

Board Bias, Information, and Investment Efficiency*

Martin Gregor[†]

Beatrice Michaeli[‡]

November 26, 2022

*For helpful comments and suggestions we thank Aleksander Aleszczyk, Tim Baldenius, Snehal Banerjee, Jeremy Bertomeu, Anne Beyer, Jonathan Bonham, Judson Caskey, Jonathan Craske, George Drymiotis, Ron Dye, Yonca Ertimur, Elia Ferracuti, Fabrizio Ferri, Paul Fischer, Henry Friedman, Jonathan Glover, Robert Göx, Wayne Guay, Ilan Guttman, Mirko Heinle, Zeqiong Huang, John Hughes, Xu Jiang, Navin Kartik, Rick Lambert, Volker Laux, Christian Leuz, Mark Maffett, Ivan Marinovic, Xiaojing Meng, Anya Nakhmurina, Stefan Reichelstein, Sander Renes, Joshua Ronen, Katherine Schipper, Anastasiya Shamshur, Kevin Smith, Philip Stocken, Jacob Thomas, Seung Yeon Yoo, Ronghuo Zheng, and workshop participants at Wharton, Stanford, University of Chicago, Kellogg, Duke, LBS, Yale, Dartmouth, NYU, USC, UCLA, UT Austin, Rice, UCSD, Purdue, CUNY, Bristol, ARW, EAA, EFMA, and the Czech Economics Society. Martin Gregor acknowledges financial support from the Czech Science Foundation (GAČR 20-03517S) and the Cooperatio Program at Charles University.

[†]Martin Gregor is at the 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.

Board Bias, Information, and Investment Efficiency

Abstract: We study how interest alignment between CEOs and corporate boards affects investment efficiency and identify a novel force behind the benefit of misaligned preferences. The model entails a CEO who encounters an investment project, collects investment-relevant information and decides either or not to present the project implementation for approval by a board of directors. The CEO can strategically choose the properties of the collected information only if the project is “novel” in the sense that it explores a new technology, business concept, or market. We find that only directors with sufficiently high anti-approval bias (e.g., with career concerns or conservative/antagonistic attitude toward corporate investments) can discipline the CEO’s empire-building tendency and opportunistic information collection and reporting. Directors with this bias, however, underinvest in projects that are not novel. From the shareholders’ perspective, the board that maximizes firm value is either biased against approval or neutral (has investment interests aligned with those of the shareholders) and always overinvests in innovations. Boards with greater expertise are more likely to be biased but their bias is less severe.

Keywords: empire-building, biased board, underinvestment, overinvestment, endogenous information

1 Introduction

Corporate boards are often tasked with the approval of significant investment opportunities such as cash acquisitions of companies or major product launches (Useem 2006). Perhaps in an attempt to ensure that boards can successfully curb the CEOs’ empire-building tendencies, current regulations either explicitly require a minimum number of directors to be outsiders (e.g., board independence standards) or facilitate activist interventions like those aiming to elect directors with cost-cutting agenda (e.g., SEC’s 2021 Universal Proxy Rules for Director Elections). However, the effect of such regulations on the efficiency of corporate investments is unclear: while several studies find empirical evidence consistent with improvement (e.g., Maffett, Nakhmurina, and Skinner 2020; Rim and Sul 2020), others find evidence consistent with deterioration (e.g., Lu and Wang 2015). Thus, the interplay between boards and CEOs in investment decisions persists.

To appreciate the complex tradeoffs at work, one needs to keep in mind that directors—even outsiders—may face private benefits and costs or have behavioral biases that affect the board’s investment preferences. Furthermore, approval decisions depend on the investment-relevant information. While such information can be found by directors, its collection is often initiated by firm management immediately after encountering a business opportunity and before involving the board. Examples include gathering of information about an acquisition target and conducting tests to evaluate customer demand or technological feasibility. In these cases, CEOs might be able to strategically determine the collection process (e.g., choose criteria for inclusion in due diligence team or focus group) in a way that maximizes the chance for board approval.

The objective of this study is to account for the features described above and study how information and investment efficiency are affected by the alignment of interests between CEOs and corporate boards. We build a model in which a CEO (“she”) finds an investment project and decides either or not to present it for approval to a board of directors.

The preferences of the shareholders, the CEO, and the board are such that each favors investments with value exceeding a player-specific level (“player’s type/bias/threshold”). The players’ alignment of interests is thus determined by the difference in their thresholds which represents private benefits/costs or investment attitudes. In line with classic agency theory (Jensen 1986), we assume that the CEO is an empire-builder and her investment threshold is below that of the shareholders.¹ The board’s threshold can assume any value: In comparison to the shareholders, the board can be classified as “pro-approval” (e.g., when directors are from the CEO’s social circle and have a positive attitude toward investments and company expansion), “anti-approval” (e.g., when directors have career concerns and conservative/antagonistic attitude toward investments and company expansion), or neutral/unbiased and aligned with the shareholders.

The CEO collects information about the encountered project before bringing it to the board for approval. However, the nature of the project could be such that the CEO has no control over the properties of the collected information, potentially because the project is similar to prior operations or because the directors are familiar with the industry. We define this type of project as “routine” and for simplicity assume that its value is fully revealed. Otherwise, the CEO can select the properties of the information she collects, possibly because the project relates to new technology, business concept, or market. We call this type of project “novel.”

In the case of a novel project, the CEO’s decision to collect and report information is modeled as a Bayesian persuasion problem: Before learning any information, the CEO commits to an observable reporting system that maps the project value into a report.

¹Early work emphasizes the incentives of CEOs to overinvest in projects that require specific human capital and make them more indispensable to the firm (Shleifer and Vishny 1989), increase their personal status and prestige (Stulz 1990), increase their compensation if tied to sales growth or firm size (Jensen and Murphy 1990) and diversify risks. Most empirical evidence of empire-building relates to expansion of foreign sales (Hope and Thomas 2008) and underperforming projects in the U.S. oil and gas industry that yield unambiguous monetary benefits to CEOs (Décaire and Sosyura 2021). In many settings, contracting either cannot fully eliminate empire-building or full elimination is too costly for the shareholders (Shleifer and Vishny 1989; Gregor and Michaeli 2021). In this paper, we take contracts as given.

To fit the examples that we have in mind (and as standard in the persuasion literature) we assume that the report is verifiable and can not be withheld.² Within the confines of our model, the CEO’s persuasion problem simplifies to a choice of a reporting cutoff, whereby projects with value above (below) the cutoff are reported as high (low). Setting a higher reporting cutoff increases the expected value of the projects with either report. After observing the CEO’s report, the board can learn the project value but at a cost.

To clearly illustrate the benefit from anti-approval boards, we begin our analysis with a simplified setting where learning by the board is prohibitively costly and the CEO takes into consideration the first business opportunity that she encounters. First, let consider the case where the CEO encountered a novel project. Despite her control over the reporting cutoff and her ability not to present certain projects, the CEO can persuade the board to approve her favored novel projects only if the board is not sufficiently biased against approval. To see why, suppose that the board’s preferences align perfectly with those of the CEO, i.e., their thresholds coincide. In this case, the CEO can simply set the reporting cutoff at her threshold and present only the novel projects reported as having high value—the board then ratifies these projects. This is also the equilibrium outcome when the preferences of the board and the CEO are only mildly misaligned (i.e., the board’s threshold is similar to that of the CEO) or when the board is strongly biased in favor of approval (i.e., the board’s threshold is significantly lower than that of the CEO). However, adopting the same strategy when the board is strongly biased against approval results in the rejection of all presented projects. To avoid this outcome, the CEO optimally increases the reporting cutoff—so that the expected value of novel projects reported as high is just sufficient for the board to approve. Faced with an opportunistically-set reporting cutoff, the board may under- or overinvest in innovations from the shareholders’ perspective. Notably, a board with significantly high anti-approval

²This assumption is not crucial for any of the results. As we explain in Section 3, allowing the CEO to withhold the report or send additional cheap talk messages, leads to an identical outcome as long as the CEO can select the properties of the collected information about novel projects.

bias elicits a cutoff choice that results in efficient novel investments.³ As we explain in Section 2, this benefit from misaligned preferences is driven by a force that has not been studied in prior literature.

When the CEO encounters a routine project, the value of the project is estimated by the existing information system and the CEO loses the ability to persuade the board via a strategically constructed report. In particular, a routine project is approved if its value exceeds the board’s threshold. Thus, while a board with severe anti-approval bias invests efficiently in innovations, it can reject a routine project against the shareholders’ interest. So which board type maximizes firm value? In our setting, a pro-approval board can never be optimal because it distorts decisions about both types of projects. A board that is only mildly biased toward rejection is also suboptimal—it not only distorts decisions about routine investments but also fails to discipline the CEO’s opportunistic reporting about novel projects. We find that only one of two board types can be optimal: neutral or strongly biased against project approval. Notably, the latter type is not biased enough to fully undo the CEO’s empire-building. In equilibrium, this board still approves some (but not all) empire-building projects. Therefore, we predict that firms *overinvest* in innovations but may or may not *underinvest* in routines.

Which of the two board types is optimal depends on the relative magnitude of the expected gain from improved approvals of novel projects and the expected loss from distorted approvals of routine ones. When the CEO is more likely to find a routine project opportunity, the expected loss is relatively large and outweighs the expected gain—thus the optimal board is more likely neutral. In addition, when the CEO has a severe tendency to overinvest, a biased board that disciplines the CEO is preferable from the shareholders’

³The result that anti-approval boards reduce the amount but increase the efficiency of investments comports with empirical evidence that shareholder activism (which leads to nomination of directors who are biased against approval in order to cut costs) is associated with less corporate investment but higher profitability (Maffett, Nakhmurina, and Skinner 2020). Our result is also consistent with the findings about hedge-fund activism of Brav, Jinag, Ma, and Tian (2018) and about director independence requirements of Rim and Sul (2020).

perspective. Our results predict that in environments with heterogeneous projects the distribution of optimal boards is bimodal. Companies managed by mildly biased CEOs have neutral boards and those managed by extreme empire-builders have anti-approval boards. The primary inefficiency in the former case is overinvestment in innovations and in the latter case—underinvestment in routines.

We extend our results by allowing the CEO to search for a project of specific type (novel or routine). Relative to the shareholders, the CEO is more inclined to prefer innovations because they allow her to control the reported information and persuade the board. Thus, when novel projects are more profitable than routine ones, the CEO searches for a novel one. In anticipation of this search, it is best for the shareholders if a severely anti-approval board is nominated. However, when routine projects are more profitable than novel ones, the CEO’s choice crucially depends on the board type. The shareholders then have two choices: they can acquiesce to the CEO’s desire to pursue innovations and accordingly nominate an anti-approval board. Alternatively, they can make the routine opportunity more attractive to the CEO by strategically nominating the board.

Lastly, we study the full-fledged case where learning by the board is not prohibitively costly. Because perfectly informed directors disregard any additional information, the CEO prepares a report that discourages the board from learning.⁴ This additionally restricts the CEO’s ability to get all her favored novel projects approved. Despite this restriction, our main finding that the optimal board is either neutral or strongly biased against approval continues to hold. We find that when learning costs are low (e.g., because directors are experienced or qualified), the optimal board is more likely biased (complementarity between bias and expertise) but also that the optimal level of bias decreases (substitution between bias and expertise). Thus, we predict that boards biased against approval are more frequently observed in industries with experienced professionals. How-

⁴Our result that boards do not learn is consistent with observations that they “depend largely on the chief executive and the company’s management for information” (The Economist 2001).

ever, the higher the directors' expertise, the lower the optimal level of board's bias.

In summary, by accounting for heterogeneity of investment projects and possible managerial control over the properties of the collected information, our model predicts that the optimal boards might have a conservative/antagonistic attitude toward corporate investments. As we elaborate in Section 2, this result departs from prior corporate governance literature and is driven by a novel force.

2 Related Literature

Our paper contributes to several strands of literature. We predict that an intermediary (board of directors) having preferences that *severely differ* from those of a sender (CEO) elicits precise information that benefits a principal (shareholders). This is a significant departure from the predictions in a strand of prior literature that makes the case for *aligned* preference in settings where the sender is exogenously endowed with perfect information and can misrepresent it at no cost (e.g., Dessein 2002; Mitusch and Strausz 2005; Adams and Ferreira 2007; Harris and Raviv 2008; Baldenius, Melumad, and Meng 2014; Chakraborty and Yilmaz 2017).⁵ The findings differ primarily because in our model the sender may choose the properties and the precision of the collected information. We believe this assumption is descriptive of many practical situations, especially those pertaining to new technologies, business concepts or markets where estimating feasibility and safety as well as predicting customer demand require new data the collection process of which could be determined by management.

Like us, several prior studies also call for misalignment but their predictions are driven

⁵A few corporate governance studies consider models with verifiable disclosure (Malenko 2014) or costly misreporting (Gregor 2020; Chen and Laux 2021) and do not study optimal interest alignment. Among other corporate governance aspects studied by the literature are: CEO turnover (Laux 2014; Meng 2020), performance manipulation (Drymiotes 2009), ability to monitor (Drymyotes 2007), board commitment to decision rule (Baldenius, Meng, and Qiu 2020), incentive compensation (Qiu 2020; Gregor and Michaeli 2021, Feng, Luo, and Michaeli 2022), and expertise (Chen, Guay, and Lambert 2020).

by different forces than the ones in our model—we contribute to this strand of literature by studying a novel force behind the benefit from diverged preferences. In Dewatripont and Tirole (1999) there are multiple agents who acquire information and their preferences are determined by the monetary incentives that they are offered. Multitasking in their model leads to competition for information acquisition among the agents and this is best utilized by the principal when the agents have different preferences. Che and Kartik (2009) study a different dimension of misalignment: a difference in prior beliefs (arising only under uncertainty) whereas we study difference in preferred policy (arising even under certainty). In their disclosure model, the greater the difference between the priors of the sender and the receiver, the more precise costly information the sender acquires because he expects that it will change the receiver’s beliefs. In our persuasion model, misalignment doesn’t affect the cost-benefit ratio of information as its acquisition is cost-free for the CEO. Notably, the authors demonstrate that disagreement over priors in their disclosure game works differently than divergent preferences: if the sender and the receiver in their model were to have different preferences but a common prior, then the receiver would always prefer an unbiased sender which is different than our findings.

In Baldenius, Meng, and Qiu (2019) friendly directors receive more precise information from the CEO whereas antagonistic ones search for information on their own. Therefore, antagonistic boards can be optimal when the information from outsiders is more valuable. In our model this channel is absent as the CEO has access to perfectly precise information (i.e., the board has no informational advantage) and chooses report properties that prevent the board from learning (i.e., the board’s information acquisition problem disappears). In Aghamola and Hashimoto (2020) a less friendly intermediary is more likely to fire the CEO. To avoid this outcome, the CEO achieves a boost in productivity by reducing the bias in her report. In our model this friction is missing as we abstract from CEO turnover. In Ball and Gao (2021) the benefit from misalignment arises due to interplay between the

agent’s bias and a restriction on the available policies. There, a biased agent needs to carefully examine whether to select an extreme policy (as a fine-tuned one is not available) and this encourages information acquisition. In our model, this channel is absent because the binary board’s action (approve or reject the project) cannot be restricted.

Misalignment can also be beneficial when the agent interacts with third parties such as suppliers, business partners and competitors. In oligopolies, delegating the product decisions to an agent who competes aggressively serves as a pre-commitment device (Fershtman and Judd 1987) and can shape the managers’ disclosure choices (Bagnoli and Watts 2015). In static bargaining, appointing a less interested agent forces the bargaining partner to reduce her share of surplus, which effectively transfers the principal’s share (Segendorff 1998).

Our paper also relates to the capital budgeting literature where a principal commits to only invest in highly-profitable projects (Antle and Eppen 1985; Harris and Raviv 1990; Baldenius, Dutta and Reichelstein 2007; Baiman, Heinle and Saouma 2013; Bastian-Johnson, Pfeiffer and Schneider 2013; 2017). In this literature there is typically an equilibrium underinvestment whereas in our model there is equilibrium overinvestment because of the CEO’s opportunistic reporting about novel projects.

Lastly, we contribute to the growing literature studying Bayesian persuasion models.⁶ In contrast to this literature, our study focuses on the optimal alignment of interests between senders and receivers from the perspective of another player. Because the board has a learning option, our study relates to the literature incorporating a receiver’s information acquisition.⁷ In our model the sender has unrestricted control over the report

⁶The model was established by Kamenica and Gentzkow (2011) and has been extended to various settings. For example, models with multiple receivers (e.g., Michaeli 2017), multiple senders (e.g., Gentzkow and Kamenica 2017a), interaction between design of public information and disclosure of private information (e.g., Friedman, Hughes and Michaeli 2020, 2021), 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 (Jiang and Stocken 2019), and asset pricing (Cianciaruso, Marinovic and Smith 2020). Earlier studies have also considered information design (e.g., Arya, Glover and Sivaramakrishnan 1997; Göx and Wagenhofer 2009).

⁷More broadly, this is related to dissemination decisions in the presence of external information (Ebert,

and therefore, unlike Huang (2016), we find that the receiver never learns in equilibrium.⁸ This result relates to findings of Matysková and Montes (2021) and Caplin, Dean, and Leahy (2019). Both studies consider discrete states of nature with entropy (variable) costs, whereas we illustrate this result in a continuous state space with fixed costs of learning. Furthermore, unlike all three of these studies, our focus is on the optimal misalignment between the interests of the sender and the receiver as well as the effect of learning costs on this misalignment.

3 Economic Setting

We consider a risk-neutral CEO (“she”) and a risk-neutral board of directors running a firm on behalf of a group of risk-neutral shareholders. The CEO finds a significant investment opportunity (“project”) and decides either to present it ($d = 1$) or not ($d = 0$) to the board for consideration. The board approves ($a = 1$) or rejects ($a = 0$) the implementation of the presented project.⁹

Preferences. The ex-post utility of player $j \in \{S, C, B\}$ (shareholders, CEO, board) is given by

$$v_j(a, d, \theta) = a \cdot d \cdot (\theta - \theta_j), \quad (1)$$

where $\theta \in [\theta_{min}, \theta_{max}]$ is the project value with $-\infty \leq \theta_{min} < 0 < \theta_{max} \leq +\infty$. The player-specific parameter $\theta_j \in [\theta_{min}, \theta_{max}]$ is common knowledge and captures misalignment of interests in a parsimonious way: each player prefers that projects with $\theta \geq \theta_j$ are

Schäfer and Schneider 2019; Frankel, Guttman and Kremer 2020; Michaeli and Wiedman 2021).

⁸The threat of board’s learning disciplines the CEO similar to other models with monitoring threat, e.g., costly state verification in Townsend (1979).

⁹The assumption that the CEO finds a project and presents it to the board for ratification is consistent with empirical evidence that, in most companies, the management is tasked with search and corporate boards with the approval of significant business opportunities (Useem 2006). We briefly discuss the outcome when the decision to undertake the project can be delegated to either player in footnote 22.

undertaken (presented *and* approved) rather than not. The parameter θ_j could reflect the players’ pecuniary/non-pecuniary benefits/costs and investment attitudes.¹⁰ We refer to it as the “player’s type/bias/threshold” and assume the following:

- (i) Without loss of generality, the shareholders’ threshold θ_S is zero (so that the shareholders naturally benefit from projects with positive value and lose from projects with negative value) but, to facilitate comparison across players, we continue to refer to it in the text as θ_S .
- (ii) The CEO’s threshold is lower than that of the shareholders, $\theta_C < \theta_S$. That is, the CEO is an empire-builder and prefers that not only value-enhancing projects but also those with value $\theta \in [\theta_C, \theta_S]$ are presented and approved—occasionally, we refer to the latter as “empire-building projects.”¹¹
- (iii) The board’s threshold can assume any value so that, compared with the shareholders, the board can be neutral/unbiased ($\theta_B = \theta_S$), biased in favor of approval ($\theta_B < \theta_S$), or biased against approval ($\theta_B > \theta_S$). The directors could be “pro-approval” due to close relationship with the CEO (e.g., belong to the same social circle), positive attitude toward investments and company expansion, or private perks. They could be “anti-approval” due to career concerns, conservative/skeptical/antagonistic attitude toward investments and company expansion, or personal disutility related to environmental and social impact.

Project type and information structure. The project could relate to a new technology, business concept or market. In such case, the type of the project is labeled “novel.” Otherwise, it is labeled “routine.” An example of a novel project could be acquisition of a

¹⁰We focus on preference misalignment for given compensation—studying how contracting can shape players’ preferences is beyond the scope of this paper. Concurrent research finds that it is often not optimal to eliminate agents’ biases contractually even if strong monetary incentives are available (Gregor and Michaeli 2021).

¹¹The assumption $\theta_C < \theta_S$ reflects the commonly observed empire-building tendency of CEOs as explained in footnote 1. We discuss the outcome when $\theta_C > \theta_S$ in footnote 26.

start-up that develops autonomous vehicles when the acquirer has no experience with such vehicles' safety and market demand; and an example of a routine project—acquisition of a peer that produces standard cars. The project type is observable. A fraction $p \in (0, 1)$ of the investment opportunities in the economy are routine and the rest are novel.

To capture any potential underlying differences between the two types of projects, we assume that the value of a routine project is drawn from a differentiable cumulative distribution G (with a corresponding probability density function g) and that of a novel project—from a differentiable cumulative distribution F (with a corresponding probability density function f). The joint probability density function of project type and project value is therefore $p \cdot g(\theta)$ if the type is routine and $(1 - p) \cdot f(\theta)$ if the type is novel. Finding a project means drawing a project from the pool of investment opportunities characterized by the joint probability distribution function.

Upon finding a project, the CEO does not observe the value θ and initiates collection of information about the project feasibility. However, if the project is routine, the CEO can not control the properties of the collected information. This could be due to company and/or directors' experience with similar projects. For simplicity, we just say that the value of routine projects is perfectly revealed by the existing information system and is observed by all parties.¹²

If the project is novel, the CEO controls the properties of the collected information—this could be due to the novelty of the technology and the lack of board's familiarity with the industry.¹³ In particular, the CEO chooses a report structure, i.e., a distribution of report realizations and a distribution of the project value conditional on any given report realization. For example, the CEO could start a due diligence process to estimate

¹²Our results qualitatively hold in a setting where the value of a routine project is only estimated with sufficient precision.

¹³For simplicity, we assume that the CEO has no control over the properties of the collected information if the project is routine and full control if the project is novel. In footnote 20, we discuss the outcome under an alternative assumption where the board learns the novel project value with some probability. This could happen either when the board seizes control over the information collection in the company or when the board has a costless private source of information.

the target value, create focus group to evaluate market demand, run an experiment to determine technological feasibility and product safety, or design a medical trial to evaluate drug effectiveness.¹⁴ The CEO’s problem is essentially a persuasion problem, where it is sufficient to consider binary reports with realizations that are supported by disjoint intervals of project values.¹⁵ Specifically, the CEO’s choice of report structure is fully characterized by the choice of a reporting cutoff $\theta_R \in \mathbb{R}$ such that a low report $r = l$ is generated if $\theta \leq \theta_R$ and a high report $r = h$ is generated if $\theta \geq \theta_R$.

As standard in the Bayesian persuasion literature, the report is verifiable and cannot be withheld if the project is presented. This assumption fits well our setting and is *not crucial* for any of the results. In particular, a proposal for an investment project of significant importance/size has to be supported by convincing evidence of the project’s feasibility—e.g., the results of due diligence process, market test, experiment, or drug trial—before being brought to the board for approval (Useem 2006). Once collected, this evidence is available within the company. Thus it is hard for the CEO to conceal or misrepresent it. Data omission or misrepresentation can also be associated with prohibitively harsh legal consequences.¹⁶ To this end, the assumption that the report is verifiable arises naturally. However, it is not critical. In particular, allowing the CEO to privately observe r and withhold it would yield an identical outcome (Gentzkow and Kamenica 2017b; Kamenica 2019). Furthermore, as we show in Appendix A, allowing the CEO to send non-verifiable messages in addition to the optimally constructed public report does not change our results—in equilibrium, the post-report cheap talk communication is uninformative and has no effect on the CEO’s reporting strategy. Interestingly, our analysis in Appendix A implies that the CEO at least weakly prefers committing to verifiable messages and

¹⁴As we elaborate below, the board can collect additional information at a cost later on, after being presented with the project.

¹⁵Under the model assumptions, the CEO is at least weakly better off choosing a binary report structure. See also Friedman, Hughes, and Michaeli (2020, 2021).

¹⁶For example, the former CEOs of Kmart and Kentucky aluminum company faced significant legal charges for providing misleading information to their boards (Peterson 2003; Associated Press 2020).

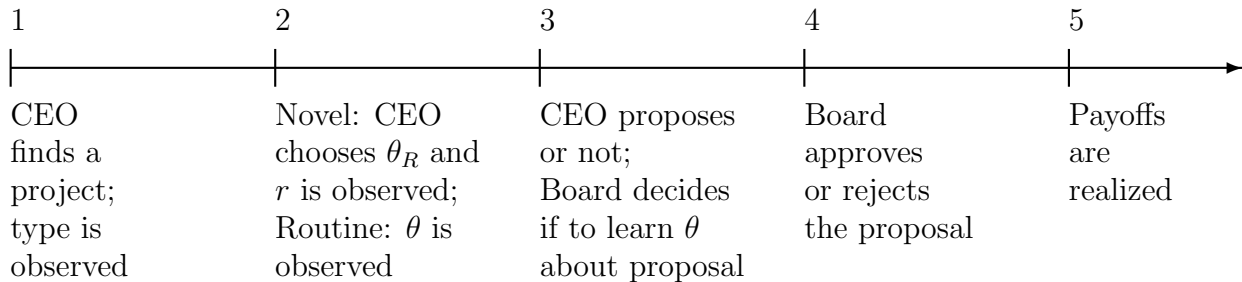


Figure 1: Timeline of the events

opting out of subsequent non-verifiable communication.

Based on the observed information, the CEO decides either or not to present the project to the board. There are no payoff consequences for the CEO for not presenting a project. If a novel project is presented, the board decides either or not to learn θ at a fixed cost $\kappa > 0$.¹⁷

Timeline. Figure 1 presents the timeline of the events. At date 1, the CEO finds a project with observable type (novel or routine). At date 2, if the project is novel, the CEO chooses the reporting cutoff θ_R and an observable report r is generated. If the project is routine, the project value θ is revealed by the existing information system and is publicly observed. At date 3, the CEO decides either or not to present the project to the board. If the novel project is presented for approval, the board decides either or not to learn θ . At date 4, the board approves or rejects the project. At date 5, the payoffs are realized.

4 Analysis Without Board's Learning

Before solving the full-fledged model, we consider a simplified setting where the collection of information by the board is prohibitively costly ($\kappa \rightarrow \infty$). Our aim is to cleanly

¹⁷ We assume that κ is known by the CEO. In footnote 31, we discuss the outcome if the CEO only knows the distribution of boards' learning costs in the economy.

illustrate that anti-approval boards can be beneficial when CEO's information collection is endogenous.

4.1 Project Encounter

We first analyze the case where finding a project means drawing a single project from the pool of investment opportunities characterized by the joint probability distribution function over the project type and value. In other words, the CEO takes into consideration only the first investment opportunity that she encounters.

4.1.1 Board's Approval Decision

The board is willing to approve an encountered and presented routine project if and only if its (revealed by the existing information system) value θ exceeds the board's threshold θ_B . If the presented project is novel, the board grants an approval if its interim (expected at date 4) payoff from an undertaken project, $\mathbb{E}[\theta \mid r, \theta_R] - \theta_B$, exceeds the zero-payoff from rejection. To characterize the report-specific decision a_r , we define the reporting cutoffs $L(\theta_B)$ and $H(\theta_B)$ at which the board is indifferent (between approving and rejecting a presented project) after low and high reports, respectively.¹⁸

Lemma 1 (Board's approval decision for given reporting cutoff). *Let $\kappa \rightarrow \infty$.*

- (i) *When $\theta_B \leq \mathbb{E}[\theta]$, the board rejects the presented novel project if and only if the report is low and the reporting cutoff is below $L(\theta_B)$; that is, $a_h = 1$ and $a_l = \mathbb{1}_{\theta_R \geq L(\theta_B)}$.*
- (ii) *When $\theta_B \geq \mathbb{E}[\theta]$, the board approves the novel project if and only if the report is high and the reporting cutoff exceeds $H(\theta_B)$; that is, $a_l = 0$ and $a_h = \mathbb{1}_{\theta_R \geq H(\theta_B)}$.*

¹⁸Formally, $L(\theta_B)$ is uniquely defined by $\mathbb{E}[\theta \mid r = l, \theta_R = L(\theta_B)] - \theta_B = 0$ and $H(\theta_B)$ by $\mathbb{E}[\theta \mid r = h, \theta_R = H(\theta_B)] - \theta_B = 0$. Because $\mathbb{E}[\theta \mid r = l, \theta_R] \leq \mathbb{E}[\theta]$ and $\mathbb{E}[\theta \mid r = h, \theta_R] \geq \mathbb{E}[\theta]$, the cutoff $L(\theta_B)$ is relevant when $\theta_B \leq \mathbb{E}[\theta]$ and $H(\theta_B)$ —when $\theta_B \geq \mathbb{E}[\theta]$. We provide further details in the appendix.

Figure 2 graphically illustrates Lemma 1 and distinguishes between four problem regions. A board with relatively low threshold, $\theta_B \leq \mathbb{E}[\theta]$, approves the novel project for any report when the reporting cutoff θ_R is large (region \mathcal{P}_1) and approves it only in the event of high report when the cutoff is small (region \mathcal{P}_2). Intuitively, a board with low threshold is more easily convinced to ratify the project. When the reporting cutoff is high, the expected value of the novel project, conditional on either report realization, is high and exceeds the board's threshold—thus the board always approves. However, when the reporting cutoff is small, the expected value of the project, conditional on $r = l$, is too low for approval even by a board with low threshold—as a result, the board ratifies the novel project only if the report is high.

Furthermore, a board with relatively high threshold, $\theta_B \geq \mathbb{E}[\theta]$, approves the novel project if $r = h$ and the cutoff θ_R is high (region \mathcal{P}_3) and rejects it otherwise (region \mathcal{P}_4). All else equal, a board with high threshold is not easily convinced to approve investment opportunities. It ratifies novel projects with expected value that is sufficiently high—this happens only following $r = h$ with a sufficiently high θ_R .

4.1.2 CEO's Reporting and Presentation Decisions

The CEO presents an encountered routine project if and only if its (revealed by the existing information system) value exceeds her threshold θ_C . For a novel project, the problem of the CEO is to choose a reporting cutoff θ_R and report-specific presentation decisions (d_l, d_h) that maximize her payoff. The best possible outcome from the CEO's perspective is when all novel projects with value $\theta \geq \theta_C$ are presented and approved and the ones with value $\theta < \theta_C$ are either not presented or rejected. It can be implemented by setting $\theta_R = \theta_C$ and making sure that (i) $a_l \cdot d_l = 0$ and (ii) $a_h \cdot d_h = 1$ at this reporting cutoff. Ensuring (i) is straightforward—all the CEO needs to do is not present the novel




Figure 2: Board's report-specific approval of a presented novel project for $\kappa \rightarrow \infty$

project if the report is low, i.e., $d_l = 0$.¹⁹ Ensuring (ii) is more challenging. The first necessary condition is that the CEO presents after observing a high report, $d_h = 1$. The second necessary condition is that the board approves after high report, $a_h = 1$. While this is the case for problem regions \mathcal{P}_1 – \mathcal{P}_3 in Figure 2, it is not for \mathcal{P}_4 . In this last region, making sure that the board approves a presented novel project with high report requires increasing the reporting cutoff to the board's indifference point, $\theta_R = H(\theta_B)$.²⁰

Lemma 2 (Reporting cutoff, presentation and approval of novel projects). *Let $\kappa \rightarrow \infty$. At date 2, the CEO chooses a reporting cutoff $\theta_R^* = R(\theta_B)$ where $R(\theta_B) \equiv \max\{\theta_C, H(\theta_B)\}$.*

¹⁹More specifically, when θ_B is in problem regions \mathcal{P}_2 – \mathcal{P}_4 of Figure 2, the CEO is indifferent between $d_l = 1$ and $d_l = 0$ because in any case $a_l = 0$. However, in region \mathcal{P}_1 , the CEO *strictly* prefers not to present projects reported to have value below $\theta_R = \theta_C$ as otherwise the board will approve them.

²⁰In equilibrium, the board approves the presented novel project. This is due to the CEO's full control over the collected information about innovations. If the board could discover the value of a novel project with some probability q then, whenever the board observes $\theta \in (\theta_B, \theta_R)$, it would reject the project—that is, the board would not simply “rubber-stamp” the CEO's proposals. If the probability q is independent of the CEO's signal, the rest of our results remain qualitatively similar.




Figure 3: Optimal reporting cutoff of a novel project for $\kappa \rightarrow \infty$

At date 3, the CEO presents the novel project only if the report is high. At date 4, the board approves the presented project.

A board with threshold $\theta_B < H^{-1}(\theta_C)$ faces a reporting cutoff of $\theta_R^* = \theta_C$. Such board approves all novel projects favored by the CEO. As θ_B increases beyond $H^{-1}(\theta_C)$, the optimal reporting cutoff also increases: the board approves fewer novel projects and is less likely to overinvest. Let θ_H be the board bias (graphically illustrated in Figure 3) associated with reporting cutoff $\theta_R^* = \theta_S$. It is easy to see that this board is biased against approval (because in this region $\theta_R^* = H(\theta_B)$ and $H^{-1}(\theta_S) > \theta_S$) but invests efficiently. Because the optimal reporting cutoff is increasing in θ_B , a board with bias/threshold $\theta_B < \theta_H$ is associated with a reporting cutoff below the shareholders' threshold and overinvests. In contrast, a board with bias $\theta_B > \theta_H$ underinvests. The CEO's empire-building tendency can be completely undone by a board with severe anti-approval bias, $\theta_B = \theta_H > \theta_S$.

Corollary 1. *Let $\kappa \rightarrow \infty$. The CEO's optimal reporting cutoff about novel projects is (weakly) increasing in θ_B and θ_C and is independent of θ_S . There exists a unique value $\theta_H = H^{-1}(\theta_S) > \theta_S$, such that a board with bias $\theta_B < \theta_H$ approves some value-destroying novel projects and a board with bias $\theta_B > \theta_H$ rejects some value-enhancing ones.*

Before we analyze the board bias that maximizes firm value, we briefly summarize the outcomes for both types of projects. Given that both players have veto over the project, the firm undertakes all routine projects with value $\theta \geq \max\{\theta_C, \theta_B\}$ and all innovations with value $\theta \geq \theta_R^* = R(\theta_B) = \max\{\theta_C, H(\theta_B)\}$. Efficient investments for routine projects are thus achieved when the board is neutral/unbiased ($\theta_B = \theta_S > \theta_C$) and for novel ones—when the board is severely biased against approval ($\theta_B = \theta_H > \theta_S > \theta_C$).^{21,22} As we elaborate in Section 2, this result identifies a novel force behind the benefit from misaligned preferences.

4.1.3 Value-Maximizing Board

Taking into account the reporting, presentation and approval decisions, the optimal board bias from the shareholders' perspective, $\tilde{\theta}_B^*$, maximizes the shareholders' welfare $W(\theta_B) \equiv pW^{routine}(\theta_B) + (1-p)W^{novel}(\theta_B)$, i.e., a convex combination of the firm value in case of routine projects, $W^{routine}(\theta_B) \equiv \int_{\max\{\theta_C, \theta_B\}}^{\theta_{max}} (\theta - \theta_S)g(\theta)d\theta$, and the firm value in case of novel projects, $W^{novel}(\theta_B) \equiv \int_{\theta_R^*}^{\theta_{max}} (\theta - \theta_S)f(\theta)d\theta$. It is easy to see from our discussion in the preceding section that, from an ex ante perspective, the optimal board

²¹In line with corporate practice, we assume that there is only one board responsible for both routine projects and innovations and that the board's preferences are consistent across projects. If the board's preferences were to depend on the project type (i.e., one bias/threshold for routine projects and another for innovations), then the optimal board would be biased against approval of novel projects and aligned with shareholders about routine ones. If the shareholders could nominate two specialized boards or even entirely spin off the innovations to a separate entity (intrapreneurship), then it would be optimal to have a neutral board for routine projects and a board that is highly biased against approval for innovations.

²²In this paper, we take delegation to the board as given. It is suboptimal to (i) delegate the decisions about both types of projects to the CEO or (ii) delegate the decisions about novel projects to the CEO and those on routine projects to the board. Under certain conditions (e.g., sufficiently small likelihood that the project is routine), it may be optimal for the shareholders to delegate only novel projects to the board and leave the rest with the CEO. Detailed analysis is available upon request.

has a threshold between θ_S and θ_H . Any other bias is associated with prohibitively large investment inefficiencies: a pro-approval board with $\theta_B < \theta_S$ overinvests and an extremely anti-approval board with $\theta_B > \theta_H$ underinvests *in both types of projects*. This observation allows us to focus solely on the interval $[\theta_S, \theta_H]$ and streamline the analysis in this section.

We first observe that, in the relevant interval, $W^{routine}(\theta_B)$ is a decreasing function. Intuitively, an increase in the board's threshold beyond θ_S leads to underinvestment in routine projects and thereby a decrease in firm value. To analyze the shape of $W^{novel}(\theta_B)$, it is instructive to classify boards into two subsets, depending on the intensity of their conflict of interest with the CEO.

Definition 1 (Conflict of interest). *When the project is novel and board's learning is prohibitively costly ($\kappa \rightarrow \infty$), the conflict of interest between a board with bias θ_B and a CEO with bias θ_C is weak if $\theta_B < H^{-1}(\theta_C)$ and strong otherwise.*

An increase of the board's threshold in the region of weak conflict (where $R(\theta_B)$ is flat) has no effect on the reporting cutoff and the approval decision—thus, in this region, the firm value from a novel project, $W^{novel}(\theta_B)$, remains constant. In contrast, an increase of the threshold in the region of strong conflict (where $R(\theta_B)$ is increasing) increases the reporting cutoff, reduces overinvestment, and increases $W^{novel}(\theta_B)$.²³ An increase of θ_B in the region of weak conflict has no effect on the reporting cutoff and the project decision: the CEO continues to set $\theta_R = \theta_C$ and the board continues to approve all projects with value above this cutoff. As a result, the firm value in the region of weak conflict remains constant. In contrast, an increase of the board's threshold in the regions of strong conflict increases the reporting cutoff, reduces overinvestment, and increases $W^{novel}(\theta_B)$ for any $\theta_B < \theta_H$. However, an increase beyond θ_H leads to underinvestment and decreases $W^{novel}(\theta_B)$. Hence $W^{novel}(\theta_B)$ is maximized when the board has a severe anti-approval

²³Specifically, $\frac{dW^{routine}(\theta_B)}{d\theta_B} = -\theta_B g(\theta_B) < 0$ for $\theta_B > \theta_S$ whereas $\frac{dW^{novel}(\theta_B)}{d\theta_B} = -\theta_C f(\theta_C) \frac{d\theta_C}{d\theta_B} = 0$ when the conflict is weak so that $\theta_R^* = \theta_C$ and $\frac{dW^{novel}(\theta_B)}{d\theta_B} = -H(\theta_B) f(H(\theta_B)) \frac{dH(\theta_B)}{d\theta_B} > 0$ when the conflict is strong so that $\theta_R^* = H(\theta_B)$.

bias, $\theta_B^* = \theta_H$, and is in a strong conflict with the CEO. Then, the information reported by the CEO leads to efficient investment levels.

Our preceding discussion implies that an increase of θ_B in the region where the conflict between the board and the CEO over novel projects is weak is associated with a decrease of welfare $W(\theta_B)$ because the shareholders only incur a “cost” from deterioration in decisions about routine operations. However, in the region where the conflict is strong, the shape of $W(\theta_B)$ depends on the relative magnitude of deterioration in routine projects and a “benefit” from improved decisions about innovations. As a result, the shareholders’ welfare is not necessarily single-peaked. This significantly complicates the identification of the optimal board. We proceed in two steps. First, in Lemma 3, we show that only two types of boards could be optimal: neutral and biased against approval. Second, in Proposition 1, we describe necessary conditions on primitives for either of these types to be optimal.

Lemma 3 (Candidates for optimal board type). *Let $\kappa \rightarrow \infty$. The optimal board type is either (i) neutral and in a weak conflict with the CEO, i.e., $\tilde{\theta}_B^* = \theta_S$ with $R(\theta_B^*) = \theta_C < \theta_S$, or (ii) strongly biased against project approval and in a strong conflict with the CEO, i.e., $\tilde{\theta}_B^* \in (\theta_S, \theta_H)$ with $R(\tilde{\theta}_B^*) \in (\theta_C, \theta_S)$.*

In particular, Lemma 3 finds that a severely biased board ($\theta_B = \theta_H$) is never optimal whereas a neutral board ($\theta_B = \theta_S$) is potentially optimal. To explain, both boards (neutral and severely biased) invest efficiently in one type of project and inefficiently in the other. In both cases, a change in the board type leads to zero marginal distortion of the efficient investment. However, for severely biased board, the change mitigates underinvestments whereas, for neutral board, the change does not mitigate overinvestments whenever the conflict between the CEO and neutral board is weak. This property makes neutral board, in contrast to the severely biased board, locally optimal.

Three empirical implications of Lemma 3 stand out. First, optimal neutral boards are

associated with weak conflicts and optimal biased boards with strong ones. Second, the shareholders always face false approvals of novel projects and may face false rejections of routine ones. Put differently, in equilibrium, investments in routine projects are either efficient or insufficient, but investments in novel projects are always excessive. Third, our model predicts that the distribution of optimal boards is *bimodal* in exogenous parameters and companies can be classified into those with optimally neutral boards and those with optimally anti-approval ones.

The exact conditions for optimality critically depend on the parameters, especially on the shape of the distribution functions F and G . Even without imposing additional restrictions, we can formulate necessary conditions on primitives for the optimality of each of the two candidate board types.

Proposition 1 (Value-maximizing board). *Let $\kappa \rightarrow \infty$.*

(i) *A necessary condition for the optimal board to be neutral is that $R(\theta_S) = \theta_C$.*

(ii) *A necessary condition for the optimal board to be biased against project approval is that $(1-p) [W^{novel}(\theta_H) - W^{novel}(\theta_S)] \geq p [W^{routine}(\theta_S) - W^{routine}(\max\{\theta_S, H^{-1}(\theta_C)\})]$.*

When condition (i) is violated (a neutral board is in a strong conflict with the CEO), condition (ii) is met.²⁴ However, when condition (i) is met (a neutral board is in a weak conflict with the CEO), condition (ii) may or may not be violated.²⁵ Thus, three scenarios exist. The right hand side of the inequality in condition (ii) represents the expected minimal loss from distortions in routine projects committed by an anti-approval board—accordingly, we refer to it as “the minimal cost of establishing a strong conflict with the CEO.” The left hand side of the inequality represents the ex ante gain from alleviated distortions in novel projects generated by a severely biased board that can

²⁴If $R(\theta_S) > \theta_C$, then $\max\{\theta_S, H^{-1}(\theta_C)\} = \theta_S$, and so the right hand side of the inequality is zero while the left hand side is positive.

²⁵Now $R(\theta_S) = \theta_C$ implies $\max\{\theta_S, H^{-1}(\theta_C)\} = H^{-1}(\theta_C) > \theta_S$.

undo the CEO’s empire-building ($\theta_B = \theta_H$), compared with a neutral board. Because this is the maximum gain that can be achieved by an anti-approval board, we refer to it as “the maximal benefit from establishing a strong conflict with the CEO.” For a biased board to ever be optimal, the maximal benefit that shareholders can gain has to exceed the minimal cost. An alternative way to describe this condition is to say that the cost of switching from a weak to a strong conflict, measured by the loss from distorted decisions on routine operations, should not be prohibitively large.

To gain further intuition about condition (ii), consider a neutral board that is in a weak conflict with the CEO. A small increase in the threshold θ_B does not improve the quality of approved novel projects, as the CEO continues to set the reporting cutoff at θ_C , but distorts the quality of approved routine projects, as some projects with value $\theta > \theta_S$ are rejected. As long as $\theta_B < H^{-1}(\theta_C)$, the shareholders only incur costs from distorted approvals of routine operations without gaining any benefits related to novel projects. However, an increase in the board’s threshold beyond $H^{-1}(\theta_C)$ —resulting in a switch from a weak to a strong conflict—raises the quality of approved projects.

For the optimal board to be biased, it is necessary that the cumulative cost of achieving a strong conflict (minimal cost) be smaller than the benefit in eliminating all distortions in novel projects (maximal benefit). Figure 4 panel (a) illustrates a case where the minimal cost is prohibitively large so that the inequality in condition (ii) is violated. In this case, the shareholders’ welfare (in blue) for a board with $\theta_B \neq \theta_S$ is lower than the welfare with a neutral board—thus the optimal board is neutral. In contrast, panel (b) presents a case where the maximal benefit outweighs the minimal cost—the inequality in condition (ii) is satisfied. Because this is only a necessary condition, the fact that it is satisfied still does not imply that the optimal board is biased. For this to happen, the total effect on shareholders’ welfare (in blue) has to exceed the welfare with a neutral board for some θ_B . In the situation of panel (b), this is true so that the optimal board is strongly (but




Figure 4: Evaluation of the necessary condition (ii) in Proposition 1 for biased board

not severely) biased against approval.²⁶

Our next result identifies a condition on the frequency of routine projects, p , that determines whether the necessary condition for the optimal board to be biased is satisfied.

Corollary 2. *Let $\kappa \rightarrow \infty$. There exists a unique value $p^* \in (0, 1]$ such that condition (ii) in Proposition 1 is violated when $p > p^*$ and satisfied otherwise.*

Intuitively, when routine projects occur with sufficiently high frequency, the expected cost from distorted decisions about routine operations is large and outweighs any expected benefit from improved decisions about novel projects. The opposite holds when routine projects are less likely—then the expected benefit exceeds the expected cost.

Condition (ii) of Proposition 1 depends also on the CEO's type, θ_C . While we can say that the condition is more likely met when the CEO has stronger empire-building tendency, little beyond that can be said for a general distribution functions F and G .²⁷ Our next corollary briefly characterizes the optimal board (in closed form) when the values of both routine and novel projects are uniformly distributed.

²⁶In relation to footnote 11, if the CEO were to incur a private cost instead of a benefit (that is, if $\theta_C > \theta_S$) the optimal board bias depends on whether the CEO can veto (not present) projects or not. In the former case the optimal board is always neutral and in the latter is either neutral or has a bias with an opposite sign to the one in our current findings.

²⁷While the left hand side of condition (ii) in Proposition 1 is independent of θ_C , the right hand side is (weakly) increasing in θ_C for any distribution functions F and G . Thus lower θ_C (stronger empire-building) makes it easier for the inequality to be satisfied.

Corollary 3 (Optimal board under a uniform distribution). *Let $\kappa \rightarrow \infty$, $\theta_S = 0$, $\theta_{max} - \theta_{min} = 1$ and suppose that the values of routine and novel projects are uniformly distributed, $f(\theta) = g(\theta) = 1$. The optimal board is anti-approval with $\tilde{\theta}_B^* = \frac{2(1-p)}{4-3p}\theta_{max} \in (\theta_S, \theta_H)$ if $\theta_C \leq -\sqrt{\frac{p}{4-3p}}\theta_{max}$. Otherwise, the optimal board is neutral with $\tilde{\theta}_B^* = \theta_S = 0$.*

In line with our findings in Lemma 3 and Proposition 1, the optimal board when project values are uniformly distributed is neutral, $\tilde{\theta}_B^* = \theta_S$, or strongly biased against project approval, $\tilde{\theta}_B^* = \frac{2(1-p)}{4-3p}\theta_{max} \in (\theta_S, \theta_H) = (0, \frac{\theta_{max}}{2})$. The latter optimal board is more strongly biased (high $\tilde{\theta}_B^*$) when the share of novel projects is large (low p) and when project values can assume large values (high θ_{max}). Under these conditions, mitigating distortions in novel projects is all the more important—and a more biased board is more capable of doing so. Furthermore, Corollary 3 shows that the optimal board is biased if the CEO is sufficiently keen on project adoption, $\theta_C \leq -\sqrt{\frac{p}{4-3p}}\theta_{max}$. This is intuitive: a CEO with severe empire-building tendency can persuade the board to approve more value-destroying novel projects. This increases the expected benefit from alleviating distortions in novel projects and makes the biased board more attractive from the shareholders' perspective.

4.2 Project Search

We now consider a scenario where the CEO is not restricted to the first encountered project. Formally, at date 1 (before the report is generated or the project value is observed), the CEO has a costless option to make additional independent draws from the pool of investment opportunities. As a consequence, the CEO searches until she finds the type of project that she prefers.²⁸ In general, the CEO considers two aspects: (i) For any

²⁸It is also possible that the CEO prefers a specific project type but her search for it is limited. For example, the CEO is allowed only a limited number of draws from the pool of opportunities. This is equivalent to the scenario with a single draw and a shift in the marginal probability p in the joint distribution function in the CEO's preferred direction. Our scenario with an unlimited number of draws is thus equivalent to the scenario with a single draw and the CEO's selection of the marginal probability p to be either zero (when novel project is preferred) or one (when routine project is preferred).

board, the novel type is more manipulable and leads to a lower threshold for acceptance than the routine type, $\theta_R = R(\theta_B) < \theta_B$. (ii) The project types may differ in profitability. To introduce the difference in profitability independently on the threshold of acceptance, we rank the distributions by (first-order) stochastic dominance.

The board's approval and the CEO's reporting and presentation decisions remain as before. At date 1, the CEO searches for a novel project if her margin from novel project, $M(\theta_B) \equiv \mathbb{E}[(\theta - \theta_C)\mathbb{1}_{\theta \geq R(\theta_B)} \mid \text{novel}] - \mathbb{E}[(\theta - \theta_C)\mathbb{1}_{\theta \geq \max\{\theta_C, \theta_B\}} \mid \text{routine}]$, is positive and searches for a routine project one otherwise.

Proposition 2 (Optimal board with first-order stochastic dominance). *Let $\kappa \rightarrow \infty$.*

- (i) *If $F(\theta) \leq G(\theta)$ for any θ , then $\theta_B^* = \theta_H$ and the CEO pursues a novel project.*
- (ii) *If $F(\theta) > G(\theta)$ for any θ and $M(\theta_S) \leq 0$, then $\theta_B^* = \theta_S$ and the CEO pursues a routine project.*
- (iii) *If $F(\theta) > G(\theta)$ for any θ and $M(\theta_S) > 0$, then $\theta_B^* \neq \theta_S$ and $\theta_B^* \notin [\theta_{\min}, \theta_C]$.*

Because of her ability to opportunistically choose the reporting cutoff of novel but not routine projects, the CEO is typically more willing (than the shareholders) to pursue innovations.²⁹ As a result, there is an asymmetry: When novel projects are at least as profitable as routine ones (in a sense of stochastic dominance), the CEO always pursues the former. But when routine projects are more profitable than novel ones, the CEO's choice is unclear.

Specifically, when the values of the projects follow the same distribution or when the value of novel projects (first-order) stochastically dominates that of routines, as is the case of Proposition 2 (i), the preferences of the board and the CEO about project type

²⁹To see why, note that the CEO's margin can be simplified to $M(\theta_B) = W^{\text{novel}}(\theta_B) - W^{\text{routine}}(\theta_B) + (\theta_S - \theta_C)[G(\max\{\theta_B, \theta_C\}) - F(R(\theta_B))]$. The preference misalignment about project selection arises only because of the third term, which reflects the difference in probability of obtaining a private benefit by the CEO.

are aligned at $\theta_B = \theta_H$. Thus the optimal board has a threshold of $\theta_B^* = \theta_H$, and the CEO pursues a novel project. Because the board can fully undo the CEO's empire-building, the equilibrium investment is efficient. When the value of routine projects stochastically dominates that of novel ones, two scenarios may arise. If the players' preferences at $\theta_B = \theta_S$ about project selection are aligned, as is the case of Proposition 2 (ii), the CEO decides to search for a routine project, the optimal board is neutral, $\theta_B^* = \theta_S$, and the equilibrium investment is efficient. However, if the preferences at $\theta_B = \theta_S$ are misaligned, as in part (iii) of the proposition, the optimal board can never be neutral. This is intuitive: when faced with a neutral board, the CEO pursues a novel project. However, conditional on the project being novel, a neutral board is suboptimal for the shareholders.

So what is the optimal board in the scenario of Proposition 2 (iii)? There are two options. First, the shareholders could acquiesce to the CEO's tendency to pursue a novel project and nominate a board with anti-approval bias to counteract empire-building; unless $M(\theta_H) < 0$, the board has severe anti-approval bias, $\theta_B = \theta_H$. Second, the shareholders could persuade the CEO to search for a routine project by nominating a board whose approval threshold makes routine projects attractive from the CEO's perspective. The best option depends on the relative magnitude of the shareholders' benefit from pursuing a routine project (stochastically dominating) project, rather than a novel (stochastically dominated) one, and the shareholders' loss from investment distortions by a board that is biased enough to influence the CEO's project selection. Due to the generality of our model, there is little we can say about the board threshold that sways the CEO into pursuing a routine project, apart from it not being neutral (as argued above) or more pro-approval than the CEO herself (as this introduces investment distortions that can be mitigated by a less pro-approval board).³⁰

³⁰One could design conditions under which the optimal board is even mildly biased in favor of approval to sway the CEO into selecting a routine project. This, for example, can occur if (i) $M(\theta) > 0$ for any $\theta \in [\theta_S, \theta_H]$; and (ii) $W^{routine}(\theta_C) > \max\{W^{novel}(\theta_H), W^{routine}(\theta_H)\}$ hold simultaneously. In this case, the optimal board has $\theta_B^* \in (\theta_C, \theta_S)$, and the CEO selects a routine project in equilibrium. To see why, note that the imposed condition (i) rules out the possibility that the CEO selects a routine project when

5 Full-Fledged Model with Board's Learning

We now reintroduce the full-fledged model where learning by the board is not prohibitively costly but return to the assumption that the CEO takes into consideration only the first investment opportunity that she encounters.

At date 4, a board that has learned θ (either because the project is novel and the existing information system revealed its value or because the board choose to learn at cost κ) approves a presented project as long as $\theta \geq \theta_B$. A board that does not know θ (because the project is novel and the board decided not to learn) makes a project decision that is described in Lemma 1 and graphically illustrated in Figure 2, where we distinguish between four problem regions.

At date 3, after observing r and prior to approving or rejecting the presented novel project, the board decides either or not to learn θ at cost $\kappa > 0$. (Clearly, there is nothing to be learned about a routine project.) In each problem region, learning is beneficial only for one of the two reports because, from the board's perspective, the decision following the other report is error-free. In other words, only one report is relevant. The board learns θ if and only if the expected benefit of learning, conditional on the relevant report, exceeds the cost of learning:

$$\mathbb{E}[X(\theta, \theta_B) \mid \theta \in [\underline{\theta}, \bar{\theta}], r] \geq \kappa, \quad (2)$$

where $[\underline{\theta}, \bar{\theta}]$ is an interval of project values at which the original (without learning) decision a_r is revised into the alternative (with learning) decision $1 - a_r$, and $X(\theta, \theta_B) \equiv |\theta - \theta_B|$ is the board's benefit of revision at θ . Table 1 summarizes the relevant values for each of the problem regions.

faced with $\theta_B \in [\theta_S, \theta_H]$. Moreover, the imposed condition (ii) rules out the optimality of any board with $\theta_B > \theta_H$ but also ensures that the board prefers to sway the CEO into selecting a routine project. Lastly, by Proposition 2, the board is never (weakly) more pro-approval than the CEO.

Region	Board bias	Reporting cutoff	Relevant report	a_r	$[\underline{\theta}, \bar{\theta}]$
\mathcal{P}_1	$\theta_B \in [\theta_{min}, \mathbb{E}[\theta]]$	$\theta_R \in (L(\theta_B), \theta_{max}]$	$r = l$	1	$[\theta_{min}, \theta_B]$
$\mathcal{P}_2, \theta_R > \theta_B$	$\theta_B \in [\theta_{min}, \mathbb{E}[\theta]]$	$\theta_R \in (\theta_B, L(\theta_B)]$	$r = l$	0	$[\theta_B, \theta_R]$
$\mathcal{P}_2, \theta_R < \theta_B$	$\theta_B \in [\theta_{min}, \mathbb{E}[\theta]]$	$\theta_R \in [\theta_{min}, \theta_B)$	$r = h$	1	$[\theta_R, \theta_B]$
$\mathcal{P}_3, \theta_R > \theta_B$	$\theta_B \in (\mathbb{E}[\theta], \theta_{max}]$	$\theta_R \in (\theta_B, \theta_{max}]$	$r = l$	0	$[\theta_B, \theta_R]$
$\mathcal{P}_3, \theta_R < \theta_B$	$\theta_B \in (\mathbb{E}[\theta], \theta_{max}]$	$\theta_R \in [H(\theta_B), \theta_B)$	$r = h$	1	$[\theta_R, \theta_B]$
\mathcal{P}_4	$\theta_B \in (\mathbb{E}[\theta], \theta_{max}]$	$\theta_R \in [\theta_{min}, H(\theta_B))$	$r = h$	0	$[\theta_B, \theta_{max}]$

Table 1: Board revisions when the board learns

Two observations arise. First, the board decides about learning only if the CEO presents the novel project. In the absence of board’s ability to learn, the CEO presents the novel project if the report is high and vetoes the project if the report is low. This property describes an optimal CEO’s presentation strategy also when the board is able to learn. If the CEO doesn’t veto the project for a low report and lets the board learn, all novel projects with $\theta \geq \theta_B$ are approved. But the CEO can achieve this (board’s most preferred) outcome also by setting a reporting cutoff $\theta_R = \theta_B$ and not presenting the project if the report is low. As a result, when multiple CEO’s optimal reporting and presentation strategies exist, we can select the one in which the CEO never lets the board learn when the relevant report is low. This happens whenever $\theta_R > \theta_B$, which defines a nonlearning region \mathcal{N}_A , graphically illustrated in Figure 5. (We label the areas where the board does not learn—either due to the CEO’s project veto or due to the board’s unwillingness to bear the cost of learning—“nonlearning regions.”)

Second, when $\theta_R \leq \theta_B$, we distinguish between board’s learning decisions in \mathcal{P}_4 and otherwise (\mathcal{P}_2 or \mathcal{P}_3). The existence of nonlearning region in \mathcal{P}_4 , denoted \mathcal{N}_C , depends on the level of the learning cost. If κ is sufficiently small, the board always learns, i.e., \mathcal{N}_C is empty. Otherwise, the board does not learn when θ_B is in the neighborhood of θ_{max} and θ_R is in the neighborhood of θ_{min} , as there the expected benefit of learning is relatively low; i.e., $\mathcal{N}_C \subseteq \mathcal{P}_4$.

Outside of \mathcal{P}_4 where $\theta_R \leq \theta_B$, the board approves the novel project after a high




Figure 5: Nonlearning regions for novel projects with a moderate cost κ

report—thus, absent learning, the board approves all projects with a value that exceeds θ_R . The board’s learning decision is thus a decision about whether to follow the distorted high report characterized by cutoff θ_R or learn at a cost and shift the approval cutoff up to θ_B . The benefit of learning arises for novel projects with values $\theta \in [\theta_R, \theta_B]$. Lemma 4 proves that a cutoff $\hat{H}(\theta_B)$ exists on a subinterval of $\theta_B \in [\theta_{min}, \theta_{max}]$ such that the board is just indifferent between learning and not learning. It also proves that the board does not learn when its threshold is close to the reporting cutoff, $\theta_B \in [\theta_R, \hat{H}^{-1}(\theta_R)]$. This nonlearning region is denoted \mathcal{N}_B . It includes also the diagonal $\theta_R = \theta_B$, where the expected benefit of learning is exactly zero. This is intuitive—if the reporting cutoff coincides with the approval threshold of the board, then knowing the true value (and not just whether it is above or below that cutoff) will not change the board’s decision.

Lemma 4 (Nonlearning regions). *When the CEO’s presentation strategy is optimal, the board does not learn if $\theta_B < \theta_R$ (nonlearning region \mathcal{N}_A) or if $\theta_B \geq \theta_R$ and either (i) $\theta_B \geq$*

$\min\{H^{-1}(\theta_R), \hat{N}^{-1}(\theta_R)\}$ (nonlearning region \mathcal{N}_C) or (ii) $\theta_B \leq \min\{H^{-1}(\theta_R), \hat{H}^{-1}(\theta_R)\}$ (nonlearning region \mathcal{N}_B), where $\hat{H}(\theta_B)$ and $\hat{N}(\theta_B)$ are defined such that $\theta_R = \hat{H}(\theta_B)$ and $\theta_R = \hat{N}(\theta_B)$ satisfy condition (2) in the respective problem region with equality.

As in Section 4 (where we analyze the optimal reporting cutoff in the absence of learning option), the CEO chooses the reporting cutoff $\hat{\theta}_R^*$ as close as possible to θ_C subject to the constraint that the board approves the presented project after a high report. The difference is that now the CEO additionally ensures the board does not learn—this may shift the reporting cutoff up. Lemma 5 characterizes these changes, and Figure 6 graphically illustrates the reporting cutoff (in bold blue).

Lemma 5 (CEO’s choice of reporting cutoff when the board has a learning option). *At date 2, the CEO chooses a reporting cutoff $\hat{\theta}_R^* = \hat{R}(\theta_B)$ and presents the novel project if $r = h$, where $\hat{R}(\theta_B) \equiv \max\{\theta_C, H(\theta_B), \hat{H}(\theta_B)\} \geq R(\theta_B)$. At date 3, the board does not learn. At date 4, the board approves the presented project.*

Why does the CEO discourage the board from learning? The reason is very intuitive. A board informed about the value of presented project cannot be persuaded by the CEO to approve empire-building projects: no matter what the reporting cutoff is, such a board approves only projects with value $\theta > \theta_B$. The same outcome can be achieved with an uninformed board if the CEO sets the reporting cutoff at $\theta_R = \theta_B$. However, this is suboptimal, as the CEO can do better by shifting the reporting cutoff closer to θ_C while preventing learning by the board; this option is always possible since $\hat{H}(\theta_B) < \theta_B$.³¹

Ensuring that the board does not learn forces the CEO to change the reporting cutoff. Specifically, the CEO (weakly) increases the reporting cutoff, $\hat{\theta}_R^* = \hat{R}(\theta_B) \geq R(\theta_B) = \theta_R^*$. As a consequence, the board’s ability to learn also affects the shareholders’ optimal board. From the shareholders perspective, the most desired outcome for novel projects is when

³¹Our result that the board never learns depends on two assumptions: (i) the CEO knows κ and (ii) the CEO can choose any reporting threshold. If either of these assumptions is not satisfied, the board may end up learning and rejecting some of the proposed projects.




Figure 6: Optimal reporting cutoff when the board has a learning option

the board is biased against approval and has a threshold

$$\hat{\theta}_H = \hat{R}^{-1}(\theta_S) = \min\{H^{-1}(\theta_S), \hat{H}^{-1}(\theta_S)\} \leq H^{-1}(\theta_S) = \theta_H. \quad (3)$$

This threshold is (weakly) increasing in κ . When the board bias is given, the introduction of a learning option thus weakly reduces the incidence of false approvals (mitigates overinvestment) when $\theta_B \leq \hat{\theta}_H$ as then the cutoff for project approvals shifts closer to the optimal cutoff θ_S , $R(\theta_B) \leq \hat{R}(\theta) \leq \theta_S$. However, the board's learning option results in weakly more false rejections (introduces or enhances underinvestment) if $\theta_B > \hat{\theta}_H$, or equivalently if $\hat{R}(\theta) > \theta_S$.

Lemma 6. *For given θ_B , the presence of board's learning option mitigates investment distortion if $\theta_B \in [\theta_{\min}, \hat{\theta}_H]$. Otherwise, it amplifies investment distortion.*

Before concluding, we comment on how board's learning affects the optimal board type. Now the interval of relevant board types is narrower, $[\theta_S, \hat{\theta}_H] \subseteq [\theta_S, \theta_H]$. In addition,

the firm value in the case of a novel project changes from $W^{novel}(\theta_B) = \int_{R(\theta_B)}^{\theta_{max}} (\theta - \theta_S) f(\theta) d\theta$ to $\widehat{W}^{novel}(\theta_B) \equiv \int_{\widehat{R}(\theta_B)}^{\theta_{max}} (\theta - \theta_S) f(\theta) d\theta$. These changes do not affect qualitatively our findings that there are only two candidates for the optimal board and that there is an association between a neutral/biased board and a weak/strong conflict of interest in equilibrium. Notably, a higher learning cost increases the minimal cost of establishing a conflict while keeping the maximal benefit of a biased board constant. Thus the condition for a biased board is more likely violated, and the board is more likely neutral.

Proposition 3 (Learning cost, optimal board type, and welfare). *An increase in the learning cost, κ , increases the likelihood that the optimal board is neutral and decreases the shareholders' welfare.*

The level of κ determines the ability of the board to learn. For example, a low cost implies that the board can easily uncover investment-relevant information—this could be either because the firm operates in a more mature industry with a lot of available data or because the directors are more experienced. And the reverse is also true: a board facing high cost has a difficult time uncovering information on its own. Our result in Proposition 3 therefore predicts that boards of firms operating in more mature industries or boards with more expertise are more likely biased.

The effect of learning cost on the optimal level $\widehat{\theta}_B^*$ when the board is already biased is hard to sign as the shape of the nonlearning constraint (ensuring that the board does not learn) depends on the distribution of the project value. Nevertheless, the asymptotic effects are clear: as $\kappa \rightarrow \infty$, the (biased) board type converges to the level without learning, $\lim_{\kappa \rightarrow \infty} \widehat{\theta}_B^* = \theta_B^* \geq \theta_S$, whereas as $\kappa \rightarrow 0^+$ and learning is cost-free, the board is always informed and so it is best for it to be neutral, $\lim_{\kappa \rightarrow 0^+} \widehat{\theta}_B^* = \theta_S$. To summarize, as learning cost decreases, (i) the board is *more likely* biased against project approval (expertise and anti-approval bias are complements), but (ii) if the board is already biased, its optimal level of bias *decreases* asymptotically (expertise and anti-approval bias are

substitutes). Together, observations (i) and (ii) imply that the optimal board type might be *nonmonotonic*.³²

The effect of board’s learning on the CEO’s welfare (at the optimal board type) is complex. On the one hand, higher learning cost relaxes the nonlearning constraint on the CEO’s optimal report, and so the CEO is more likely to persuade the board to approve the projects she prefers. On the other hand, when κ is high, the optimal board is more biased against approval, which hurts the CEO due to lower acceptance rate of routine projects. Whether the CEO benefits from higher learning costs depends on the relative magnitude of the two effects, which in turn depends on the distributions F and G . Our last result illustrates that, when project values are uniformly distributed, learning by the board is beneficial even for the CEO if, paradoxically, the CEO has a severe tendency toward empire-building.

Corollary 4 (CEO’s preference over board’s learning with uniform distribution). *Suppose that $p \in (0, 1)$, $\theta_{max} - \theta_{min} = 1$, $\theta_S = 0$ and $f(\theta) = g(\theta) = 1$. The CEO strictly prefers board’s cost-free learning, $\kappa = 0$, over a prohibitively large learning cost if $\theta_C \leq -\sqrt{\frac{p}{4-3p}}\theta_{max}$.*

6 Concluding Remarks

Our paper studies the optimal bias of corporate boards tasked with the approval of investment opportunities proposed by empire-building CEOs. In line with the empirical evidence (e.g., Maffett, Nakhmurina, and Skinner 2020), we find that reducing the alignment of interests between CEOs and boards may improve investment efficiency and identify a novel force behind the benefits of misaligned preferences. Accounting for heterogeneity of project types, we also predict a bimodal distribution of boards in the

³²To see why, consider a set of parameters such that the optimal board is neutral without learning (for $\kappa \rightarrow \infty$) but, because of observation (i), is biased with learning (for interior κ). By observation (ii), as $\kappa \rightarrow 0^+$, the optimal board type converges back to neutral.

economy: a peak where boards are neutral (aligned with shareholders) and a peak where they are strongly biased against approval (e.g., consisting of directors with conservative/antagonistic attitude). We expect that boards in firms with a large share of novel projects and managed by CEOs with strong tendencies toward empire-building are less likely neutral. Our model also predicts that firms overinvest in novel projects and sometimes underinvest in routine operations. The effect of directors' expertise is nontrivial: even though our results imply that boards rely on the information provided by management, directors' expertise disciplines the reporting opportunism of CEOs and improves investment efficiency. We anticipate that boards with higher expertise are more likely but less severely biased. These results provide new testable predictions about the links between board composition and investment efficiency.

Our paper also illuminates the ongoing debate about board independence. Prior analytical research suggests that CEOs are less forthcoming about exogenously acquired non-verifiable information when faced with directors whose preferences are not aligned with theirs—a finding that may raise concerns about unintended consequences of board independence requirements and regulations that facilitate shareholder activism (for nomination of cost-cutting directors). We show that CEOs prefer to commit to communicating verifiable information. Then, nominating a neutral or anti-approval board is optimal for shareholders so our study highlights positive effect of such requirements.

Appendix

A Robustness: Cheap Talk

In the main analysis, we solely focus on the CEO's collection of publicly observable information. In practice, after preparing a formal report for a board meeting, CEOs may come across additional information and informally communicate it to the board. In this extension, we briefly demonstrate that our main results are robust to accounting for such possibility. In particular, we consider a setting where, after the public report is observed and the project is proposed, the CEO privately learns the novel project value and sends a non-verifiable message to the board (cheap talk communication).³³

Proposition 4 (Robustness of the results with cheap talk). *Suppose that, after the public report is observed and the project is proposed, the CEO privately learns the novel project value and sends a non-verifiable message to the board. Then, the CEO's equilibrium non-verifiable message is uninformative and the equilibrium reporting cutoff is $\theta_R^* = R(\theta_B)$ as defined in Lemma 2.*

In equilibrium, the post-report cheap talk communication is uninformative and has no effect on the CEO's reporting strategy.³⁴ Put differently, the presence of cheap talk option preserves the communicated hard information without conveying additional soft information. The intuition is simple: any further refinement of the information contained in the message $r = h$ (i.e., information $\theta \geq \theta_R = R(\theta_B)$) increases the payoff of the board and reduces that of the CEO. Our model predicts that the CEO at least weakly prefers committing to verifiable messages and essentially opting out of subsequent non-verifiable communication.³⁵ This is because commitment creates credibility of the communicated information. Given Proposition 4, our findings remain intact in a setting where the CEO privately learns the novel project value (after setting the reporting cutoff) and sends a non-verifiable message to the board.³⁶

B Proofs

Proof of Lemma 1: To begin, let $f_r(\theta)$ be the probability density function after $r \in \{l, h\}$ is observed. This probability is obtained by truncation of the prior distribution F at the

³³The timing of the public report realization does not matter.

³⁴Essentially, the CEO's message only confirms that the project value belongs to the subinterval associated with the realized report.

³⁵If the conflict of interests between the board and the CEO is weak ($\theta_B < H^{-1}(\theta_C)$ and thus $\theta_R = \theta_C$), the CEO is indifferent between (i) setting $\theta_R = R(\theta_B) = \theta_C$ and $(d_l, d_h) = (0, 1)$ and opting out from subsequent cheap talk; or (ii) setting $\theta_R \in [\theta_{min}, \theta_{max}]$ and $(d_l, d_h) = (1, 1)$ and leaving all communication to the cheap talk. If the conflict is strong, option (i) is strictly preferred by the CEO.

³⁶If the CEO cannot choose the properties of the publicly observable report and can only send a non-verifiable message to a board that takes a binary action (accept/reject), the optimal board from shareholders' perspective would be neutral/unbiased. Analysis available upon request.

reporting cutoff θ_R , because after $r = h$ the board realizes that the project value is in $[\theta_R, \theta_{max}]$, and after $r = l$ the board realizes that the project value is in $[\theta_{min}, \theta_R]$:

$$f_h(\theta) = \begin{cases} 0, & \text{if } \theta \in [\theta_{min}, \theta_R), \\ \frac{f(\theta)}{1-F(\theta_R)}, & \text{if } \theta \in [\theta_R, \theta_{max}]. \end{cases}$$

$$f_l(\theta) = \begin{cases} \frac{f(\theta)}{F(\theta_R)}, & \text{if } \theta \in [\theta_{min}, \theta_R], \\ 0, & \text{if } \theta \in (\theta_R, \theta_{max}]; \end{cases}$$

Now note that the expected at date 4 project value (after observing a report r with cutoff θ_R) is as follows:

$$\mathbb{E}[\theta \mid r, \theta_R] = \int_{\theta_{min}}^{\theta_{max}} \theta f_r(\theta) d\theta.$$

The board's interim (expected at date 4) payoff from approval is then,

$$\mathbb{E}[v_B(a = 1, \theta) \mid r, \theta_R] = \int_{\theta_{min}}^{\theta_{max}} \theta f_r(\theta) d\theta - \theta_B. \quad (4)$$

When deciding whether to approve the project, the board compares the interim payoff in (4) with the zero-payoff from project rejection. Thus, the approval decision of the board is report-specific and is given by

$$a_r = \mathbb{1}_{\int_{\theta_{min}}^{\theta_{max}} \theta f_r(\theta) d\theta - \theta_B \geq 0}, \quad \forall r \in \{l, h\}$$

where the indicator function equals 1 if the board's interim payoff from approval is non-negative and equals zero otherwise. Because the board's interim payoff from approval is increasing in θ_R and decreasing in θ_B (for any report r), a board facing a higher (given) reporting cutoff θ_R (i.e., a higher expected project value conditional on the report) and/or a board with lower bias θ_B is more likely to approve the project proposed by the CEO.

For given θ_R , the range of the expected project value when $r = l$ is $[\theta_{min}, \mathbb{E}[\theta]]$ (see $\mathbb{E}[\theta \mid r = l, \theta_R] \leq \mathbb{E}[\theta]$), whereas the range of the expected project value when $r = h$ is $[\mathbb{E}[\theta], \theta_{max}]$ (see $\mathbb{E}[\theta \mid r = h, \theta_R] \geq \mathbb{E}[\theta]$). The two intervals overlap only in the expected value of the project, $\mathbb{E}[\theta]$. This implies that $\theta_B \lesseqgtr \mathbb{E}[\theta]$ determines the range of the indicator function.

Exploiting this observation, we characterize the board's indifference over the project approval in the two-dimensional space of cutoffs, $(\theta_B, \theta_R) \in [\theta_{min}, \theta_{max}]^2$. First, consider the low report realization. The function $L : [\theta_{min}, \mathbb{E}[\theta]] \rightarrow [\theta_{min}, \theta_{max}]$ that yields cutoff θ_R such that the board is indifferent for a low report, is implicitly characterized by $\int_{\theta_{min}}^{\theta_{max}} \theta f_l(\theta) d\theta = \theta_B$ such that $\theta_R = L(\theta_B)$. Equivalently, $L^{-1}(\theta_B) = \int_{\theta_{min}}^{\theta_{max}} \theta f_l(\theta) d\theta$. Similarly, consider the high report realization case and specify a function $H : [\mathbb{E}[\theta], \theta_{max}] \rightarrow [\theta_{min}, \theta_{max}]$ that yields a cutoff θ_R such that the board is indifferent for a high report. The function is implicitly characterized by $\int_{\theta_{min}}^{\theta_{max}} \theta f_h(\theta) d\theta = \theta_B$ such that $\theta_R = H(\theta_B)$. Equivalently, $H^{-1}(\theta_B) = \int_{\theta_{min}}^{\theta_{max}} \theta f_r(\theta) d\theta$.

To characterize a_r , we exploit that the board's interim payoff from approval is increasing in θ_R . For $\theta_B \leq \mathbb{E}[\theta]$, consider a low report realization. The board is indifferent at $\theta_R = L(\theta_B) \in [\theta_{min}, \theta_{max}]$. We have $a_l = 0$ if $\theta_R \leq L(\theta_B)$ and $a_l = 1$ if $\theta_R \geq L(\theta_B)$. For a high report realization, observe that $a_h = 1$ because

$$H^{-1}(\theta_R) - \theta_B \geq H^{-1}(\theta_R) - \mathbb{E}[\theta] \geq H^{-1}(\theta_{min}) - \mathbb{E}[\theta] = 0.$$

For $\theta_B \geq \mathbb{E}[\theta]$, consider a high report realization. The board is indifferent at $\theta_R = H(\theta_B) \in [\theta_{min}, \theta_{max}]$. We have $a_h = 0$ if $\theta_R \leq H(\theta_B)$ and $a_h = 1$ if $\theta_R \geq H(\theta_B)$. For a low report realization, observe that $a_l = 0$ because

$$L^{-1}(\theta_R) - \theta_B \leq L^{-1}(\theta_R) - \mathbb{E}[\theta] \leq L^{-1}(\theta_{max}) - \mathbb{E}[\theta] = 0.$$

Proof of Lemma 2: The proof follows from the discussion in the text.

Proof of Corollary 1: The properties of the CEO's optimal reporting cutoff θ_R^* are given by the properties of R function in Lemma 2. Uniqueness of θ_H follows from the fact that H function (and also its inverse H^{-1} function) is increasing. The observation that any $\theta_B \neq \theta_H$ distorts novel investments follows from the fact that a conflict between shareholders and CEO, $\theta_S > \theta_C$, implies that R function is in its increasing part, $R(\theta_H) = H(\theta_H) > \theta_C$.

Proof of Lemma 3: Three separate claims eliminate all other candidates for the optimal board type.

Claim 1. *If the optimal board induces a weak conflict over novel project, it must be neutral: $R(\theta_B^*) = \theta_C \Rightarrow \tilde{\theta}_B^* = \theta_S$*

Suppose not: $\theta_S < \tilde{\theta}_B^*$. Then, from the shape of $R(\theta_B)$, a weak conflict is induced for any $\theta_B \in [\theta_S, \tilde{\theta}_B^*]$. (By Definition 1 and Lemma 2, the weak conflict exists for relevant board types, $\theta_B \in [\theta_S, \theta_H]$, on an interval $\theta_B \in [\theta_S, H^{-1}(\theta_C)]$.) On this interval, however, $W^{routine}$ is decreasing in θ_B whereas W^{novel} is constant in θ_B , and therefore the maximizer is the minimal board type θ_S , which contradicts $\theta_S < \tilde{\theta}_B^*$.

Claim 2. *If the optimal board induces a strong conflict over novel project, it cannot be a severely biased board: $R(\tilde{\theta}_B^*) > \theta_C \Rightarrow \tilde{\theta}_B^* < \theta_H$*

To begin, note that a severely biased board induces a strong conflict always, $R(\theta_H) = \theta_S > \theta_C$; by continuity of $R(\theta_B)$ and $\theta_S - \theta_C > 0$, it is inducing a strong conflict also on a left neighborhood of θ_H . We will now analyze $W(\theta_B)$ on this neighborhood. We know that the marginal benefit for novel projects decreases to zero when $\tilde{\theta}_B^*$ approaches θ_H , $\lim_{\theta_B \rightarrow \theta_H^-} \frac{dW^{novel}(\theta_B)}{d\theta_B} = 0$. At the same time, the marginal cost for routine projects is negative, $\lim_{\theta_B \rightarrow \theta_H^-} \frac{dW^{routine}(\theta_B)}{d\theta_B} = -\theta_H g(\theta_H) < 0$. To combine, the shareholders' welfare is decreasing when all value-destroying novel projects are eliminated, $\lim_{\theta_B \rightarrow \theta_H^-} \frac{dW(\theta_B)}{d\theta_B} = -p\theta_H g(\theta_H) < 0$, and therefore the optimal board is $\tilde{\theta}_B^* < \theta_H$.

Claim 3. *If the optimal board induces a strong conflict over novel project, it cannot be neutral: $R(\theta_B^*) > \theta_C \Rightarrow \tilde{\theta}_B^* > \theta_S$*

Suppose not: $\tilde{\theta}_B^* = \theta_S$ and $R(\theta_S) > \theta_C$. Then, by the shape of $R(\theta_B)$, a strong conflict is induced for any relevant board type, $\theta_B \in [\theta_S, \theta_H]$. At $\theta_B = \theta_S$, the marginal cost for routine projects is zero, $\left. \frac{dW^{routine}(\theta_B)}{d\theta_B} \right|_{\theta_B=\theta_S} = -\theta_S g(\theta_S) = 0$. In contrast, the marginal benefit for novel projects is positive for the strong conflict (where $R(\theta_S) = H(\theta_S) > \theta_C$ when the conflict is strong), $\left. \frac{dW^{novel}(\theta_B)}{d\theta_B} \right|_{\theta_B=\theta_S} = 1 - F(H(\theta_S)) > 0$. By combining the two effects, the ex ante shareholders' payoff is increasing at $\theta_B = \theta_S$, $\left. \frac{dW(\theta_B)}{d\theta_B} \right|_{\theta_B=\theta_S} = (1-p)[1 - F(H(\theta_S))] > 0$. As a result, the optimal board has $\tilde{\theta}_B^* > \theta_S$, which is a contradiction.

Proof of Proposition 1: For each condition, we construct a separate claim.

Claim 4. *If $\tilde{\theta}_B^* = \theta_S$ then $R(\theta_S) = \theta_C$.*

Suppose not: $\tilde{\theta}_B^* = \theta_S$ and $R(\theta_S) > \theta_C$. Then, by Claim 3 from Proof of Lemma 3, $\tilde{\theta}_B^* > \theta_S$, which is a contradiction.

Claim 5. *If $\tilde{\theta}_B^* > \theta_S$ then it holds that $(1-p)[W^{novel}(\theta_H) - W^{novel}(\theta_S)] \geq p[W^{routine}(\theta_S) - W^{routine}(\max\{\theta_S, H^{-1}(\theta_C)\})]$*

Recall the left-hand side of the inequality (LHS) is the maximal benefit of a strong conflict (novel projects side) and the right-hand side of the inequality (RHS) is the minimal cost of a strong conflict (routine projects side). We prove by contradiction: Suppose the optimal board is biased but the maximal benefit is below the minimal cost (LHS below RHS). Three cases exist:

- $H^{-1}(\theta_C) \leq \theta_S$. Then, a strong conflict is induced for any relevant board type $\theta_B \geq \theta_S$. The minimal cost is zero but the maximal benefit is positive, $W^{novel}(\theta_H) - W^{novel}(\theta_S) > 0$. This contradicts that the maximal benefit is below the minimal cost.
- $H^{-1}(\theta_C) > \theta_S$ and $\tilde{\theta}_B^* \in [\theta_S, H^{-1}(\theta_C)]$ (weak conflict induced). By Claim 1 in Proof of Lemma 3, the optimal board is neutral if the conflict is weak; this is a contradiction.
- $H^{-1}(\theta_C) > \theta_S$ and $\tilde{\theta}_B^* \in [H^{-1}(\theta_C), \theta_H]$ (strong conflict induced). On this interval, the shareholders' payoff is bounded by \bar{W} from above,

$$\begin{aligned} W(\theta_B) &= p W^{routine}(\theta_B) + (1-p) W^{novel}(\theta_B) \\ &< p W^{routine}(H^{-1}(\theta_C)) + (1-p) W^{novel}(\theta_H) \equiv \bar{W}, \end{aligned}$$

because on this interval, $W^{routine}$ is maximized at $\theta_B = H^{-1}(\theta_C)$ and W^{novel} is maximized at θ_H . Now, we use that the condition that the maximal benefit exceeds the

minimal cost can be rearranged into $W(\theta_S) \leq \overline{W}$. Therefore, since we suppose the maximal benefit is below the minimal cost, we have $W(\theta_S) > \overline{W}$. This contradicts that \overline{W} is an upper bound of $W(\theta_B)$.

Proof of Corollary 2: Like in Proof of Proposition 1, LHS denotes the left-hand side of the inequality in condition (ii) (the maximal benefit of a strong conflict) and RHS denotes the right-hand side of the inequality (the minimal cost of a strong conflict). The proof follows from the following three observations:

- LHS is decreasing in p whereas RHS is increasing in p ; therefore, a unique cutoff $p^* \in [0, 1]$ exists.
- At the limit, as $p \rightarrow 0$, LHS becomes $W^{novel}(\theta_H) - W^{novel}(\theta_S) > 0$ whereas RHS becomes 0 so that the condition is satisfied. Therefore, the cutoff is positive, $p^* > 0$.
- At the limit, as $p \rightarrow 1$, LHS becomes 0 whereas RHS becomes $W^{routine}(\theta_S) - W^{routine}(\max\{\theta_S, H^{-1}(\theta_C)\}) \geq 0$. If $\theta_S < H^{-1}(\theta_C)$, RHS is positive, the condition is violated and the cutoff is $p^* < 1$. Otherwise, both LHS and RHS are zero, the condition is satisfied, and $p^* = 1$.

Proof of Corollary 3: We use $F(\theta) = \theta - \theta_{max} + 1$ and $f(\theta) = 1$ if $\theta \in [\theta_{max} - 1, \theta_{max}]$, which implies $H^{-1}(\theta_R) = \int_{\theta_R}^{\theta_{max}} \frac{1}{1-F(\theta_R)} \theta f(\theta) d\theta = \frac{1}{2}(\theta_{max} + \theta_R)$ if $\theta_R \in [\theta_{max} - 1, \theta_{max}]$. The inverse function yields that the board is indifferent at a high report if the reporting cutoff satisfies $\theta_R = H(\theta_B) = 2\theta_B - \theta_{max}$ such that $\theta_B \in [\theta_{max} - \frac{1}{2}, \theta_{max}]$. Note that the constraint $\theta_R = H(\theta_B)$ is linear.

To find the optimal board type, recall that the optimal board is biased only if the induced conflict is strong, $R(\tilde{\theta}_B^*) = H(\tilde{\theta}_B^*)$. Therefore, if the optimal board is biased, the reporting cutoff is set as $\theta_R = R(\theta_B) = H(\theta_B)$. Using this observation, we will proceed in the following three steps: (i) We will find a (unique) pair $(\theta_B^\dagger, \theta_R^\dagger)$ that maximizes the shareholders' expected payoff $W(\theta_B, \theta_R)$ subject to the constraint $\theta_R = H(\theta_B)$, where $(\theta_B, \theta_R) \in \mathbb{R}^2$. Here we will use that $W(\theta_B, \theta_R)$ is strictly quasiconcave for $(\theta_B, \theta_R) \in [\theta_S, H^{-1}(\theta_S)] \times [\theta_C, \theta_S]$. As a result, the upper contour sets are convex. Therefore, constrained optimization of the shareholder's expected payoff on a linear constraint will yield a unique optimum. (ii) We will compare $W(\theta_B^\dagger, \theta_R^\dagger)$ and $W(\theta_S, \theta_C)$; the latter corresponds to the neutral board. If $W(\theta_S, \theta_C) > W(\theta_B^\dagger, \theta_R^\dagger)$, then the neutral board is optimal because (θ_S, θ_C) is always feasible. (iii) If not, $W(\theta_S, \theta_C) \leq W(\theta_B^\dagger, \theta_R^\dagger)$, then we will check feasibility of $(\theta_B^\dagger, \theta_R^\dagger)$. Specifically, we will check if $\theta_B^\dagger \in [\theta_S, H^{-1}(\theta_S)]$ and $\theta_R^\dagger \in [\theta_C, \theta_S]$. If $(\theta_B^\dagger, \theta_R^\dagger)$ is not feasible, we will check if it implies that the optimal board is neutral.

Recall $\theta_S = 0$. In the first step, we obtain: $(\theta_B^\dagger, \theta_R^\dagger) = \left(\frac{2(1-p)}{4-3p} \theta_{max}, \frac{-p}{4-3p} \theta_{max} \right)$. In the second step, we derive the shareholder's expected payoff, $W(\theta_B, \theta_R) = pW^{routine}(\theta_B, \theta_R) + (1-p)W^{novel}(\theta_B, \theta_R) = \frac{1}{2}(\theta_{max}^2 - p\theta_B^2 - (1-p)\theta_R^2)$. The condition $W(\theta_S, \theta_C) \leq W(\theta_B^\dagger, \theta_R^\dagger)$ rearranges into

$$\theta_C \leq -\sqrt{\frac{p}{4-3p}} \theta_{max}. \quad (5)$$

In the third step, we evaluate feasibility of $(\theta_B^\dagger, \theta_R^\dagger)$. It is easy to check that $\tilde{\theta}_B^* \in [\theta_S, H^{-1}(\theta_S)] = [0, \frac{\theta_{max}}{2}]$. Also, $\theta_R^\dagger < \theta_S = 0$. Finally, we must verify a feasibility condition $\theta_R^\dagger \geq \theta_C$, which is equivalent to $\theta_C \leq \frac{-p}{4-3p}\theta_{max}$. This condition is however implied by the condition in (5). As a result, if the pair $(\theta_B^\dagger, \theta_R^\dagger)$ is not feasible, it is worse for the shareholders than the pair (θ_S, θ_C) . And since $(\theta_B^\dagger, \theta_R^\dagger)$ is the optimal pair on the linear constraint $\theta_R = H(\theta_B)$, any other feasible pair on the linear constraint is also worse for the shareholders than the pair (θ_S, θ_C) .

Proof of Proposition 2: For a given project type t , let $W^t(\theta_B)$ be the shareholders welfare when the board is of type θ_B and the CEO is willing to select project type t facing board of type θ_B . Let \bar{W}^t denote the maximal welfare for each project type, i.e., $\bar{W}^{routine} = W^{routine}(\theta_S)$ and $\bar{W}^{novel} = W^{novel}(\theta_H)$.

Part (i). We first prove that if $F(\theta) \leq G(\theta)$ for any θ , then $\bar{W}^{novel} \geq \bar{W}^{routine}$. We know that the shareholders' optimal decision cutoff for each project type is at θ_S . Applying this decision cutoff (i.e., for any $\theta < \theta_S$, the probability mass is shifted to an atom at θ_S where the shareholders' welfare is zero), we obtain adjusted distributions F^{adj} and G^{adj} that account for the rejections of projects with negative value. Namely, $F^{adj}(\theta) = G^{adj}(\theta) = 0$ if $\theta < \theta_S$, and $F^{adj}(\theta) = F(\theta)$ and $G^{adj}(\theta) = G(\theta)$ if $\theta \geq \theta_S$. Since the adjusted distributions preserve the first-order stochastic dominance, $F^{adj}(\theta) \leq G^{adj}(\theta)$ for any θ , novel projects are on average more profitable when the shareholders can implement the optimal decision cutoff θ_S for any project type, $\bar{W}^{novel} = \int_{\theta_S}^{\theta_{max}} \theta f(\theta) d\theta \geq \int_{\theta_S}^{\theta_{max}} \theta g(\theta) d\theta = \bar{W}^{routine}$.

Thus, if $F(\theta) \leq G(\theta)$ for any θ , the shareholders primarily seek to obtain value \bar{W}^{novel} . This value is attainable by the shareholders. Namely, when the board is extreme, $\theta_B = \theta_H$, inequalities $W^{novel}(\theta_H) = \bar{W}^{novel} \geq \bar{W}^{routine} > W^{routine}(\theta_H)$ and $G(\theta_H) - F(R(\theta_H)) \geq F(\theta_H) - F(R(\theta_H)) > 0$ jointly imply $M(\theta_H) > 0$. In other words, CEO's and shareholders' preferences over project selection at $\theta_B = \theta_H$ are aligned and this value is attainable.

Part (ii). By analogy to the claim above, if $F(\theta) \geq G(\theta)$ for any θ , then $\bar{W}^{novel} \leq \bar{W}^{routine}$. The shareholders primarily seek to attain value $\bar{W}^{routine}$. If $M(\theta_S) \leq 0$, this value is attainable when $\theta_B = \theta_S$ because CEO's and shareholders' preferences are aligned.

Part (iii). Like in Part (ii), the shareholders primarily seek to attain value $\bar{W}^{routine}$. But this value is not feasible since $\theta_B = \theta_S$ yields $M(\theta_S) > 0$; the CEO's and shareholders' preferences are not aligned. First, by contradiction, we prove $\theta_B^* \neq \theta_S$. Suppose $\theta_B^* = \theta_S$. Since $M(\theta_S) > 0$, it means that a combination of neutral board and novel project is optimal for the shareholders.

- When $R(\theta_S) > \theta_C$ (strong conflict), $R(\theta_B)$ is increasing if $\theta_B \in [\theta_S, \theta_H]$ and therefore also $W^{novel}(\theta_B)$ is increasing. For the neutral board to be constrained optimal for novel projects, we must have that $M(\theta_B) \leq 0$ if $\theta_B \in (\theta_S, \theta_H]$ (i.e., any increase in θ_B forces the CEO to select the routine project), but this is impossible due to continuity of M function and $M(\theta_S) > 0$.

- When $R(\theta_S) = \theta_C$ (weak conflict), $R(\theta_B)$ is constant if $\theta_B \in (\theta_S, H^{-1}(\theta_C)]$. Therefore, on this interval, the CEO's expected payoff from a novel project is also constant, whereas her expected payoff from a routine project is decreasing. As a result, $M(\theta_B) > M(\theta_S) > 0$ if $\theta_B \in (\theta_S, H^{-1}(\theta_C)]$. By continuity of M function, we must therefore have $M(\theta_B) > 0$ also for $\theta_B \in [H^{-1}(\theta_B), H^{-1}(\theta_B) + \varepsilon)$ where R function is increasing. But if $R(\theta_B)$ is increasing, also $W^{novel}(\theta_B)$ is increasing, and the shareholders will strictly prefer decisions on the novel project with a biased board $\theta_B \in (H^{-1}(\theta_B), H^{-1}(\theta_B) + \varepsilon)$ to decisions on the novel project with a neutral board $\theta_B = \theta_S$. This is a contradiction.

Second, we prove $\theta_B^* > \theta_C$. We start by observing that a board aligned with the CEO, $\theta_B = \theta_C$, convinces the CEO to select a routine project. This is because $R(\theta_C) = \theta_C$, and therefore $M(\theta_C) = W^{novel}(\theta_C) - W^{routine}(\theta_C) < 0$. The strict inequality follows from strict stochastic dominance of the routine project. At the same time, recall $M(\theta_S) > 0$. By continuity of M function, $M(\theta_B) < 0$ also on the neighborhood $\theta_B \in (\theta_C, \theta_C + \epsilon)$ for some small $\epsilon > 0$. But shareholders welfare, $W^{routine}(\theta_B)$, is increasing on the neighborhood and thus $\theta_B^* \neq \theta_C$.

Next, suppose $\theta_B^* < \theta_C$. If the CEO selects a routine, the outcome with θ_B^* is equivalent to the one with $\theta_B = \theta_C$ because the decision cutoff for routine projects is at $\max\{\theta_B^*, \theta_C\} = \theta_C$. If the CEO selects an innovation, θ_B^* is suboptimal for the shareholders because a higher welfare is achieved for $\theta_B = \theta_S$ where the CEO also selects an innovation.

Proof of Lemma 4: The expected benefit of learning conditional on the report r is $\int_{\underline{\theta}}^{\bar{\theta}} X(\theta, \theta_B) \frac{f(\theta)}{1-F(\theta_R)} d\theta = \frac{Y_j(\theta_B, \theta_R)}{1-F(\theta_R)}$, where $Y_j(\theta_B, \theta_R) \equiv \int_{\underline{\theta}}^{\bar{\theta}} X(\theta, \theta_B) f(\theta) d\theta = \int_{\underline{\theta}}^{\bar{\theta}} |\theta - \theta_B| f(\theta) d\theta$ is the unconditional expected benefit from learning in region \mathcal{P}_j and $[\underline{\theta}, \bar{\theta}]$ is summarized in Table 1. From discussion in the text, the board never learns when the CEO plays her optimal presentation strategy and $\theta_R > \theta_B$ because the relevant report is $r = l$ and it is then (weakly) optimal for the CEO to veto the project and eliminate board's option to learn. Therefore, nonlearning region \mathcal{N}_A includes all pairs (θ_B, θ_B) located above the diagonal $\theta_R = \theta_B$.

To obtain nonlearning region \mathcal{N}_B , we analyze the board's learning decision in regions \mathcal{P}_2 and \mathcal{P}_3 where the pairs are below the diagonal, $\theta_R \leq \theta_B$. The unconditional benefit from learning, $Y(\theta_B, \theta_R) = \int_{\theta_R}^{\theta_B} (\theta_B - \theta) f(\theta) d\theta$, is zero at the diagonal $\theta_B = \theta_R$ and is increasing in θ_B . (It is increasing even for $\theta_B > \theta_{max}$, where $f(\theta_B) = 0$.) Therefore, for any given θ_R , there exists a *unique* cutoff $\hat{\theta}_B > \theta_R$ such that $Y(\hat{\theta}_B, \theta_R) = [1 - F(\theta_R)]\kappa > 0$. The board is indifferent between learning and non-learning at the cutoff, doesn't learn if her type is below the cutoff $\hat{\theta}_B$ and learns if her type is above the cutoff $\hat{\theta}_B$.

To analyze properties of the cutoff, we introduce an implicit function $I(\theta_B, \theta_R) \equiv Y(\theta_B, \theta_R) - [1 - F(\theta_R)]\kappa$. Notice I is increasing in θ_B since Y is increasing in θ_B . However, the effect of θ_R on the value of the implicit function is non-monotonic. Precisely, the marginal effect is $\frac{\partial I}{\partial \theta_R} = f(\theta_R)(\theta_R - \theta_B + \kappa)$, so the implicit function is (weakly) decreasing in θ_R when the reporting cutoff is sufficiently far from the board's optimal cutoff, $\frac{\partial I}{\partial \theta_R} < 0$

if $\theta_R < \theta_B - \kappa$, and (weakly) increasing in θ_R when the reporting cutoff is sufficiently close to the board's optimal cutoff, $\frac{\partial I}{\partial \theta_R} > 0$ if $\theta_R \in (\theta_B - \kappa, \theta_B)$. Both effects are strict if $\theta_R \in [\theta_{min}, \theta_{max}]$; otherwise, the marginal effect is zero. Next, observe that the implicit function is negative at the diagonal, $I(\theta_B, \theta_B) < 0$, and therefore by continuity, it is negative also whenever the reporting cutoff is closely below the board's optimal cutoff.

Therefore, $I(\hat{\theta}_B, \theta_R) = 0$ holds only for parameters when the board's optimal cutoff is sufficiently far from the reporting cutoff, $\hat{\theta}_B > \theta_R + \kappa$. Consequently, at $\theta_B = \hat{\theta}_B$, the implicit function is decreasing in the reporting cutoff θ_R if $\hat{\theta}_B \leq \theta_{max}$; otherwise, it is constant in θ_R . Finally, we combine these properties of the implicit function at $\hat{\theta}_B$ with the implicit function theorem, and observe that $\hat{\theta}_B$ is weakly increasing in θ_R ,

$$\left. \frac{d\theta_B}{d\theta_R} \right|_{I(\theta_B, \theta_R)=0} = -\frac{\frac{\partial I}{\partial \theta_R}}{\frac{\partial I}{\partial \theta_B}} \geq 0.$$

As long as $\theta_B \leq \theta_{max}$ and $\theta_R \geq \theta_{min}$, we inversely have $\frac{d\theta_R}{d\theta_B} \big|_{I(\theta_B, \theta_R)=0} > 0$, and therefore for these values of (θ_B, θ_R) we may introduce an *increasing* function $\hat{H}(\theta_B)$ that yields θ_R such that the board is indifferent, $[1 - F(\hat{H}(\theta_B))]\kappa = Y(\hat{H}(\theta_B), \theta_R)$. Its inverse function, $\hat{H}^{-1}(\theta_R)$, is also increasing; moreover, the domain of the inverse function can be extended to any $\theta_R \in [\theta_{min}, \theta_{max}]$ because a unique $\hat{\theta}_B$ exists for any $\theta_R \in [\theta_{min}, \theta_{max}]$. Finally, we use that the cutoff $\hat{\theta}_B = \hat{H}^{-1}(\theta_R)$ is unique, and also that location of pairs (θ_B, θ_R) in regions \mathcal{P}_2 and \mathcal{P}_3 and below the diagonal implies $\theta_B \in [\theta_R, \hat{H}^{-1}(\theta_R)]$. As a consequence, non-learning region \mathcal{N}_B contains any pair of $(\theta_B, \theta_R) \in [\theta_{min}, \theta_{max}]^2$ such that $\theta_B \geq \theta_R$ and $\theta_B \leq \min\{H^{-1}(\theta_R), \hat{H}^{-1}(\theta_R)\}$.

To obtain nonlearning region \mathcal{N}_C , we analyze the board's learning decision in region \mathcal{P}_4 where $\theta_R \leq H(\theta_B)$. In this case, it is convenient to fix $\theta_B \in [\mathbb{E}[\theta], \theta_{max}]$ and analyze the board's decision when θ_R changes. The unconditional benefit from learning, $Y(\theta_B, \theta_R) = \int_{\theta_B}^{\theta_{max}} (\theta_B - \theta)f(\theta)d\theta$, is weakly positive (strictly if $\theta_B < \theta_{max}$) and constant in θ_R . The unconditional cost from learning, $[1 - F(\theta_R)]\kappa$, is decreasing in θ_R to zero at $\theta_R = \theta_{max}$. Therefore, an increase in θ_R only reduces the cost side of the board's tradeoff, and if the board doesn't learn, it is for θ_R in the neighborhood of θ_{min} . When $Y(\theta_B, \theta_{min}) \geq \kappa$, the board learns for any θ_R . More precisely, when $Y(\theta_B, \theta_{min}) < \kappa$, there exists a unique cutoff $\hat{\theta}_R > \theta_{min}$, where the board is indifferent at the cutoff, $Y(\theta_B, \hat{\theta}_R) = [1 - F(\hat{\theta}_R)]\kappa$, learns if $\theta_R \geq \hat{\theta}_R$, and doesn't learn if $\theta_R < \hat{\theta}_R$. Therefore, we may introduce a cutoff function $\hat{N}(\theta_B)$ which describes the board's indifference, $\hat{\theta}_R = \hat{N}(\theta_B)$, and the domain of the function is any θ_B such that $Y(\theta_B, \theta_{min}) < \kappa$.

Proof of Lemma 5: Like in Proof of Lemma 2, the best possible outcome for the CEO is when all projects with value $\theta \geq \theta_C$ are presented and approved, and all other projects are not presented or rejected. This outcome is implemented by $\theta_R = \theta_C$ and $(d_l, d_h) = (0, 1)$ if $\theta_R = \theta_C \geq \max\{H(\theta_B), \hat{H}(\theta_B)\}$ (i.e., with high report, the board approves the project without learning).

Suppose $\theta_C < \max\{H(\theta_B), \hat{H}(\theta_B)\}$ and thus the CEO's best possible outcome cannot be implemented. (Notice board that learns implement it if and only if $\theta_B = \theta_C$, but

then $\theta_B = \theta_C > \max\{H(\theta_B), \hat{H}(\theta_B)\}$, which we have covered above.) We prove that the optimal report is $\hat{\theta}_R = \max\{H(\theta_B), \hat{H}(\theta_B)\} < \theta_B$. First, motivating board's learning by $\theta_R < \hat{H}(\theta_B)$ implies that the decision cutoff is θ_B , but the decision cutoff $\hat{\theta}_R$ is closer to the CEO's optimal cutoff, since $\theta_C < \hat{\theta}_R < \theta_B$. Second, motivating board's rejection after high report by $\theta_R < H(\theta_B)$ implies that the decision cutoff is θ_{min} and the CEO's value zero. But, the decision cutoff $\hat{\theta}_R$ gives the CEO a positive expected value. Third, with any reporting cutoff $\theta_R > \hat{H}(\theta_B)$ and positive recommendations for any r leads (for relevant board types $\theta_B \geq \theta_S$) to project rejection with $r = l$ and the decision cutoff θ_R , but the decision cutoff $\hat{\theta}_R$ is closer to CEO's optimal cutoff, since $\theta_C < \hat{\theta}_B < \theta_R$. The same suboptimal decision cutoff is achieved for any reporting cutoff $\theta_R > \hat{H}(\theta_B)$ and positive recommendation only when $r = h$.

Proof of Lemma 6: The proof follows from the discussion in the text.

Proof of Proposition 3: First, we analyze the effect of a higher learning cost on shareholders' welfare. The shareholders' optimal board problem is to maximize $W(\theta_B)$ subject to the constraint $\theta_R = \hat{R}(\theta_B)$. The relevant board types are $\theta_B \in [\theta_S, \hat{\theta}_H]$; for these types, $\theta_R \in [\theta_C, \theta_S]$. Intuitively, the shareholders want as low θ_B as possible (to not distort routine projects) and also as high θ_R as possible (to not distort novel projects). A higher κ reduces $\hat{H}(\theta_B)$ function and consequently (weakly) reduces the reporting cutoff function, $\hat{R}(\theta_B) = \max\{\theta_C, H(\theta_B), \hat{H}(\theta_B)\}$. A higher κ thus makes the shareholder's constraint, $\theta_R = \hat{R}(\theta_B)$ tighter, which implies that the maximized shareholders' welfare is (weakly) reduced.

Second, we analyze the effect of a higher learning cost on the type of the board (neutral or biased). A higher κ (weakly) reduces $\hat{R}(\theta_B)$. We know that the shareholder's optimal board is chosen from two conditionally optimal boards: (i) for low θ_B such that $\hat{R}(\theta_B) = \theta_C$ (weak conflicts), the conditionally optimal board is neutral, whereas for (ii) high θ_B such that $\hat{R}(\theta_B) > \theta_C$ (strong conflicts), the conditionally optimal board is biased. The shareholders welfare associated with the neutral board in a weak conflict is constant in the cost κ , since the pair $(\theta_B, \theta_R) = (\theta_S, \theta_C)$ is invariant in the cost, whereas the welfare associated with the biased board in a strong conflict is (weakly) decreasing in the cost κ as $\hat{R}(\theta_B)$ (weakly) decreases (see above). This implies that the shareholders may switch from the biased board to the neutral board, but not vice versa.

To complete the argument, we also have to analyze the effect of a higher learning cost on the existence of a weak conflict for a neutral board which is a necessary condition for a neutral board. This effect is positive; a decrease in $\hat{R}(\theta_B)$ implies that it is more likely that a weak conflict begins to exist for a neutral board, i.e., it is more likely that we switch from a case with $\hat{R}(\theta_S) > \theta_C$ to a case with $\hat{R}(\theta_S) < \theta_C$. To sum up, with a higher learning cost, it is more likely that a weak conflict begins to exist for a neutral board, and also more likely that the board switches from the biased board to the neutral board (conditional on existence of weak conflict for the neutral board); the two effects thus jointly imply that with a higher learning cost, the board is more likely neutral.

Proof of Corollary 4: By Corollary 3, we have two parametrical cases:

- In absence of learning, suppose the shareholders' optimal board is neutral. Then, the CEO is strictly better off when $\kappa \rightarrow \infty$ because the decision cutoff for routine projects is θ_S for both cost-free and prohibitive learning, but the decision cutoff for innovations is θ_C for prohibitive learning and $\theta_S > \theta_C$ for cost-free learning.
- In absence of learning, suppose the shareholders' optimal board is biased. Let $(\theta_B^\dagger, \theta_R^\dagger)$ be the optimal biased board type and resulting reporting cutoff; this is also the pair of decision cutoffs for routine projects and innovations. This pair must be compared with a pair (θ_S, θ_C) . In our example, for any pair of cutoffs (θ_B, θ_R) , the CEO's ex ante expected payoff is

$$W_C(\theta_B, \theta_R) \equiv \frac{\theta_{max}^2}{2} - \theta_C \theta_{max} + p \left(-\frac{\theta_B^2}{2} + \theta_C \theta_B \right) + (1-p) \left(-\frac{\theta_R^2}{2} + \theta_C \theta_R \right).$$

The CEO is better off with costless board's learning than without board's learning if $W_C(\theta_S, \theta_S) > W_C(\theta_B^\dagger, \theta_R^\dagger)$. Rearranging and simplifying, $W_C(\theta_S, \theta_S) - W_C(\theta_B^\dagger, \theta_R^\dagger) = \frac{p(1-p)(\theta_{max}-2\theta_C)}{2(4-3p)} > 0$, because $p \in (0, 1)$, $\theta_{max} > 0$ and $\theta_C < 0$.

Proof of Proposition 4: We proceed in two steps. In the first step, we take reporting strategy $\theta_R = R(\theta_B)$ and also presentation strategy $(d_l, d_h) = (0, 1)$ as given and prove that then only a babbling cheap talk equilibrium exists. In the second step, we prove that the CEO doesn't deviate from the reporting strategy $R(\theta_B)$ and presentation strategy $(d_l, d_h) = (0, 1)$.

Step 1: We analyze only report-specific cheap talk for $r = h$, because the project is not presented for $r = l$. Conditional on high report, the project value is $\theta \in M_h \equiv [\theta_R, \theta_{max}]$. There exists a babbling equilibrium where the CEO sends a single message that θ belongs to M_h : in this case the project is approved. It is easy to see that there does not exist a two-message equilibrium. In such equilibrium the interval M_h is partitioned into (M_h^-, M_h^+) . Because the board approves the project when the message is that $\theta \in M_h$, it also approves the project when the message is that $\theta \in M_h^+$. However, if the CEO observes $\theta \in M_h^-$, she communicates that $\theta \in M_h^+$ because she wants the project to be approved. Thus the communication degenerates into babbling.

Step 2: In the previous step we showed that, when the reporting cutoff is $\theta_R = R(\theta_B)$, the cheap talk is uninformative and the board's information is characterized by a partition at θ_R . Is there any other partition that is better for the CEO from an ex ante perspective? By Lemma 2, the partition at the cutoff $R(\theta_B)$ is the best two-partition for the CEO that satisfies that the board is willing to approve the project after message $r = h$. Also, a finer partitioning cannot increase the CEO's expected payoff. To do so, such partition must result in project approval for some $\theta \in (\theta_C, R(\theta_B))$, which is impossible since $\theta = R(\theta_B)$ is the lowest project type that the board approves in any partition.

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. *Journal of Accounting Research*, 50, 437–486.
- [3] Antle, R., Eppen, G. (1985). Capital rationing and organizational slack in capital budgeting. *Management Science* 31(2), 163–174.
- [4] Arya, A., Glover, J., Sivaramakrishnan, K., (1997). Commitment issues in budgeting. *Journal of Accounting Research*, 35, 273–278.
- [5] Associated Press (2020): Report: former Aluminum Mill CEO misled investors, board. April 27, 2020.
- [6] Bagnoli, M., Watts, S. G. (2015). Delegating disclosure and production choices. *The Accounting Review*, 90, 835–857.
- [7] Baiman, S., Heinle, M., Saouma, R. (2013). Multi-stage capital budgeting with delayed consumption of slack *Management Science* 59(4).
- [8] Baldenius, T., Dutta, S., Reichelstein, S. (2007). Cost allocations for capital budgeting. *The Accounting Review* 82(4), 837–867.
- [9] Baldenius, T., Melumad, N., Meng, X. (2014). Board composition and CEO power. *Journal of Financial Economics*, 112(1), 53–68.
- [10] Baldenius, T., Meng, X., Qiu, L. (2019). Biased boards. *The Accounting Review*, 94(2), 1–27.
- [11] Baldenius, T., Meng, X., Qiu, L. (2020). The value of board commitment. *Review of Accounting Studies*, forthcoming.
- [12] Ball, I., Gao, X. (2022). Benefiting from Bias. Working paper.
- [13] Bastian-Johnson, N., Pfeiffer, T., Schneider, G. (2013). Multistage capital budgeting for shared investments. *Management Science* 59, 1213–1228.
- [14] Bastian-Johnson, N., Pfeiffer, T., Schneider, G. (2017). Two-stage capital budgeting under resource constraints. *Review of Accounting Studies* 22, 933–963.
- [15] Bertomeu, J., Cheynel, E., (2015). Asset measurement in imperfect credit markets. *Journal of Accounting Research* 53 (5), 965–984.
- [16] Brav, A. Jiang, W., Ma, S., Tian, X. (2018). How does hedge fund activism reshape corporate innovation? *Journal of Financial Economics* 130, 237–264.

- [17] Caplin, A., Dean, M., Leahy, J. (2019). Rational inattention, optimal consideration sets, and stochastic choice. *Review of Economic Studies*, 86(3), 1061–1094.
- [18] Chakraborty, A., Yilmaz, B. (2017). Authority, consensus, and governance. *Review of Financial Studies*, 30(12), 4267–4316.
- [19] Che, Y., Kartik, N. (2009). Opinions as incentives. *Journal of Political Economy*, 117(5), 815–860.
- [20] Chen, K., Guay, W., Lambert, R. (2020). Multidimensional corporate governance: board design, executive compensation, and information sharing. Working paper.
- [21] Chen, Y., Laux, V. (2021). Overconfidence, manipulation, and board monitoring. Working paper.
- [22] Cianciaruso, D., Marinovic, I., Smith, K. (2020). Information design in financial markets. Working paper.
- [23] Decaire, H., Sosyura, D. (2021) Self-Dealing in Corporate Investment. Working paper.
- [24] 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>
- [25] Dessein, W. (2002). Authority and Communication in Organizations. *Review of Economic Studies*, 69, 811–838.
- [26] Dewatripont, M., Titole, J. (1999). Advocates. *Journal of Political Economy*, 107(1), 1–39.
- [27] Dordzhieva, A., Laux, V., Zhang, R. (2020). Signaling private information via accounting system design. Working Paper.
- [28] Drymiotis, G., (2007). The monitoring role of insiders. *Journal of Accounting and Economics*, 44 (3), 359–377.
- [29] Drymiotis, G., (2009). Managerial influencing of boards of directors. *Journal of Management Accounting Research* 20 (Special Issue), 19–45.
- [30] Ebert, M., Schäfer, U. and Schneider, G.T. (2019). Information Leaks and Voluntary Disclosure. Working paper.
- [31] Feng, F., Luo, Y., Michaeli, B. (2022). In search of a unicorn. Working paper.
- [32] Fershtman, C., Judd K. (1987). Equilibrium incentives in oligopoly. *The American Economic Review*, 77(5), 927–940.

- [33] Frenkel, S., Guttman, I., and Kremer, I. (2019). The effect of exogenous information on voluntary disclosure and market quality. *Journal of Financial Economics*, forthcoming.
- [34] Friedman, H., Hughes, J., Michaeli, B., (2020). Optimal reporting when additional information might arrive. *Journal of Accounting and Economics*, 69 (2–3), 101276.
- [35] Friedman, H., Hughes, J., Michaeli, B., (2021). A rationale for imperfect reporting standards. *Management Science*, forthcoming.
- [36] Gentzkow M., Kamenica E., (2017a). Competition in persuasion. *Review of Economic Studies*, 84, 300–322.
- [37] Gentzkow M., Kamenica E., (2017b). Disclosure of endogenous information. *Economic Theory Bulletin*, 5, 47–56.
- [38] Gregor, M. (2020). Board directors under CEO’s pressure: misreading, misreporting, and dissent risk. Working paper.
- [39] Gregor, M., Michaeli, B. (2021). Boards and executive compensation: another look. Working paper.
- [40] Göx, R., Michaeli, B., (2019). Optimal information design and incentive contracts with performance measure manipulation. Working paper.
- [41] Göx, R.F., Wagenhofer, A. (2009). Optimal impairment rules. *Journal of Accounting and Economics*, 48(1), 2–16.
- [42] Harris, M., Raviv, A. (2008). A theory of board control and size. *Review of Financial Studies*, 21(4), 1797–1832.
- [43] Hope, O., Thomas, W. (2008). Managerial empire building and firm disclosure. *Journal of Accounting Research* 46(3), 591–626.
- [44] Huang, Z., (2016). Optimal reporting systems with investor information acquisition. Working paper.
- [45] Jensen, M. (1986). Agency costs of free cash flow, corporate finance, and takeovers. *American Economic Review* 76(2), 323–329.
- [46] Jensen, M., Murphy, K. (1990). Performance pay and top-management incentives. *Journal of Political Economy* 98(2), 225–264.
- [47] Jiang, X., Stocken, P. (2019). Public communication between managers and analysts. Working paper.
- [48] Jiang, X., Yang, M. (2017). Properties of optimal accounting rules in a signaling game. *Journal of Accounting and Economics*, 63, 499–512.

- [49] Kamenica, E., Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6), 2590–2615.
- [50] Kamenica, E. (2019). Bayesian persuasion and information design. *Annual Review of Economics* 11, 249–272.
- [51] Laux, V. (2014). Corporate governance, board oversight, and CEO turnover. *Foundations and Trends in Accounting* 8(1), 1–73.
- [52] Lu, J., Wang, W. (2015). Board independence and corporate investments. *Review of Financial Economics*, 24, 52–64.
- [53] Maffett, M., Nakhmurina, A., Skinner, D. (2020). Importing activists: determinants and consequences of increased cross-border shareholder activism. Working paper.
- [54] Malenko, N. (2014). Communication and decision-making in corporate boards. *Review of Financial Studies*, 27(5), 1486–1532.
- [55] Matysková, L., Montes, A. (2021). Bayesian persuasion with costly information acquisition. Working paper.
- [56] Meng, X. (2020). Information, incentives and CEO replacement. Working paper.
- [57] Michaeli, B. (2017). Divide and inform: rationing information to facilitate persuasion. *The Accounting Review*, 92(5), 167–199.
- [58] Michaeli, B., Wiedman, E. (2020). Investor beliefs and corporate disclosure timing in the presence of external sources of information. Working paper.
- [59] Mitusch, K., Strausz, R. (2005). Mediation in situations of conflict and limited commitment. *Journal of Law, Economics, and Organization*, 21(2), 467–500.
- [60] Peterson, M. (2003). Kmart, saying ex-chief misled board, wants payments returned. *The New York Times*, February 26, 2003.
- [61] Qiu, L. (2020). Say on pay and the “quiet life” of boards. Working paper.
- [62] Rim, H, Sul, E. (2020). Director independence and corporate investment efficiency: evidence from board reforms worldwide. Working paper.
- [63] Segendorff, B. (1998). Delegation and threat in bargaining. *Games and Economic Behavior*, 23(2), 266–283.
- [64] Shleifer, A., Vishny, R. (1989). Management entrenchment: The case of manager-specific investments. *Journal of Financial Economics* 25(1), 123–139.
- [65] Stulz, R. (1990). Managerial discretion and optimal financing policies. *Journal of Financial Economics* 26(1), 3–27.

- [66] The Economist, (2001). The fading appeal of the board room. March 31, 2001.
- [67] Townsend, R. (1979). Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory* 21(2), 265–293.
- [68] Useem, M. (2006). How well-run boards make decisions. *Harvard Business Review*, November 2016.