{2}} && \text{(A9)}\\ &= 2\mu_s C_{N-1}^{\frac{N-1}{2}}\left(\varrho_1(q,\mu_s)\,(1-\varrho_1(q,\mu_s))\right)^{\frac{N-1}{2}}, && \text{(A10)}\end{aligned}$$

where

$$\begin{aligned}\varrho_1(q,\mu_s) &= q+(1-q)\,\pi(q,\mu_s) && \text{(A11)}\\ \varrho_0(q,\mu_s) &= (1-q)\,\pi(q,\mu_s) && \text{(A12)}\end{aligned}$$

are the equilibrium probabilities that a shareholder votes "for" conditional on $\theta=1$ and $\theta=0$, respectively (and on s), and $\pi(q,\mu_s)$ is the equilibrium probability that a non-subscribing shareholder votes "for," given by (7).

Proof of Lemma 1: The proof is provided in Section III.A of the Internet Appendix.

Proof of Proposition 4: We start by noting that since the equilibrium in the voting subgame depends on the recommendation only via $\Pr(Piv|q,\mu_s)$, the advisor's problem (14) can be solved using concavification of the function $\Pr(Piv|q,\mu_s)$, which is a function of a single variable μ_s. The concave closure of this function is a combination of either the function itself or a linear function connecting two points. Hence, it is without loss of generality to restrict attention to binary recommendations that induce at most two different posteriors.

We first prove part 1 (applying to small μ). It is easy to show that $\Pr(Piv|q,\mu_s)$ is increasing in μ_s if and only if $\mu_s<\frac{1}{2}$. Consider $\mu<\underline{\mu}$ from Lemma 1. Since $\Pr(Piv|q,0)=0$ and $\Pr(Piv|q,\mu_s)$ is strictly increasing and strictly convex in μ_s for $\mu_s\in\left(0,\underline{\mu}\right)$, then the concave closure of $\Pr(Piv|q,\mu_s)$ is linear in the neighborhood of $\mu_s=0$, and the optimal recommendation design for $\mu\in\left(0,\underline{\mu}\right)$ takes the following form: signal $s=0$ induces belief $\mu_0=0$ and signal $s=1$ induces belief μ_1, where μ_1 is one that maximizes

$$\Pr(Piv|q,\mu_1)\,\tau_1+\Pr(Piv|q,0)\,\tau_0 \tag{A13}$$

subject to the Bayes plausibility constraint $\mu_1\tau_1 + \mu_0\tau_0 = \mu$. Since $\mu_0 = 0$, the latter implies $\tau_1 = \frac{\mu}{\mu_1}$, and since $\Pr(Piv|q,0) = 0$, the average probability of a split vote is $\Pr(Piv|q,\mu_1)\frac{\mu}{\mu_1}$, and μ_1 solves

$$\mu_1 = \arg\max_m \frac{\Pr(Piv|q,m)}{m}\mu. \tag{A14}$$

The point μ_1 is one where the linear function that starts at $(0,0)$ is tangent to the function $\Pr(Piv|q,\cdot)$. To find μ_1, we substitute (A10) into (A14) and get

$$\frac{\Pr(Piv|q,m)}{m}\mu = 2C_{N-1}^{\frac{N-1}{2}}\left(\varrho_1(q,m)\left(1-\varrho_1(q,m)\right)\right)^{\frac{N-1}{2}}\mu, \tag{A15}$$

where $\varrho_1(q,m)$ is given by (A11). Hence, μ_1 solves

$$\max_m \left(\varrho_1(q,m)\left(1-\varrho_1(q,m)\right)\right)^{\frac{N-1}{2}}. \tag{A16}$$

Thus, if feasible, the optimal posterior μ_1 is such that $\varrho_1(q,\mu_1) = \frac{1}{2}$. In other words, the probability of a shareholder voting "for" conditional on a recommendation against the prior $(s=1)$ and the state being in line with the recommendation $(\theta = 1)$, is 50%. If $m = \frac{1}{2}$, then $\pi(q,m) = \frac{1}{2}$, so $\varrho_1(q,m) = \frac{1+q}{2} > \frac{1}{2}$. Hence, $m = \frac{1}{2}$ does not solve (A16) if there exists $m \in \left(\mu_1, \frac{1}{2}\right)$ that satisfies $\varrho_1(q,m) = \frac{1}{2}$. Consider $m \neq \frac{1}{2}$. Then, using (7), $\varrho_1(q,m) = \frac{1}{2}$ is equivalent to

$$\begin{gathered}(1-q)(1-\pi(q,m)) = \frac{1}{2} \Leftrightarrow \frac{1}{2} + \frac{2q - \sqrt{(z-1)^2 + 4q^2 z}}{2(z-1)} = \frac{1}{2} \\ \Leftrightarrow z^2 - 2\left(1-2q^2\right)z + 1 - 4q^2 = 0 \Leftrightarrow (z-1)\left(z-1+4q^2\right) = 0,\end{gathered} \tag{A17}$$

where $z \equiv \left(\frac{m}{1-m}\right)^{\frac{2}{N-1}}$. Since $m \neq \frac{1}{2}$, the only root is $z = 1-4q^2$. Equating it to $\left(\frac{\mu_1}{1-\mu_1}\right)^{\frac{2}{N-1}}$ implies that the optimal posterior is

$$\mu_1(q) \equiv \frac{\left(1-4q^2\right)^{\frac{N-1}{2}}}{1+\left(1-4q^2\right)^{\frac{N-1}{2}}}. \tag{A18}$$

Note that $\mu_1 < \frac{1}{2}$. It follows that the optimal recommendation induces beliefs $\mu_0 = 0$ and $\mu_1(q) \in \left(\mu, \frac{1}{2}\right)$ given by (A18). The average probability of a shareholder being pivotal given this recommendation design is $2C_{N-1}^{\frac{N-1}{2}}\left(\frac{1}{4}\right)^{\frac{N-1}{2}}\mu = \mu C_{N-1}^{\frac{N-1}{2}}2^{2-N}$. To show that this recommendation is indeed optimal for any $\mu \leq \mu_1(q)$, consider any other recommendation with posteriors $\mu_l < \mu_h$. The Bayes plausibility constraint implies that $\mu_l < \mu < \mu_h$ and that μ_l and μ_h are induced with probability $\frac{\mu_h - \mu}{\mu_h - \mu_l}$ and $\frac{\mu - \mu_l}{\mu_h - \mu_l}$, respectively. Note that for any μ_s, (A10) implies that

$$\Pr\left(Piv|q,\mu_s\right)=2\mu_s C_{N-1}^{\frac{N-1}{2}}\left(\varrho_1\left(q,\mu_s\right)\left(1-\varrho_1\left(q,\mu_s\right)\right)\right)^{\frac{N-1}{2}}\leq 2\mu_s C_{N-1}^{\frac{N-1}{2}}\left(\frac{1}{4}\right)^{\frac{N-1}{2}}=\mu_s C_{N-1}^{\frac{N-1}{2}}2^{2-N}. \tag{A19}$$

Hence, the average probability of being pivotal given (μ_l,μ_h) is

$$\begin{aligned}&\tfrac{\mu_h-\mu}{\mu_h-\mu_l}\Pr\left(Piv|q,\mu_l\right)+\tfrac{\mu-\mu_l}{\mu_h-\mu_l}\Pr\left(Piv|q,\mu_h\right)\\ \leq\ &\tfrac{\mu_h-\mu}{\mu_h-\mu_l}\mu_l C_{N-1}^{\frac{N-1}{2}}2^{2-N}+\tfrac{\mu-\mu_l}{\mu_h-\mu_l}\mu_h C_{N-1}^{\frac{N-1}{2}}2^{2-N}=\mu C_{N-1}^{\frac{N-1}{2}}2^{2-N},\end{aligned} \tag{A20}$$

as required.

Next, we prove part 2 (applying to large μ). Given the symmetry of the problem in μ around $\mu=\frac{1}{2}$, this proof is almost identical to the proof of part 1 and hence is relegated to Section III.B of the Internet Appendix. In that proof, we show that for any $\mu\geq\mu_0(q)$, the optimal recommendation induces beliefs $\mu_0(q)$ and $\mu_1=1$, and that the average probability of a shareholder being pivotal given this recommendation design is $(1-\mu)\,C_{N-1}^{\frac{N-1}{2}}2^{2-N}$.

Proof of Corollary 1: The statement directly follows from Proposition 4 and the fact that function $\mu_0(q)$ increases in q, $\lim_{q\to 0}\mu_0(q)=\frac{1}{2}$, and $\lim_{q\to\frac{1}{2}}\mu_0(q)=1$, and function $\mu_1(q)$ decreases in q, $\lim_{q\to 0}\mu_1(q)=\frac{1}{2}$, and $\lim_{q\to\frac{1}{2}}\mu_1(q)=0$.

Proof of Proposition 5: By Lemma 1, $\Pr\left(Piv|q,\mu_s\right)$ is strictly concave in $\left(\frac{1}{2}-\varepsilon,\frac{1}{2}+\varepsilon\right)$. Moreover, $\Pr\left(Piv|q,\mu_s\right)$ is increasing for $\mu_s<\frac{1}{2}$ and decreasing for $\mu_s>\frac{1}{2}$, and hence has a maximum at $\frac{1}{2}$. Therefore, the concavification approach implies that the optimal recommendation is uninformative.

Proof of Proposition 6: Since the equilibrium in the voting subgame depends on the recommendation only via $\Pr\left(Piv|q,\mu_s\right)$, we can without loss of generality focus on binary recommendations (for the same reasons as in the proof of Proposition 4).

To prove the first statement of the proposition, we first note that for any μ,

$$\begin{aligned}&\max_q\Pr\left(Piv|q,\mu\right)=\max_q 2\mu C_{N-1}^{\frac{N-1}{2}}\left(\varrho_1\left(q,\mu\right)\left(1-\varrho_1\left(q,\mu\right)\right)\right)^{\frac{N-1}{2}}\\ &\leq 2\mu C_{N-1}^{\frac{N-1}{2}}\max_{\varrho_1\in[0,1]}\left(\varrho_1\left(1-\varrho_1\right)\right)^{\frac{N-1}{2}}=\mu C_{N-1}^{\frac{N-1}{2}}2^{2-N},\end{aligned} \tag{A21}$$

where the first equality follows from (A10). Now consider an arbitrary binary recommendation, such that signal $s=0$ induces belief μ_l and signal $s=1$ induces belief μ_h, and an arbitrary choice of q. The Bayes plausibility constraint requires that $\mu_l\leq\mu\leq\mu_h$, and for any $\mu_l\in[0,\mu]$

and $\mu_h \in [\mu, 1]$, the Bayes plausibility constraint requires that $\tau_0 = \frac{\mu_h - \mu}{\mu_h - \mu_l}$ and $\tau_1 = \frac{\mu - \mu_l}{\mu_h - \mu_l}$. Then the expected profit of the advisor, $\Pi(q, \mu_l, \mu_h)$, satisfies

$$\Pi(q, \mu_l, \mu_h) \le q^* H^{-1}(1 - q^*)\, \mu C_{N-1}^{\frac{N-1}{2}} 2^{1-N} N. \tag{A22}$$

To see this, note that

$$\frac{2}{N}\Pi(q, \mu_l, \mu_h) = qH^{-1}(1-q)\left(\frac{\mu_h - \mu}{\mu_h - \mu_l}\Pr(Piv|q, \mu_l) + \frac{\mu - \mu_l}{\mu_h - \mu_l}\Pr(Piv|q, \mu_h)\right) \tag{A23}$$

$$\le \max_{q,\mu_l,\mu_h}\left\{\left(qH^{-1}(1-q)\right)\left(\frac{\mu_h - \mu}{\mu_h - \mu_l}\Pr(Piv|q, \mu_l) + \frac{\mu - \mu_l}{\mu_h - \mu_l}\Pr(Piv|q, \mu_h)\right)\right\} \tag{A24}$$

$$\le \max_q\left\{qH^{-1}(1-q)\right\}\max_{\mu_l,\mu_h}\left\{\frac{\mu_h - \mu}{\mu_h - \mu_l}\max_q \Pr(Piv|q, \mu_l) + \frac{\mu - \mu_l}{\mu_h - \mu_l}\max_q \Pr(Piv|q, \mu_h)\right\}$$

$$\le \max_q\left\{qH^{-1}(1-q)\right\}\max_{\mu_l,\mu_h}\left\{\frac{\mu_h - \mu}{\mu_h - \mu_l}\mu_l C_{N-1}^{\frac{N-1}{2}} 2^{2-N} + \frac{\mu - \mu_l}{\mu_h - \mu_l}\mu_h C_{N-1}^{\frac{N-1}{2}} 2^{2-N}\right\} \tag{A26}$$

$$= \max_q\left\{qH^{-1}(1-q)\right\}\mu C_{N-1}^{\frac{N-1}{2}} 2^{2-N} \le q^* H^{-1}(1-q^*)\,\mu C_{N-1}^{\frac{N-1}{2}} 2^{2-N}. \tag{A27}$$

The first equality follows from (10) and (11). The first inequality is trivial. The second inequality follows from the fact that the maximum of a product cannot be higher than the product of the maxima. The third inequality follows from (A21). The last equality follows algebraically. Finally, the last inequality follows from the fact that q^* maximizes $qH^{-1}(1-q)$.

Suppose that $\mu \le \mu_1(q^*)$ and $q^* < \frac{1}{2}$. Then, the proof of Proposition 4 shows that choosing $q = q^*$ and the information policy with $\mu_l = 0$ and $\mu_h = \mu_1(q^*)$ yields the expected profit of $q^* H^{-1}(1-q^*)\,\mu C_{N-1}^{\frac{N-1}{2}} 2^{1-N} N$. Hence, (A22) implies that $(\mu_l, \mu_h, q) = (0, \mu_1(q^*), q^*)$ is optimal. The optimal fee that induces q^* then follows directly from (11).

By symmetry of $\Pr(Piv|q, \mu)$ around $\mu = \frac{1}{2}$, the proof of the second statement of the proposition is analogous.

Proof of Proposition 7: The statement holds for any distribution H for which $qH^{-1}(1-q)$ is increasing on $[0, q^*]$ and $q^* > \frac{1}{2}$. The assumption of an increasing hazard rate is sufficient but not necessary: as we show at the end of this proof, an increasing hazard rate of $H(\cdot)$ guarantees that $qH^{-1}(1-q)$ is increasing on $[0, q^*]$.

Consider any $\mu \in (0, \frac{1}{2}]$ (the proof for $\mu \ge \frac{1}{2}$ is similar). Note that $\mu_1(q)$ defined in Proposition 4 decreases in q, $\lim_{q \to 0} \mu_1(q) = \frac{1}{2}$, and $\lim_{q \to \frac{1}{2}} \mu_1(q) = 0$. Hence, for any $\mu \in (0, \frac{1}{2}]$, we can define $\hat{q} \equiv \mu_1^{-1}(\mu) \in [0, \frac{1}{2})$, which is the highest q for which $\mu_1(q) \ge \mu$. For any $q < \hat{q}$, Proposition 4 implies that the optimal recommendation is partially informative, with the probability of signal $s = 1$ being $\frac{\mu}{\mu_1(q)}$, and the average probability of being pivotal being

$\mu C_{N-1}^{\frac{N-1}{2}} 2^{2-N}$. By continuity, when $q = \hat{q}$ and the recommendation becomes uninformative, the probability that a shareholder is pivotal is also $\mu C_{N-1}^{\frac{N-1}{2}} 2^{2-N}$. Since $\hat{q} < \frac{1}{2} \leq q^*$ and $qH^{-1}(1-q)$ is increasing on $[0, q^*]$, we have $qH^{-1}(1-q) < \hat{q}H^{-1}(1-\hat{q})$ for any $q < \hat{q}$. Since the advisor's expected profit is the average probability of being pivotal multiplied by $\frac{N}{2} H^{-1}(q) q$, the advisor is strictly better off picking an uninformative recommendation and the fee that induces $\hat{q}$, rather than picking some $q < \hat{q}$ and $\mu_0 = 0$, $\mu_1(q)$ (which is the optimal recommendation design for this q by Proposition 4), as required.

It remains to prove that an increasing hazard rate of $H(\cdot)$ implies that $\lambda(q) \equiv qH^{-1}(1-q)$ is increasing on $[0, q^*]$. Denoting $\eta(q) \equiv H^{-1}(1-q)$, we get

$$\lambda'(q) = H^{-1}(1-q) - \frac{q}{H'(H^{-1}(1-q))} = \eta(q) - \frac{1 - H(\eta(q))}{H'(\eta(q))}. \tag{A28}$$

Since the hazard rate $\frac{H'(\eta)}{1-H(\eta)}$ is increasing in η, the function $\eta - \frac{1-H(\eta)}{H'(\eta)}$ is increasing in η as well. Since $\eta(q)$ is decreasing, $\lambda'(q)$ is decreasing in q. Finally, note that when $q = 0$, $\lambda'(0) > 0$. It follows that if q^* is the maximum of $\lambda(q)$, then $\lambda(q)$ increases on $[0, q^*]$.

Proof of Proposition 8: The complete proof is provided in Section III.C of the Internet Appendix, and we only provide a sketch of the proof here.

Fix any realization of public signal s inducing posterior μ_s. Then, the set of signals R can be divided into three subsets: R_0, R_1, R_m, defined, respectively, as all signals that induce posteriors (conditional on r and s) strictly below $\frac{1}{2}$ (inducing subscribers to vote against); strictly above $\frac{1}{2}$ (inducing subscribers to vote in favor); and equal to $\frac{1}{2}$. Because shareholders are indifferent between voting for and against if their posterior belief is exactly $\frac{1}{2}$, the value of the report is zero conditional on the third set of signals, and hence it is optimal to not induce such beliefs. Moreover, because the probability of being pivotal is the same for all signals from R_0 (and for all signals from R_1), we can combine such signals into one. Thus, the problem reduces to a report with binary signals: $R = \{0, 1\}$. We next repeat the arguments in the proof of Proposition 2 and show that for any $\mu_s \in (0, 1)$, non-subscribing shareholders vote for the proposal with probability $\pi(q, \tilde{z}_s) \in (0, 1)$, where $\pi(q, \cdot)$ is given by (7) and

$$\tilde{z}_s = \left(\frac{p_1 - \frac{1}{2}}{p_1 - \mu_s} \frac{p_0 + \mu_s - 1}{p_0 - \frac{1}{2}} \right)^{\frac{2}{N-1}}, \tag{A29}$$

where $p_0 \equiv \Pr(\theta = 0 | r = 0, \mu_s)$ and $p_1 \equiv \Pr(\theta = 1 | r = 1, \mu_s)$. Finally, we repeat the arguments in Section II.B and, using (A29), derive the value of the report as a function of q, μ_s, p_0, and p_1. We complete the proof by showing that for any μ_s, this value is maximized if $p_0 = p_1 = 1$.

REFERENCES

Admati, Anat R., and Paul Pfleiderer, 1986, A monopolistic market for information, *Journal of Economic Theory* 39, 400–438.

Admati, Anat R., and Paul Pfleiderer, 1990, Direct and indirect sale of information, *Econometrica* 58, 901–928.

Alexander, Cindy R., Mark A. Chen, Duane J. Seppi, and Chester S. Spatt, 2010, Interim news and the role of proxy voting advice, *Review of Financial Studies* 23, 4419–4454.

Alonso, Ricardo, and Odilon Camara, 2016, Persuading voters, *American Economic Review* 106, 3590–3605.

Alonso, Ricardo, and Konstantinos E. Zachariadis, 2021, Persuading large investors, Working paper, London School of Economics.

Austen-Smith, David, and Jeffrey S. Banks, 1996, Information aggregation, rationality, and the Condorcet Jury Theorem, *American Political Science Review* 90, 34–45.

Bardhi, Arjada, and Yingni Guo, 2018, Modes of persuasion toward unanimous consent, *Theoretical Economics* 13, 1111–1149.

Bar-Isaac, Heski, and Joel D. Shapiro, 2020, Blockholder voting, *Journal of Financial Economics* 136, 695–717.

Bebchuk, Lucian A., and Alma Cohen, 2005, The costs of entrenched boards, *Journal of Financial Economics* 78, 409–433.

Bergemann, Dirk, and Alessandro Bonatti, 2019, Markets for information: An introduction, *Annual Review of Economics* 11, 85–107.

Bew, Robyn, and Richard Fields, 2012, Voting decisions at US mutual funds: How investors really use proxy advisers. IRRC Institute.

Black, Bernard S., and John C. Coffee Jr., 1994, Hail Britannia: Institutional investor behavior under limited regulation, *Michigan Law Review* 92, 1997–2087.

Blonien, Patrick, Alan D. Crane, Kevin Crotty, and David De Angelis, 2024, Proxy advice and errors in shareholder voting, Working paper.

Bolton, Patrick, Tao Li, Enrichetta Ravina, and Howard Rosenthal, 2020, Investor ideology, *Journal of Financial Economics* 137, 320–352.

Boone, Audra L., Gillan, Stuart, and Mitch Towner, 2020, The role of proxy advisors and large passive funds in shareholder voting: Lions or lambs? Working paper.

Bouton, Laurent, Aniol Llorente-Saguer, Antonin Macé, and Dimitrios Xefteris, 2021, Voting in shareholders meetings, NBER Working Paper 29005.

Bouvard, Matthieu, and Raphael Levy, 2018, Two-sided reputation in certification markets, *Management Science* 64, 4755–4774.

Brav, Alon, Wei Jiang, Tao Li, and James Pinnington, 2023, Shareholder monitoring through voting: New evidence from proxy contests, *Review of Financial Studies* 37, 591–638.

Brav, Alon, and Richmond D. Mathews, 2011, Empty voting and the efficiency of corporate governance, *Journal of Financial Economics* 99, 289–307.

Bubb, Ryan, and Emiliano Catan, 2020, The party structure of mutual funds, *Review of Financial Studies* 35, 2839–2878.

Buechel, Berno, Lydia Mechtenberg, and Alexander F. Wagner, 2024, When two experts are better than one: The example of shareholder voting, Working Paper, University of Fribourg.

Cabezon, Felipe, 2022, Executive compensation: The trend toward one size fits all, Working Paper, Virginia Tech.

Cespa, Giovanni, 2008, Information sales and insider trading with long-lived information, *Journal of Finance* 63, 639–672.

Chan, Jimmy, Seher Gupta, Fei Li, and Yun Wang, 2019, Pivotal persuasion, *Journal of Economic Theory* 180, 178–202.

Chang, Briana, and Martin Szydlowski, 2020, The market for conflicted advice, *Journal of Finance* 75, 867–903.

Choi, Stephen, Jill Fisch, and Marcel Kahan, 2009, The power of proxy advisors: Myth or reality, *Emory Law Journal* 59, 869-918.

Crawford, Vincent P., and Joel Sobel, 1982, Strategic information transmission, *Econometrica* 50, 1431–1451.

Cremers, KJ Martijn, Lubomir P. Litov, and Simone M. Sepe, 2017, Staggered boards and long-term firm value, revisited, *Journal of Financial Economics* 126, 422–444.

Cvijanovic, Dragana, Moqi Groen-Xu, and Konstantinos E. Zachariadis, 2020, Free-riders and underdogs: Participation in corporate voting, ECGI Working Paper No. 649/2020.

Dikolli, Shane S., Mary Margaret Frank, Zhe Michael Guo, and Luann J. Lynch, 2022, Walk the talk: ESG mutual fund voting on shareholder proposals, *Review of Accounting Studies* 27, 864–896.

Dimmock, Stephen G., William C. Gerken, Zoran Ivković, and Scott J. Weisbenner, 2018, Capital gains lock-in and governance choices, *Journal of Financial Economics* 127, 113–135.

Ertimur, Yonca, Fabrizio Ferri, and David Oesch, 2013, Shareholder votes and proxy advisors: Evidence from say on pay, *Journal of Accounting Research* 51, 951–996.

Ertimur, Yonca, Fabrizio Ferri, and David Oesch, 2018, Understanding uncontested director elections, *Management Science* 64, 3400–3420.

Esö, Peter, Stephen Hansen, and Lucy White, 2014, A theory of vote trading and information aggregation, Working Paper, University of Oxford.

Fahlenbrach, Rüdiger, Nicolas Rudolf, and Alexis Wegerich, 2024, Leading by example: Can one universal shareholder's voting pre-disclosure influence voting outcomes? Working paper.

Faleye, Olubunmi, 2007, Classified boards, firm value, and managerial entrenchment, *Journal of Financial Economics* 83, 501–529.

Feddersen, Timothy, and Wolfgang Pesendorfer, 1998, Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting, *American Political Science Review* 92, 23–35.

Fishman, Michael J., and Kathleen M. Hagerty, 1995, The incentive to sell financial market information, *Journal of Financial Intermediation* 4, 95–115.

Fos, Vyacheslav, and Clifford G. Holderness, 2023, The distribution of voting rights to shareholders, *Journal of Financial and Quantitative Analysis* 58, 1878–1910.

Frenkel, Sivan, 2015, Repeated interaction and rating inflation: A model of double reputation, *American Economic Journal: Microeconomics* 7, 250–280.

Fulghieri, Paolo, Günter Strobl, and Han Xia, 2014, The economics of solicited and unsolicited credit ratings, *Review of Financial Studies* 27, 484–518.

Gallagher, Daniel M., 2013, Remarks at 12th European Corporate Governance & Company Law Conference.

Gallagher, Daniel M., 2014, Outsized power and influence: The role of proxy advisers, Critical Legal Issues Working Paper Series, Washington Legal Foundation.

Gantchev, Nickolay, and Mariassunta Giannetti, 2021, The costs and benefits of shareholder democracy: Gadflies and low-cost activism, *Review of Financial Studies* 34, 5629–5675.

Garcia, Diego, and Francesco Sangiorgi, 2011, Information sales and strategic trading, *Review of Financial Studies* 24, 3069–3104.

Goel, Anand M., and Anjan V. Thakor, 2015, Information reliability and welfare: A theory of coarse credit ratings, *Journal of Financial Economics* 115, 541–557.

Goldstein, Itay, and Chong Huang, 2016, Bayesian persuasion in coordination games, *American Economic Review* 106, 592–596.

Goldstein, Itay, and Yaron Leitner, 2018, Stress tests and information disclosure, *Journal of Economic Theory* 177, 34–69.

Guo, Yingni, and Eran Shmaya, 2019, The interval structure of optimal disclosure, *Econometrica* 87, 653–675.

Hayne, Christie, and Marshall Vance, 2019, Information intermediary or de facto standard setter? Field evidence on the indirect and direct influence of proxy advisors, *Journal of Accounting Research* 57, 969–1011.

Hu, Edwin, Nadya Malenko, and Jonathon Zytnick, 2024, Custom proxy voting advice, ECGI Working Paper No. 975/2024.

Iliev, Peter, Jonathan Kalodimos, and Michelle Lowry, 2021, Investors' attention to corporate governance, *Review of Financial Studies* 34, 5581–5628.

Iliev, Peter, and Michelle Lowry, 2015, Are mutual funds active voters? *Review of Financial Studies* 28, 446–485.

Iliev, Peter, and Svetla Vitanova, 2023, Investors' voting preferences, Working paper.

Inostroza, Nicolas, 2021, Persuading multiple audiences: An information design approach to banking regulation, Working Paper, University of Toronto.

Inostroza, Nicolas, and Alessandro Pavan, 2020, Persuasion in global games with application to stress testing, Working paper, Northwestern University.

Jackson, Robert, and Jonathon Zytnick, 2022, Individual investor ideology, Working paper, New York University.

Johnson, William C., Jonathan M. Karpoff, and Sangho Yi, 2015, The bonding hypothesis of takeover defenses: Evidence from IPO firms, *Journal of Financial Economics* 117, 307–332.

Kamenica, Emir, and Matthew Gentzkow, 2011, Bayesian persuasion, *American Economic Review* 101, 2590–2615.

Kerman, Toygar, P. Jean-Jacques Herings, and Dominik Karos, 2020, Persuading strategic voters, Working paper, Corvinus University of Budapest.

Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li, 2017, Persuasion of a privately informed receiver, *Econometrica* 85, 1949–1964.

Larcker, David F., Allan L. McCall, and Gaizka Ormazabal, 2015, Outsourcing shareholder voting to proxy advisory firms, *Journal of Law and Economics* 58, 173–204.

Leitner, Yaron, and Basil Williams, 2023, Model secrecy and stress tests, *Journal of Finance* 23, 1055–1095.

Leitner, Yaron, and Bilge Yilmaz, 2019, Regulating a model, *Journal of Financial Economics* 131, 251–268.

Levit, Doron, and Nadya Malenko, 2011, Nonbinding voting for shareholder proposals, *Journal of Finance* 66, 1579–1614.

Levit, Doron, and Anton Tsoy, 2022, A theory of one-size-fits-all recommendations, *American Economic Journal: Microeconomics* 14, 318–347.

Li, Tao, 2018, Outsourcing corporate governance: Conflicts of interest within the proxy advisory industry, *Management Science* 64, 2951–2971.

Ma, Shichao and Yan Xiong, 2021, Information bias in the proxy advisory market, *Review of Corporate Finance Studies* 10, 82–135.

Malenko, Andrey and Nadya Malenko, 2019, Proxy advisory firms: The economics of selling information to voters, *Journal of Finance* 74, 2441–2490.

Malenko, Andrey and Nadya Malenko, 2024, Voting choice, ECGI Finance Working Paper No. 910/2023.

Malenko, Nadya, and Yao Shen, 2016, The role of proxy advisory firms: Evidence from a regression-discontinuity design, *Review of Financial Studies* 29, 3394–3427.

Marquardt, Blair B., Brett W. Myers, and Xu Niu, 2018. Strategic voting and insider ownership, *Journal of Corporate Finance* 51, 50–71.

Matsusaka, John G., and Chong Shu, 2021, A theory of proxy advice when investors have social goals. Working Paper, University of Southern California.

Matvos, Gregor, and Michael Ostrovsky, 2010, Heterogeneity and peer effects in mutual fund proxy voting, *Journal of Financial Economics* 98, 90–112.

Maug, Ernst, 1999, How effective is proxy voting? Information aggregation and conflict resolution in corporate voting contests, Working paper, Duke University.

Maug, Ernst, and Kristian Rydqvist, 2009, Do shareholders vote strategically? Voting behavior, proposals screening, and majority rules, *Review of Finance* 13, 47–79.

McCahery, Joseph A., Zacharias Sautner, and Laura T. Starks (2016). Behind the scenes: The corporate governance preferences of institutional investors, *Journal of Finance* 71, 2905–2932.

Michaeli, Beatrice, 2017, Divide and inform: Rationing information to facilitate persuasion, *Accounting Review* 92, 167–199.

Orlov, Dmitry, Pavel Zryumov, and Andrzej Skrzypacz, 2023, The design of macro-prudential stress tests, *Review of Financial Studies* 36, 4460–4501.

Puchniak, Dan W., and Umakanth Varottil, 2023, Rethinking acting in concert: Activist ESG stewardship is shareholder democracy, ECGI Law Working Paper No. 731/2023.

Rayo, Luis, and Ilya Segal, 2010, Optimal information disclosure, *Journal of Political Economy* 118, 949–87.

Sangiorgi, Francesco, and Chester S. Spatt, 2017, The economics of credit rating agencies, *Foundations and Trends in Finance* 12, 1–116.

Sarvary, Miklos, and Philip M. Parker, 1997, Marketing information: A competitive analysis, *Marketing Science* 16, 24–38.

Spatt, Chester S., 2021, Proxy advisory firms, governance, market failure, and regulation, *Review of Corporate Finance Studies* 10, 136–157.

Shu, Chong, 2024, The proxy advisory industry: Influencing and being influenced, *Journal of Financial Economics* 154, 103810.

About ECGI

ECGI is an international non-profit membership organisation that provides a platform for debate and dialogue among academics, policymakers, and business leaders, with a focus on major corporate governance, ESG, and stewardship issues.

The views expressed in this working paper are those of the authors, not those of the ECGI or its members.