How CEO-Friendly Should Boards With Limited Attention Be?*

Martin Gregor† Beatrice Michaeli‡

January 8, 2021

*We thank Tim Baldenius, Gregory Burke, Brian Cadman, Judson Caskey, Jonathan Craske, Mark DeFond, Yonca Ertimur, Elia Ferracuti, Fabrizio Ferri, Henry Friedman, Jonathan Glover, Robert Göx, Ilan Guttman, John Heater, Mirko Heinle, Shane Heitzman, John Hughes, Xu Jiang, Navin Kartik, Anya Kleymenova, Clive Lennox, Tracie Majors, Xiaojing Meng, Katherine Schipper, Richard Sloan, Lorien Stice-Lawrence, Philip Stocken, Rahul Vashishtha, T.J. Wong, and workshop participants at Duke Fuqua Business School, Dartmouth Tuck School of Business, University of Southern California, and UCLA Anderson School of Management for helpful comments and suggestions. Martin Gregor acknowledges financial support from the Czech Science Foundation (GAČR 20-03517S). Beatrice Michaeli acknowledges financial support from UCLA Anderson School of Management.

†Martin Gregor is at Institute of Economic Studies, Faculty of Social Sciences, Charles University, and can be reached at martin.gregor@fsv.cuni.cz.

‡Beatrice Michaeli is at UCLA Anderson School of Management, and can be reached at beatrice.michaeli@anderson.ucla.edu.

How CEO-Friendly Should Boards With Limited Attention Be?

Abstract: A CEO who is an empire-builder reports information about an investment opportunity ("project"). Before approving or rejecting the project, a board of directors decides whether and how much additional information to collect, i.e., whether to remain rationally inattentive. We show that the CEO prepares and communicates a report that is just sufficiently precise so as to persuade the board not to learn any additional information and to approve some value-destroying projects (type-I error). The informativeness of the report is increasing in the misalignment of interests between the board and the CEO. Because more informative reports reduce the probability of approval error, the shareholders may optimally assemble a board that is "unfriendly" to the CEO. Our model predicts that (i) board-dependence regulations lead to decrease in corporate investments but an increase in return to shareholders; (ii) the return on investment is lower in companies with busy directors operating in less mature industries; (iii) the conflict of interest between boards and CEOs is either mild or moderately strong and the likelihood of strong conflict is lower in companies with busy directors and companies operating in less mature industries or industries with attractive outside opportunities for directors.

Keywords: Bayesian persuasion, rational inattention, empire building, board of directors, friendly boards, alignment of interests

JEL: D71, D72, D81, D82, D83, G34

1 Introduction

Prior literature shows that CEOs facing friendly boards have stronger incentives to communicate any private information with which they are exogenously endowed (Adams and Ferreira, 2007; Baldenius, Meng and Qiu, 2019). Because independent boards are perceived as less CEO-friendly, this result sheds light on some of the anticipated effects of board-dependence regulations. However, the information available to CEOs is often endogenous: as part of their exploration of investment opportunities and preparation for presentation to the board, CEOs decide what type and how much information to collect. It is unclear whether facing friendly boards increases or decreases the incentives of CEOs to acquire information in the first place. Furthermore, directors may also learn additional decision-relevant information on their own but learning information could be costly and time-consuming. Thus, boards may choose to be rationally inattentive and this choice endogenously depends on the information communicated by CEOs and the boards' attitude and preferences. In this paper, we consider a scenario in which the information collected by CEOs and boards is endogenous and study the effects of boards' alignment of interests with CEOs on the board's rational inattention, the total available decision-relevant information and the investment efficiency. We also study the optimal board nomination within the confines of the model by solving for the alignment of interests between boards and CEOs that maximizes shareholders' value.

We develop a model in which a board of directors approves or rejects an investment opportunity ("project") proposed by a CEO ("she").[1] The preferences of the shareholders and the CEO are not perfectly aligned. Specifically, the CEO's payoff from undertaking the project is higher than the shareholders' payoff due to some private benefits. Thus, the CEO favors not only projects that increase firm value but also some projects that destroy it. We say that the CEO has "empire-building tendencies." The alignment of

[1] "[The board's] most important function is to approve or send back for amendment management's recommendations about the future direction of the corporation" (Wommack, 1979).

interests of the board with those of the shareholders and the CEO determines the board's attitude towards the project proposed by the CEO. Specifically, a board with preferences that are more aligned with those of the CEO (e.g., because the directors are insiders and are loyal to the CEO or because of some other private benefits) is more CEO-friendly in the sense that the board has a higher tendency to approve not only the value-increasing projects but also some value-destroying projects that the CEO favors. And vice versa: a board with preferences that are more misaligned with those of the CEO (e.g., because the directors are outsiders, have career concerns or interest in competing companies) is more unfriendly to the CEO and has a higher tendency to reject some of the projects favored by the CEO.

The CEO prepares a report with investment-relevant information and presents it to the board.[2] Besides the report, the board can learn additional information. Because the board's time and attention are costly, it may choose to learn imperfect or even no information at all, i.e., be rationally inattentive. In the process of nominating directors, the shareholders strategically determine the attitude and alignment of interests between the CEO and the board (e.g., by including insiders, outsiders and directors with career concerns or interest in other companies).

As a first step, we study the properties of the CEO's report for a given misalignment of interests between the board and the CEO. We show that the CEO prepares a report that discourages learning by the board. The reason is as follows. Any information that the board can learn can also be collected and reported by the CEO. However, while the CEO controls the properties of the report, she has no control over the additional information collected by the board. Thus, if the board learns, the CEO faces additional uncertainty regarding the beliefs used as a basis for the project approval decision. To avoid this additional uncertainty, the CEO prepares a sufficiently informative report. This renders

[2] Agrismart and Plymouth Tube cases focus on a CEO who collects and presents convincing information to the board for approval of a business opportunity (Wouters and Davila, 2018; Ward and Zsolnay, 2016).

the learning of additional information by the board unnecessary. Notably, while the board's learning option influences the optimal properties of the report, this option is never exercised in equilibrium. Our result predicts that boards remain rationally inattentive and rely on the information provided by CEOs without searching for additional information on their own.[3]

Furthermore, we find that the board never rejects value-increasing projects (i.e., never commits type-II error) but may approve value-decreasing projects (i.e., may commit type-I error). The reason is that the CEO has no benefit when the board rejects a value-increasing project but benefits if the board approves some value-destroying projects. Thus, the CEO maximizes her expected payoff by preparing an imperfectly precise report that encourages the board to always approve high type (value-increasing) projects and sometimes approve intermediate type (value-destroying) projects. This is achieved by reporting the lowest type projects in a precise manner and pooling high-type with intermediate projects.[4] Interestingly, the optimal report is more informative than the one needed for mere project approval.[5] This is because the CEO needs to ensure that the board does not learn additional information which requires communicating more precise information.

We also find that the CEO prepares a more informative report when faced with a less friendly board. This is because a board with preferences that differ substantially from those of the CEO is less likely to approve the project. In fact, such board approves without learning only when the CEO's report implies that the project is value-increasing with sufficiently high probability. Because an unfriendly board receives more precise reports,

[3]Our result is consistent with references in the press that boards "depend largely on the chief executive and the company's management for information" (The Economist, 2001). Active board's learning requires a friction that is external to our setting such as a limit on the set of available report properties.

[4]For example, the CEO could solicit an expert to take soil samples and determine the contamination levels in an area of potential business opportunity. Our model assumes that perfect precision is achievable. In practice, information precision could be limited, which could explain empirical evidence on the existence of type-II errors in decision making.

[5]This result is in contrast to findings in persuasion models without receiver's learning (e.g., Kamenica and Gentzkow, 2011).

it commits fewer errors and so approves fewer but more profitable projects. Thus, an exogenous increase in the severity of the conflict of interest between the board and the CEO reduces the firm's total investments. Our findings therefore provide new insights and intuition for the declines in investment rates following board-dependence regulations (Frydman and Hilt, 2017; Jagannathan, Jiao, and Krishnamurthy, 2020). At the same time, we predict that such regulations will be associated with higher returns on corporate investments.

Next, we study the optimal board nomination within the confines of our model by solving for the value-maximizing alignment of interests between the CEO and the board of directors.[6] Because unfriendly boards commit fewer approval errors (as a result of the more precise information reported by the CEO), the shareholders benefit from appointing directors who have more misaligned interests with the CEO. Such directors also benefit from the more precise information but at the same time they also receive a lower payoff from an approved project.[7] Therefore, all else being equal, more unfriendly directors might be less willing to participate in the board and would rather undertake an outside opportunity (e.g., accept a position at another firm). This limits the degree of the conflict of interest between the board and the CEO that the shareholders can invoke by nominating unfriendly directors. We show that the optimal alignment of interests depends on a surprisingly simple condition. Specifically, there exists a critical board type which determines the equilibrium severity of the conflict between the CEO and the board. If the critical type is willing to participate, the shareholders assemble an unfriendly board with an interior level of interest misalignment. However, if the critical type is not will-

[6]In our model, the misalignment of the CEO's interests with those of the shareholders is taken as given (either because the CEO is already hired, or because all available CEO candidates have some degree of private interests). While the shareholders cannot align directly the CEO's private interests with their own interests, they can strategically nominate directors with certain characteristics and, by doing so, determine the alignment of interests between the CEO and the board.

[7]In particular, the board's expected payoff is concave in the interest misalignment. For small increase in the misalignment, the benefit from higher precision outweighs the reduction in payoff from project approval. The opposite is true, however, for a large increase in the misalignment.

ing to participate, the shareholders are indifferent between a range of possible levels of interest misalignment. In this case, the board is CEO-friendly (to the degree allowed by board-dependence regulations). We thus predict the existence of two distinctly different board types.

A comparative statics analysis shows that shareholders are more likely to assemble a CEO-friendly board when learning is particularly costly. In our setting, learning is costlier when the board members are more inattentive (e.g., because they serve on several boards—such directors are often referred to as "busy") or when there are fewer available information sources (which is the case in less mature industries). In addition, we predict that, when directors have better outside opportunities, the shareholders assemble a more friendly board and the equilibrium conflict between the CEO and the board is less likely to be severe. This is because friendly directors are more likely to accept a position with the board.

Our paper contributes to several strands of literature. In terms of employed paradigms, our paper belongs to the growing literature studying Bayesian persuasion models.[8] As in Huang (2016), the sender designs a report and a receiver decides whether to collect additional information. However, in Huang (2016), the receiver incurs a fixed cost for perfectly learning the underlying state whereas in our setup, in line with the stochastic discrete choice theory (Matějka and McKay, 2015), we assume that the receiver is rationally inattentive and the learning cost is entropy-based so that learning more precise information is costlier. Unlike Huang (2016), we find that the receiver never learns additional information in equilibrium and study the optimal misalignment of interests between the sender

[8]The Bayesian persuasion model was established by Kamenica and Gentzkow (2011) and has been extended to settings with multiple receivers (e.g., Michaeli, 2017), multiple senders (e.g., Gentzkow and Kamenica, 2017), information acquisition by receivers (e.g., Huang, 2016; Matysková, 2018; Caplin, Dean and Leahy, 2019), interaction between ex ante design of public information and ex post disclosure of private information (e.g., Friedman, Hughes and Michaeli, 2020), agency problems (e.g., Göx and Michaeli, 2019), liquidation decisions (e.g., Bertomeu and Cheynel, 2015), signaling (e.g., Jiang and Yang, 2017; Dordzhieva, Laux and Zheng, 2020), mutual persuasion (e.g., Jiang and Stocken, 2019), and asset pricing (e.g., Cianciaruso, Marinovic and Smith, 2020). Earlier studies have also considered ex ante information design (e.g., Arya, Glover and Sivaramakrishnan, 1997; Göx and Wagenhofer, 2009).

and the receiver. Like us, Bloedel and Segal (2020) study a persuasion model with an inattentive receiver. However, Bloedel and Segal (2020) assume that the receiver interprets the reports imprecisely and find that the sender partitions the states into three intervals with pooling at the bottom and top, and separation in the middle. We assume instead that the board interprets the report precisely but can learn additional information. In contrast to Bloedel and Segal (2020), we find that low types are always revealed but some or all intermediate types are pooled with high types. Similar to Matysková (2018) and Caplin et. al (2019), we find that the report designed by the sender (the CEO in our case) discourages the receiver (the board in our case) from learning additional information in equilibrium. Different from Matysková (2018) and Caplin et. al (2019), we study how the receiver's characteristics and alignment of interests with the sender influence the ex ante design of the report. We also study the optimal choice of these characteristics by a third party (the shareholders). Furthermore, in an extension, we consider the ability of the CEO to communicate subsequently observed private information via cheap talk. Similar to our approach in the extension, Jain (2020) studies a model in which the sender engages in both persuasion and cheap talk. However, Jain (2020) finds that the sender designs a less informative signal in the presence of cheap talk. In contrast, the informativeness of the signal in our setting remains unaffected by cheap talk.

In our model, the board is characterized by a certain degree of interest misalignment. Hence, our paper also relates to the literature on directors' conservative bias. Jiang, Wan, and Zhao (2015) find that career concerns induce conservative bias. Employing cheap talk models, a stream of the literature shows that board's conservatism is an impediment to communication and a lower congruence of interests diminishes the credibility of precise messages. Hence, the equilibrium cheap talk is less informative (Chakraborty and Yilmaz, 2017; Harris and Raviv, 2008; Baldenius, Melumad, and Meng, 2014; Baldenius, Meng and Qiu 2019). Related, Malenko (2014) studies how conformity pressure affects verifiable disclosures by board members and Gregor (2020) studies how board friendliness affects misreporting by CEOs. Our paper also contributes to the literature on optimal board composition. A common topic in this literature is the existence of informational frictions that give rise to a trade-off between insiders' opportunism and outsiders' ignorance. Adams and Ferreira (2007) attribute this trade-off to the existence of monitoring and advisory tasks of the board—this is absent from our model because the CEO's report is unrestricted (and can be perfectly informative). Levit and Malenko (2015) study how directors' reputational concerns affect board structure. Levit (2012) studies the optimal board structure with an emphasis on directors' expertise.[9]

2 Model setup

The model entails a CEO, a board of directors, and shareholders of a firm.[10] The CEO discovers an investment opportunity (project) of an unknown value v_θ, where $\theta \in \Theta \equiv \{1, 2, 3\}$ is a random variable representing the project type with $v_1 < v_2 < 0 < v_3$.[11] The CEO proposes the project to the board. The board approves ($a = 1$) or rejects ($a = 0$) the project. We assume that, when indifferent, the board approves the project. The payoff of the shareholders is given by

$$\pi(a, \theta) = av_\theta,$$

[9]Prior literature has studied additional aspects. Baldenius, Meng and Qiu (2020) study the effect of board commitment to a decision rule on the its communication with the CEO. In a model with CEO's disclosure of private information, Chen, Guay and Lambert (2020) show that boards with greater independence also have higher expertise. Banerjee and Szydlowski (2020) study the effects of VCs friendliness. Drymiotes (2007) shows that friendly boards can be more effective at monitoring. Drymiotes (2009) studies managers' ability to influence their performance evaluation. Several studies have analyzed issues related to CEO turnover (Laux, 2014; Aghamola and Hashimoto, 2020; Meng, 2020). Qiu (2020) finds that boards pursue a "quiet life."

[10]In our model, the board of directors and the shareholders are collective entities making decisions that maximize the aggregate preferences of its members. The individual preferences of the members and the various ways a collective decision can be made are beyond the scope of this study. Our results extend to a setting in which collective decisions are based on a majority rule—in such case the decision is determined by the preferences of the median member.

[11]To facilitate the exposition and intuition provision, our model considers a state with three possible realizations but our results extend to a setting with multiple state realizations.

which implies that the shareholders receive a strictly positive payoff only if the project yielding the highest value, v_3, is approved (i.e., only if $\theta = 3$ and $a = 1$). The payoff of the CEO is given by

$$w(a, \theta) = a(v_\theta + \beta),$$

where β is a private "empire-building" benefit $\beta \in (-v_2, -v_1)$, i.e., $v_1 + \beta < 0 < v_2 + \beta$.[12] This assumption implies that the CEO receives a positive payoff not only if the project of highest value, v_3, is approved but also if the one of intermediate value, v_2, is approved (i.e., when $\theta \in \{2, 3\}$ and $a = 1$). Put differently, the CEO is biased in favor of adopting the intermediate project.

The board's payoff is given by

$$u(a, \theta) = a(v_\theta - \gamma),$$

where $\gamma \in (v_2, v_3)$ determines the degree of misalignment of interests between the board, the shareholders and the CEO.[13,14,15] The smaller γ, the more aligned are the preferences

[12] The CEO's private benefit, β, and the sensitivity of her objective to the firm value are normalized. In the absence of normalization, the CEO's payoff could be expressed as $w_0(a, \theta) = xav_\theta + a\beta_0$ where $x \in (0, 1)$ is the sensitivity of the CEO's payoff to the firm value and β_0 represents the non-normalized private benefit of undertaking the project. It is immediate that the objective with a non-normalized benefit is equivalent to the objective with a normalized benefit because $w_0(a, \theta) \propto w(a, \theta)$ when $\beta \equiv \beta_0/x$.

[13] Note that γ can assume both positive and negative values. The board receives a private approval benefit when $\gamma < 0$ and incurs an approval cost when $\gamma > 0$. Approval benefits could stem from close relationship with the CEO or private perks from the project. An example of approval cost could be a situation where directors own shares in or serve on the board of a competing company. Alternatively, directors at early stages of their careers who are concerned about being perceived as overly agreeable and thus less competent could also face approval costs. Lastly, an approval cost could also be due to conservative bias: "[t]he board has taken a risk-averse view ... a very good reason for boards to focus on risk was to avoid the stigma of becoming high-profile failures" (Deloitte, 2015). See also Jiang, Wan, and Zhao (2015). If $\gamma = 0$, the preferences of the board and the shareholders perfectly align.

[14] If $\gamma > v_3$, the board always rejects the project, and the information conveyed by the CEO is irrelevant. If $\gamma \leq v_1$, the board approves all projects. Hence, a board with $\gamma \notin (v_1, v_3]$ can never be optimal for the shareholders. Furthermore, a board with $\gamma \in (v_1, v_2)$ is also not optimal for the shareholders because the board, like the CEO, is willing to implement projects of intermediate value. Then, in the absence of the conflict of interest, the CEO chooses a perfectly informative report and the board approves all intermediate (value-destroying) projects.

[15] As in the case of the CEO (see footnote 12 for details), the board's approval cost/benefit from approving the project, γ, and its sensitivity to the firm value are also normalized.

of the board with those of the CEO. We say that the board is more "friendly" to the CEO. A friendly board could be composed of insider directors (or directors from the CEO's social circle) who have an established relationship with the CEO or directors with private benefits from project approval. The larger γ, the more "unfriendly" the board is to the CEO. An unfriendly board could consist of outside directors owning shares in a competing company and directors who are more conservative in their approval decisions. Occasionally, we refer to γ as the board's bias that determines the severity of the conflict of interest between the CEO and the board.

The prior probability distribution of the random variable θ is common knowledge. Formally, the prior belief is given by $\boldsymbol{\mu}^o \equiv (\mu_1^o, \mu_2^o, \mu_3^o) \in \Delta(\Theta)$, where $\mu_\theta^o \equiv \Pr(\theta)$, $\theta \in \Theta$. The CEO chooses the properties of a verifiable public report R about the project value from a finite set $\mathbb{S}$ of possible realizations that cannot be misreported. Our assumption that the CEO has control over the properties of the report reflects the fact that CEOs choose the type and the precision of the information that they collect.[16] Once the CEO collects information with specific precision (e.g., by soliciting an advice from a specific advisor or purchasing a specific type of data collection software), she cannot change the precision of the information. That is, the CEO commits to the properties of the information included in the report in advance. The assumption that CEOs cannot misreport the collected verifiable information can be rationalized as follows. Because projects requiring boards' approval are usually of significant importance to the firm (Useem, 2006), CEOs need to collect and include in the report convincing data supported by evidence (e.g., expert advice). Such evidence is available within the company, can be verified and thus it

[16]Even in companies with guidelines about the information that CEOs have to provide to the board (e.g., profitability analysis of the project), CEOs have discretion over the specifics and the precision of the collected information. For example, a CEO can decide to collect and include in the report very precise evidence that the project is not value-destroying or that the project is value-enhancing. The collection could involve purchase and analysis of data or receipt of advice from an expert.

is hard for the CEO to conceal or misrepresent it.[17,18] In the main part of the paper, we focus solely on the commitment to the precision of verifiable evidence. In Section 4, we extend our analysis by allowing the CEO to privately observe additional soft information that she can misreport.

For our analysis, it is useful to denote $\phi^r \equiv \Pr(R = r)$ the probability that the realized report is r. A report realization $r \in \mathbb{S}$ induces an interim belief $\boldsymbol{\mu}^r \equiv (\mu_1^r, \mu_2^r, \mu_3^r)$, where $\mu_\theta^r \equiv \Pr(\theta \mid r)$, $\theta \in \Theta$. To avoid clutter, we occasionally suppress the superscript r and simply use $\boldsymbol{\mu}$. Only when the realized report r becomes relevant for the analysis, we use $\boldsymbol{\mu}^r$. Because every report realization r is associated with a specific belief $\boldsymbol{\mu}^r$, the distribution $\boldsymbol{\phi} = (\phi^r)_{r \in \mathbb{S}}$ also describes the (probability) distribution of the interim beliefs.[19] The CEO can select any $\boldsymbol{\phi}$ with a finite number of report realizations as long as it satisfies the martingale (Bayes-plausibility) property, $\mathbb{E}_{\boldsymbol{\phi}}[\boldsymbol{\mu}] = \boldsymbol{\mu}^o$. The distribution $\boldsymbol{\phi}$ and the posterior beliefs $\boldsymbol{\mu}$ jointly determine the precision of the report.[20]

After observing the realization r and forming interim belief $\boldsymbol{\mu}^r$, the board may obtain an additional signal T from a set $\mathbb{T}$ of finite number of possible realizations.[21] Examples of such learning include searching for industry economic projections and reading financial statements. We denote $\tau^t \equiv \Pr(T = t)$ the probability that the additional signal realization is t. The properties of the additional signal are characterized by $\boldsymbol{\tau} = (\tau^t)_{t \in \mathbb{T}}$.

[17] The assumption that the CEO cannot conceal the report is not critical for our results. As long as the board is aware that the CEO attempted to collect information (e.g., because the CEO purchased a data analysis software), the information contained in a withheld report will unravel. Analysis is available upon request.

[18] Data misrepresentation can also be associated with prohibitively harsh legal consequences. For example, the former CEOs of Kmart and Kentucky aluminum company faced significant legal charges for providing misleading information to their boards of directors (Peterson, 2003; Associated Press, 2020).

[19] Specifically, ϕ^r is also the probability that the interim belief is $\boldsymbol{\mu}^r$ and so there is a one-to-one mapping between the distribution of the report and the distribution of the beliefs that the report induces.

[20] An alternative way of describing the precision of the system is by considering the probability of a report realization conditional on the state θ. This, however, can easily be derived from the distribution $\boldsymbol{\phi}$ and the posterior beliefs $\boldsymbol{\mu}$ using Bayes rule: $\Pr(r \mid \theta) = \frac{\Pr(\theta|r)\Pr(s)}{\Pr(\theta)} = \frac{\mu_\theta^r \phi^r}{\mu_\theta^o}$ for all $\theta \in \Theta$ and $r \in \mathbb{S}$. A perfectly informative report is one in which every state is always mapped to only one of the report realizations.

[21] For our analysis, it is irrelevant whether the additional signal is private or not.

For every realization $t \in \mathbb{T}$, there are corresponding final beliefs $\mu^t \equiv (\mu_1^t, \mu_2^t, \mu_3^t)$, where $\mu_\theta^t \equiv \Pr(\theta \mid r, t)$, $\theta \in \Theta$.[22] The martingale property, $\mathbb{E}_\tau[\mu^t] = \mu$, must hold.

Learning additional information is costly to the board. This assumption could reflect the fact that directors have limited time and attention.[23] Thus, learning additional information, beyond that contained in the report, is personally costly. Additionally, the costly nature of the signal T could be due to the board's disadvantage in access to firm-relevant information, relative to the CEO. In line with the rational inattention literature, we assume that the board's cost of obtaining the additional signal is proportional to the reduction of the (expected) Shannon entropy (Sims, 2003; Matějka and McKay, 2015) so that learning a more informative signal is costlier for the board.[24] Formally, the Shannon entropy of the $|\Theta|$-dimensional interim belief μ (for given r and ϕ before observing t) is given by: $H(\mu) = -\sum_{\theta\in\Theta} \mu_\theta \ln \mu_\theta$, where $0 \ln 0 = 0$ holds by convention. The total entropy-based cost of a signal distributed by τ over a support $\mathbb{T}$ is $\sum_{t\in\mathbb{T}} \tau^t H(\mu^t)$. Then, the board's personal (entropy-based) cost of the signal T is linear in the reduction of the expected entropy, $J(\mu, \tau) = \lambda \left\{H(\mu) - \sum_{t\in\mathbb{T}} \tau^t H(\mu^t)\right\}$, where $\lambda \geq 0$ is the marginal cost of reducing the entropy.[25] The lower λ, the easier it is for the board to learn information. One possible interpretation could be that busy boards (i.e., boards consisting of directors with multiple positions) have higher marginal cost of learning. Alternatively, the cost of learning could be related to the industry maturity in a sense that more information sources are available in more mature industries.

When nominating the board, the shareholders optimally choose its preference align-

[22] The finite belief μ^t is again a function of the realized report r. We suppress it to avoid clutter.

[23] It is not uncommon for directors to have other occupations or serve on several boards.

[24] The results extend to any specification with a posterior-separable cost function (Matysková, 2018).

[25] To fix ideas, consider an example with intermediate beliefs $\mu^r = (1/3, 1/3, 1/3)$ so that the Shannon entropy is $H(\mu^r) = -(1/3 \ln 1/3 + 1/3 \ln 1/3 + 1/3 \ln 1/3) = \ln 3$. A perfectly informative additional signal eliminates all uncertainty and thereby results in Shannon entropy of zero. Thus, the board's cost of acquiring such signal is $J(\mu, \tau) = \lambda(\ln 3 - 0) = \lambda \ln 3$. Alternatively, a completely uninformative private signal does not reduce any uncertainty and so the reduction in (expected) Shannon entropy is zero and the cost incurred by the board is $J(\mu, \tau) = \lambda(\ln 3 - \ln 3) = 0$. Any imperfectly informative additional signal will be associated with a cost $J(\mu, \tau) \in (0, \lambda \ln 3)$.

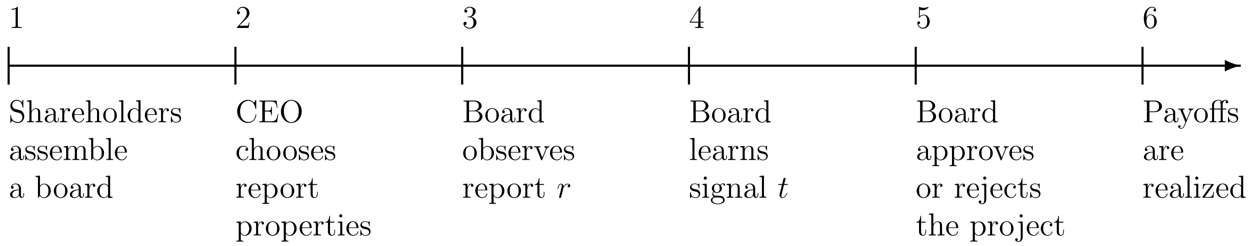

Figure 1: Timeline of the events

ment parameter γ (e.g., by deciding how many insiders, outsiders and directors with private approval costs or benefits to include).[26] For the directors to participate (i.e., accept the position as members of the board), their individual rationality constraint has to be satisfied. As participation in the board requires time and attention that can be productively employed elsewhere, we assume the board's reservation utility is strictly positive but not prohibitively large, $\underline{U} \in (0, \mu_3^o(v_3 - v_2))$.[27]

Figure 1 presents the timeline of the events. At date 1, the shareholders assemble the board of directors characterized by γ. If the board's reservation utility is not met, the board rejects participation and the game ends. Otherwise, the game proceeds. At date 2, the CEO chooses the properties of a report of the project value. At date 3, the board observes the report. At date 4, the board decides how much additional information to learn by choosing the properties of an additional signal. At date 5, the additional signal is realized and the board approves or rejects the project. At date 6, the payoffs are realized.

[26] Accounting for CEO's influence over the nomination of some directors requires an analysis of the board's collective decision making process. Both of these frictions are beyond the scope of this study.

[27] As we elaborate in detail later, a board with $\gamma \in (v_2, v_3)$ does not participate if $\underline{U} \geq \mu_3^o(v_3 - v_2)$.

3 Analysis

3.1 First-best benchmark

To see how the board's reservation utility constrains the shareholder's nomination choices, we first briefly discuss the optimal board's bias γ in the absence of informational frictions. If, at date 3, the board is perfectly informed about the project value, it only approves the value-enhancing project, $\theta = 3$. Thus, from the shareholders' point of view any $\gamma \in (v_2, v_3)$ is optimal as long as the board agrees to participate.[28] With anticipated perfect information the board's participation constraint can be presented as

$$\gamma \leq \overline{\gamma} \equiv v_3 - \frac{\underline{U}}{\mu_3^o} \in (v_3, v_2).$$

Thus, in the first-best benchmark, the shareholders choose any bias $\gamma \in (v_2, \overline{\gamma}]$. An interesting question is whether the first-best board incurs an approval benefit ($\gamma < 0$), i.e., is biased in the same direction as the CEO. This is always the case if $\overline{\gamma} < 0$, which holds when the board's reservation utility is sufficiently high, $\overline{U} \geq \mu_3^o v_3$.[29] Put differently, in industries with very good outside opportunities, boards are inevitably CEO-friendly because only directors with approval benefits are willing to participate.

3.2 Board's project approval and learning

We solve the model by backward induction. At date 5, for given report $r \in \mathbb{S}$ and additional signal $t \in \mathbb{T}$, the board approves the project if approval yields a higher expected payoff than rejection:

$$\mathbb{E}\left[u(1,\theta) \mid r,t\right] \geq \mathbb{E}\left[u(0,\theta) \mid r,t\right] = 0. \tag{1}$$

[28] The indifference between several levels of board's bias is driven by the binary nature of the project approval decision.

[29] This is feasible as $\mu_3^o v_3 \in (0, \mu_3(v_3 - v_2))$.

Because the board is only moderately biased, the approval decision depends on the report and the signal.[30] The information from the report and the signal is summarized in the final belief μ^t, so we let $a(\mu^t)$ denote the board's approval decision for the final belief μ^t. Formally, $a(\mu^t) = \mathbb{1}\{\sum_{\theta\in\Theta} \mu_\theta^t v_\theta - \gamma \geq 0\}$, where $\mathbb{1}$ is an indicator function.

At date 4, after observing report $r \in \mathbb{S}$ and anticipating its approval strategy, the board decides whether and how much additional information to learn. Technically speaking, the board's problem is to find a distribution τ that maximizes the board's conditionally expected payoff net of personal learning (entropy-based) costs,

$$\mathbb{E}\left[u(a(\mu^t), \theta) \mid r, \tau\right] - J(\mu, \tau),$$

subject to the martingale property. For every report realization r, and corresponding intermediate belief μ^r, there is an optimal corresponding signal distribution $\tau(\mu^r)$ of the additional signal learned by the board.

As we show below, of specific interest for our analysis are the report realizations (and the associated beliefs) for which the board chooses *not to learn* any additional information (or, to put differently, learns an uninformative signal). To study these report realizations we use a graphical argument. Because the state can assume three possible values, the board's beliefs are elements of a two-dimensional simplex $\Delta(\Theta)$. We denote its corners as $\{A, B, C\}$ with

$$\mu^A = (1, 0, 0); \quad \mu^B = (0, 1, 0); \quad \mu^C = (0, 0, 1).$$

Upon observing $r = A$, the board is certain that the project value is v_1. Similarly, after $r = B$ ($r = C$) the board believes that the project value is v_2 (v_3). For any possible report, the board's interim belief is within the two dimensional simplex and thus is a convex combination of the corner beliefs. We label the simplex $\Delta^{ABC} \equiv \Delta(\Theta)$. In addition, let $NL^a \subset \Delta^{ABC}$ represent the set of interim beliefs for which the board chooses not to learn

[30] If $\gamma \notin (v_1, v_3)$ the board either rejects or approves the project, regardless of r and t.

any additional information (and thus incurs zero learning entropy-based costs) and takes an action a. We refer to NL^a as the "nonlearning region of action a."

Lemma 1. *There exist two non-empty nonlearning regions, $NL^1 \subset \Delta^{ABC}$ and $NL^0 \subset \Delta^{ABC}$, such that:*

(i) If $\mu^r \in NL^1$, the board does not learn additional information and approves the project. For any properties of the report, the nonlearning region NL^1 of the board shrinks in the board's bias γ and expands in the learning cost parameter λ.

(ii) If $\mu^r \in NL^0$, the board does not learn additional information and rejects the project. For any properties of the report, the nonlearning region NL^0 expands in the board's bias γ and the learning cost parameter λ.

We first describe the non-learning region of approval, NL^1. Our arguments are graphically illustrated in Figure 2. In the proximity of point C, the board is fairly certain that the project yields the highest value, v_3, and therefore approves the project without learning additional information.[31] Thus, NL^1 has to contain the neighborhood of the point C. To find the precise boundaries of this neighborhood, we follow Caplin et al. (2019) and search for the interim beliefs in Δ^{ABC} at which the optimal signal T converges to a degenerate signal with a single realization in which the board approves the project outright without learning. (Beyond the boundaries of the non-learning region, the board does not approve the project without additional information but may approve it if it were to learn additional favorable information.) One such point is $D \in \Delta(\Theta)$ at which the proposed project yields the lowest value, v_1, with zero probability and the board's interim beliefs satisfy

$$\mu^D = \left(0, 1 - \mu_3^D, \mu_3^D\right).$$

The conditional probability that the project value is high, $\mu_3^D \in (0,1)$, is derived in the

[31]Because $\gamma \in (v_2, v_3)$, the only project approved immediately when the type is certain is $\theta = 3$.

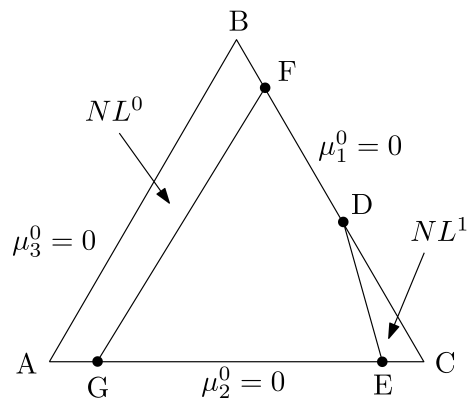

Figure 2: Interim beliefs and nonlearning regions in the simplex $\Delta(\Theta)$

proof to Lemma 1 in the Appendix. Because $\mu_1^D = 0$, point D is graphically positioned on the line BC. Larger μ_3^D implies that the project is more likely to have high, rather than intermediate, value. Graphically, this shifts point D away from point B in the direction of point C.

Similarly, we can find a point $E \in \Delta(\Theta)$, at which the project yields the intermediate value, v_2, with zero probability ($\mu_2^E = 0$, which implies that the point is positioned on the AC-line), the optimal additional signal T converges to a degenerate signal, and the board approves the project without learning:

$$\mu^E = \left(1 - \mu_3^E, 0, \mu_3^E\right).$$

The conditional probability that the project yields high value, $\mu_3^E \in (0,1)$, is again formally derived in the Appendix. Larger μ_3^E implies that the project has more likely high value and graphically shifts point E further away from point A towards point C.

It is easy to see that the nonlearning region when the board approves the project, NL^1, coincides with the area CDE graphically illustrated in Figure 2.[32] The magnitudes of μ_3^D and μ_3^E perfectly characterize the size of the nonlearning region NL^1. Notably, μ_3^D

[32]For $\gamma \in (v_2, v_3)$, points D and E never coincide with point C, hence the nonlearning region NL^1 is non-empty.

and μ_3^E are *independent* of the properties ϕ but crucially depend on the board's bias γ and the learning cost λ.

From a technical standpoint, because smaller λ and larger γ increase μ_3^E and μ_3^D, the outright approved project has a higher expected value. Thus, when the nonlearning region NL^1 shrinks (as points D and E shift closer to the point C), the expected value goes up. The intuition is as follows. All else being equal, boards with more severe conflict of interest with the CEO (high γ) are less inclined to approve the empire-building projects that the CEO proposes. Furthermore, if it is easier for a board to learn additional information (low λ), there will be fewer instances in which the board will be willing to forego learning.

We next consider the nonlearning region of project rejection, NL^0. Because $\gamma \in (v_2, v_3)$, the board rejects projects that yield the lowest and intermediate values with certainty. Hence, NL^0 includes the line AB.[33] In the neighborhood of AB, the project yields low or intermediate value with such a high probability that the board rejects the project outright without learning additional information. We find the points $F \in \Delta(\Theta)$ and $G \in \Delta(\Theta)$ at which the optimal additional signal T converges to an uninformative signal at which the project is rejected outright:

$$\mu^F = \left(0, 1 - \mu_3^F, \mu_3^F\right) \text{ and } \mu^G = \left(1 - \mu_3^G, 0, \mu_3^G\right),$$

with both $\mu_3^F \in (0,1)$ and $\mu_3^G \in (0,1)$ again formally provided in the Appendix. Graphically, the nonlearning region NL^0 coincides with the area $ABFG$.[34] The magnitudes of μ_3^F and μ_3^G perfectly characterize the size of the nonlearning region of rejection. As before, both μ_3^F and μ_3^G are independent of the properties ϕ but crucially depend on the board's bias γ and the learning cost parameter λ.

The intuition for the comparative statics with respect to λ is the same as the one

[33] Any point on this line implies that the project has either low, or intermediate value.

[34] For $\gamma \in (v_2, v_3)$, points G and F never coincide with points A and B, hence the nonlearning region NL^0 is non-empty.

related to part (i) of Lemma 1: The costlier it is to learn additional information, the more often the board will choose to forego learning. However, now the nonlearning region is *expanding* in the board's bias γ. Technically, this is because larger γ decreases μ_3^F and μ_3^G, which makes it less likely that the project has high value, and shifts points F and G closer to points B and A, respectively. Intuitively, more confrontational boards (high γ) are *more* inclined to reject the project without learning additional information.

3.3 Report properties

At date 2, the CEO chooses the properties of the report. Technically speaking, the CEO's problem is to find a distribution ϕ that maximizes her expected payoff subject to the martingale property. Kamenica and Gentzkow (2011) show that the optimal distribution can be obtained by concavification of the CEO's expected payoff function conditional on the observed report, i.e., the CEO's payoff function over the interim beliefs invoked by the report. Characterizing the CEO's payoff function in our setting, in the presence of the board's learning option, is nontrivial. Nevertheless, the solution of the problem can be presented in an elegant way.

Similar to the findings of recent literature on generalized Bayesian persuasion of a rationally inattentive receiver (Matysková, 2018; Caplin et al., 2019), the CEO's problem in our model can be simplified to (i) ensuring that the interim beliefs are such that the board does not learn additional information and (ii) finding the distribution of interim beliefs that maximizes the probability of board's approval of intermediate and high-type projects. To understand why the CEO wants to avoid learning by the board note that, while at date 2 the CEO is uncertain which report will be realized, she controls the set of possible supported realizations and the probability with which each one of these realizations is generated. Because the CEO is unconstrained in her choice, any information that the board can possibly learn can also be directly provided by the CEO. However, if the board learns, the CEO faces an additional uncertainty related to the signal t which cannot benefit her. Therefore, the CEO prefers that the board does not learn. This happens whenever the report invokes interim beliefs in the board's nonlearning regions. In such case, the CEO controls the set of final posterior beliefs used as a basis for the board's approval decision, as well as the ex ante distribution of these beliefs.[35,36]

Because the CEO needs to make sure the board does not learn additional information, the candidate report realizations are in the nonlearning regions NL^0 and NL^1. Our problem can be further simplified by focusing only on the extreme points of these nonlearning regions (Matysková, 2018). An extreme point of NL^a is a point with belief $\mu \in NL^a$, which does not lie on any open line segment joining two points of NL^a. Focusing on the extreme points means considering only report realizations that imply the project is one of two types, i.e., realizations that rule out one state with certainty.[37]

In our case, the set of extreme points of NL^0 is $\{A, B, F, G\}$, and the set of extreme points of NL^1 is $\{C, D, E\}$. Thus, the set of candidate realizations that could be supported by the optimal report is

$$\widehat{\mathbb{S}} = \{A, B, C, D, E, F, G\}.$$

The properties of the report on the reduced set $\widehat{\mathbb{S}}$ are represented by a distribution

$$\Phi = \{\phi^A, \phi^B, \phi^C, \phi^D, \phi^E, \phi^F, \phi^G\},$$

where, as before, $\phi^r \equiv \Pr(R = r)$, $r \in \widehat{\mathbb{S}}$. The optimal report is characterized by the op-

[35]Technically speaking, as discussed below in relation to panel (b) of Figure 4, the concave closure is above the CEO's expected payoff function conditional on the observed report for any posterior belief associated with board's learning.

[36]Caplin et al. (2019) demonstrate that the board's optimal signal mixes between final beliefs that are borders of non-learning regions. The CEO can invoke them directly by the choice of the reporting technology. In addition, the CEO may invoke other than border points from the nonlearning regions. We elaborate on this point in relation to Figure 4.

[37]Technically, because the CEO's expected payoff function is linear within non-learning regions, the concave closure cannot involve only points in the interior, and any value in the interior (implying that the project can be of any type) can be achieved through a lottery on the extreme points.

timal distribution $\widetilde{\boldsymbol{\phi}}$ over the reduced set $\widehat{\mathbb{S}}$, where $\widetilde{\boldsymbol{\phi}} = \left(\widetilde{\phi}^r\right)_{r\in\widehat{\mathbb{S}}} = \arg\max_{\boldsymbol{\phi}\in\Delta(\widehat{\mathbb{S}})} \mathbb{E}[w(\cdot)]$, subject to the martingale property, $\boldsymbol{\mu}^o = \sum_{r\in\widehat{\mathbb{S}}} \phi^r \boldsymbol{\mu}^r$. Lemma 2 below reduces the dimensionality of the problem by establishing that the optimal report properties do not support realizations $\{E, F, G\}$.

Lemma 2. *The optimal report is not supported by interim beliefs that lie on the points $\{E, F, G\}$, i.e., $\widetilde{\phi}^E = \widetilde{\phi}^F = \widetilde{\phi}^G = 0$.*

To understand the intuition behind Lemma 2, consider a fully informative report with three realizations: $r \in \{A, B, C\}$. Such report is not optimal for the CEO because she favors intermediate value projects and those are not approved with a fully informative report. Adding the realization $r = D$ increases the probability of approval of the intermediate value projects because this report is sent for both high and intermediate value projects. In contrast, adding the realizations $r \in \{F, G\}$ does not increase the approval probability as these reports are associated with project rejection. Finally, adding the realization $r = E$ is not helpful either—this report is sent only for low and high value projects, and there is no conflict of interest between the CEO and the board over these projects.

It is worthwhile to point that, while the board's learning option is never exercised in equilibrium, it influences the CEO's choice of report properties. This is because the optimal properties in our model encourage the board not only to approve the project but also not to learn any additional information. In particular, because $\boldsymbol{\mu}^D$ is optimally set to render learning by the board unnecessary and ensure approval, the board strictly prefers approval at point D, i.e., $\mathbb{E}[u(a, \theta) \mid r = D] = \sum_{\theta\in\Theta} \mu_\theta^D v_\theta - \gamma > 0$. Our observation is in contrast with the properties of reports in persuasion models without receiver's learning where an optimal report of a biased persuader makes the decision-maker exactly indifferent between approval and rejection.

We graphically illustrate this point in Figure 3. In the figure, the approval region of

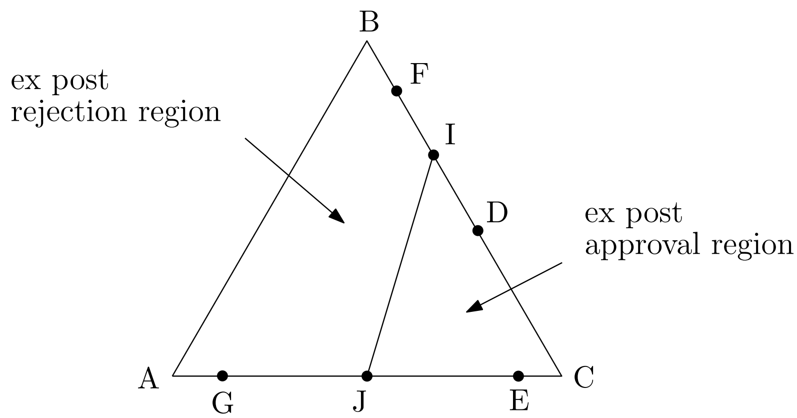

Figure 3: Ex post approval and rejection regions in the simplex $\Delta(\Theta)$

final beliefs for which the board (weakly) prefers approval is CJI. The rejection region of final beliefs for which the board (strictly) prefers rejection is $ABIJ$ and it is a complement to the approval region.[38] The relative size of the non-learning regions within the approval and rejection regions depends on the board's learning cost. When learning is prohibitively costly ($\lambda \to \infty$), points D and F converge to I, and points E and G converge to H; we obtain a classic persuasion problem (Kamenica and Gentzkow, 2011). In the absence of learning costs ($\lambda = 0$), points D and E converge to C, point F converges to B, and point G converges to A; the board then anyway learns the project type so the optimal report is also fully informative. For any interior level of learning cost, $\lambda \in (0, \infty)$, the nonlearning region does not coincide with the respective action region. Hence, when a non-learning board selects an action (approval or rejection), the alternative action is perceived as strictly inferior to the selected action.

To gain further intuition for our results, it is instructive to consider the expected payoffs of the CEO and the board as a function of the beliefs on the BC-line.[39] The solid line in panel (a) of Figure 4 represents the board's conditionally expected payoff

[38] Formally, the approval region of final beliefs is $AR^1 \equiv \{\mu \in \Delta(\Theta) : \sum_{\theta\in\Theta} \mu_\theta v_\theta - \gamma \geq 0\}$, and the rejection region is $AR^0 \equiv \Delta(\Theta) \setminus AR^1$.

[39] Because in our model the state can assume three values, this is not a precise representation. Yet, it provides helpful intuition for our results.

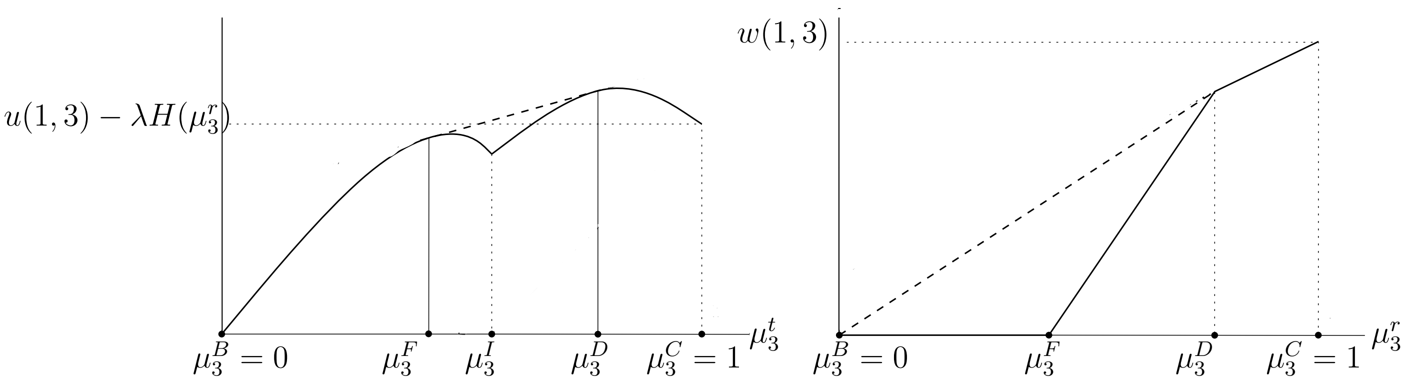

(a) Board's payoff as function of μ_3^t (b) CEO's payoff as function of μ_3^r

Figure 4: Expected payoff of the board and the CEO as functions of beliefs
Here, $w(1,3) = v_3 + \beta$ is the ex post CEO's payoff when a high type project is approved. Similarly, $u(1,3) - \lambda H(\mu_3^r) = v_3 - \gamma - \lambda(H(\mu_3^r) - H(1))$ is the ex post board's payoff when the board learns that the project is high type with certainty ($\mu_3^t = \mu_3^C = 1$) and approves the project.

after observing the report r and the additional signal t. For any final belief outside the interval (μ_3^F, μ_3^D) this payoff is concave so that the board does not learn. Intuitively, the board is fairly certain that the project is either intermediate type and will anyway reject it, or fairly certain that the project is high type and will anyway accept it. Thus, for the extreme beliefs, spending time and attention on learning is not beneficial. The board strictly prefers to learn only if the final belief is in the interior interval (μ_3^F, μ_3^D). The concave closure (in dashed line) then is strictly above the board's conditionally expected payoff and the board's net payoff emerges here as a result of the concavification. Namely, because the action is binary (approve or reject the project), the board's optimal learning strategy is to induce a lottery of the final beliefs at the border of the nonlearning region of rejection, μ_3^F, and the border of the nonlearning region of approval, μ_3^D (Caplin et al., 2019).

The solid line in panel (b) of Figure 4 represents the CEO's conditionally expected payoff after preparing the report r. For any intermediate belief below μ_3^F, the board rejects the project (without learning) and so the CEO's payoff is zero. For any intermediate belief above μ_3^F, the CEO's payoff is strictly positive because the project is approved with a positive probability (the board in this region will either learn and approve with some probability, or not learn and approve outright). Depending on the magnitude of the prior belief μ_3^o, it is optimal for the CEO to prepare a report which induces a lottery of beliefs μ_3^D and μ_3^B or a lottery of beliefs μ_3^D and μ_3^C. We elaborate on these choices in further detail after formally establishing the optimal report properties in Proposition 1. For now, what matters is that it is always optimal for the CEO to choose report properties that support $r = D$.

The result in Lemma 2 implies that the optimal report can only be supported by realizations $r \in \{A, B, C, D\}$. Clearly, upon observing $r \in \{A, B, C\}$, the board is fully informed and commits no approval errors. As a consequence, the only realization supported by the optimal report that leaves the board not fully informed is $r = D$. This report implies that the project either yields high, or intermediate value. For our subsequent analysis, it is instructive to describe the point D through its likelihood ratio

$$\rho(\gamma) \equiv \frac{\Pr(\theta = 3 \mid r = D)}{\Pr(\theta = 2 \mid r = D)} = \frac{\mu_3^D}{1 - \mu_3^D}.$$

The larger $\rho(\gamma)$, the larger is the probability that the project is high type, conditional on observing $r = D$. Because, by construction, the board approves the project at point D and the only approval error occurs at point D, we label $\rho(\gamma)$ the "approval precision" of the board.

Corollary 1 (Precision of approval)**.** *The board's approval precision $\rho(\gamma)$ is decreasing in the cost of learning, λ, and increasing in the board bias, γ.*

A board with lower learning costs or stronger conflict of interest with the CEO is less likely to approve value decreasing projects. For future reference, we additionally label the likelihood ratio at the prior beliefs, $\rho^o \equiv \mu_3^o / \mu_2^o$, the "guaranteed" level of approval precision. We provide further details about ρ^o in Section 3.4.

As a next step, we will classify all possible situations into two types depending on the intensity of the conflict between the CEO and the board. Let Δ_{ABD} be the convex hull of $\{A, B, D\}$ and Δ_{ACD} be the convex hull of $\{A, C, D\}$. When the prior belief μ^o is in Δ_{ABD}, the project is more likely to be intermediate type. Because this project is value-destroying for the firm and its shareholders, we label this case "strong conflict." It is easy to see that strong conflict is equivalent to a situation in which the approval precision exceeds the guaranteed level of precision, $\rho(\gamma) > \rho^o$. Because the approval precision $\rho(\gamma)$ is continuously increasing in γ and the guaranteed approval precision ρ^o is independent of γ, we can also present the strong conflict as a situation in which $\gamma > \gamma^o$. Here, $\gamma^o \in (v_2, v_3)$ is a cutoff level satisfying $\rho^o = \rho(\gamma^o)$.[40,41] Intuitively, the larger γ, the more unfriendly is the board to the CEO and the larger is their conflict of interests. Similarly, when the prior belief μ^o is in Δ_{ACD}, the project is more likely to be high type. Because this project adds value to the firm, we label this case "weak conflict" of interests.[42] Following a similar argument, weak conflict is equivalent to $\rho(\gamma) < \rho^o \Leftrightarrow \gamma < \gamma^o$.

We are now ready to solve for the optimal properties $\widetilde{\phi}$. Proposition 1 shows that the system depends on the severity of the conflict defined above, and is not fully informative as long as the projects could be of intermediate value, $\mu_2^o > 0$.

Proposition 1 (Optimal properties of the CEO's report)**.**

(i) For a strong conflict between the board and the CEO ($\gamma > \gamma^o$), the report is characterized by the optimal distribution $\widetilde{\phi} = \phi_{ABD} \equiv \left(\mu_1^o, \mu_2^o - \frac{\mu_3^o(1-\mu_3^D)}{\mu_3^D}, 0, \frac{\mu_3^o}{\mu_3^D}, 0, 0, 0\right)$. The board approves the project after observing $r = D$ and rejects the project otherwise.

(ii) For a weak conflict between the board and the CEO ($\gamma < \gamma^o$), the report is character-

[40] To see that $\gamma^o > v_2$, we use that $\rho(v_2) = 0$ and $\rho^o > 0$. To see that $\gamma^o < v_3$, we use that $\rho(v_3) \to \infty > \rho^o$.

[41] Notably, because μ_3^D does not depend on the properties ϕ, the cutoff γ^o is also independent of ϕ.

[42] Note that $\Delta_{ABD} \cup \Delta_{ACD} = \Delta(\Theta) = \Delta_{ABC}$.

ized by the optimal distribution $\widetilde{\phi} = \phi_{ACD} \equiv \left(\mu_1^o, 0, \mu_3^o - \frac{\mu_2^o \mu_3^D}{1-\mu_3^D}, \frac{\mu_2^o}{1-\mu_3^D}, 0, 0, 0\right)$*. The board rejects the project after observing* $r = A$ *and approves the project otherwise.*

When the conflict between the board and the CEO is strong ($\gamma > \gamma^o$), only realizations $r \in \{A, B, D\}$ are supported. The project with the lowest (highest) value is always reported as $r = A$ ($r = D$). However, the intermediate type project is sometimes reported as $r = B$ and sometimes as $r = D$.[43] Because, by construction, only $r = D$ is associated with project approval, all high type and *some* intermediate type projects are approved. When the conflict is weak ($\gamma < \gamma^o$), the optimal report only supports $r \in \{A, C, D\}$. The project with the lowest (intermediate) value is always reported as $r = A$ ($r = D$). The high type project is sometimes reported as $r = C$ and sometimes as $r = D$.[44] Now the board approves the project following two reports, $r \in \{C, D\}$, and all high as well as *all* intermediate type projects are approved. The reason is that, when the conflict is weak, the prior beliefs about the project value are high, and so persuading the board to approve all intermediate projects by mixing them with some high type projects is feasible.

To gain further intuition for the difference in the optimal report properties with weak and strong conflict, it is instructive to recall our discussion related to panel (b) of Figure 4. When the conflict is strong, the prior belief μ_3^o is between μ_3^B and μ_3^D. Thus, the optimal report induces a lottery of those beliefs by supporting realizations $r \in \{B, D\}$. In contrast, when the conflict is weak, the prior belief μ_3^o is between μ_3^D and μ_3^C. The optimal report then induces a lottery of those beliefs by supporting realizations $r \in \{D, C\}$.

In summary, for a mild conflict of interest between the CEO and the board associated with low γ, the board approves all intermediate and high type projects. Once the level

[43]Using Bayes rule, $\Pr(r = A|\theta = 1) = 1$, and $\Pr(r = B|\theta = 1) = \Pr(r = D|\theta = 1) = 0$. Furthermore, $\Pr(r = A|\theta = 3) = \Pr(r = B|\theta = 3) = 0$ and $\Pr(r = D|\theta = 3) = 1$. Lastly, $\Pr(r = A|\theta = 2) = 0$, $\Pr(r = B|\theta = 2) = 1 - \frac{\mu_3^o}{\mu_2^o} \cdot \frac{1-\mu_3^D}{\mu_3^D} \in (0, 1)$ and $\Pr(r = D|\theta = 2) = \frac{\mu_3^o}{\mu_2^o} \cdot \frac{1-\mu_3^D}{\mu_3^D} \in (0, 1)$.

[44]Technically, $\Pr(r = A|\theta = 1) = 1$, and $\Pr(r = C|\theta = 1) = \Pr(r = D|\theta = 1) = 0$. Furthermore, $\Pr(r = A|\theta = 2) = \Pr(r = C|\theta = 2) = 0$ and $\Pr(r = D|\theta = 2) = 1$. Lastly, $\Pr(r = A|\theta = 3) = 0$, $\Pr(r = C|\theta = 3) = 1 - \frac{\mu_3^o}{\mu_2^o} \cdot \frac{1-\mu_3^D}{\mu_3^D} \in (0, 1)$ and $\Pr(r = D|\theta = 3) = \frac{\mu_3^o}{\mu_2^o} \cdot \frac{1-\mu_3^D}{\mu_3^D} \in (0, 1)$.

of bias reaches the critical cutoff γ^o, the board begins to reject some of the intermediate type projects. Any increase in γ beyond that point further decreases the chance that the intermediate type project is approved.

Corollary 2 (Probability of approval). *When the conflict is weak ($\gamma < \gamma^o$), the ex ante probability of project approval is* $\Pr(a = 1) = \Pr(r = D) + \Pr(r = C) = \mu_2^o + \mu_3^o$. *When the conflict is strong ($\gamma > \gamma^o$), the ex ante probability of project approval is given by* $\Pr(a = 1) = \Pr(r = D) = \frac{\mu_3^o}{\mu_3^D} \in (\mu_3^o, \mu_2^o + \mu_3^o)$ *and it is decreasing in γ.*

When the conflict of interest between the board and the CEO is more severe, the board is less likely to approve the project. Because outside directors have preferences that are less aligned with CEOs (high γ), our result in Corollary 2 may help explain the declines in investment rates following board-dependence regulations (Frydman and Hilt, 2017; Jagannathan, Jiao, and Krishnamurthy, 2020). At the same time, because boards characterized with higher γ approve fewer value-decreasing projects (see also the comparative statics in Corollary 1), we predict that board-dependence regulations are associated with higher return on corporate investments. Furthermore, because directors serving on multiple boards have more limited time and less mature industries have fewer information sources to learn from, our results in Corollaries 1 and 2 predict lower return on corporate investments in companies with busy boards and companies operating in less mature industries.

Regardless of the severity of the conflict, the optimal report properties do not distort the approval decision for the projects about which the CEO and the board agree. The lowest-value projects are always rejected, and the highest-value projects are always approved. The distortion is that, some (or all) intermediate projects are also approved. Similar to the classic binary case with a persuader who is biased in favor of project adoption, we observe no type-II errors (no rejections of high type projects) and some type-I errors (approvals of some value-destroying projects) in the board's equilibrium action

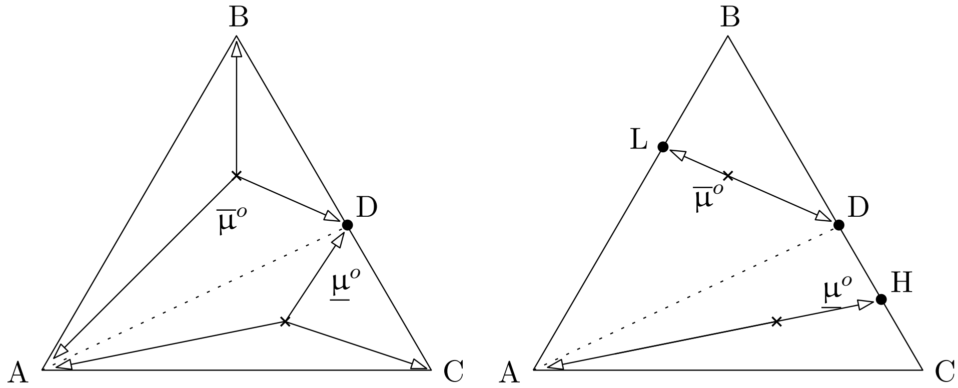

(a) 3D reports (b) 2D reports

Figure 5: Equivalent characterizations of the optimal report
A weak conflict is represented by $\underline{\mu}^o$, and a strong conflict is represented by $\overline{\mu}^o$.

strategy. The intuition for this result is that the CEO can persuade the board to approve intermediate value projects only by pooling them with the highest value projects. Thus, the CEO maximizes her expected payoff by preparing a report that only allows type-I errors but never allows type-II errors.

Because the board's action is binary (approve or reject), it is immediate that the optimal report can be simplified into one with only two realizations. Corollary 3 formally states this observation and Figure 5 graphically illustrates the two equivalent characterizations of the optimal system in a simplex $\Delta(\Theta)$.

Corollary 3 (2D optimal signal)**.**

(i) For a strong conflict ($\gamma > \gamma^o$), the optimal distribution $\widetilde{\phi}$ is equivalent to a distribution ϕ_{LD} over a subset $\{L, D\}$. The board approves the project after observing $r = D$ and rejects the project otherwise.

(ii) For a weak conflict ($\gamma < \gamma^o$), the optimal distribution $\widetilde{\phi}$ is equivalent to a distribution ϕ_{AH} over a subset $\{A, H\}$. The board approves the project after observing $r = H$ and rejects the project otherwise.

For a strong conflict, we can coarsen the optimal 3D-report realizations with interim beliefs μ^A and μ^B at points A and B leading to rejection into a single realization with an interim belief μ^L at point L on the line AB. Report $r = L$ implies that the project has either low or intermediate value with non-zero probability. For a weak conflict, we can coarsen the optimal 3D-report realizations with interim beliefs μ^C and μ^D at points C and D leading to approval into a single realization with an interim belief μ^H at point H on the line CD. Report $r = H$ implies that the project has either intermediate or high value with non-zero probability. Notice that $\mu^L \in NL^0$ and $\mu^H \in NL^1$; hence, following the coarse report L or H, the board indeed does not learn and selects the action outright.

3.4 Board nomination

3.4.1 Shareholders' expected payoff

Let $\overline{\Pi}$ be the expected payoff that the shareholders receive when the board is fully informed. A fully informed board approves the project only if $\theta = 3$, and therefore

$$\overline{\Pi} = \Pr(\theta = 3) \cdot v_3 = \mu_3^o v_3.$$

However, our discussion in the preceding section reveals that this maximum payoff is not attainable to the shareholders because the board is not fully informed and, in addition to high type projects, may approve intermediate projects (i.e., may commit type-I errors). Let $\Pi(\gamma)$ denote the expected shareholders' payoff (conditional on board's participation):

$$\begin{aligned} \Pi(\gamma) &= \Pr(\theta = 3) \cdot v_3 + \Pr(a = 1 \cap \theta = 2) \cdot v_2 \\ &= \overline{\Pi} + v_2 \cdot \frac{\mu_3^o}{\max\{\rho(\gamma), \rho^o\}} < \overline{\Pi}. \end{aligned} \tag{2}$$

Because the intermediate type project is value-destroying ($v_2 < 0$), the expected payoff of the shareholders is always lower than the maximum payoff. Furthermore, since the strong conflict exists if and only if $\rho(\gamma) > \rho^o$, this representation illustrates that the shareholder's payoff under weak conflict corresponds to the minimum payoff under the strong conflict, i.e., to the payoff associated with the approval precision ρ^o. Thus, from the shareholders perspective, we can interpret the weak conflict as a special case of the strong conflict with the minimal ("guaranteed") level of approval precision ρ^o. If the board participates, the shareholders receive at least the minimal (guaranteed) payoff $\Pi(\gamma^o)$ corresponding to the cutoff level of bias γ^o. Even if the bias is low and the conflict between the CEO and the board is weak, the outcome from the shareholders point of view is the same as if the bias was artificially increased to the cutoff level γ^o. However, to obtain a higher than the guaranteed payoff, the shareholders must increase the board's bias above the cutoff level, γ^o, thereby invoking a strong conflict.

3.4.2 Optimal board bias

By Corollary 1, the approval precision is increasing in γ, so the shareholders benefit from nominating a board that is as unfriendly as possible. However, a board with sufficiently high γ could be better off with the outside option yielding a utility of $\underline{U}$ and could prefer not to participate. Formally accounting for the board's individual rationality constraint reveals the levels of bias for which the payoff $\Pi(\gamma)$ is not implementable. In the following, we identify these prohibitively large levels of γ. As weak conflict and strong conflict involve structurally different reports, we first consider separately the cases of weak ($\gamma < \gamma^o$) and strong conflict ($\gamma \geq \gamma^o$) and then present the equilibrium board's bias which ultimately determines the level of conflict in the equilibrium.

When $\gamma < \gamma^o$, the board's individual rationality constraint can be rewritten as

$$\gamma \leq \gamma^w \equiv v_3 - \frac{\mu_2^o(v_3 - v_2) + \underline{U}}{\mu_2^o + \mu_3^o}. \tag{3}$$

Our assumption that the reservation utility is not prohibitively large, $\underline{U} < \mu_3^o(v_3 - v_2)$, implies that the maximal implementable level of board bias exceeds the value of the intermediate type project, $\gamma^w > v_2$, hence the individual rationality constraint is satisfied on a non-empty interval. For the values of bias on this interval, $\gamma \in (v_2, \min\{\gamma^o, \gamma^w\}]$, the shareholders achieve their guaranteed payoff $\Pi(\gamma^o)$ and are thereby indifferent between any level of γ from this interval.

When $\gamma \geq \gamma^o$, it is possible to rewrite the participation condition as

$$\rho(\gamma)[\mu_3^o(v_3 - \gamma) - \underline{U}] \geq \mu_3^o(\gamma - v_2) > 0. \tag{4}$$

Because the approval precision $\rho(\gamma)$ is positive and the bias γ is moderate ($\gamma \geq \gamma^o > v_2$), a necessary condition for the board to participate is that the left-hand side of (4) is positive, and therefore $\underline{U} < \mu_3^o(v_3 - \gamma)$. We can rewrite this condition as an upper bound on the implementable levels of the board bias,

$$\gamma < \gamma^s \equiv v_3 - \frac{\underline{U}}{\mu_3^o}. \tag{5}$$

This bound serves as a necessary (but not sufficient) condition for participation.[45] When the board's bias exceeds the cutoff γ^s, the board does not participate. When the bias is below the cutoff γ^s, participation is possible but is not guaranteed. In other words, the restriction $\gamma < \gamma^s$ helps to narrow down the interval of implementable biases but is not sufficient to identify the highest value of this interval. To obtain this value, we transform the participation constraint into

$$\rho(\gamma) - \psi(\gamma) \geq 0.$$

[45]Because we assume throughout that the reservation utility is not prohibitively large, $\underline{U} < \mu_3^o(v_3 - v_2)$, we obtain $\gamma^s > v_2$, hence a non-empty interval of biases $\gamma \in (v_2, \gamma^s)$ complies with this upper bound.

It represents the difference between the approval precision, $\rho(\gamma)$, and

$$\psi(\gamma) \equiv \frac{\mu_3^o(\gamma - v_2)}{\mu_3^o(v_3 - \gamma) - \underline{U}},$$

which we call the "minimum approval precision required by the board."

Lemma 3. *The difference between the true and minimum approval precisions, $\rho(\gamma) - \psi(\gamma)$, is concave in γ for $\gamma \in (v_2, \gamma^s)$, and decreasing in γ in the left neighborhood of γ^s.*

To provide intuition, the effect of γ on the board's willingness to participate depends on two countervailing forces. On the one hand, larger γ disciplines the CEO to prepare a more informative report. As a result, the board's approval precision $\rho(\gamma)$ increases, thereby affecting positively the board's willingness to participate. On the other hand, for given report informativeness, larger γ reduces the board's benefit from project approval and thereby increases the minimum approval precision required by the board, $\psi(\gamma)$. As a result, the board is less willing to participate. In essence, Lemma 3 shows that the first effect dominates for low levels of bias whereas the second effect dominates for high levels of bias.

We are now ready to show that the equilibrium board composition depends on a surprisingly simple condition: the participation decision of the cutoff type γ^o.[46]

Proposition 2 (Optimal board composition)**.** *If the board type γ^o is not willing to participate, the shareholders assemble a CEO-friendly board characterized with any $\gamma \in (v_2, \gamma^w]$, and the equilibrium conflict is weak. If the board type γ^o is willing to participate, the shareholders nominate an unfriendly board with $\gamma^* \in (\gamma^o, \gamma^s)$, where $\rho(\gamma^*) - \psi(\gamma^*) = 0$, and the equilibrium conflict is strong.*

If the board type γ^o is not willing to participate, the shareholders can induce only a weak conflict by nominating a friendly board with bias $\gamma \in (v_2, \gamma^w]$. For any bias level

[46] In the proof to Proposition 2, we demonstrate that this condition boils down to whether $\gamma^o \leq \gamma^w$ or not, and also whether $\rho(\gamma^o) - \psi(\gamma^o) \leq 0$ or not.

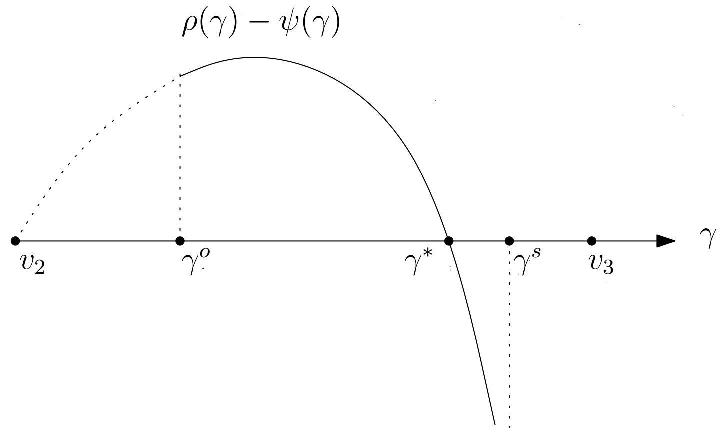

Figure 6: Optimal board bias when type γ^o is willing to participate

from this interval, the CEO prepares a report that convinces the board to approve all conflicting projects. The shareholders earn only their guaranteed payoff, $\Pi(\gamma^o)$. Corrolary 4 conducts comparative statics of the cutoff value γ^w.

Corollary 4. *The cutoff γ^w is increasing in μ_3^o and decreasing in μ_2^o and $\underline{U}$.*

The analysis implies that the shareholders assemble a more friendly board, when the board's reservation utility and frequency of intermediate type (conflicting) projects is larger, and the frequency of high type projects is smaller.

If the board type γ^o is willing to participate, both weak and strong conflicts can be induced in the process of board nomination. This case is depicted in Figure 6. All types $\gamma < \gamma^*$ are willing to participate. However, because the shareholders benefit from larger γ, they optimally assemble an unfriendly board with the maximal implementable bias $\gamma^* \in (\gamma^o, \gamma^s)$, where the participation constraint is just binding. The equilibrium conflict is strong and the CEO prepares a report that convinces the board to approve only some of the conflicting projects.[47] The shareholders earn more than the guaranteed payoff.

At this point, it is instructive to compare the optimal board's bias with the first-best benchmark. As shown in Section 3.1, the shareholders are indifferent between any bias

[47]Note that $\gamma^* > \gamma^o$, that is, when type γ^o is willing to participate the equilibrium conflict is stronger.

$\gamma \in (v_2, \overline{\gamma}]$ when the board is anticipated to observe perfect information about the project type. In contrast, Proposition 2 implies that the shareholders optimally assemble a board with bias $\gamma < \gamma^s = \overline{\gamma}$ or $\gamma < \gamma^w < \overline{\gamma}$. Thus, the presence of informational frictions restricts the implementable levels of bias below $\overline{\gamma}$. The reason is that in the second-best scenario the board is never fully informed in equilibrium which, all else equal, reduces the expected payoff from participation in the board. Ensuring the board's participation then comes at the expense of invoking a weaker conflict between the board and the CEO.

Proposition 3. *The conflict is more likely strong and the optimal board's bias γ^* increases if the prior probability μ_3^o that the project value is high increases, and the board's reservation utility $\underline{U}$, the learning cost λ, and the prior probability μ_2^o that the project value is intermediate decrease.*

The comparative statics in Proposition 3 implies that firms with less busy directors operating in more mature industries (low γ) are associated with stronger conflict between CEOs and boards. In these cases, because learning is easier and cheaper, convincing the board not to learn additional information requires a more informative report. This increases the board's approval precision $\rho(\gamma)$. At the same time, λ has no effect on the minimal level of approval precision $\psi(\gamma)$ required by the board. As a result, the board is more willing to participate and the shareholders can assemble a less friendly board that just breaks even.

Our results in Proposition 3 and Corollary 4 additionally imply that when directors have better outside opportunities, the shareholders assemble a more friendly board and the equilibrium conflict between the CEO and the board is less likely to be strong. Formally, there always exists an interior critical level $\widehat{U} \in (0, \mu_3^o(v_3 - v_2))$ such that, if the board's reservation utility exceeds this threshold, the shareholders assemble a board that has a weak conflict with the CEO. Otherwise, the shareholders assemble a board that has a strong conflict with the CEO. Hence, an exogenous improvement in outside opportunities can force firms to assemble more friendly boards that have weaker conflicts of interests with CEOs.

4 Discussion of robustness with private information

Over time, CEOs can privately learn value-relevant information and communicate it to their boards. It is immediate that ex-post nonverifiable (cheap talk) communication does no better than ex ante commitment to a report from the CEO's point of view. The reason is that, because the CEO is unconstrained in her choice of a report, she can always commit to a report that induces the same beliefs as those induced by the cheap talk communication. Thus, by revealed preference, when the CEO can choose either cheap talk communication or commit ex ante to the properties of a report, it has to be the case that the CEO is at least weakly better off with ex ante commitment.[48] It is unclear, however, whether the prospect of subsequent receipt and communication of private information at a later date affects the ex ante choice of report properties and the optimal board composition (i.e., when the two channels of communication are simultaneously available to the CEO). To study this question, we extend our model by allowing the CEO to privately observe the project value and send a costless nonverifiable (cheap talk) message to the board at date 3.

As the board's action is binary, the CEO's communication at date 3 is effectively also binary, involving mixing over recommendations (to approve or not approve). We let $(m_\theta(\mu))_{\theta\in\Theta}$ represent the cheap talk communication strategy at date 3—this strategy is conditional on the realization of the report (or, equivalently, on the corresponding interim belief μ). The cheap talk communication strategy induces a lottery of interim beliefs

[48]This observation may not hold if choosing the report properties is costly. When the conflict is weak, cheap talk communication of private information is credible and can perfectly replicate the distribution of beliefs induced by the optimal report. As a result, if the report is associated with a small cost, the CEO prefers to rely on cheap talk communication instead. Analysis is available upon request.

(μ_+, μ_-), where the former belief is for a positive recommendation, and the latter is for a negative recommendation. The induced distribution is:[49]

$$(\phi_+, \phi_-) \equiv \left(\sum_{\theta \in \Theta} m_\theta(\mu)\mu_\theta, \sum_{\theta \in \Theta} (1 - m_\theta(\mu))\mu_\theta \right).$$

In Section 3.3, we identified the structure of the report that the CEO prepares when she does not expect to receive subsequent private information. The CEO's problem becomes more complex when she receives and communicates private information at date 3. Then, the CEO's expected payoff (conditional on the signal realization) reflects the cheap talk communication. While we could solve the concavification problem by mechanically updating the (conditionally) expected payoff, there is a more elegant way to address the problem. It rests on the fact that cheap talk communication is a nested signal to the report. When selecting report properties at date 2, the CEO effectively chooses a compound lottery (distribution) of beliefs at date 4 instead of a simple lottery (distribution) of beliefs as it was the case in the absence of cheap talk. For the CEO, it is however irrelevant whether a particular lottery is generated as a simple lottery or as a reduced lottery that represents a compound lottery. The only relevant difference for the CEO's persuasion problem is that the cheap talk option at date 3 may affect the set of feasible distributions.

In particular, in the original game in Section 3.3, all distributions that satisfy the martingale property are feasible. In the extended game studied in this section, some of these distributions may not be feasible because the CEO may deviate at some report realizations by giving an informative recommendation. Hence, it is convenient to think of the cheap talk option as a collection of constraints on the feasibility of reports in the

[49] All these variables are functions of the belief μ, which we suppress to save on notation. The beliefs satisfy Bayes rule and the martingale property $\mu = \phi_+\mu_+ + \phi_-\mu_-$. In this notation, we disregard the minor issue that with any type-biased pure cheap talk strategy (i.e., $m_1 = m_2 = m_3$, and $m_\theta \in \{0, 1\}$), one of the beliefs μ_+ or μ_- is an out-of-equilibrium belief that is realized with zero probability and hence is not defined by Bayes rule.

original persuasion problem. Then, the optimal report properties are affected only if the constraints imposed by the cheap talk are actually binding. In our case, however, these constraints *do not bind* in equilibrium (for both 2D and 3D reports). Thus, as our result below shows, in our setting the subsequent receipt and cheap talk communication of private information is irrelevant for the choice of report properties.[50]

Lemma 4. *Subsequent receipt and communication of private information do not affect the CEO's ex ante choice of report properties, the board's information acquisition, and the board's action strategy.*

Our result implies that the distribution $\widetilde{\Phi}$ from Proposition 1 is always implementable through a compound lottery with nested equilibrium cheap talk messages. Put differently, the same distribution of the board's interim beliefs is implemented by a report and a collection of cheap talk strategies that do not add additional information to any report realization. Given that the distribution of final beliefs $\widetilde{\Phi}$ remains unchanged in the presence of private information the following result follows immediately without a formal proof.

Proposition 4. *The optimal board composition with receipt and nonverifiable (cheap talk) communication of private information is as described in Proposition 2.*

5 Concluding remarks

We study the interaction between an empire-building CEO and a board of directors about an investment opportunity. We show that the CEO prepares an imperfectly informative report that encourages the board not to gather additional information and approve some value-destroying projects (i.e., commit type-I errors). We show that the precision of the

[50] For a setting that affects the ex ante choice of report properties, see Jain (2020). For a setting in which the subsequent receipt and verifiable disclosure of private information affect the ex ante design of public reports, see Friedman, Hughes and Michaeli (2020).

CEO's report is decreasing in the friendliness of the board. Thus, unfriendly boards observe more precise information and commit fewer approval errors. That is why shareholders benefit from nominating as unfriendly board as possible in the presence of an empire-building CEO. However, directors with such conflicting interests may not be willing to serve on the board when their antagonism implies also lower benefits from undertaking the investments. Taking into account the board's participation constraint, we establish the existence of two distinctly different optimal board types: one that has a weak conflict with the CEO and one that has a strong conflict with the CEO. Our results predict that, in more mature industries and in industries with poorer outside opportunities, shareholders are more likely to assemble unfriendly boards. Our model takes the compensation contracts and the collective decision mechanisms as given. Future work may consider the ability of shareholders to further shape the preferences of CEOs and directors through incentive contracting and implementing various protocols of collective decision-making.

Appendix

Proof of Lemma 1: When characterizing the extreme points of NL^0 and NL^1, we apply Proposition 2 in Caplin et al. (2019). It characterizes the optimal strategy of a rationally inattentive receiver (i.e., board) that satisfies the martingale property. In particular, we use the fact that, if the project decision is binary (approve or reject) and an interim belief is located in a learning region, the optimal board information acquisition strategy is a lottery over two final beliefs. The first final belief, μ^t, is in the border of NL^1 (a point t) and leads to the project approval, and the second final belief $\mu^{t'}$ is in the border of NL^0 (a point t') and leads to the project rejection. To derive these borderline posterior beliefs, we apply the 'Invariant Likelihood Ratio (ILR) Equations for Chosen Options' property. Its geometric interpretation in 2D-space is that the slopes of the board's net payoff function at points t and t', $\sum_{\theta\in\Theta}\mu_\theta^t u(1,\theta)+\lambda H(\boldsymbol{\mu}^t)$ and $\sum_{\theta\in\Theta}u(0,\theta)+\lambda H(\boldsymbol{\mu}^{t'})$, must be identical. Intuitively, this property uses the fact that the board's net payoff function emerges as a result of concavification and, after concavification, the payoff function is by definition linear in the learning region and strictly concave in non-learning regions. Figure 5 in Caplin et al. (2019) illustrates this property.

We begin with deriving the extreme points D and F on the BC-line where $\mu_1 = 0$ and $\mu_2+\mu_3 = 1$. In the borderline point D, the board approves the project, and the board's net payoff function is $\mu_3^D(v_3-\gamma-\lambda\ln\mu_3^D)+(1-\mu_3^D)(v_2-\gamma-\lambda\ln(1-\mu_3^D))$; its slope in μ_3-dimension is $(v_3-\gamma-\lambda\ln\mu_3^D)-(v_2-\gamma-\lambda\ln(1-\mu_3^D))$. In the borderline point F, the board rejects the project, and the board's net payoff function is $-\mu_3^F\lambda\ln\mu_3^F-(1-\mu_3^F)\lambda\ln(1-\mu_3^F))$. Its slope in μ_3-dimension is $-\lambda\ln\mu_3^F+\lambda\ln(1-\mu_3^F)$. Like in Caplin et al. (2019), we use that the slopes are equal if the state-specific components are equal across states, $v_3-\gamma-\lambda\ln\mu_3^D = -\lambda\ln\mu_3^F$ and $v_2-\gamma-\lambda\ln(1-\mu_3^D) = -\lambda\ln(1-\mu_3^F)$. After rearranging, $e^{\frac{1}{\lambda}(v_3-\gamma)} = \frac{\mu_3^D}{\mu_3^F}$ and $e^{\frac{1}{\lambda}(v_2-\gamma)} = \frac{1-\mu_3^D}{1-\mu_3^F}$. The unique solution to this system of two linear equations is

$$(\mu_3^D, \mu_3^F) = \left(\frac{e^{\frac{1}{\lambda}(v_3 - v_2)} - e^{\frac{1}{\lambda}(v_3 - \gamma)}}{e^{\frac{1}{\lambda}(v_3 - v_2)} - 1}, \frac{e^{\frac{1}{\lambda}(\gamma - v_2)} - 1}{e^{\frac{1}{\lambda}(v_3 - v_2)} - 1} \right).$$

Note that for $\gamma \in (v_2, v_3)$, it always holds that $\mu_3^D \in (0,1)$ and $\mu_3^F \in (0,1)$. That is, NL^1 is non-empty.

The extreme points E and G are on the AC-line where $\mu_2 = 0$ and $\mu_1 + \mu_3 = 1$. In the borderline point E, the board approves the project, and the board's net payoff function is $\mu_3^E(v_3 - \gamma - \lambda \ln \mu_3^E) + (1 - \mu_3^E)(v_1 - \gamma - \lambda \ln(1 - \mu_3^E))$. Its slope in μ_3-dimension is $(v_3 - \gamma - \lambda \ln \mu_3^E) - (v_1 - \gamma - \lambda \ln(1 - \mu_3^E))$. In the borderline point G, the board rejects the project, and the board's net payoff function is $-\mu_3^G \lambda \ln \mu_3^G - (1 - \mu_3^G)\lambda \ln(1 - \mu_3^G))$. Its slope in μ_3-dimension is $-\lambda \ln \mu_3^G + \lambda \ln(1 - \mu_3^G)$. Imposing that the state-specific components are equal across states and rearranging yields $e^{\frac{1}{\lambda}(v_3 - \gamma)} = \frac{\mu_3^E}{\mu_3^G}$ and $e^{\frac{1}{\lambda}(v_1 - \gamma)} = \frac{1 - \mu_3^E}{1 - \mu_3^G}$. The unique solution to this system of two linear equations is

$$(\mu_3^E, \mu_3^G) = \left(\frac{e^{\frac{1}{\lambda}(v_3 - v_1)} - e^{\frac{1}{\lambda}(v_3 - \gamma)}}{e^{\frac{1}{\lambda}(v_3 - v_1)} - 1}, \frac{e^{\frac{1}{\lambda}(\gamma - v_1)} - 1}{e^{\frac{1}{\lambda}(v_3 - v_1)} - 1} \right).$$

Note that for $\gamma \in (v_2, v_3)$, it always holds that $\mu_3^E \in (0,1)$ and $\mu_3^G \in (0,1)$. That is, NL^0 is non-empty. □

Proof of Lemma 2: The proof is by contradiction. Let $W(\mu) \equiv \mathbb{E}[w(\cdot)|r]$ be the expected payoff of the CEO at the interim stage for given report realization (and the interim belief μ that it invokes). We refer to $W(\mu)$ as the CEO's "indirect value function."[51] The optimal report is characterized by the distribution $\widetilde{\phi}$ over the reduced set $\widehat{\mathbb{S}}$, where

$$\widetilde{\phi} = \left(\widetilde{\phi}^r\right)_{r \in \widehat{\mathbb{S}}} = \arg \max_{\phi \in \Delta(\widehat{\mathbb{S}})} \sum_{r \in \widehat{\mathbb{S}}} \phi^r W(\mu^r), \tag{6}$$

[51]Because the report discourages additional information acquisition, the board's action at the interim stage (i.e., for given belief μ) is a deterministic variable, characterized by a function, $\sigma : \Delta(\Theta) \to \{0,1\}$. With this observation, we can easily express the expected payoff of the CEO at the interim stage: $W(\mu) = \sigma(\mu) \left(\sum_{\theta \in \Theta} \mu_\theta v_\theta + b\right)$.

subject to the martingale property, $\upmu^o = \sum_{r\in\widehat{\mathbb{S}}} \phi^r \upmu^r$.[52] We show that any distribution with a positive ϕ^E, ϕ^F, or ϕ^G is not optimal.

(i) Suppose $\widetilde{\phi}^E > 0$: Construct a refined distribution $\widehat{\upphi}$ by redistributing all probability mass from point E to points A and C. The probability mass is divided into $(1-\mu_3^E, \mu_3^E) = (\mu_1^E, \mu_3^E)$ shares such that the martingale property is satisfied, $(1-\mu_3^E)\upmu^A + \mu_3^E \upmu^C = \upmu^E$. That is, $\widehat{\phi}_A = \widetilde{\phi}^A + \mu_1^E \widetilde{\phi}^E < 1$, $\widehat{\phi}_C = \widetilde{\phi}^C + \mu_3^E \widetilde{\phi}^E < 1$, $\widehat{\phi}_E = 0$, and $\widehat{\upphi} = \widetilde{\upphi}$ otherwise. The redistribution increases the CEO's objective,

$$\sum_{r\in\widehat{\mathbb{S}}} \widehat{\upphi} W(\upmu) - \sum_{r\in\widehat{\mathbb{S}}} \widetilde{\upphi} W(\upmu) = \widetilde{\phi}^E[\mu_1^E W(\upmu^A) + \mu_3^E W(\upmu^C) - W(\upmu^E)] = -\widetilde{\phi}^E \mu_1^E W(\upmu^A) > 0,$$

which follows from $\mu_1^E > 0$, $\widetilde{\phi}^E > 0$, and $W(\upmu^A) = v_1 + b < 0$.

(ii) Suppose $\widetilde{\phi}^F > 0$: By analogy, construct a refined distribution $\widehat{\upphi}$ by redistributing all probability mass from point F to points B and C. The probability mass is divided into $(1-\mu_3^F, \mu_3^F) = (\mu_2^F, \mu_3^F)$ shares such that the martingale property is satisfied, $(1-\mu_3^F)\upmu^B + \mu_3^F \upmu^C = \upmu^F$. That is, $\widehat{\phi}_B = \widetilde{\phi}^B + \mu_2^F \widetilde{\phi}^F < 1$, $\widehat{\phi}_C = \widetilde{\phi}^C + \mu_3^F \widetilde{\phi}^F < 1$, $\widehat{\phi}_F = 0$, and $\widehat{\upphi} = \widetilde{\upphi}$ otherwise. The redistribution increases the CEO's objective,

$$\sum_{r\in\widehat{\mathbb{S}}} \widehat{\upphi} W(\upmu) - \sum_{r\in\widehat{\mathbb{S}}} \widetilde{\upphi} W(\upmu) = \widetilde{\phi}^F[\mu_2^F W(\upmu^B) + \mu_3^F W(\upmu^C) - W(\upmu^F)] = \widetilde{\phi}^F \mu_3^F W(\upmu^C) > 0,$$

which follows from $\mu_3^F > 0$, $\widetilde{\phi}^F > 0$, and $W(\upmu^C) = v_3 + b > 0$.

(iii) Suppose $\widetilde{\phi}^G > 0$: Like in the case of point E, construct a refined distribution $\widehat{\upphi}$ by redistributing all probability mass from point G to points A and C. The probability mass is divided into $(1-\mu_3^G, \mu_3^G) = (\mu_1^G, \mu_3^G)$ shares such that the martingale property is satisfied, $(1-\mu_3^G)\upmu^A + \mu_3^G \upmu^C = \upmu^G$. That is, $\widehat{\phi}_A = \widetilde{\phi}^A + \mu_1^G \widetilde{\phi}^G < 1$, $\widehat{\phi}_C =$

[52]More precisely, the problem is characterized by 18 constraints, where 14 inequalities and one equality are due to the existence of a simplex, $0 \leq \phi^r \leq 1, r \in \widehat{\mathbb{S}}$, $\sum_{r\in\widehat{\mathbb{S}}} \phi^r = 1$, and three equalities are the martingale properties: $\mu_\theta^o = \sum_{r\in\widehat{\mathbb{S}}} \phi^r \mu_\theta^r$, $\theta = 1, 2, 3$.

$\widetilde{\phi}^C + \mu_3^G \widetilde{\phi}^G < 1$, $\widehat{\phi}_G = 0$, and $\widehat{\phi} = \widetilde{\phi}$ otherwise. The redistribution increases the CEO's objective,

$$\sum_{r \in \widehat{\mathbb{S}}} \widehat{\phi} W(\mu) - \sum_{r \in \widehat{\mathbb{S}}} \widetilde{\phi} W(\mu) = \widetilde{\phi}^G [\mu_1^G W(\mu^A) + \mu_3^G W(\mu^C) - W(\mu^G)] = \widetilde{\phi}^G \mu_3^G W(\mu^C) > 0,$$

which follows from $\mu_3^G > 0$, $\widetilde{\phi}^G > 0$, and $W(\mu^C) = v_3 + b > 0$. □

Proof of Proposition 1: Let, as in the proof of Lemma 2, $W(\mu) \equiv \mathbb{E}[w(\cdot)|r]$ be the indirect value function. We analyze whether an alternative convex hull exists such that the prior lies in the hull and a signal can be constructed from extreme points of the hull. By Carathéodory's theorem, concavification over a two-dimensional simplex is based on at most three linearly biased points. In our case, convex hulls created by triplets of linearly biased points are $\{ABD, ACD, ABC\}$. In addition, we have convex hulls constructed by pairs of points, i.e., lines. We proceed in two steps. First, we suppose that the prior is not on any line between points $\{A, B, C, D\}$, which is equivalent to be in the interior of Δ_{ABD} or in the interior of Δ_{ACD}. This eliminates convex hulls constructed by pairs of points. Second, we analyze the boundaries of Δ_{ABD} and Δ_{ACD}.

- For any prior in the interiors, $\mu_0 \in \text{int}(\Delta_{ABD}) \cup \text{int}(\Delta_{ACD})$, the only alternative convex hull is Δ_{ABC}. Replacing $\widetilde{\phi}$ (where $\widetilde{\phi} = \phi_{ABD}$ or $\widetilde{\phi} = \phi_{ACD}$) by ϕ_{ABC} is equivalent to redistributing all probability mass $\widetilde{\phi}^D$ from point D to points B and C. The probability mass is divided into $(1 - \mu_3^D, \mu_3^D) = (\mu_2^D, \mu_3^D)$ shares such that the martingale property is satisfied, $(1 - \mu_3^D)\mu^B + \mu_3^D \mu^C = \mu^D$. That is, $\phi_{ABC}^B = \widetilde{\phi}^B + \mu_2^D \widetilde{\phi}^D < 1$, $\phi_{ABC}^C = \widetilde{\phi}^C + \mu_3^D \widetilde{\phi}^D < 1$, $\phi_{ABC}^D = 0$, and $\phi_{ABC} = \widetilde{\phi}$

otherwise. This redistribution decreases the CEO's objective,

$$\sum_{r\in\widehat{\mathbb{S}}} \phi_{ABC} W(\mu) - \sum_{r\in\widehat{\mathbb{S}}} \widetilde{\phi} W(\mu) = \widetilde{\phi}^D[\mu_2^D W(\mu^B) + \mu_3^D W(\mu^C) - W(\mu^D)]$$
$$= -\widetilde{\phi}^D(1-\mu_3^D)W(\mu^B) < 0,$$

because $\mu_2^o > 0$ and $\mu_3^o > 0$ (hence, $\widetilde{\phi}^D > 0$), and $W(\mu^B) = v_2 + b > 0$.

- Consider boundaries of the simplex. For $\mu_1^o = 0$, we replicate the argument from above; redistribution of the probability mass from point D to points B and C decreases the CEO's objective as $W(\mu^D) > \mu_2^D W(\mu^B) + \mu_3^D W(\mu^C)$. For $\mu_2^o = 0$, all feasible signal distributions based on $\{A, B, C, D\}$ are exactly equivalent, $(\phi^A, \phi^B, \phi^C, \phi^D) = (\mu_1^o, 0, \mu_3^o, 0)$. The same holds for $\mu_3^o = 0$, where $(\phi^A, \phi^B, \phi^C, \phi^D) = (\mu_1^o, \mu_2^o, 0, 0)$.

- The remaining case is when μ_0 is located on AD line but not on the boundary. Then, $\frac{\mu_2^o}{1-\mu_3^D} = \frac{\mu_3^o}{\mu_3^D}$, and $\phi_{ABD} = \phi_{ACD}$. No other distribution based on $\{A, B, C, D\}$ is feasible.

To sum up, none of the alternative concavifications over $\widehat{\mathbb{S}}$ is a solution to the CEO's linear programming problem in equation (6). □

Proof of Corollary 2: Follows directly from Proposition 1 and is omitted. □

Proof of Corollary 3: For a strong conflict ($\mu_0 \in \Delta_{ABD}$) we can coarsen the report realizations with interim beliefs μ^A and μ^B at points A and B leading to rejection, into a single realization with interim belief μ^L at point L, where

$$\mu^L = \frac{\widetilde{\phi}^A}{\widetilde{\phi}^A + \widetilde{\phi}^B}\mu^A + \frac{\widetilde{\phi}^B}{\widetilde{\phi}^A + \widetilde{\phi}^B}\mu^B.$$

The optimal 3D-distribution then is equivalent to a 2D-distribution ϕ_{LD} over a subset $\{L, D\}$ where

$$(\phi_{LD}^{L}, \phi_{LD}^{D}) \equiv \left(\phi_{ABD}^{A} + \phi_{ABD}^{B}, \phi_{ABD}^{D}\right).$$

For a weak conflict ($\mu_0 \in \Delta_{ACD}$) we can coarsen the report realizations with interim beliefs μ^C and μ^D at points C and D leading to approval, into a single realization with interim belief μ^H at point H, where

$$\mu^H = \frac{\widetilde{\phi}^C}{\widetilde{\phi}^C + \widetilde{\phi}^D}\mu^C + \frac{\widetilde{\phi}^D}{\widetilde{\phi}^C + \widetilde{\phi}^D}\mu^D.$$

The optimal 3D-distribution is equivalent to a 2D-distribution ϕ_{AH} over a subset $\{A, H\}$, where

$$(\phi_{AH}^{A}, \phi_{AH}^{H}) \equiv \left(\phi_{ACD}^{A}, \phi_{ACD}^{C} + \phi_{ACD}^{D}\right).$$

$\square$

Proof of Lemma 3: We derive the first and second derivatives of the functions $\rho(\gamma)$ and $\psi(\gamma)$, to be denoted $\rho_\gamma, \rho_{\gamma\gamma}, \psi_\gamma, \psi_{\gamma\gamma}$. To determine the signs of ψ_γ and $\psi_{\gamma\gamma}$, notice that $\gamma < \gamma^s$ and $\gamma - v_2 > 0$.

$$\rho_\gamma = \frac{e^{\frac{1}{\lambda}(v_3-\gamma)}(e^{\frac{1}{\lambda}(v_3-v_2)} - 1)}{\lambda(e^{\frac{1}{\lambda}(v_3-\gamma)} - 1)^2} > 0, \psi_\gamma = \frac{v_3 - v_2 - \frac{\underline{U}}{\mu_3^o}}{(v_3 - \gamma - \frac{\underline{U}}{\mu_3^o})^2} > 0$$

$$\rho_{\gamma\gamma} = -\frac{1}{\lambda}\rho_\gamma < 0, \psi_{\gamma\gamma} = \frac{2\psi_\gamma}{v_3 - \gamma - \frac{\underline{U}}{\mu_3^o}} > 0$$

By inspection of the derivatives, we cannot determine the slope of Δ_γ unambiguously, but we can observe that the function is concave for $\gamma \in (v_2, \gamma^s)$, $\rho_{\gamma\gamma} - \psi_{\gamma\gamma} < 0$. Next, we evaluate the function at the lowest possible bias v_2 and also at the level where it approaches the positivity constraint γ^s (from below). Comparing the two values reveals that the function has decreased:

$$\lim_{\gamma \to v_2} (\rho(\gamma) - \psi(\gamma)) = 0 > -\infty = \lim_{\gamma \to \gamma^s} (\rho(\gamma) - \psi(\gamma)).$$

To see the latter, notice

$$\lim_{\gamma \to \gamma^s} \rho(\gamma) = \rho(\gamma^s) = \frac{e^{\frac{1}{\lambda}(v_3 - v_2)} - e^{\frac{1}{\lambda}\underline{U}/\mu_3^o}}{e^{\frac{1}{\lambda}\underline{U}/\mu_3^o} - 1} > -\infty = \lim_{\gamma \to \gamma^s} -\psi(\gamma).$$

As the function is concave, continuous, and non-increasing for some $\gamma \in (v_2, \gamma^s)$, it must be either decreasing everywhere or increasing first and then decreasing. In both cases, the function is decreasing on the left neighborhood of γ^s.

Proof of Proposition 2: The proof proceeds in two main steps. In the first step, we begin with a taxonomy; we will characterize a pair of outcomes for the weak conflict and a pair of outcomes for the strong conflict. In the second step, we demonstrate that the pairs are equivalent. Hence, a single condition determines which of the outcomes is selected under the weak (respectively strong) conflict. We also show that this condition also characterizes whether the shareholders implements the weak conflict or the strong conflict.

Step 1. Taxonomy of outcomes. For the weak conflict, we will speak of an unconstrained outcome if $\gamma^w \geq \gamma^o$ and a constrained outcome if $\gamma^w < \gamma^o$. For the strong conflict, by concavity, the function $\rho(\gamma) - \psi(\gamma)$ is either (i) increasing everywhere, or (ii) increasing and then decreasing, or (iii) decreasing everywhere. By Lemma 3, the first scenario is eliminated. Knowing that the function is concave, at most two solutions of $\rho(\gamma) - \psi(\gamma) = 0$ exist. We have already identified one solution at the lowest feasible bias, $\gamma = v_2$ (but it is obviously not admissible for a strong conflict). Therefore, two outcomes may exist:

- The second solution doesn't exist or is not feasible for a strong conflict. By Lemma 3, $\rho(\gamma) - \psi(\gamma)$ is decreasing is a lower neighborhood of γ^s. If the second solution doesn't exist (because the function is decreasing everywhere) or is not feasible for a strong conflict (as the value of the solution is below γ^o), the function is decreasing and negative on the entire interval $\gamma \in [\gamma^o, \gamma^s)$, and hence the strong conflict is not implementable (i.e., the participation constraint binds everywhere in the interval).

- The second solution exists and is feasible for a strong conflict. We will denote the solution $\gamma^* \in [\gamma^o, \gamma^s) : \rho(\gamma^*) - \psi(\gamma^*) = 0$. As the function is concave, it is decreasing in the neighborhood of the second (higher) solution, and by recalling $\rho(v_2) - \psi(v_2) = 0$, we obtain $\Delta_\gamma > 0$ for any $\gamma \in (v_2, \gamma^*]$. Therefore, the strong conflict is implementable for any $\gamma \in [\gamma^o, \gamma^*)$.

Step 2. Equivalence between outcomes for weak and strong conflicts. Next, we examine the link between the taxonomy for the weak conflict and the strong conflict. Not surprisingly, the link is very tight as (i) each taxonomy depends on the willingness to participate at the cutoff γ^o and (ii) the participation constraints in (3) and (4) for the cutoff bias, γ^o, are identical. Formally, if the board of the type γ^o participates under the weak conflict (an unconstrained outcome), then it also participates under the strong conflict, $\rho(\gamma^o) - \psi(\gamma^o) \geq 0$, and the second solution under the strong conflict exists. And if the board of the type γ^o doesn't participate under the weak conflict (a constrained outcome), then the board doesn't participate under the strong conflict either, $\rho(\gamma^o) - \psi(\gamma^o) < 0$, and the second solution under the strong conflict doesn't exist or is not feasible.

Finally, we exploit this condition to derive the optimal level of bias.

- If the γ^o-type board participates, any $\gamma \in [v_2, \gamma^*]$ is implementable. As $\Pi(\gamma)$ is increasing in γ (because the payoff is increasing in $\rho(\cdot)$, and $\rho(\cdot)$ is increasing in γ), the optimal board is at the highest implementable level of bias, $\gamma = \gamma^*$, and the shareholders earn payoff $\Pi(\gamma^*) > \Pi(\gamma^o)$.

- If the γ^o-type board doesn't participates, any $\gamma \in [v_2, \gamma^w *]$ is implementable. Here, the shareholders are indifferent as they earn only a guaranteed payoff $\Pi(\gamma^o)$.

Proof of Proposition 3: We begin with comparative statics of the shareholder's decision to adopt the weak conflict or strong conflict. The critical condition, i.e., whether the board of the cutoff type is willing to participate or not, can be written simply as $\gamma^w - \gamma^o \geq 0$. We thus introduce and analyze a function

$$G(\underline{U}, \mu_2^o, \mu_3^o) \equiv \gamma^w - \gamma^o = \frac{\mu_2^o v_2 + \mu_3^o v_3 - \underline{U}}{\mu_2^o + \mu_3^o} - v_3 + \lambda \ln \frac{\mu_2^o e^{\frac{1}{\lambda}(v_3 - v_2)} + \mu_3^o}{\mu_2^o + \mu_3^o}.$$

As $G(\cdot)$ is decreasing in $\underline{U}$, the board is more likely friendly if her reservation utility is high. For the analysis of the change in the frequency of irrelevant projects, μ_1^o, given a constant approval precision under prior beliefs (ρ^o unchanged), we obtain that the function $G(\cdot|\rho^o) = \frac{1}{1+\rho^o}\ v_2 + \frac{\rho^o}{1+\rho^o}\ v_3 - \frac{\underline{U}}{1-\mu_1^o} - \gamma^o$ is clearly decreasing in μ_1^o. Thus, the board is more likely friendly if the frequency of the irrelevant projects increases without affecting the board optimal decision for a signal that mixes all conflicting and profitable projects. We also obtain that $G(\cdot)$ is increasing in μ_3^o, given a constant share of conflicting projects, μ_2^o, and decreasing in μ_2^o, given a constant share of profitable projects, μ_3^o. (Here, in the derivation, we exploit that the reservation utility is not prohibitive.) To get that the function is decreasing in λ, we will use that γ^o is implicitly defined by $\rho(\lambda) - \rho^o = 0$, and use $\rho_\gamma > 0$ and $\rho_\lambda < 0$ from Corollary 1.

Second, we conduct comparative statics for the optimal board's bias under the strong conflict, using the Implicit Function Theorem. For instance, for board reservation utility,

$$\frac{d\gamma^*}{d\underline{U}} = -\frac{\rho_{\underline{U}} - \psi_{\underline{U}}}{\rho_\gamma - \psi_\gamma}.$$

At $\gamma = \gamma^*$, we have $\rho_\gamma - \psi_\gamma < 0$; hence, in the comparative statics, it is always sufficient to examine the sign of the numerator. Specifically for the board reservation utility, $\rho_{\underline{U}} - \psi_{\underline{U}} = -\psi_{\underline{U}} < 0$, hence the bias is decreasing under the strong conflict. The bias is also decreasing in the frequency of irrelevant projects (given fixed approval likelihood) and the frequency of conflicting projects (given fixed profitable projects). In contrast, the frequency of profitable projects (given fixed agency conflict) and maturity of the industry increase the optimal value of bias.

Proof of Lemma 4: For the optimal 3D-signal, we only need to check four posteriors for the optimal persuasion, $\{A, B, C, D\}$. At any belief μ^j, where $j \in \{A, B, C\}$, the CEO's type is revealed and the cheap talk is uninformative. At μ^D, there is no cheap talk either. The CEO prefers approval (for any type), and hence always recommends.

For the optimal 2D-signals, we additionally need to check $\{L, H\}$. At μ^H, the incentives are exactly as the ones at μ^D. At μ^L, for low or intermediate type, the board prefers to reject, and hence no CEO type can benefit from any other message. $\square$

References

[1] Adams, R. B., Ferreira, D. (2007). A theory of friendly boards. *Journal of Finance*, 62(1), 217–250.

[2] Aghamola, C., Hashimoto, T. (2020). Aggressive Boards and CEO Turnover. Working paper.

[3] Arya, A., Glover, J., Sivaramakrishnan, K., (1997). Commitment issues in budgeting. *Journal of Accounting Research*, 35, 273–278.

[4] Associated Press (2020): Report: former Aluminum Mill CEO misled investors, board. April 27, 2020.

[5] Baldenius, T., Melumad, N., Meng, X. (2014). Board composition and CEO power. *Journal of Financial Economics*, 112(1), 53–68.

[6] Baldenius, T., Meng, X., Qiu, L. (2019). Biased boards. *The Accounting Review*, 94(2), 1–27.

[7] Baldenius, T., Meng, X., Qiu, L. (2020). The Value of Board Commitment. *Review of Accounting Studies*, forthcoming.

[8] Banerjee, S., Szydlowski, M. (2020). Friends Don't Lie: Monitoring and Communication With Risky Investments. Working paper.

[9] Bloedel, A., Segal, I., (2020). Persuading a Rationally Inattentive Agent. Working Paper.

[10] Bertomeu, J., Cheynel, E., (2015). Asset measurement in imperfect credit markets. *Journal of Accounting Research* 53.

[11] Caplin, A., Dean, M., Leahy, J. (2019). Rational inattention, optimal consideration sets, and stochastic choice. *Review of Economic Studies*, 86(3), 1061–1094.

[12] Chakraborty, A., Yilmaz, B. (2017). Authority, consensus, and governance. *Review of Financial Studies*, 30(12), 4267–4316.

[13] Chen, K., Guay, W., Lambert, R., (2020). Multidimensional Corporate Governance: Board Design, Executive Compensation, and Information Sharing. Working Paper.

[14] Cianciaruso, D., Marinovic, I., Smith, K. (2020). Information Design in Financial Markets. Working Paper.

[15] Deloitte (2015). Directors' Cut: Board Effectiveness. Available at: https://www2.deloitte.com/content/dam/Deloitte/au/Documents/risk/deloitte-au-risk-director-cut-310815.pdf

[16] Dordzhieva, A., Laux, V., Zhang, R. (2020). Signaling entrepreneurial vision via information system design. Working Paper.

[17] Drymiotes, G., (2007). The monitoring role of insiders. *Journal of Accounting and Economics*, 44 (3), 359-377.

[18] Drymiotes, G., (2009). Managerial Influencing of Boards of Directors. *Journal of Management Accounting Research* 20 (Special Issue), 19-45.

[19] Friedman, H., Hughes, J., Michaeli, B., (2020). Optimal reporting when additional information might arrive. *Journal of Accounting and Economics*, 69 (2-3).

[20] Frydman, C., Hilt, E. (2017). Investment banks as corporate monitors in the early twentieth century United States. *American Economic Review*, 107(7), 1938–70.

[21] Garcia, A. (2018). Director skills: diversity of thought and experience in the boardroom. ISS analytics, institutional shareholder services, September 14, 2018. Available at: https://www.issgovernance.com.

[22] Gentzkow M., Kamenica E., (2017). Competition in persuasion. *Review of Economic Studies*, 84, 300–322.

[23] Gregor, M. (2020). Board directors under CEO's pressure: misreading, misreporting, and dissent risk. Mimeo.

[24] Göx, R., Michaeli, B., (2019). Optimal information design and incentive contracts with performance measure manipulation. Working paper.

[25] Göx, R.F., Wagenhofer, A. (2009). Optimal impairment rules. *Journal of Accounting and Economics*, 48(1), 2–16.

[26] Harris, M., Raviv, A. (2008). A theory of board control and size. *Review of Financial Studies*, 21(4), 1797-1832.

[27] Harris, M., Raviv, A. (2010). Control of corporate decisions: Shareholders vs. management. *Review of Financial Studies*, 23(11), 4115-4147

[28] Huang, Z., (2016). Optimal reporting systems with investor information acquisition. Mimeo.

[29] Jagannathan, M., Jiao, W., Krishnamurthy, S. (2020). Missing them yet? Investment banker directors in the 21st century. *Journal of Corporate Finance*, forthcoming.

[30] Jain, V. (2020). Bayesian persuasion with cheap talk. Working Paper.

[31] Jensen, M. (1993). The modern industrial revolution, exit, and the failure of internal control systems. *Journal of Finance*, 48, 831-880.

[32] Jiang, X., Yang, M. (2017). Properties of optimal accounting rules in a signaling game. *Journal of Accounting and Economics*, 63, 499-512.

[33] Jiang, X., Stocken, P. (2019). Public Communication between Managers and Analysts. Working Paper.

[34] Jiang, W., Wan, H., Zhao, S. (2015). Reputation concerns of biased directors: evidence from individual director voting. *Review of Financial Studies*, 29(3), 655–696.

[35] Kamenica, E., Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6), 2590–2615.

[36] Laux, V. (2014). Corporate governance, board oversight, and CEO turnover. *Foundations and Trends in Accounting* 8(1), 1-73.

[37] Levit, D. (2012). Expertise, structure, and reputation of corporate boards. Available at SSRN 1866384.

[38] Levit, D., Malenko, N. (2016). The Labor Market for Directors and Externalities in Corporate Governance. *The Journal of Finance*, 71, 775-808.

[39] Malenko, N. (2014). Communication and decision-making in corporate boards. *Review of Financial Studies*, 27(5), 1486–1532.

[40] Matějka, F., McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1), 272–98.

[41] Matysková, L. (2018). Bayesian persuasion with costly information acquisition. CERGE-EI Working Paper Series (614).

[42] Meng, X. (2020). Information, incentives and CEO replacement. Working paper.

[43] Michaeli, B. (2017). Divide and inform: rationing information to facilitate persuasion. *The Accounting Review*, 92(5), 167–199.

[44] Peterson, M. (2003). Kmart, saying ex-chief misled board, wants payments returned. The New York Times, February 26, 2003.

[45] Qiu, L. (2020). Say on pay and the “quiet life” of boards. Working paper.

[46] Sims, C.A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3), 665–690.

[47] The Economist, (2001). The fading appeal of the board room. March 31, 2001.

[48] Useem, M. (2006). How Well-Run Boards Make Decisions. *Harvard Business Review*, November 2016.

[49] Ward, J., Zsolnay, C. (2016) *Plymouth Tube Company: How to Build Ownership Consensus*. Case KEL961. Kellogg School of Business.

[50] Wommack, W. (1979). The Board's Most Important Function. *Harvard Business Review*, September 1979.

[51] Wouters, M., Davila, K. (2018). *AgriSmart: Funding New Corporate Ventures*. Case No. W18185. Ivey Publishing.