

When do proxy advisors improve corporate decisions?*

Berno Buechel[†], Lydia Mechtenberg[‡] and Alexander F. Wagner[§]

April 23, 2026

Forthcoming, Economic Inquiry

Abstract

This paper studies the impact of proxy advisors on shareholder decision-making. We posit two assumptions: (i) the board is at least as well informed as any individual shareholder; (ii) shareholders can condition their information acquisition on the proxy advisor's recommendation. If only (i) holds, shareholders optimally rubber-stamp the board's proposal and do not invest in private research. If (ii) additionally holds, disagreement between the board and the proxy advisor can trigger private information acquisition by some shareholders. We identify the conditions under which this mechanism restores value-of-information logic and improves the informational quality of shareholder voting outcomes.

Keywords: Collective decision-making, information acquisition, proxy advisors, strategic voting

JEL classification: G23, D72, D83

*We thank the editor (Tim Salmon) and two anonymous referees, as well as Bruno Biais, Alessandra Casella, Romain Ducret, Christian Ewerhart, Alia Gizatulina, Thorsten Hens, Dusan Isakov, Nick Netzer, Shaoting Pi, and Jean-Charles Rochet and the audiences of the CEPS Political Economy Workshop at Paris-Saclay 2024, the conference of the Swiss Society of Economics and Statistics 2024, the annual congress of the Verein für Socialpolitik 2024, the Voting and Governance workshop at Yale University 2023, the conference of the theoretical committee of the Verein für Socialpolitik 2023, a workshop organized under the grant by NBIM (“Evolution and determinants of institutional investor preferences and their impact on portfolio companies”), the Swiss Theory Day 2023, the SFI research days 2021 as well as seminar participants at the University of Zurich and Bielefeld for their helpful comments. We are particularly indebted to Navin Kartik, Andrey Malenko, Nadya Malenko, Ernst Maug, Adam Meirowitz, Stefan Klössner, Maximilian Janisch, and Thomas Lehéricy who all provided detailed feedback on earlier versions. Moreover, we thank Bénédicte Droz, Enrico Fenoaltea, Karin Loetscher, and Andrea Sommer for great research assistance. We acknowledge financial support from the University of Zurich Research Priority Program “Financial market regulation.” Buechel: University of Fribourg, Department of Economics, Bd. de Pérolles 90, CH-1700 Fribourg, Switzerland; Mechtenberg: University of Hamburg, Department of Economics, von-Melle-Park 5, D-20146 Hamburg, Germany; Wagner: University of Zurich, Department of Banking and Finance, Plattenstrasse 14, CH-8032 Zurich, Switzerland. Some earlier versions of this paper circulated under the title “When two experts are better than one: The example of shareholder voting.” During the preparation of this manuscript, we used AI-based language assistance tools to improve clarity and grammar. All content was reviewed and edited by the authors, who take full responsibility for the substance and interpretation of the paper. We declare that we have no relevant or material financial interests that relate to the research described in this paper. Until August 2019, Wagner was chairman of a non-profit proxy advisor, the Swipra Foundation.

[†]University of Fribourg. Email: berno.buechel@unifr.ch.

[‡]University of Hamburg. Email: Lydia.Mechtenberg@uni-hamburg.de.

[§]University of Zurich, CEPR, ECGI, Swiss Finance Institute. Email: alexander.wagner@df.uzh.ch.

1 Introduction

Shareholders vote on a variety of important issues, including director elections, executive compensation, and certain aspects of mergers and acquisitions. During the past two decades, shareholders' decision making has changed due to the rise of a new business: Proxy advisory firms (such as ISS and Glass Lewis) provide voting recommendations to shareholders. These recommendations have substantial impact on shareholder voting outcomes.¹ There is an ongoing public and scientific debate about the effects of proxy advisors (PAs in what follows) on the quality of decision making in shareholder meetings.² The regulation of PAs is highly contentious, as evidenced, for example, by frequent and substantial changes of rules applied by the U.S. Securities and Exchange Commission (SEC) within the last few years. A key point of contention is that PAs may crowd out shareholders' incentives to invest in own research. Shareholders who rely on a PA's recommendation as a substitute for own research save costs individually, but negatively affect the collective by not contributing new information into the decision-making process (Malenko and Malenko, 2019).

In this paper, we show that under two practically relevant assumptions, the presence of a PA actually leads to either more, or, at least, not fewer shareholders investing into private information acquisition, and hence improves corporate decision quality. The basic intuition is that shareholders who otherwise, without a PA, "rubber-stamp" any board proposal, use the PA as a filter to identify controversial issues that deserve further investigation. Proving that this intuition holds generally, even when voters are strategic, and bringing to light its hidden assumptions, necessitates a formal model. By solving the model, we provide the conditions under which disagreement of public signals endogenously restores the acquisition of private information in voting. From a conceptual perspective, the mechanism we study can be understood as a form of sequential value-of-information logic: shareholders optimally acquire a costly private signal only when earlier public signals leave posterior beliefs sufficiently close to indifference.

In our model, there are three types of agents: Shareholders, the firm's board of directors, and a PA. Shareholders and the board both care about firm value, whereas the PA cares for its profit. There is a proposal on an issue (e.g., a director election). Agents are imperfectly informed about the correct decision on the issue, i.e., about which decision will increase firm

¹See, for example, Alexander et al. (2010); Choi, Fisch, and Kahan (2010); Ertimur, Ferri, and Oesch (2013); Iliev and Lowry (2015); Larcker, McCall, and Ormazabal (2013, 2015); Li (2018); Malenko and Shen (2016); McCahery, Sautner, and Starks (2016); Matsusaka and Shu (2021). PAs can also have an impact by increasing shareholder engagement with firms (Dey, Starkweather, and White, 2022).

²Spatt (2021) provides a survey of the literature with a focus on regulatory issues.

value most. The board and the PA receive a private imperfect and independently distributed signal about the correct decision. The board recommends a decision based on its own signal. For brevity, we refer to this as the “board’s proposal.” Then, each shareholder individually decides whether to buy the PA’s vote recommendation, i.e., the PA’s signal, and whether to invest in own research, i.e., to obtain a private signal. Finally, shareholders vote, and the simple majority rule determines the outcome.

Two key assumptions set this model apart from the existing literature. First, we posit that the board is better-informed than any single shareholder alone. This assumption, “BIB” (for better-informed board), is in line with a long tradition of studies in corporate governance arguing that insiders (the board and management) have information about the company that may be superior to that of shareholders (Jensen and Meckling, 1976). The second key assumption is “PAF” (proxy advice first): *after* receiving proxy advice, a shareholder can decide upon additional research about the issue at hand. This assumption is more likely to hold in regulatory settings that provide shareholders with sufficient time to conduct research.

We solve our game-theoretic model for pure Perfect Bayesian Nash equilibria, mainly focusing on equilibria that are not Pareto-dominated by other equilibria.³ We find that the presence of a PA increases the shareholders’ incentives to invest in own research or leaves these incentives unchanged relative to the case without a PA.

The underlying logic builds on standard value-of-information reasoning, but its application in a voting environment with strategic agents is non-trivial. Assume first that there is no PA. Each shareholder only needs to consider the case in which her vote decides whether the proposal passes or not, i.e., where she is pivotal. If all shareholders vote on the board’s proposal based on their own information, then a shareholder being pivotal must mean that the affirmative and dissenting signals of other shareholders are equally frequent. In this situation, only the information of the pivotal shareholder and the board’s information are crucial. Even if the pivotal shareholder has invested in her own research and this signal is not in favor of the board’s proposal, she is still better off voting in favor of the proposal since the board, by Assumption BIB, is better informed than the shareholder. Therefore, the shareholder would not invest in her own research in the first place and would prefer to always endorse the proposal. This behavior is typically inefficient. Substituting the own research with the informativeness of the board’s proposal is individually rational but collectively harmful, as it leads to correlated mistakes.⁴ Thus, although sequential value-of-information logic would

³Pareto-dominated equilibria are based on coordination failure.

⁴This result holds as long as shareholders do not anticipate the board’s proposal to be so conflicted as to be less informative than even a single shareholder’s own research and, thus, close to uninformative, a model

prescribe private research in controversial cases, it breaks down in the benchmark voting environment due to free-riding and pivotality considerations.

Compared to this benchmark, the presence of a PA leads to higher decision quality in our model. In this sense, the presence of an early second public signal restores the conditions under which standard sequential value-of-information logic becomes operative despite the voting externality. Intuitively, for a shareholder it pays off to invest in own research when there is sufficient controversy about whether the proposal should be accepted. Thus, a PA does not just contribute an additional signal into the decision-making process, but triggers shareholders to conditionally generate their own signals. This complementarity effect is strongest when the board and the PA are similarly well informed such that their contradicting signals indicate strong controversy. We show these results for the whole parameter space and address the cases where each of the two baseline assumptions is violated--providing the exact conditions for endogenous signal acquisition. We also discuss how the core mechanism extends to settings with non-trivial proxy advisor fees or conflicts of interest, clarifying which aspects of the results are robust to monopolistic pricing or biased advice and which require stronger assumptions.

We further extend the analysis to asymmetric equilibria. Asymmetric equilibria allow shareholders to specialize in different strategies, even though they are *ex ante* identical.⁵ It turns out that the crucial strategy of conditionally investing in research is also pervasive in asymmetric equilibria and is, in fact, used in a much larger area of the parameter space than in symmetric equilibria. In addition, depending on the parameters, there are shareholders who always invest in an own signal without subscribing to the PA, shareholders who rubber-stamp the board's proposals or shareholders who "robo-vote," i.e., always vote according to the PA's recommendations. With asymmetric equilibria, too, PAs make the voting outcome more informed.

Although our model is necessarily stylized, it provides important perspectives on some recent regulatory developments, which we discuss in Section 6. For example, some recent proposals would arguably lead to a failure of Assumption PAF, because timelines around shareholder meetings would be adjusted such that investors simply would not have time to do research *contingent on* PA advice. Such interventions would, therefore, inadvertently lead to lower quality corporate decision-making.

variation violating BIB that we discuss in Section 5.

⁵Of course, another reason for shareholders to behave heterogeneously in practice is that they differ *ex ante*. Studies that feature shareholder heterogeneity include Levit, Malenko, and Maug (2021) and Levit, Malenko, and Maug (2024).

Our paper contributes to the literature on common interest voting and on voting in corporations. The literature on strategic voting in a setting of common interest started with Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1996, 1997, 1998).⁶ Informational inefficiencies of voting occur when the symmetry assumptions of the standard Condorcet model are violated, either with regard to the signal technology (Austen-Smith and Banks, 1996) or with regard to the information-transmission process (Gerardi and Yariy, 2007; Iaryczower, Shi, and Shum, 2018; Buechel and Mechtenberg, 2019). Inefficiencies are also generated when private information becomes costly (Persico, 2004; Gershkov and Szentes, 2009), or when public information is provided that is not of sufficiently high quality (Kawamura and Vlaseros, 2017; Jeong, 2019; Liu, 2019; Malenko and Malenko, 2019). Generally speaking, inefficiencies can be generated by correlation of private beliefs across voters, either through public information or through information-transmission processes between voters that are not optimally tailored to the signal (quality) distribution.⁷

A related strand is the literature on voting in corporations. Seminal works in this literature that specifically elucidate the role of strategic voting include Maug (1999), Maug and Yilmaz (2002), and Maug and Rydqvist (2009). A relatively recent innovation has been to enhance the analysis by considering the presence of proxy advisors (Malenko and Malenko, 2019; Malenko, Malenko, and Spatt, 2025). Other work on (strategic) voting in the corporate finance context includes Matvos and Ostrovsky (2006), Brav and Mathews (2011), Levit and Malenko (2011), Van Wesep (2014), Bar-Isaac and Shapiro (2020), Ma and Xiong (2021), Meirowitz and Pi (2022), and Parlasca and Voss (2023), among others.

We contribute to both strands of literature: First, for common-interest voting environments, we characterize when standard sequential value-of-information logic fails and when it is restored once an additional early public signal is introduced. Addressing the whole parameter space, we find the boundaries of signal precision where information acquisition and information aggregation breaks down. Second, we show that proxy advice given early enough can foster shareholders' own investments in research. This has the policy implication that shareholders need to have sufficient time to conduct such research *after* receiving the advice from PAs. Under these conditions, PAs are likely to improve decision quality and social welfare.

⁶For studies of informational efficiency in games with strategic complements and strategic substitutes, see Angeletos and Pavan (2007) and Hellwig and Veldkamp (2009).

⁷Accordingly, Levy and Razin (2015) show that correlation neglect can enhance informational efficiency.

2 Model Setup

2.1 Basic Ingredients

We model voting on corporate decisions as strategic voting under uncertainty (Austen-Smith and Banks, 1996; Feddersen and Pesendorfer, 1996, 1997, 1998) and follow frameworks such as Kawamura and Vlaseros (2017), Malenko and Malenko (2019), Bar-Isaac and Shapiro (2020) and Ma and Xiong (2021). A firm is owned by $N > 1$ shareholders, where N is odd. The firm faces uncertainty with respect to a binary decision. Examples vary by jurisdiction and include director elections, dividends, shareholder proposals, compensation-related matters, etc. Making the ex post correct decision will increase firm value by an amount normalized to 1, while the wrong decision leaves it unchanged.

More formally, there are two states of the world $\theta \in \{A, B\}$ with equal prior probability. Slightly abusing notation, we assume that the firm has to decide on a binary issue $\{A, B\}$ that yields value 1 if and only if the decision matches the true state.⁸

The board of directors receives a binary signal s_B regarding the issue to be voted on. The signal takes on values a or b . The signal quality is $q_B \in (\frac{1}{2}, 1)$, i.e., $Pr[s_B = a | \theta = A] = Pr[s_B = b | \theta = B] = q_B$. We assume that the board then recommends either action A or B . We call this the “board’s proposal.”

A profit-maximizing proxy advisor (PA) offers advice to shareholders at fee $f > 0$. The PA receives a signal about the true state as well. The quality of that signal is $q_P \in (\frac{1}{2}, 1)$. The PA provides a vote recommendation *for* or *against* the board’s proposal to subscribing shareholders.

Before that, shareholders decide whether to subscribe to the PA’s offer. If a shareholder subscribes, she receives the PA’s recommendation.⁹ A shareholder *then* decides whether to invest $c > 0$ in own research about the issue at hand. If a shareholder expends own research costs, this leads to a private signal of quality $q_S \in (\frac{1}{2}, 1)$. Hence, a shareholder knows the board’s signal, can subscribe to get the PA’s signal, and then can invest into an own private signal. When the shareholder meeting is held, each shareholder votes *yes* or *no*. Abstentions are excluded.¹⁰ For simplicity, each shareholder holds one share of the firm and each share provides one vote. The decision that receives a majority of votes is implemented.

⁸As long as the vote we consider remains a decision between two different actions, with one Pareto-dominating the other conditional on the unobservable state of the world, the main argument of the paper would still apply when introducing a continuous state space or a continuous signal technology.

⁹Hence, the PA sells information directly in the sense of Admati and Pfleiderer (1990).

¹⁰In practice, shareholders may also abstain. However, according to most institutional settings abstentions are counted (either as *yes* or *no*) and hence shareholders’ voting action is essentially binary.

Conditional on state θ , all signals are independent, and precision levels q_B , q_P , and q_S are common knowledge.

Starting from this standard framework, our first key assumption is that the board knows better than any single shareholder what is good for the company.

Assumption 1 (BIB). *The board is at least as well informed as a single shareholder, i.e.,*

$$q_S \leq q_B.$$

“BIB” stands for better-informed board. One interpretation is that the proposal itself already contains valuable information that would be hard to match for a single shareholder alone. For the quality of the PA q_P we do not make any analogous assumption that restricts it to be above or below the other agents’ qualities.

In the course of the analysis it will come in handy to transform signal qualities $q \in (0.5, 1)$ into log-odds $\log(\frac{q}{1-q}) \in (0, \infty)$. We denote the log-odds of the board being correct as $\ell_B := \log(\frac{q_B}{1-q_B})$ and likewise $\ell_S := \log(\frac{q_S}{1-q_S})$ for the shareholders and $\ell_P := \log(\frac{q_P}{1-q_P})$ for the PA. Then Assumption BIB reads $\ell_S \leq \ell_B$.¹¹ This notation is convenient since it allows us to aggregate signal qualities by summation. To see this, consider the board’s signal b and assume, for instance, that both the PA and one shareholder have received signals a and there is no further information. Then, the board’s signal is rather correct than not if and only if $q_B(1 - q_P)(1 - q_S) \geq (1 - q_B)q_Pq_S$, which is equivalent to $\ell_B \geq \ell_P + \ell_S$.

Our second key assumption is that shareholders can condition their research investment on the PA’s recommendation.

Assumption 2 (PAF). *Subscribing shareholders decide upon own research investment after they have received the PA’s recommendation.*

“PAF” stands for “proxy advice first” and means that the proxy advisor’s recommendation arrives early enough to still acquire private information on the specific issue.¹²

¹¹Nitzan and Paroush (1982) show that among voters with idiosyncratic signal precision the optimal voting weights would be according to these log-odds.

¹²Relaxing this assumption would change the timing of our model such that shareholders have to decide simultaneously about subscribing to the PA and about investing in own research. We discuss the consequences of making this assumption in our model in Section 6.2 and provide the corresponding results in Supplementary Online Material (SOM) Section 1.3.

2.2 Simplification and Timeline

It turns out that we can substantially simplify the exposition by fixing the signal and behavior of the board and the behavior of the PA. The board receives a signal and then makes a recommendation. We let the board's signal be always b (for board).¹³ We fix the board's behavior by assuming it makes the proposal according to its signal, i.e., it has received signal b and now proposes action B .¹⁴ Likewise, we fix the PA's behavior to set fee $f > 0$ and recommend according to its signal, i.e., it recommends *for* if it has received signal b (for board) and it recommends *against* if it has received signal a (against board).

These simplifications do not affect the substance of our analysis and results. Since the board, like the shareholders, aims at maximizing firm value, revealing its signal to help shareholders decide for the optimal policy is in its own interest. The proxy advisor, in turn, can generate profits only from helping individual shareholders to make a decision that is even more informed than it would be without the PA. The reason is that the shareholders' willingness to pay for proxy advice depends only on the PA's contribution to the informedness of the vote.¹⁵ In this sense, the fee f should be interpreted as reflecting shareholders' willingness to pay for an incremental public signal, rather than as the outcome of strategic pricing, market power, or contractual frictions on the part of the proxy advisor. Hence, the optimal strategy of the PA is to reveal its true signal whenever asked to do so, and to set the fee f equal to the shareholder's willingness to pay for this revelation. Hence, explicitly including a strategic PA and an endogenous fee does not add interesting results for the informational mechanism we study, such that we can focus on the main implications derived from the simplified setup. The role of non-trivial fees and research costs for equilibrium behavior is discussed explicitly in Section 5.3.

The timeline, which is illustrated by Figure 1, summarizes the simplified setup. At $t = 0$ nature draws a state of the world and signals for all potential recipients of signals. The players learn the board's proposal and hence its signal. At $t = 1$ each shareholder decides whether to pay the fee for the PA's report. Those who pay the fee receive the truthful vote recommendation which is equivalent to learning the PA's signal. At $t = 2$ each shareholder decides whether to invest costs c to receive an own independently and identically distributed

¹³This will exclude strategies that depend on the label of the alternative, such as always voting *yes* for alternative A and *no* for alternative B independent of which alternative the board has proposed.

¹⁴In Section 5.4, we discuss to which extent, and with which implications, a re-interpretation of our model covers conflicts of interest between board and shareholders.

¹⁵In Malenko, Malenko, and Spatt (2021), the PA biases a costless public recommendation to enhance shareholders' perceived pivotality since their willingness to pay for the PA's costly signal increases in pivotality. Heterogeneity in voters, which allows these manipulations to be effective, is created by subjective values.

signal of quality q_S . At $t = 3$ shareholders vote. At $t = 4$ the proposal passes if a majority approves it and payoffs are realized.

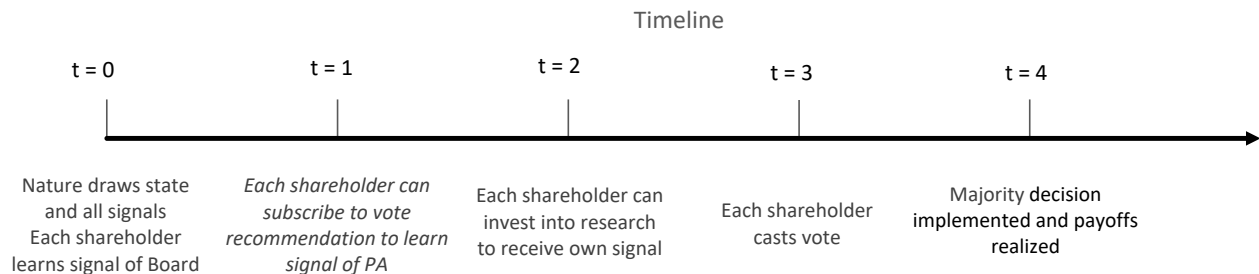


Figure 1: Timeline. For simplicity, the board’s and PA’s behavior is fixed. (Actions in italics only apply if there is a PA.)

2.3 Strategies

The most important strategic aspects concern the shareholders. They have several strategies both on the information acquisition stages ($t = 1$ and $t = 2$, respectively) and on the voting stage ($t = 3$). On the information acquisition stages, there are six strategies: A shareholder who does not subscribe may invest in own research (NotSubscribe-Invest) or not (NotSubscribe-NotInvest); a shareholder who does subscribe may unconditionally invest in research (Subscribe-Invest) or not (Subscribe-NotInvest) or, else, may invest in research only if the recommendation is *for* (Subscribe-InvestIFF*for*) or only if the recommendation is *against* (Subscribe-InvestIFF*against*).

In the voting stage, any shareholder chooses *yes* or *no*. The set of voting strategies depends on the acquired information which may include the PA’s signal and the own signal. For instance, for a shareholder who acquired both kinds of information (e.g., with Subscribe-Invest), a voting strategy is a mapping $v_i : \{for, against\} \times \{a, b\} \rightarrow \{yes, no\}$. Slightly abusing notation, we write σ_i for the information acquisition *and* voting strategy of a shareholder i , and we use $\sigma = (\sigma_1, \dots, \sigma_N)$ to denote a strategy profile of shareholders.

We study Perfect Bayesian Nash equilibria, i.e., players best respond to their beliefs and update their beliefs according to Bayes’ rule wherever possible. We focus on pure strategies, but analyze both symmetric and asymmetric strategy profiles.

Asymmetric strategy profiles capture heterogeneous behavior of shareholders better than symmetric mixed strategies would; and they are more efficient in common interest voting games (see, e.g., Kawamura and Vlaseros, 2024). Existence of equilibria is not an issue, as a symmetric equilibrium always exists, e.g., every shareholder *rubber-stamps* (i.e., votes *yes*

without subscribing to the PA nor investing into an own signal). Multiplicity of equilibria we handle by selecting those that are Pareto-efficient among them. This eliminates equilibria that are due to miscoordination.

To analyze the model we take the perspective of a regulator who compares a market with a PA, as in the game defined above, with a market in which no PA is admitted (i.e., the game above without actions at $t = 1$). The regulator maximizes welfare which requires maximizing decision quality in our setup. We will assume that costs of information acquisition, be it fee f or costs c , are relatively small compared to benefits of decision quality. Hence, when shareholders have to trade off costs of information acquisition with benefits of higher firm value, we will assume that the latter dominates.¹⁶ When there are two strategies with the same decision quality, then shareholders strictly prefer the one with lower costs, as we assume that costs are strictly positive. The quality of corporate decisions is measured by $\Pi(\sigma)$, the ex ante probability that the decision will match the true state.¹⁷ In what follows, all propositions of the main text are proven in Appendix A in this document. Some further arguments used in the discussion in the main text are formally shown in the Supplementary Online Material (SOM), which is available here: <https://bit.ly/proxy-SOM>.

3 Main Results

Benchmark: Without Proxy Advisor. Consider first the benchmark situation that no PA is admitted. Thus, posit that in the timeline of Figure 1 actions at $t = 1$ are suppressed. Then a shareholder’s information acquisition decision reduces to whether to acquire an own signal or not in $t = 2$. Suppose for a moment that all shareholders do acquire such a signal and vote according to it. We call this strategy profile *UNIS*, for “UNconditional Investment in own Signal,” where the term “unconditional” will be justified later, when shareholders could potentially condition their investment in own research on the PA’s vote recommendation.¹⁸ In this strategy profile the decision quality amounts to $\Pi(\sigma^{UNIS}) = \pi(N)$, where $\pi(N) := \sum_{i=\frac{N+1}{2}}^N \binom{N}{i} q_S^i (1 - q_S)^{N-i}$ is the probability that a majority decision of N shareholders is correct.

¹⁶We discuss the relaxation of this assumption in Section 5.3.

¹⁷This is also called *informational efficiency*, which can be distinguished from *economic efficiency* (see, e.g., Buechel and Mechtenberg, 2019). Economic efficiency means welfare, which here can be defined as $\Pi(\sigma)$ net of the investment costs in own research since the prices paid to the PA are transfers. When investment costs c become arbitrarily small, the two concepts coincide.

¹⁸For simplicity, we use the same labels for strategies as for the symmetric strategy *profiles* composed of these strategies.

While the decision quality of such voting behavior is usually very high (De Caritat, 1785), it is unfortunately not an equilibrium under Assumption BIB. The intuition is straightforward once spelled out. A single shareholder i can improve by deviating to not acquire a signal and vote *yes*. When this shareholder i is pivotal, the signals of all $N - 1$ other shareholders are split: there are as many a -signals as there are b -signals among them. Now, even if i 's signal points against the board's proposal, Assumption BIB, i.e., the assumption that the board is at least as well informed as i , makes it beneficial to vote *yes*, i.e., for the board's proposal, and not to acquire own information in the first place. We call this latter strategy and its corresponding strategy profile “*Rubber-stamping*”.¹⁹ This insight generalizes.

Proposition 1 (SYM without PA). *Suppose no PA is admitted. If Assumption BIB holds, then there does not exist a symmetric equilibrium in which shareholders invest in own research. Hence, decision quality in symmetric equilibria is bounded by: $\Pi(\sigma) \leq q_B$. The Pareto-efficient²⁰ symmetric equilibrium is Rubber-stamping and leads to decision quality $\Pi(\sigma^{Rubber}) = q_B$.*

Proposition 1 shows that without a PA the quality of decision making is bounded by the quality of the board whose proposal is rubber-stamped by the shareholders. This result is related to the first theorem in Austen-Smith and Banks (1996), which shows that it is not an equilibrium to vote according to the private signal when a proposal is ex ante more likely to be correct. The result is also related to the substitution effect of Malenko and Malenko (2019), but occurs on a different level: For the scenario that the board's proposals are uninformative, Malenko and Malenko (2019) obtain over-reliance of shareholders on the PA's recommendations. For the scenario that the board's proposals are informative, we obtain over-reliance on this proposal without any PA, as a new benchmark.²¹

Including Proxy Advice. The presence of a proxy advisor substantially increases a shareholder's set of information-acquisition strategies. One of them, *Subscribe-InvestIFFagainst*, gives rise to the following symmetric strategy profile, which we denote by $\hat{\sigma}$ and call “*CAIS* (Conditional on Advice Invest in Signal):” All shareholders subscribe to proxy advice; if the recommendation is *for*, they vote *yes*; if the recommendation is *against*, they invest in own

¹⁹Assumption BIB, $q_S \leq q_B$, is in fact necessary and sufficient for the conclusion of Proposition 1. Since we have Assumption BIB as a leading assumption, we only show sufficiency in the proof of Proposition 1.

²⁰Recall that the criterion of Pareto-efficiency is applied within the set of symmetric equilibria.

²¹The result that no shareholder invests in own research is stylized. When studying asymmetric equilibria without PA, we find the same message in a less stylized form: Without PA, there always exist shareholders who do not invest in own research, while the number of shareholders who do is bounded.

research and vote according to their own signal, i.e., vote *yes* if the signal is *b* and *no* if it is *a*. CAIS is illustrated in Table 1.

In this strategy profile shareholders use the PA's recommendation as a filter: *for* recommendations are followed without being challenged; *against* recommendations trigger further investigation of the issue.

		Own Signal	
		<i>b</i> (for board)	<i>a</i> (against board)
PA's recommendation	<i>for</i>	yes	
	<i>against</i>	yes	no

Table 1: Strategy CAIS: Invest in own research if and only if vote recommendation is *against*; and after *for* recommendation vote *yes*, after *against* recommendation vote *yes* if and only if signal is for board.

It turns out that based on this strategy profile the negative result of Proposition 1 can be mitigated by the presence of a PA.

Proposition 2 (SYM with PA). *Let Assumptions BIB and PAF hold. Let costs c be arbitrarily small and let fee f be sufficiently smaller. Suppose there is a PA with $\ell_P \in (\ell_B - \ell_S, \ell_B + \ell_S)$. Then there exists a symmetric equilibrium in which shareholders conditionally invest in own research. The Pareto-efficient equilibrium is CAIS and leads to decision quality $\Pi(\hat{\sigma}) > q_B$. Otherwise (i.e., if $\ell_P \notin (\ell_B - \ell_S, \ell_B + \ell_S)$), the Pareto-efficient equilibrium is Rubber-stamping with $\Pi(\sigma^{Rubber}) = q_B$.*

In the proof of Proposition 2 (Appendix A.2) we proceed as in the proof of Proposition 1 (Appendix A.1), by providing first the full characterization of all equilibria as a lemma and then selecting those that are Pareto-efficient among them. Comparing Proposition 2 with Proposition 1, we conclude that the presence of a PA either strictly improves decision quality or leaves it unchanged, compared to the setting without a PA. The condition for the strict improvement can be rewritten as $|\ell_B - \ell_P| < \ell_S$, which has the following interpretation: the difference in quality of board and PA is smaller than the information quality of one shareholder. If this conditions is satisfied, we have a new equilibrium (CAIS) with a PA in which all shareholders conditionally invest in own research.

The first intuition for the conditions of the PA being beneficial as stated in Proposition 2 can be seen from their violations. Consider the symmetric strategy profile CAIS. If $\ell_P \leq \ell_B - \ell_S$, we have $\ell_S + \ell_P \leq \ell_B$, i.e., the board is better informed than the PA and one shareholder together. Then there is a deviation from CAIS to Rubber-stamping. Intuitively,

the board is sufficiently well informed that it does not individually pay off to acquire any information, even if it were costless. If $\ell_P \geq \ell_B + \ell_S$, i.e., the PA is better informed than the board and one shareholder together, then there is a deviation from CAIS, e.g., to not investing and to voting against the board’s proposal. Indeed, the deviating shareholder’s vote is only pivotal if board and PA disagree and voting *no* improves decision quality, given that the PA is so well informed. If costs c or f are not small enough, there is again a beneficial deviation, e.g., to Rubber-stamping, which saves costs. Finally, if the PA’s fee f is not sufficiently smaller than the costs c , then deviating to UNIS saves costs without affecting the outcome.²²

In Section 1 of the Supplementary Online Material (SOM), we extend the analysis to the complete characterization of all symmetric equilibria in pure strategies, with and without Assumptions BIB and PAF. Most importantly, that analysis shows that the two key assumptions Assumption BIB and PAF are not only sufficient but also necessary for the conclusion, as we will discuss below.

4 Illustration and Discussion

Numerical Example. Example 1 numerically illustrates Propositions 1 and 2.

Example 1 (Symmetric Equilibria). *Let $q_B = 0.75$, $q_P = 0.7$, and $q_S = 0.6$. Then $\ell_B = 0.477$, $\ell_P = 0.368$, and $\ell_S = 0.176$ such that the condition $\ell_P \in (\ell_B - \ell_S, \ell_B + \ell_S)$ of Proposition 2 is satisfied, as $0.368 \in (0.477 - 0.176, 0.477 + 0.176)$. Table 2 illustrates the implications of Propositions 1 and 2 for decision quality. First, not admitting a PA leads to Rubber-stamping and hence to a decision quality of $q_B = 0.75$, independent of the number of shareholders N (Proposition 1). Second, when a PA is admitted, CAIS is the Pareto-efficient symmetric equilibrium, which delivers a strictly higher decision quality (by Proposition 2). Its decision quality is further increasing in the number of shareholders N and approaching $0.925 < 1$ for large N . The decision quality differs from one by the probability that the board’s and the PA’s signal agree on the wrong state. Finally, Table 2 shows the hypothetical case in which all shareholders play UNIS, i.e., invest in own research. This is not an equilibrium but a classic benchmark capturing the quality of majority decisions by N independently informed*

²²The assumption c small enough assures that shareholders who can improve decision quality by investing in own research would not shy away due to the high costs. The assumption that the costs are larger than zero matters when deviations that do not affect decision quality are considered. The assumption that fees f are sufficiently smaller than c means that the results answer the question whether there is a fee f such that a PA can profitably be active in the market.

voters, as already pointed out by the Marquis de Condorcet (*De Caritat*, 1785). In this benchmark case, decision quality may start low, but becomes larger than in equilibrium for a sufficiently large number of voters.

Setting	Decision quality	$N = 3$	$N = 5$	$N = 21$	$N = 101$	$N = 1,001$
No PA	$\Pi(\sigma^{Rubber}) = q_B$	0.75	0.75	0.75	0.75	0.75
With PA	$\Pi(\hat{\sigma}) = q_B q_P + p^{dis} \pi(N)$	0.784	0.798	0.855	0.917	0.925
Hypothetical	$\Pi(\sigma^{UNIS}) = \pi(N)$	0.648	0.683	0.826	0.979	1.0

Table 2: Decision quality in Example 1, illustrating Propositions 1 and 2 for $q_B = 0.75$, $q_P = 0.7$, and $q_S = 0.6$. The table considers the two Pareto-efficient symmetric equilibria, Rubber-stamping and CAIS, and strategy profile UNIS, which is not an equilibrium. $p^{dis} := (1 - q_B)q_P + q_B(1 - q_P)$ is the probability that the board’s and the PA’s signal differ.

To illustrate how information-acquisition costs affect welfare comparisons, consider the individual costs $C(\sigma)$ of the three symmetric strategy profiles: $C(\sigma^{Rubber}) = 0$, $C(\hat{\sigma}) = f + p^{dis}c$, and $C(\sigma^{UNIS}) = c$ for each shareholder. While f is a welfare-neutral transfer from shareholders to the proxy advisor, non-trivial costs c matter for the welfare comparison, as they represent real resource expenditures. Admitting non-trivial costs, CAIS yields higher welfare than Rubber-stamping if and only if $\Pi(\hat{\sigma}) - \Pi(\sigma^{Rubber}) > Np^{dis}c$, which is equivalent to $c < \frac{1}{N}(\pi(N) - \frac{q_B(1-q_P)}{p^{dis}})$, where $p^{dis} := (1 - q_B)q_P + q_B(1 - q_P)$. It measures when the costs of the more informative equilibrium are justified by higher decision quality. For instance, in Example 1 with $N = 21$, that is $c < 0.0125$, which means that the individual private costs of information acquisition must not exceed 1.25% of the normalized value of the decision for the firm.²³

We now turn to illustrating Propositions 1 and 2 graphically, while at the same time extending our analysis to the entire parameter space.

Graphical Illustration. Figure 2 illustrates the full parameter space. An entry (x, y) in this coordinate system has the simple interpretation that the board is equally well informed as x shareholders, while the PA is equally well informed as y shareholders.²⁴

In the upper panel of Figure 2, no PA is admitted. By Proposition 1, Rubber-stamping is the Pareto-efficient equilibrium under Assumption BIB, i.e., for $q_S \leq q_B$. This is illustrated

²³We will return to non-trivial costs and study their implications for deviations and equilibria in Subsection 5.3.

²⁴“Equally well informed” means that if x shareholders have received a signal a (against the board) then both states A and B are equally likely. Hence, if more than x shareholders have received a signal a and there is no other information, then the board should be overruled.

in the area $\frac{\ell_B}{\ell_S} \geq 1$. Assumption BIB is necessary and sufficient for this conclusion as UNIS is the Pareto-efficient equilibrium for $\frac{\ell_B}{\ell_S} < 1$. Hence, when there is no PA, information acquisition occurs if and only if the board is less well informed than a single shareholder, i.e., when Assumption BIB is violated.²⁵

In the lower panel of Figure 2, there is a PA and Assumption PAF is satisfied. Proposition 2 has shown that under Assumption BIB, i.e., for $\frac{\ell_B}{\ell_S} \geq 1$, we have either CAIS or Rubber-stamping as Pareto-efficient symmetric equilibrium. Moreover, by Proposition 2 the parameter space in which CAIS is an equilibrium is given by the condition $\ell_P \in (\ell_B - \ell_S, \ell_B + \ell_S)$, which defines a corridor around the 45-degree line.²⁶ On the 45-degree line the board and the PA are exactly equally well informed, i.e., $\frac{\ell_P}{\ell_S} = \frac{\ell_B}{\ell_S}$ (or $q_B = q_P$). Note that this corridor is not bounded from the upper right. Hence, for arbitrarily well-informed board and PA, there is still an equilibrium with conditional information acquisition of all shareholders, as long as the board and the PA are roughly equally-well informed. The intuition is that whenever their signals contradict each other, there is sufficient controversy to invest in own research. Below this corridor, the board's signal is more informative than the PA's and a single shareholder's signal together such that any shareholder can deviate to Rubber-stamping. Similarly, above the corridor the PAs' signal is more informative than the board's and a single shareholder's signal together such that any shareholder has an incentive to deviate, e.g., to subscribe to the PA and follow its advice. A strategy profile consisting of only this strategy, call it *Follow* (or "robo-voting"), is however not an equilibrium, as any agent can save costs by not subscribing to the PA without affecting the outcome.²⁷ That is why Rubber-stamping is also the Pareto-efficient symmetric equilibrium in this area of the parameter space.

We now assess the comparative statics of changing information quality. Assume $\frac{\ell_B}{\ell_S} > 1$ and start with an uninformed PA: $q_P \approx 0.5$, i.e., $\frac{\ell_P}{\ell_S} \approx 0$. Decision quality remains unaffected by the PA's information quality q_P (or $\frac{\ell_P}{\ell_S}$) at first, then discontinuously increases from q_B to $\Pi(\hat{\sigma})$. Within the region where CAIS is an equilibrium, decision quality further improves as $\Pi(\hat{\sigma})$ is continuously increasing in q_P . Finally, it returns to the level q_B when Rubber-stamping is played again. Hence, there is a non-monotonic effect of a PA's information quality on the corporate decision quality with the latter being highest for a PA that is slightly better informed than the board.²⁸ Comparative-static effects of the board's information quality

²⁵The additional results, due to the violation of Assumption BIB, are provided in the Supplementary Online Material (SOM) Section 1.

²⁶When studying asymmetric equilibria, we show that CAIS can be played by a majority of shareholders far beyond this corridor. The corridor only restricts the area in which *all* shareholders play CAIS.

²⁷We will revisit this question when studying asymmetric equilibria. It turns out that Follow can be part of an equilibrium strategy profile.

²⁸A non-monotonic effect of the PA's recommendation quality on the corporate decision quality is also




Figure 2: Parameter space with Pareto-efficient symmetric equilibria. Upper panel: without a PA; lower panel: with a PA. Both panels: The areas to the left of $\frac{\ell_B}{\ell_S} = 1$ are precluded by Assumption BIB; they are still depicted here to show the effect of the assumption.

are analogous if $\frac{\ell_P}{\ell_S} > 1$, i.e., the PA is better informed than a single shareholder. Finally, increasing signal quality of the shareholders, q_S , reduces $\frac{\ell_B}{\ell_S}$ and $\frac{\ell_P}{\ell_S}$, which means graphically moving towards the origin. This improves decision quality of CAIS because shareholders base their decision on their own information when the PA's recommendation is *against*.

Assumption BIB, i.e., $\frac{\ell_B}{\ell_S} \geq 1$, rules out UNIS, the strategy profile in which all shareholders acquire information. Violating BIB while satisfying PAF, UNIS is the Pareto-efficient symmetric equilibrium in the lower left corner of the parameter space (in the lower panel of Figure 2), which is defined by the condition $\ell_S > \ell_B + \ell_P$. Hence, in the presence of a PA, UNIS requires that one single shareholder must be better informed than board and PA

predicted by the analysis of Malenko and Malenko (2019). In their model, higher PA quality always weakly reduces the shareholders' investment in private signals such that maximal research incentives are obtained for the lowest PA quality. In our model, maximal research incentives are obtained for intermediate PA quality, namely when it equals the board's quality. If proxy advice does not arrive sufficiently early or if the PA is already better informed than the board, a competence-increasing regulation of the PA may undermine shareholders' research incentives and affect decision quality negatively.

together. Interestingly, this is an even stronger condition than the condition for UNIS when no PA is admitted ($\ell_S > \ell_B$).

Let us now compare the upper panel with the lower panel. Under Assumption BIB, i.e., for $\frac{\ell_B}{\ell_S} \geq 1$, the presence of a PA weakly improves decision quality, as it replaces Rubber-stamping with CAIS if anything. When Assumption BIB is violated, there can be a different effect. Suppose that the quality of the board is not much better than a coin flip, i.e., $q_B \approx 0.5$. Then $\frac{\ell_B}{\ell_S} \approx 0$ and there is the equilibrium with full information acquisition (UNIS) and high decision quality, as long as no PA is admitted. The presence of a PA who is better informed than a single shareholder ($\frac{\ell_P}{\ell_S} > 1$) destroys this equilibrium and reduces decision quality from $\pi(N)$ to $q_B \approx 0.5$. The reason is that conditional on pivotality a shareholder prefers to follow the PA's recommendation over acquiring and using the own signal. Hence, Assumption BIB dramatically changes how admission of a PA affects decision quality when studying symmetric equilibria.

5 Extensions and Robustness

5.1 Asymmetric Equilibria

We have characterized the (Pareto-efficient) symmetric equilibria. Since the characterization of all (Pareto-efficient) asymmetric equilibria is quite cumbersome and involves many case distinctions, we relegate it to Supplementary Online Material (SOM) Section 2 and provide only the main result and the essence of the other findings here. Interestingly, although we model shareholders as ex ante homogeneous, there is specialization on different strategies in the Pareto-efficient equilibria, e.g., in one typical equilibrium, some shareholders play CAIS, some shareholders play UNIS, while others play either Rubber-stamping, or always follow the PA's recommendation, depending on whether the board or the PA is better informed.

An overview of the Pareto-efficient asymmetric equilibria is presented in Figure 3 which is proven in SOM Lemma 2.3. *Protest* is defined as the strategy in which a shareholder does not invest in research and votes *no*. *Follow* refers to the strategy where a player buys the PA's recommendation and follows it when voting (without investing in own research).

Our basic results remain similar. In particular, we can first show that without PA, the number of shareholders who invest in own research is bounded from above. That is, in the equilibrium without a PA there are always some shareholders not investing in research, given that Assumption BIB is satisfied. Second, when admitting a PA whose signal quality is not too far from the board's, the number of shareholders who invest or conditionally




Figure 3: Parameter space with Pareto-efficient (potentially asymmetric) equilibria. Upper panel without a PA based on SOM Proposition 2.1, lower panel with a PA based on SOM Lemma 2.3. CAIS is part of the equilibrium strategy profile in large areas of the parameter space.

invest weakly increases. Again, the crucial innovation is that the PA's recommendation is used as a condition to invest in own research in the strategy CAIS. While this was true for all shareholders in Proposition 2 in a certain parameter range, we now find this in much larger areas of the parameter space, while no longer all shareholders use this strategy (see again Figure 3). More precisely, if we have $\frac{|\ell_B - \ell_P|}{\ell_S} < \frac{N+1}{2}$, then $N - \frac{|\ell_B - \ell_P|}{\ell_S}$, i.e., more than half of all shareholders, invest in own information, either conditionally as in CAIS or even unconditionally as in UNIS. The above condition means that the difference between the information quality of the PA and the information quality of the board must not exceed the aggregated information quality of about half of all shareholders together, which graphically widens the corridor in the lower panel of Figure 2 from starting at 1 to starting at $\frac{N+1}{2}$ in the lower panel of Figure 3 on both axes.²⁹ Moreover, the number of investing shareholders, $N - \frac{|\ell_B - \ell_P|}{\ell_S}$, is decreasing in the difference of information quality between the board and the PA. Hence, we find the strongest research incentives for shareholders when the PA is as well informed as the board.³⁰

Finally, the question remains how the effects of a PA on equilibrium behavior translates into decision quality. Proposition 3 provides the answer.

Proposition 3 (ASYM). *Let costs c and f be sufficiently small. For any setting of signal qualities $q_B, q_P, q_S \in (\frac{1}{2}, 1)$, decision quality in any Pareto-efficient equilibrium with a PA under Assumption PAF, $\Pi(\sigma^*)$, is weakly higher than the decision quality in any strategy profile without a PA (including their Pareto-efficient equilibria), i.e., $\Pi(\sigma^*) \geq \bar{\Pi}^{no-PA}$, where $\bar{\Pi}^{no-PA}$ is the maximal decision quality for any strategy profile in the game without a PA.*

The proof of Proposition 3 considers all Pareto-efficient strategy profiles and shows that each of them is an equilibrium with maximal decision quality. It then exploits the insight that any decision quality attainable without a PA can be replicated in the presence of a PA that is simply ignored. Hence, decision quality cannot be reduced by admitting a PA. There are some technical differences between Proposition 3 on asymmetric equilibria and the propositions on symmetric equilibria. First, the simple line of argument used in this proof does not apply to symmetric equilibria. Accordingly, the inefficiency issue that arises there does not carry over to asymmetric equilibria. Second, the assumptions differ: Proposition 3 does not require the fee f to be sufficiently smaller than the cost of private research c . This makes the analysis consistent with the interpretation in Section 2.2, where the fee f reflects

²⁹Observe that the larger the number of shareholders N , the less demanding this assumption is.

³⁰Other comparative-static effects might be different for asymmetric equilibria than for symmetric, see Supplementary Online Material (SOM) Section 2.

shareholders' willingness to pay for the informational value of proxy advice, rather than strategic pricing or market power on the part of the proxy advisor.

Despite these technical differences, the main insight that PAs weakly improve decision quality carries over. This result holds throughout the entire parameter space when considering asymmetric strategy profiles. Moreover, the novel type of equilibrium behavior identified in this paper (CAIS) exists in a broad range of the parameter space.

5.2 One Dominant Shareholder

We have thus far assumed that $N > 1$. Let us now consider the case of only one shareholder $N = 1$, which applies to any company with a shareholder who holds a decisive majority of shares. We can show that both main results carry over to this case. First, without a PA, there is no incentive to invest in research under Assumption BIB, i.e., for $q_S \leq q_B$. Second, the presence of a PA with appropriate information quality improves decision quality, as it leads to a Pareto-efficient equilibrium in which the shareholder conditionally invests in research.

Does this imply that the voting context is unnecessary for our analysis? Certainly not. The single-shareholder case masks the importance of the timing. Since a single shareholder is always pivotal, the Assumption PAF is not necessary for research investment in that special case. That is, even when the subscription decision and the information acquisition decision are made simultaneously, there is an equilibrium with investment in own research for $N = 1$. In this equilibrium strategy the shareholder subscribes to the vote recommendation and invests in own research (Subscribe-Invest) and votes *yes* if and only if at least one of the two supports the board's proposal. Hence, for the case of only one shareholder, there is a complementarity between proxy advice and own research, independently of the timing of the two decisions; while our analysis of voting brings to light that complementarity relies on the timing and addresses the free-rider problem.

5.3 Non-trivial Costs

We assumed that costs of private information acquisition c are small and, for Proposition 2, that the fee for subscribing to proxy advice f is sufficiently smaller. This subsection examines how restrictive these assumptions are and clarifies how behavior changes when fees or research costs are non-trivial, starting with the role of the proxy advisor's fee.

If the PA's fee f is not sufficiently smaller than the costs c , then deviating from symmetric CAIS to UNIS saves costs without affecting the outcome. Indeed, decision quality would not change because in the case of pivotality, the potentially deviating shareholder votes according to the own signal in both strategy profiles. However, this shareholder's costs would change from $C(\hat{\sigma}) = f + p^{dis}c$, with $p^{dis} = q_B(1 - q_P) + (1 - q_B)q_P$, to $C(\sigma^{UNIS}) = c$. Hence, the symmetric equilibrium CAIS requires $f \leq \bar{f} := (1 - p^{dis})c$, which we have formulated as f sufficiently smaller than c . As $p^{dis} \in (0, \frac{1}{2})$, a sufficient condition is $f \leq \frac{c}{2}$. In words, the PA's fee must not exceed half of the individual costs of private research, an assumption that seems realistic. If the fee exceeds this bound, shareholders optimally refrain from subscribing to the proxy advisor, and the conditional information acquisition mechanism no longer operates.

Consider now a monopolistic PA that chooses the subscription fee to maximize profits. As long as shareholders are homogeneous and we focus on symmetric strategy profiles, this extension does not alter the equilibrium logic of the model. If the fee set by the PA exceeds shareholders' willingness to pay for proxy advice, demand collapses and the proxy advisor earns zero profits. Hence, a profit-maximizing monopolistic PA optimally sets the fee up to the shareholders' willingness to pay for the subscription, which we determined here as \bar{f} . Accordingly, treating the fee f as exogenous and sufficiently small can be interpreted as a reduced-form representation of monopolistic pricing under homogeneous demand.

Let us now turn to the costs of private information acquisition c . If they are not sufficiently small, a shareholder has a profitable deviation from CAIS to Rubber-stamping. This move saves private costs for the deviating shareholder but reduces expected decision quality. Computing the expected change in costs and benefits yields the following *small cost condition*, which is sufficient for CAIS to be an equilibrium, given $f \leq \frac{c}{2}$.³¹

$$c < \frac{1}{N} \binom{N-1}{\frac{N-1}{2}} (1 - q_S)^{\frac{N-1}{2}} q_S^{\frac{N-1}{2}} [(1 - q_B)q_Pq_S - q_B(1 - q_P)(1 - q_S)].$$

The first factor on the right-hand side stems from the fact that any voter owns only the share $\frac{1}{N}$ of the firm. This factor disappears when costs are expressed relative to the

³¹The derivation and some examples are provided in Supplementary Online Material (SOM) Section 1.4.

voter’s per-share stake, i.e., as $c^{rel} := \frac{c}{1/N}$. The last factor on the right-hand side, in square brackets, is the conditional change in decision quality and depends on the signal qualities. The middle factors capture the probability of being pivotal, which decreases with N and, for given symmetry, also declines as q_S moves away from 0.5. More precisely, they capture the probability that all $N - 1$ other agents’ signals are split, which decreases quickly with N when q_S is not close to 0.5. How restrictive this small cost condition is depends on the signal qualities and on the number of voters N .

In Example 1 in the main text, the condition requires private costs to be below 0.18% of the per-share stake for $N = 21$ and below 0.52% for $N = 5$. These values are around 1% for $N = 21$ and 3% for $N = 5$ in some other examples that also satisfy all conditions of Proposition 2.

This illustrates that the symmetric equilibrium with conditional private research by all shareholders is most plausible as a benchmark for settings with a limited number of actively engaged voters or with very low marginal research costs, rather than as a literal description of broad shareholder participation in large firms. Thus, the small-cost assumption is restrictive and corresponds to situations in which much is at stake for the firm in the decision or when shareholders indeed face very low marginal research costs. Otherwise, private research equilibria, despite being potentially efficient, unravel due to a free-rider problem that emerges with non-trivial costs.

In reality, costs of private research are heterogeneous across shareholders. There are large investors with economies of scale and small costs per issue, while private costs for other shareholders may be prohibitively high. Our analysis therefore highlights that proxy advice is most likely to improve decision quality when at least a subset of shareholders faces sufficiently low research costs.

5.4 Conflicts of Interest

In a re-interpretation of our model the board has a partial conflict of interest with the shareholders. Suppose the effect of conflicted interests is that it reduces the likelihood that the board’s proposal is correct from q_B to some $\tilde{q}_B < q_B$. More technically, suppose the board’s bias is a random variable that is drawn and private information. Shareholders know the distribution of the bias, but not its realization. The distribution of the bias is such that either the board’s proposal is determined by the bias or that it is determined by the signal. Moreover, suppose that the board’s bias is symmetrically distributed around zero.³² This

³²Asymmetry of the bias distribution would make one type of proposal more informative than the other.

introduces noise into the informativeness of the board’s proposal as shareholders put positive probability on the case that the proposal is independent of the signal. Then the assumption of a high quality board $\tilde{q}_B \geq q_S$ thus means that the board is not only better informed, but also that the board’s agency problem is limited. Conversely, a low \tilde{q}_B means either that the board has a low expertise or that it has a high agency problem such that its proposal is not very informative.

Reconsidering the comparative-statics on q_B (cf., e.g., Figure 2), we can thus also address how the agency problem affects the decision quality. Start with a very well informed board and a small agency problem: $q_B > \tilde{q}_B > q_P > q_S$. Reducing \tilde{q}_B first fosters the shareholders’ research incentives up to the point $\tilde{q}_B \approx q_P$, then reduces them. This non-monotonicity makes it possible that agency problems may even increase the quality of corporate decision making, as boards whose proposals are less informative may incentivize shareholders to (conditionally) invest in own research.

If we further increase the agency problem up to a point where the board’s proposal becomes close to uninformative (i.e., $\tilde{q}_B \approx \frac{1}{2}$), a proxy advisor is indeed detrimental for the reasons elaborated in the literature and illustrated at the end of Section 4. However, a second PA might now challenge the first PA’s recommendation. This could induce the controversy needed to make shareholders invest in own research.

Now, consider instead a biased PA. The implications are analogous to the discussion above. If the bias is symmetric across states and private information to the PA, it affects the model as a reduction in the effective precision of the PA’s signal q_P . We discussed those comparative-static effects already in Section 4. If, instead, the PA or the board is biased in a non-symmetric way, one type of recommendation becomes more informative than the other, which leads to a non-trivial extension beyond the scope of the present analysis.

6 Practical and Regulatory Implications

6.1 Shareholder Behavior in Practice

Our theoretical findings are consistent with observed heterogeneity in how institutional investors use proxy advisors. Proxy advisory reports typically include both a vote recommendation and detailed background analysis. Many large asset managers state that they use these materials as an input into their own decision process rather than mechanically following the recommendation (Swipra, 2018), and bespoke voting policies are increasingly common. This reflects the optionality that our model identifies when shareholders decide whether to

acquire private information in addition to public signals from the board and a proxy advisor.

At the same time, empirical studies find that a nontrivial share of investors often vote in line with proxy advisor recommendations without additional discernible research, especially for routine proposals (Shu, 2021). In the context of our model, such “robo-voting” can be interpreted in two distinct ways. First, even conditionally informed shareholders often vote with advice on routine items. Second, when private information acquisition costs vary across investors, some investors optimally choose strategies that resemble rubber-stamping or following advice without further research, as in the asymmetric equilibria we characterize. Related evidence at the issuer level is provided by Ertimur, Kim, and Pawliczek (2024), who document that firms sometimes respond to negative proxy advisor recommendations by issuing detailed supplemental proxy filings. These filings are effective only when they lead to a revision of the proxy advisor’s recommendation, highlighting the role of proxy advisors as informational gatekeepers and as focal points in the voting process.

6.2 Regulation of Timing of Proxy Advice

Consider the situation when proxy advice arrives *after* the shareholders’ decision to invest in own research, i.e., when Assumption PAF is violated. All actions occur as illustrated in the timeline (Figure 1), but proxy advice arrives at the end of period $t = 2$. We consider the cases where Assumption BIB holds and where it does not hold. If Assumption BIB holds, the Pareto-efficient symmetric equilibrium is Rubber-stamping and hence decision quality is bounded by q_B . Hence, there is no positive effect of having a PA, as decision quality with or without a PA is bounded by the quality of the board. If Assumption BIB does not hold, i.e., if the board does not have the best information regarding what is good for the company, we find that UNIS is an equilibrium and Pareto-efficient if and only if $\ell_S \geq \ell_B + \ell_P$; otherwise, Rubber-stamping is the Pareto-efficient equilibrium. This condition is the same as in our model with early proxy advice (see bottom right corner of the lower panel of Figure 2). It is more demanding than the condition in the setting without a PA, which was $\ell_S > \ell_B$. Specifically, the condition $\ell_S \geq \ell_B + \ell_P$ means that a single shareholder has to be better informed, not only than the board, but than both the board and the PA together. Hence, the introduction of a PA whose information arrives late, if anything, weakens the shareholders’ research incentives.

Overall, a key insight from our model is that the informational value of proxy advice critically depends on its timing relative to shareholders’ decisions to acquire private information. In the model, “sufficiently early” has a precise meaning: proxy advice must arrive

before shareholders decide whether to incur the cost c of private research. Only in this case can shareholders condition their information acquisition on the advice, which is the margin through which proxy advice affects equilibrium behavior.

To interpret this condition in practice, it is useful to think in terms of an *effective research window*, that is, the time between the arrival of proposal-specific information and investors' internal voting cutoffs. Proxy advice is “early” in the sense of the model if it expands this window and thereby allows investors to further update their information. By contrast, advice that arrives only after investors have effectively committed to their information acquisition strategy can influence final votes but has very different consequences for voting strategies.

Recent empirical evidence supports this interpretation. Hirst and Robertson (2022) show that investors often do not know the content of proposals sufficiently in advance, implying that the relevant window for informed voting may open relatively late in the process. Iliev, Kalodimos, and Lowry (2021) document that investor attention to governance is limited and varies systematically over time, consistent with constraints on processing large volumes of information. Complementing this, Calluzzo and Kedia (2025) show that the concentration of shareholder meetings during the proxy season leads to substantial workload compression for proxy advisors and is associated with lower-quality recommendations.

These findings suggest that the effectiveness of proxy advice depends not only on its informational content but also on whether it arrives at a point in time when investors have the capacity to process it and adjust their behavior. This interpretation aligns with Hu, Malenko, and Zytneck (2024), who show that institutional voting follows a two-stage structure in which investors decide, after observing proposals, whether to conduct additional research and deviate from default voting policies.

Viewed through this lens, recent SEC interventions can be understood as attempts to ensure that proxy advice is provided within the effective research window. The relevant issue is not merely whether recommendations are distributed before the shareholder meeting, but whether they are disseminated with sufficient lead time for investors to process them, compare them with issuer responses, and, if warranted, conduct additional research. In our model, this is exactly the condition under which proxy advice improves decision quality.

For example, in August 2019, the U.S. Securities and Exchange Commission (SEC) issued guidance stating that investment advisers fulfill their fiduciary duties of care and loyalty only if they exercise oversight over proxy advisory firms, including scrutiny of methodologies, engagement practices, and error-correction procedures.³³ Subsequent SEC

³³<https://www.sec.gov/rules/interp/2019/ia-5325.pdf>

proposals and rule changes focused explicitly on the timing of proxy advice and issuer responses. A proposed rule from November 2019 raised concerns that investors might receive proxy advisor recommendations too late to meaningfully conduct independent research.³⁴ While the rule adopted in July 2020 partially addressed this issue by requiring mechanisms for timely awareness of issuer responses,³⁵ a further proposal in November 2021 emphasized the importance of delivering independent proxy voting advice to investors in a timely manner.³⁶ This approach was finalized in July 2022.³⁷

Taken together, these developments highlight that the regulatory relevance of timing is more specific than simply “earlier is better”: rules matter insofar as they ensure that proxy advice arrives at a stage where investors can still revise their information acquisition decisions, rather than merely their final votes.

Recent developments in the proxy advisory industry further underscore the relevance of our analysis. Major proxy advisors have adjusted their service offerings, not by abandoning voting recommendations altogether, but by complementing them with additional analytical tools and greater customization. For the 2026 proxy season, ISS continues to issue standardized benchmark voting recommendations, while at the same time expanding governance research platforms that provide analytical content without explicit vote recommendations.³⁸ Similarly, Glass Lewis has moved away from a single uniform recommendation style and now offers multiple voting policy frameworks, allowing institutional clients to choose among different recommendation styles that better align with their own governance preferences.³⁹

Importantly, these changes do not eliminate the presence of public signals in the form of proxy advisor recommendations. Rather, they modify how such signals are generated, interpreted, and integrated into investors’ decision-making processes. From the perspective of our model, these developments can be understood as affecting the precision, framing, and salience of the proxy advisor’s signal, while preserving its role as an early public input into shareholder voting. As long as recommendations (or clearly interpretable voting perspectives) remain available sufficiently early, the core mechanism we study---namely that disagreement between public signals can trigger conditional private information acquisition---continues to apply.

³⁴<https://www.sec.gov/rules/proposed/2019/34-87457.pdf>

³⁵<https://www.sec.gov/rules/final/2020/34-89372.pdf>

³⁶<https://www.sec.gov/rules/proposed/2021/34-93595.pdf>

³⁷<https://www.sec.gov/rules/final/2022/34-95266.pdf>

³⁸<https://insights.issgovernance.com/posts/iss-stoxx-introduces-new-research-services-to-support-investors-proprietary-stewardship-programs/>

³⁹<https://www.glasslewis.com/article/personal-commitment-change-proxy-voting-practices>

More broadly, the coexistence of standardized recommendations and optional analytical tools highlights that proxy advice increasingly serves as a screening and prioritization device rather than a mechanical voting instruction. This interpretation is consistent with our theoretical framework, in which proxy advice helps shareholders identify controversial proposals that warrant further investigation, while routine proposals are more likely to be decided without additional private research.

6.3 Regulation of Proxy Advisor Competence

Regulators continue to focus on disclosure and oversight of proxy advisory methodologies and conflicts of interest. For example, U.S. regulatory actions have stressed compliance with antifraud provisions and enhanced transparency around methodologies and conflicts, while in the EU the Shareholder Rights Directive mandates disclosure of methodology and conflicts to institutional investor clients.

These competence-enhancing measures interact with the model's insight on signal quality. In our model, a proxy advisor with somewhat better information than the board provides incentives for private information acquisition, and yields the highest collective decision quality. At the same time, if the proxy advisor is sufficiently better informed than the board, or if improvements in signal precision make one public expert dominant, the incentives for shareholders to acquire private information may weaken, potentially lowering decision quality. Taken together, these observations suggest that regulatory efforts to improve proxy advisor competence and transparency should be evaluated not only in terms of signal quality per se, but in terms of how they affect both the timing of information and relative precision vis-à-vis board information.

7 Conclusion

In this paper we analyze the effects of proxy advisory firms (PAs) on corporate decisions. As a benchmark PAs are not admitted. When the board's proposals are sufficiently informative, shareholders do not have incentives to conduct their own research and simply rubber-stamp the board's proposals. Hence in the absence of PAs, decision quality is bounded by the quality of the board (Proposition 1). Introducing a PA whose information level is not too far from the board's alters this result and leads to a higher decision quality (Proposition 2). From a conceptual perspective, this effect reflects the restoration of standard sequential value-of-information logic in a voting environment: shareholders optimally acquire a costly private

signal only when earlier public signals generate sufficient disagreement. This result hinges on the timing of proxy advice: only when the PA's recommendation arrives sufficiently early can shareholders condition their private research decisions on it. Extending the analysis from symmetric equilibria to asymmetric equilibria, we find that many but not all shareholders play this conditional investment strategy. Importantly, we arrive at the same overall conclusion: PAs improve corporate decision quality (Proposition 3).

Importantly, our results do not imply that adding more public signals is always beneficial. The effect depends on whether additional signals increase disagreement and thereby trigger private information acquisition. While two (at best equally precise) signals create more controversy than one, adding another one reverts the effect: while two signals are offsetting each other, the third is provided as if there was only a single one. Hence, multiple public signals induce more information acquisition if and only if they create more controversy, e.g., if there is an even number of equally precise signals, or if signal qualities differ in a way that increases controversy. Having an informative proposal and one similarly informed advisor is a cheap and simple way of doing so.⁴⁰

While our analysis is tailored to shareholder voting, related concerns arise in other environments in which public signals may crowd out or stimulate private information acquisition. For example, Sangiorgi and Spatt (2017) argue that credit rating agencies can crowd out independent information production by investors. Our results highlight the mechanism through which public signals can either suppress or restore private research incentives in collective decision-making problems.

More broadly, the paper shows that in voting environments with endogenous information acquisition, standard value-of-information logic may fail due to free-riding and pivotality considerations, and may be restored only under specific timing and information conditions. In the context of shareholder voting, these conditions help explain when proxy advice improves decision quality and when it does not.

⁴⁰The larger the (even) number of conditionally independent informative signals of equal quality, the smaller is the probability of a tie in the signal realizations.

A Appendix: Proofs

A.1 Proof of Proposition 1

To prove Proposition 1 Lemma A.1 is helpful.

Lemma A.1 (SYM without PA: All Equilibria). *Let Assumption BIB hold. Suppose no PA is admitted.*

- i. Define Protest as the strategy profile in which no shareholder invests in research and all shareholders vote no. Protest is a symmetric equilibrium for any $q_B, q_S \in (0.5, 1)$. Its decision quality is $1 - q_B$.*
- ii. Rubber-stamping (i.e., no shareholder invests in research and all shareholders vote yes) is a symmetric equilibrium for any $q_B, q_S \in (0.5, 1)$. Its decision quality is q_B .*
- iii. There are no other symmetric equilibria.*

Proof. We address each part separately.

- i. We have $N \geq 3$ shareholders (because $N > 1$ and odd). When all shareholders vote *no*, a single shareholder is never pivotal. Hence, there is no way to increase decision quality. Deviations can thus only affect costs. Since no information is acquired in this information-acquisition strategy (NotSubscribe-NotInvest), costs are minimal. Hence, there is no unilateral improvement.

Decisions always implement the opposite of the board's proposal. By assumption of the simplified model, the board's proposal corresponds to its signal (B). Hence, the ex ante probability that the true state matches the decision equals the probability that the board's signal does not match the true state, which is $1 - q_B$.

- ii. The proof that Rubber-stamping is an equilibrium is fully analogous to part i. of Lemma A.1. With Rubber-stamping, the decision quality equals the ex ante probability that the board's signal matches the true state, which is q_B .
- iii. There are only two information-acquisition strategies. For not investing in an own signal both strategies are symmetric equilibria (see part i. and ii.). Consider now investment in an own signal: Since shareholders pay c they must condition on their own signal. Otherwise, they could improve their utility by voting the same and not investing c . Conditioning on their signal leaves two pure strategies: vote *yes* if b and *no* if a (i.e., UNIS) or the opposite (vote *yes* if a and *no* if b). If voting *yes* after a (against) was optimal, then voting *no* after a would also be so. Hence, shareholders could profitably deviate to unconditionally voting A .

We finally show that UNIS is not an equilibrium under Assumption BIB, i.e., $q_S \leq q_B$.⁴¹ Consider shareholder i deviates to Rubber-stamping. The deviation changes the outcome only if i is pivotal and the own signal is a : Under UNIS i would vote *no*, under Rubber-stamping i would vote *yes*. Pivotality implies that among the $N - 1$ other shareholders

⁴¹In fact, UNIS is a symmetric equilibrium if and only if $q_S > q_B$.

the signals are split in $\frac{N-1}{2}$ a -signals and $\frac{N-1}{2}$ b -signals. Conditional on that case, B is more likely to be true than A (such that Rubber-stamping weakly improves decision quality) if and only if

$$\begin{aligned} q_B(1 - q_S) \binom{N-1}{\frac{N-1}{2}} q_S^{\frac{N-1}{2}} (1 - q_S)^{\frac{N-1}{2}} &\geq (1 - q_B)q_S \binom{N-1}{\frac{N-1}{2}} (1 - q_S)^{\frac{N-1}{2}} q_S^{\frac{N-1}{2}} \\ q_B(1 - q_S) &\geq (1 - q_B)q_S \\ \frac{q_B}{1 - q_B} &\geq \frac{q_S}{1 - q_S} \\ \ell_B &\geq \ell_S. \end{aligned}$$

Hence, Rubber-stamping weakly improves decision quality for $q_S \leq q_B$, which is Assumption BIB. Moreover, Rubber-stamping saves costs c . Therefore, it strictly improves utility of the deviating shareholder i . □

Now, we use Lemma A.1 to prove Proposition 1. Under Assumption BIB there are only two equilibria. Equilibrium Rubber-stamping leads to the same costs as the Protest equilibrium. Rubber-stamping Pareto-dominates Protest because it leads to higher decision quality $\Pi(\sigma^{Rubber}) = q_B > 0.5 > 1 - q_B = \Pi(\sigma^{Protest})$.

A.2 Proof of Proposition 2

To prove Proposition 2 Lemma A.2 is helpful.

Lemma A.2 (SYM with PA: All Equilibria). *Let Assumptions BIB and PAF hold. Let costs c be arbitrarily small and let fee f be sufficiently smaller.*

- i. Protest (i.e., no shareholder invests in research and all shareholders vote no) is a symmetric equilibrium for any $\ell_B, \ell_S \in (0, \infty)$. Its decision quality is $1 - q_B$.*
- ii. Rubber-stamping (i.e., no shareholder invests in research and all shareholders vote yes) is a symmetric equilibrium for any $\ell_B, \ell_S \in (0, \infty)$. Its decision quality is q_B .*
- iii. CAIS is a symmetric equilibrium if and only if $\ell_P \in (\ell_B - \ell_S, \ell_B + \ell_S)$. Its decision quality is: $\Pi(\hat{\sigma}) = q_B q_P + [(1 - q_B)q_P + q_B(1 - q_P)]\pi(N)$, with $\pi(N) := \sum_{i=\frac{N+1}{2}}^N \binom{N}{i} q_S^i (1 - q_S)^{N-i}$.*
- iv. CAIS-2 is a symmetric equilibrium if and only if $\ell_P \in (\ell_B - \ell_S, \ell_B + \ell_S)$. Its decision quality is: $\Pi(\sigma^{CAIS-2}) = (1 - q_B)(1 - q_P) + [(1 - q_B)q_P + q_B(1 - q_P)]\pi(N)$.*
- v. There are no other symmetric equilibria. In particular, there is no equilibrium in which all shareholders subscribe to proxy advice and unconditionally invest in own signal (Subscribe-Invest).*

Proof. We address each part of Lemma A.2 separately.

- i. The proof is identical to the proof of Lemma A.1, part i.

ii. The proof is identical to the proof of Lemma A.1, part ii.

iii. CAIS is illustrated in Table 1. We show that CAIS is an equilibrium if and only if $\ell_P \in (\ell_B - \ell_S, \ell_B + \ell_S)$.

Suppose first that $\ell_P \notin (\ell_B - \ell_S, \ell_B + \ell_S)$, i.e., either $\ell_P \leq \ell_B - \ell_S$ or $\ell_P \geq \ell_B + \ell_S$. We show that CAIS cannot be an equilibrium. In CAIS pivotality implies that the vote recommendation is *against* and that among the $N - 1$ other shareholders the signals are split in $\frac{N-1}{2}$ *a*-signals and $\frac{N-1}{2}$ *b*-signals. (Indeed, after recommendation *for* no shareholder is pivotal.)

Let $\ell_P \leq \ell_B - \ell_S$. Consider a shareholder i who deviates to Rubber-stamping. This deviation alters the decision in comparison to CAIS if the vote recommendation is *against*, all other shareholder's signals are split, and i 's signal is *a*: In CAIS, i would vote *no*, in the deviation i would vote *yes*. This deviation weakly improves decision quality if $\ell_B \geq \ell_P + \ell_S$, which holds by assumption. Since, the deviation saves costs c , it increases i 's expected utility.

Let $\ell_P \geq \ell_B + \ell_S$. Consider a shareholder i who deviates to voting *no* without information acquisition (as in Protest). This deviation alters the decision in comparison to CAIS if the vote recommendation is *against*, all other shareholders' signals are split, and i 's signal is *b*: In CAIS, i would vote *yes*, in the deviation i would vote *no*. This deviation weakly improves decision quality if $\ell_P \geq \ell_B + \ell_S$, which holds by assumption. Since the deviation saves costs c , it increases i 's expected utility.

Hence, if $\ell_P \notin (\ell_B - \ell_S, \ell_B + \ell_S)$, CAIS is not an equilibrium.

Now suppose that $\ell_P \in (\ell_B - \ell_S, \ell_B + \ell_S)$. In order to show that CAIS is an equilibrium, we demonstrate that there is no individual deviation that improves utility. We use the following principle: if a deviation is more attractive than an other deviation in terms of utility, then excluding the former is sufficient to exclude the latter. We organize the potential deviations by information-acquisition strategy. There are six information-acquisition strategies to consider. Pivotality always implies that the vote recommendation is *against* and that among the $N - 1$ other shareholders the signals are split in $\frac{N-1}{2}$ *a*-signals and $\frac{N-1}{2}$ *b*-signals.

(1) NotSubscribe-NotInvest. Deviating to NotSubscribe-NotInvest and voting *yes* as in Rubber-stamping is not an improvement for low enough costs given $\ell_S + \ell_P > \ell_B$. This deviation only changes the outcome if the PA has recommended *against*, i has received signal *a* (against), and all other shareholders' signals are split. It would weakly improve decision quality iff

$$\begin{aligned}
q_B(1 - q_P)(1 - q_S) \binom{N-1}{\frac{N-1}{2}} q_S^{\frac{N-1}{2}} (1 - q_S)^{\frac{N-1}{2}} &\geq (1 - q_B)q_Pq_S \binom{N-1}{\frac{N-1}{2}} (1 - q_S)^{\frac{N-1}{2}} q_S^{\frac{N-1}{2}} \\
q_B(1 - q_P)(1 - q_S) &\geq (1 - q_B)q_Pq_S \\
\frac{q_B}{1 - q_B} &\geq \frac{q_P}{1 - q_P} \cdot \frac{q_S}{1 - q_S} \\
\ell_B &\geq \ell_P + \ell_S.
\end{aligned}$$

By assumption $\ell_P > \ell_B - \ell_S$, this deviation strictly decreases decision quality. It does save costs f always and c with probability $q_B(1 - q_P) + (1 - q_B)q_P$. For low enough costs f and c , Rubber-stamping does not increase utility because of its lower decision quality.

Deviation to vote *no* without information acquisition (as in Protest) is not an improvement for low enough costs given $\ell_P < \ell_B + \ell_S$.

- (2) NotSubscribe-Invest. Deviation NotSubscribe-Invest and voting according to the own signal as in UNIS does not change the outcome. Indeed, after a *for* recommendation i is not pivotal, after an *against* recommendation i votes under her deviation as she does under CAIS. Hence, this deviation is an improvement only if it saves costs. It is not an improvement if $f \leq c[q_Bq_P + (1 - q_B)(1 - q_P)]$, which is satisfied if f is sufficiently lower than c .
- (3) Subscribe-NotInvest. The deviation to buying the PA's recommendation and following it is not an improvement given $\ell_P < \ell_B + \ell_S$ and low enough c .
- (4) Subscribe-Invest. Deviation to buy both recommendation and signal. Case 1, illustrated in Table A.2, is outcome equivalent, but more costly. Case 2, illustrated in Table A.3, is not an improvement given $\ell_P < \ell_B + \ell_S$.
- (5) Subscribe-InvestIFF*for*. Consider the deviation to buying the PA's recommendation and investing in an own signal iff the recommendation is *for*. The case illustrated in Table A.4 is not an improvement given $\ell_P < \ell_B + \ell_S$. The alternative case, which differs by voting *yes* after the *against* recommendation, is not an improvement given $\ell_S + \ell_P > \ell_B$.
- (6) Subscribe-InvestIFF*against*. Consider the deviation to applying the same information-acquisition strategy as in CAIS, but a different voting strategy. The most attractive deviation is to vote *no* after the *for* recommendation. This is outcome equivalent and equally costly and, hence, not an improvement.

Hence, under the conditions assumed in part iii. of the Lemma CAIS is an equilibrium.

Finally, concerning decision quality, notice that if board and PA receive the same signal, this signal determines the decision, and if they receive a different signal, the signal that is received by a majority of shareholders determines the decision. Therefore, decision quality in CAIS is $(q_Bq_P) * 1 + (1 - q_B)(1 - q_P) * 0 + [(1 - q_B)q_P + q_B(1 - q_P)]\pi(N) * 1$, as q_Bq_P is the probability that the board and the PA both receive the same and correct signal, and $[(1 - q_B)q_P + q_B(1 - q_P)]$ is the probability that the two receive signals that are different from each other.

iv. CAIS-2 is illustrated in Table A.1.

		Own Signal	
		b (for board)	a (against)
PA	for	no	
	$against$	yes	no

Table A.1: CAIS-2: Invest in research iff vote recommendation is *against*; after *for* recommendation vote *no*, after *against* recommendation vote *yes* iff the own signal is b .

The proof that CAIS-2 is an equilibrium if and only if $\ell_P \in (\ell_B - \ell_S, \ell_B + \ell_S)$ is identical to the proof that CAIS is an equilibrium under these conditions (cf. Proof of Lemma A.2, part iii.).⁴²

Concerning decision quality, notice that if board and proxy advisor receive the same signal, the decision is contrary to this signal, and if they receive different signals, the signal that is received by a majority of shareholders determines the decision. Therefore, decision quality in CAIS-2 is $(q_B q_P) * 0 + (1 - q_B)(1 - q_P) * 1 + [(1 - q_B)q_P + q_B(1 - q_P)]\pi(N) * 1$.

v. To show that there are no additional equilibria, we exhaustively discuss all pure strategies. Again, we organize the discussion by information-acquisition strategy.

- (1) NotSubscribe-NotInvest. There are only voting strategies *yes* or *no*. Both lead to equilibria as shown in parts i and ii.
- (2) NotSubscribe-Invest. Since shareholders pay c they must condition on their own signal. Otherwise, they could improve their utility by voting the same and not investing c . Conditioning on the own signal leaves two pure strategies: vote *yes* if the signal is b and *no* if the signal is a (i.e., as in UNIS) or the opposite (vote *yes* if the signal is a and *no* if the signal is b). If voting *yes* after a (*against*) was optimal, then voting *no* after a would also be optimal. Hence, shareholders could improve their utility by unconditionally voting A . Only UNIS remains. Under Assumption BIB, NotSubscribe-NotInvest and voting *yes* as in Rubber-stamping is a profitable deviation from UNIS.
- (3) Subscribe-NotInvest. Since shareholders pay f they must condition on the PA's recommendation. For instance, they vote *yes* after *for* and *no* after *against*; or they do the opposite. In either case, no shareholder is pivotal since all vote for the same, given a particular recommendation.

A shareholder can improve her utility by not paying f and voting, e.g., *yes*. Hence, there is no symmetric equilibrium with this information-acquisition strategy.

- (4) Subscribe-Invest. Since shareholders pay both f and c they must condition their voting strategy on both the PA's vote recommendation and the own signal. Otherwise, they could improve their utility with the same voting behavior, but saving costs. This means that in fact only two voting strategies remain.

Case 1: vote *yes* except if the PA's recommendation is *against* and the own signal is a , as in Table A.2. In this case no shareholder is pivotal if the PA recommends *for*.

⁴²This is not surprising, as both strategies have the same information-acquisition strategy, Subscribe-InvestIFF*against*, and they only differ in a voting action, where no player is pivotal.

Hence, shareholder i can only be pivotal if the recommendation is *against*. If so, i would vote according to her signal. Hence, deviating to unconditionally investing in an own signal and voting accordingly, as in UNIS, would not change the outcome because either i is not pivotal or i would also vote according to the signal. However, acting as in UNIS saves fee f . Thus, this is a profitable deviation, and the strategy profile of case 1, illustrated in Table A.2, cannot be an equilibrium.

		Own Signal	
		b (for board)	a (against)
PA	for	yes	yes
	$against$	yes	no

Table A.2: A strategy profile based on acquiring both proxy advice and own signal, case 1: Subscribe-Invest and vote *yes*, except if PA's recommendation is *against* and the own signal is a .

Case 2: vote *no* except if both the PA's recommendation is *for* and the own signal is b , as in Table A.3. The analogous argument as above for case 1 applies, as follows: In this case no shareholder is pivotal if the PA recommends *against*. Hence, shareholder i can only be pivotal if the recommendation is *for*. If so, i would vote according to signal. Hence, deviating to unconditionally investing in an own signal as in UNIS would not change the outcome because either i is not pivotal or i would also vote according to the signal. Acting as in UNIS, however, saves fee f . Thus, strategy profile of case 2 cannot be an equilibrium.

		Own Signal	
		b (for board)	a (against)
PA	for	yes	no
	$against$	no	no

Table A.3: A strategy profile based on acquiring both proxy advice and own signal, case 2: Subscribe-Invest and vote *no*, except if PA's recommendation is *for* and the own signal is b .

Therefore, there cannot be a symmetric equilibrium with this information-acquisition strategy (Subscribe-Invest), in which shareholders unconditionally buy both the PA's recommendation and an own signal.

- (5) Subscribe-Invest IFF *for*. Since shareholders pay f and sometimes c they must condition their voting strategy on the recommendation and the own signal when they acquire them. In particular, after having bought the own signal on top of the recommendation *for*, shareholders must vote according to their signal in equilibrium. Voting the opposite is dominated, and not conditioning as well. This leaves two cases, which we address as Cand. 5a and Cand. 5b. We show that none of them is an equilibrium under Assumption 1.⁴³ Consider first Cand. 5a: shareholders vote

⁴³In fact, each of these strategy profiles is a symmetric equilibrium if and only if $\ell_S > \ell_B + \ell_P$.

yes except if the PA's vote recommendation is *for* and the own signal is *a* (against) as in Table A.4.

		Own Signal	
		<i>b</i> (for board)	<i>a</i> (against)
PA	<i>for</i>	yes	no
	<i>against</i>	yes	

Table A.4: Cand. 5a. A strategy profile based on acquiring an own signal iff the recommendation is *for*: Subscribe-InvestIFF*for* and vote *yes*, except if PA's recommendation is *for* and the own signal is *b*.

Consider shareholder i who deviates to NotSubscribe-NotInvest and voting *yes*, as in Rubber-stamping. This deviation only alters the outcome when the vote recommendation is *for*, all other shareholders' signals are split, and i 's signal is *a* (*against*): Under Cand. 5a, i would vote *no*, but under her deviation she votes *yes*. Decision quality improves by this deviation if $\ell_B + \ell_P > \ell_S$. This condition is satisfied by Assumption 1. Moreover, the costs are lower under this deviation than under Cand. 5a. Hence, Cand. 5a cannot be an equilibrium.

Now consider Cand. 5b. Shareholders vote *no* except if the PA's vote recommendation is *for* and the own signal is *b*. Again, no shareholder is pivotal after recommendation *against*. Hence, deviating to NotSubscribe-NotInvest and voting *yes* as in Rubber-stamping is an improvement, identical to the case of Cand. 5a above.

- (6) Subscribe-InvestIFF*against*. Since shareholders pay f and sometimes c they must condition their voting strategy on the recommendation and the own signal when they acquire them. In particular, after having bought the own signal on top of the recommendation *against*, shareholders must vote according to their signal in equilibrium. Voting the opposite is dominated, and not conditioning as well. This leaves two cases: CAIS and CAIS-2, which we have addressed. Hence, there are no further equilibria.

□

Now we can turn to the proof of Proposition 2. Suppose there is a PA with $\ell_P \in (\ell_B - \ell_S, \ell_B + \ell_S)$. To show that CAIS is an equilibrium and Pareto-efficient, we use Lemma A.2, which shows that besides CAIS there are three further equilibria in this parameter space: Rubber-stamping, Protest, and CAIS-2. It remains to show that CAIS Pareto-dominates in this area.

First, CAIS has the same costs as CAIS-2 and decision qualities are: $\Pi(\hat{\sigma}) = q_B q_P + [(1 - q_B)q_P + q_B(1 - q_P)]\pi(N)$. $\Pi(\sigma^{CAIS-2}) = (1 - q_B)(1 - q_P) + [(1 - q_B)q_P + q_B(1 - q_P)]\pi(N)$. CAIS has higher decision quality iff $q_B q_P > (1 - q_B)(1 - q_P)$, which always holds as $q_B, q_P > 0.5$. Hence, CAIS Pareto-dominates CAIS-2.

Second, decision quality of Rubber-stamping is q_B and decision quality of Protest is

$1 - q_B < q_B$. CAIS has strictly higher decision quality than both iff

$$\begin{aligned}
\Pi(\hat{\sigma}) &> q_B \\
q_B q_P + [q_B(1 - q_P) + (1 - q_B)q_P]\pi(N) &> q_B \\
q_B(1 - q_P)\pi(N) + (1 - q_B)q_P\pi(N) &> q_B(1 - q_P) \\
(1 - q_B)q_P\pi(N) &> q_B(1 - q_P)[1 - \pi(N)] \\
\frac{\pi(N)}{1 - \pi(N)} \cdot \frac{q_P}{1 - q_P} &> \frac{q_B}{1 - q_B} \\
\log\left(\frac{\pi(N)}{1 - \pi(N)}\right) + \log\left(\frac{q_P}{1 - q_P}\right) &> \log\left(\frac{q_B}{1 - q_B}\right) \\
\ell_N + \ell_P &> \ell_B, \\
\text{where } \ell_N &:= \log\left(\frac{\pi(N)}{1 - \pi(N)}\right).
\end{aligned} \tag{A.1}$$

Since $\ell_N > \ell_S$ and by assumption $\ell_P > \ell_B - \ell_S$, we have $\ell_N + \ell_P > \ell_S + \ell_P > \ell_B$. Hence, CAIS leads to strictly higher decision quality than both Rubber-stamping and Protest. It induces higher costs f and c . Thus, for low enough costs, CAIS Pareto-dominates them.

Now, suppose that $\ell_P \notin (\ell_B - \ell_S, \ell_B + \ell_S)$. To show that the Pareto-efficient equilibrium is Rubber-stamping, we use again Lemma A.2. Under Assumption BIB and for $\ell_P \notin (\ell_B - \ell_S, \ell_B + \ell_S)$, only two equilibria remain: Rubber-stamping and Protest. Rubber-stamping Pareto-dominates because it leads to higher decision quality $\Pi(\sigma^{Rubber}) = q_B > 0.5 > 1 - q_B = \Pi(\sigma^{Protest})$, while it induces the same costs.

A.3 Proof of Proposition 3

Suppose there is a PA and Assumption PAF holds.⁴⁴ Let S be the set of all pure strategy profiles.⁴⁵ Let $\Pi : S \rightarrow [0, 1]$ be the decision quality. Let $S^{max} \subset S$ be the set of all strategy profiles that maximize Π . As S is finite, S^{max} is non-empty.

A player's strategy is called *minimal* if the voting strategy conditions on all pieces of information that are acquired. For instance, consider information-acquisition strategy *Subscribe-InvestIFFagainst*: When an own signal has been acquired after vote recommendation *against*, the voting behavior must differ between signal realization a and signal realization b to be part of a minimal strategy. Observe that for any strategy that is not minimal, the voting behavior can be mimicked by a strategy with lower costs, saving c or f or both. A strategy profile(!) is called *minimal* if all players' strategies are minimal *and* if any player's reduction of information acquisition (not subscribing and/or not acquiring an

⁴⁴We are particularly thankful to Maximilian Janisch and Thomas Lehéricy who suggested this proof idea for this proposition. Interestingly, it can be applied to asymmetric equilibria, but not to symmetric equilibria. The reason is that when studying symmetric equilibria, the space under consideration changes because deviations to strategies that thus form an asymmetric strategy profile are admitted. An alternative proof approach would consider our game as a game with a potential function in the sense of (Monderer and Shapley, 1996).

⁴⁵Then there are 16 strategies for each shareholder. We restrict attention to pure strategies.

own signal) changes the decision quality.⁴⁶ Now let $S^* \subset S^{max}$ be the strategy profiles that are minimal and lead to maximal decision quality.

We first show, as Claim 1, that all $\sigma^* \in S^*$ are equilibria. A shareholder can only achieve higher utility than in σ^* by higher decision quality or lower costs. Higher decision quality is impossible per definition. Lower costs reduce decision quality because σ^* is minimal. As costs are by assumption sufficiently small, lower costs that reduce decision quality are not an improvement.

Second, we show, as Claim 2, that any Pareto-efficient strategy profile must belong to S^* . Suppose first that σ' is Pareto-efficient, but not in S^* . Then it is either not maximizing decision quality or not minimal. If it does not maximize decision quality, take another strategy profile, say σ^* , that does and every shareholder is better off. The reason is that any difference in decision quality is always larger than the difference in costs, which are by assumption sufficiently small; formally: $\Pi(\sigma^*) - \Pi(\sigma') > c + f \implies u_i(\sigma^*) > u_i(\sigma')$ for all i . If σ' is not minimal, there is a player who can save costs without affecting decision quality and utility of other shareholders.

By Claim 1 and 2 together, each Pareto-efficient strategy profile is an equilibrium with maximal decision quality. Hence, there clearly exists an equilibrium, say σ^* , with maximal decision quality. Now consider any Pareto-efficient equilibrium σ , i.e., an equilibrium that is not Pareto-dominated by any other equilibrium. This equilibrium must also maximize decision quality, i.e., $\sigma \in S^{max}$. Otherwise, it would be dominated by σ^* , as higher decision quality means strictly higher utility for every shareholder (again due to the small cost assumption). Therefore, *every Pareto-efficient equilibrium must maximize decision quality.*

Let us now turn to the model without a PA.⁴⁷ Let T be the set of all pure strategy profiles (without a PA). In full analogy to above, we define T^* as the set of strategy profiles that are maximizing decision quality and that are minimal. Now observe that any strategy profile in T (without a PA) corresponds to a strategy profile in S (with a PA) where simply no player subscribes to the PA's vote recommendation. Consequently, any decision quality obtained with a strategy profile in T can also be obtained with a strategy profile in S . Let $\tilde{\sigma} \in S$ be a strategy profile that mimicks the maximal decision quality obtainable without a PA. Let $\bar{\Pi}^{with-PA}$, respectively $\bar{\Pi}^{no-PA}$, denote the maximal decision quality in the framework with a PA, respectively without a PA (for any strategy profile in the corresponding games). Since in the Pareto-efficient Nash equilibria with a PA, decision quality is maximal, we have $\Pi(\sigma^*) = \bar{\Pi}^{with-PA} \geq \Pi(\tilde{\sigma}) = \bar{\Pi}^{no-PA}$ for any Pareto-efficient equilibrium σ^* in the game with a PA.

⁴⁶Indeed, there are strategy profiles where all strategies are minimal, but still some players can reduce their information-acquisition without affecting the decision. For instance, when $N - 1$ players play NotSubscribe-NotInvest and vote *yes* as if in Rubber-stamping and one player does not subscribe to the PA but unconditionally invests into an own signal, as in UNIS. If the latter player stops acquiring an own signal, decisions are unaffected because she is never pivotal. By definition, such strategy profiles are not minimal.

⁴⁷There are only four strategies for each player: as in Rubber-stamping, as in Protest, as in UNIS, and voting contrary to the own signal.

References

- Admati, A. R., and P. Pfleiderer. 1990. Direct and indirect sale of information. *Econometrica* 58:901--28.
- Alexander, C. R., M. A. Chen, D. J. Seppi, and C. S. Spatt. 2010. Interim news and the role of proxy voting advice. *Review of Financial Studies* 23:4419--54.
- Angeletos, G.-M., and A. Pavan. 2007. Efficient use of information and social value of information. *Econometrica* 75:1103--42.
- Austen-Smith, D., and J. S. Banks. 1996. Information aggregation, rationality, and the Condorcet Jury Theorem. *American Political Science Review* 90:34--45.
- Bar-Isaac, H., and J. Shapiro. 2020. Blockholder voting. *Journal of Financial Economics* 136:695--717.
- Brav, A., and R. D. Mathews. 2011. Empty voting and the efficiency of corporate governance. *Journal of Financial Economics* 99:289--307.
- Buechel, B., and L. Mechtenberg. 2019. The swing voter's curse in social networks. *Games and Economic Behavior* 118:241--68.
- Calluzzo, P., and S. Kedia. 2025. Managing the workload of the proxy season. *Journal of Law and Economics* 68:871--900.
- Choi, S., J. E. Fisch, and M. Kahan. 2010. The power of proxy advisors: Myth or reality? *Emory Law Journal* 59:869--907.
- De Caritat, Marie Jean Antoine Nicolas, M. d. C. 1785. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. L'imprimerie royale.
- Dey, A., A. Starkweather, and J. T. White. 2022. Proxy Advisory Firms and Corporate Shareholder Engagement. *Working Paper*.
- Ertimur, Y., F. Ferri, and D. Oesch. 2013. Shareholder votes and proxy advisors: Evidence from say on pay. *Journal of Accounting Research* 51:951--96.
- Ertimur, Y., J. Kim, and A. Pawliczek. 2024. Firm Responses to Proxy Advisor Recommendations: Evidence from Supplemental Proxy Filings. *Working Paper*.
- Feddersen, T., and W. Pesendorfer. 1996. The swing voter's curse. *American Economic Review* 408--24.
- . 1997. Voting behavior and information aggregation in elections with private information. *Econometrica* 1029--58.
- . 1998. Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting. *American Political Science Review* 23--35.
- Gerardi, D., and L. Yariv. 2007. Deliberative voting. *Journal of Economic Theory* 134:317--38.
- Gershkov, A., and B. Szentes. 2009. Optimal voting schemes with costly information acquisition. *Journal of Economic Theory* 144:36--68.
- Hellwig, C., and L. Veldkamp. 2009. Knowing What Others Know: Coordination Motives in Information Acquisition. *Review of Economic Studies* 76:223--51.
- Hirst, S., and K. Robertson. 2022. Hidden agendas in shareholder voting. *Yale Journal on Regulation* 39:1161--218.
- Hu, E., N. Malenko, and J. Zytneck. 2024. Custom proxy voting advice. NBER Working Paper 32559, National Bureau of Economic Research.

- Iaryczower, M., X. Shi, and M. Shum. 2018. Can Words Get in the Way? The Effect of Deliberation in Collective Decision Making. *Journal of Political Economy* 126:688--734.
- Iliev, P., J. Kalodimos, and M. Lowry. 2021. Investors' attention to corporate governance. *Review of Financial Studies* 34:5581--628.
- Iliev, P., and M. Lowry. 2015. Are mutual funds active voters? *Review of Financial Studies* 28:446--85.
- Jensen, M., and W. H. Meckling. 1976. Theory of the firm: Managerial behavior, agency costs, and ownership structure. *Journal of Financial Economics* 3:305--60.
- Jeong, D. 2019. Using cheap talk to polarize or unify a group of decision makers. *Journal of Economic Theory* 180:50--80.
- Kawamura, K., and V. Vlaseros. 2017. Expert information and majority decisions. *Journal of Public Economics* 147:77--88.
- . 2024. Efficient equilibria in common interest voting games. *International Journal of Game Theory* 53:475--92.
- Larcker, D. F., A. McCall, and G. Ormazabal. 2013. Proxy advisory firms and stock option repricing. *Journal of Accounting and Economics* 56:149--69.
- . 2015. Outsourcing shareholder voting to proxy advisory firms. *Journal of Law and Economics* 58:173--204.
- Levit, D., and N. Malenko. 2011. Nonbinding voting for shareholder proposals. *The Journal of Finance* 66:1579--614.
- Levit, D., N. Malenko, and E. Maug. 2021. The Voting Premium. *Working Paper*.
- . 2024. Trading and Shareholder Democracy. *The Journal of Finance* 79:257--304.
- Levy, G., and R. Razin. 2015. Correlation neglect, voting behavior, and information aggregation. *American Economic Review* 105:1634--45.
- Li, T. 2018. Outsourcing corporate governance: Conflicts of interest within the proxy advisory industry. *Management Science* 64:2473--972.
- Liu, S. 2019. Voting with public information. *Games and Economic Behavior* 113:694--719.
- Ma, S., and Y. Xiong. 2021. Information bias in the proxy advisory market. *Review of Corporate Finance Studies* 10:82--135.
- Malenko, A., and N. Malenko. 2019. Proxy advisory firms: The economics of selling information to voters. *The Journal of Finance* 74:2441--90.
- Malenko, A., N. Malenko, and C. Spatt. 2021. Creating controversy in proxy voting advice. *Working Paper*.
- . 2025. Creating controversy in proxy voting advice. *The Journal of Finance* 80:2303--54.
- Malenko, N., and Y. Shen. 2016. The role of proxy advisory firms: Evidence from a regression-discontinuity design. *Review of Financial Studies* 29:3394--427.
- Matusaka, J. G., and C. Shu. 2021. Does proxy advice allow funds to cast informed votes? *Working Paper*.
- Matvos, G., and M. Ostrovsky. 2006. Strategic Proxy Voting. *Working Paper*.
- Maug, E. 1999. How effective is proxy voting? Information aggregation and conflict resolution in corporate voting contests. *Working Paper*.
- Maug, E., and K. Rydqvist. 2009. Do shareholders vote strategically? Voting behavior, proposal

- screening, and majority rules. *Review of Finance* 13:47--79.
- Maug, E., and B. Yilmaz. 2002. Two-class voting: A mechanism for conflict resolution. *American Economic Review* 92:1448--71.
- McCahery, J. A., Z. Sautner, and L. T. Starks. 2016. Behind the scenes: The corporate governance preferences of institutional investors. *Journal of Finance* 71:2905--32.
- Meirowitz, A., and S. Pi. 2022. The Shareholder's Dilemma: Voting and Trading. *Journal of Financial Economics* 146:1073--96.
- Monderer, D., and L. S. Shapley. 1996. Potential games. *Games and Economic Behavior* 14:124--43. ISSN 0899-8256. doi:<https://doi.org/10.1006/game.1996.0044>.
- Nitzan, S., and J. Paroush. 1982. Optimal decision rules in uncertain dichotomous choice situations. *International Economic Review* 289--97.
- Parlasca, M., and P. Voss. 2023. Voting and trading on proxy advice. *Available at SSRN 4673547* .
- Persico, N. 2004. Committee design with endogenous information. *Review of Economic Studies* 71:165--91.
- Sangiorgi, F., and C. Spatt. 2017. The economics of credit rating agencies. *Foundations and Trends in Finance* 12:1--116.
- Shu, C. 2021. The proxy advisory industry: Influencing and being influenced. *Working Paper*.
- Spatt, C. 2021. Proxy advisory firms, governance, market failure, and regulation. *Review of Corporate Finance Studies* 10:136--57.
- Swipra. 2018. Corporate Governance Survey 2018. https://www.swipra.ch/uploads/news/SWI_PRA_Survey_2018_findings.pdf.
- Van Wesep, E. D. 2014. The Idealized Electoral College voting mechanism and shareholder power. *Journal of Financial Economics* 113:90--108.