Artificial Intelligence in the Boardroom*

Daniel Ferreira Jin Li

February 9, 2026

Abstract

We analyze how the use of artificial intelligence affects the monitoring and advisory relationship between CEOs and corporate boards. AI can serve as an advisor to CEOs, partially substituting for board advice, and thus reducing CEOs' incentives to share information with directors. In equilibrium, firms respond by reducing board independence to restore information flows. Lower monitoring intensity decreases CEO turnover, making CEOs more entrenched. Lower dismissal risk allows firms to reduce CEO compensation. We show that AI adoption by the board mitigates some of the negative aspects of the CEO's AI use, but cannot restore efficient levels of board monitoring. Our analysis suggests that the adoption of AI in the boardroom may have unintended consequences for corporate governance.

Keywords: Boards, CEOs, Artificial Intelligence, Corporate Governance

*We thank Mike Burkart, Luis Garicano, Dirk Jenter, and seminar participants at LSE. We thank Zehao Zhang for excellent research assistance. Ferreira: London School of Economics, CEPR and ECGI. Li: HKU Business School.

1 Introduction

Artificial intelligence systems have evolved from sophisticated chatbots to become agents that can solve problems and perform complex tasks in a semi-independent manner. Such systems can both help experts and replace them.[1] This transformation has been particularly pronounced in domains requiring analytical reasoning, pattern recognition, and strategic assessment—capabilities that were once the exclusive province of human expertise.[2] Major corporations are already deploying AI systems that can analyze vast datasets, identify market trends, and evaluate strategic options. The speed and scope of this technological shift raises questions about how organizations structure their decision-making processes and divide tasks between humans and artificial intelligence systems.[3]

In this paper, we develop a framework for studying the impact of generative artificial intelligence on the relationship between CEOs and corporate boards. As AI increasingly serves as an analytical tool for executive decision-making,[4] it potentially disrupts the traditional information flows between executives and directors. While CEOs and other executives can now access AI-generated insights for strategic planning, idea generation, and project assessment, this capability may reduce their reliance on board expertise and alter their incentives to share information.[5] We show that these changes in communication dynamics have important implications for board composition, monitoring, and executive compensation.

[1] Korinek (2025) shows how AI systems can enhance the productivity of economics researchers by both working in collaboration with them and carrying out independent research tasks.

[2] In a randomized controlled trial with 758 Boston Consulting Group consultants, Dell'Acqua et al. (2023) find that those using GPT-4 completed 12.2% more tasks, finished them 25.1% faster, and achieved over 40% higher quality output compared to controls.

[3] Babina et al. (2025) use resume and job posting data to document significant flattening of hierarchical structures with AI adoption, showing increases in junior-level workers and decreases in middle management roles.

[4] Kenny, Kowalkiewicz, and Oosthuizen (2025) present cases and examples of how CEOs can use AI for strategic planning. Csaszar, Ketkar, and Kim (2024) present evidence of the use of AI as a strategic decision-making tool in firms.

[5] Larcker et al. (2025) argue that AI has the potential to alter the relationship between executives and directors significantly. Shekshnia and Yakubovich (2025) discuss how AI helps directors.

The model is as follows. A CEO exerts effort to identify a problem (e.g., find a deal, generate an idea, etc.). Once a problem is identified, the CEO must assess it and decide whether to solve it on her own or with the help of the board. The CEO has access to an AI technology. We consider AI primarily as an *oracle*, that is, it "*serves as an advisor, offering information, predictions, or recommendations while leaving all decisions and actions to human users*" (Jiang (2025), p. 114). The AI technology has two capabilities: it enhances the CEO's problem-solving skills (a *skill-augmenting capability*), and it also solves some problems independently of the CEO's ability (a *problem-solving capability*).

We consider the problem of a group of controlling shareholders who must choose the board's monitoring intensity (i.e., board independence) and the CEO's compensation contract, which we jointly refer to as a *governance structure*. An optimal governance structure must incentivize the CEO both to exert effort and to ask the board for help when the CEO is unable to solve a problem on her own. Asking for help is perceived as a negative signal about the CEO's ability, which increases the likelihood that an independent board will replace the CEO. To induce the CEO to communicate truthfully with the board, the optimal governance structure must involve a board that is partially "CEO-friendly," that is, a board that may opt to retain the CEO even when a better replacement is available. As a consequence, the optimal structure will typically commit to some degree of ex-post inefficient CEO entrenchment to enhance the CEO's ex-ante incentives to seek the board's advice.

A CEO with access to an AI technology that has problem-solving capabilities can use it as a substitute for the board's advice. The more capable AI is, the less likely is the CEO to ask the board for help. Therefore, to restore information flows between the CEO and the board, the optimal board structure must be more CEO-friendly when such a capability exists. As the CEO faces a lower risk of dismissal under a friendlier board, the optimal compensation level is lower with AI than without it. Because the CEO's incentive constraints must be binding, the decrease in dismissal risk is exactly offset by the reduction in pay, leaving the CEO no better off. Thus, the AI-driven increase in CEO entrenchment

destroys firm value, without benefiting the CEO.

In contrast, if AI makes CEOs more skilled at solving problems, CEOs become less likely to involve the board in problem-solving. As CEOs become more productive when they have access to AI, firms find it easier to incentivize them to exert effort. Thus, CEOs can be induced to exert effort with lower pay. However, lower pay reduces CEO incentives to communicate with the board. Thus, as before, the optimal governance structure requires low pay and low monitoring intensity. We conclude that CEO access to AI increases board friendliness and reduces CEO pay. The value implications of AI adoption by CEOs are ambiguous: while AI's problem-solving capabilities are welfare-destroying, its skill-augmenting capabilities are value-enhancing.

While AI adoption by CEOs unambiguously reduces both the board's monitoring intensity and CEO pay levels, a natural follow-on question is whether AI adoption by the board can counter the effects of CEO AI adoption. We extend the model to allow the board to directly benefit from AI adoption in its advisory and monitoring roles. We show that AI adoption by the board attenuates some of the effects of CEO AI adoption on monitoring and pay levels. However, to reverse the original effects, shocks to AI technology must be significantly biased towards the board. If we instead expect AI shocks to benefit CEOs and boards in similar ways, AI adoption must decrease board monitoring and CEO pay levels.

Our paper is the first formal analysis of the impact of AI on corporate boards. Our model relates to a theoretical literature on the dual role of boards as advisors and monitors of management,[6] which includes Adams and Ferreira (2007), Harris and Raviv (2008), Baldenius, Melumad, and Meng (2014), Baldenius, Meng, and Qiu (2019), Levit (2020),

[6] A more recent literature has also focused on the "mediating" role of boards. Burkart, Miglietta, and Ostergaard (2023) present evidence that boards monitor management and mediate between different types of shareholders. Ewens and Malenko (2024) provide evidence of boards that both advise and mediate between investors and executives. Also consistent with a mediating role of boards, Ferreira, Ferreira, and Mariano (2018) show evidence that changes in creditors' power relative to shareholders affect appointments to the board.

Drymiotes and Sivaramakrishnan (2021), and Jiang and Laux (2023), among others.[7] More generally, the paper relates to the broader literature on information asymmetries in the boardroom, as in Warther (1998), Raheja (2005), Song and Thakor (2006), Baranchuk and Dybvig (2009), Inderst and Mueller (2010), Levit (2012), Malenko (2014), Levit and Malenko (2016), Chakraborty and Yilmaz (2017), Chemmanur and Fedaseyeu (2018), and Donaldson, Malenko, and Piacentino (2020).[8]

Our model follows Adams and Ferreira (2007) in that the CEO must first provide information to the board before receiving advice. That is, ours is a "two-way" communication model. In this setup, CEO-friendly boards can be optimal: The board commits to monitoring with low intensity to provide the CEO with incentives to share information. The logic is reminiscent of Burkart, Gromb, and Panunzi (1997) and Aghion and Tirole (1997), who show that ex-post intervention by principals undermines agents' ex-ante incentives.

As in Song and Thakor (2006), in our model, the CEO is responsible for generating a project idea with an unknown probability of success before deciding whether to reveal it to the board. CEOs are heterogeneous in ability, enjoy private benefits, and have career concerns. The board monitors by gathering information about the CEO and occasionally replacing a low-ability CEO. CEO compensation is endogenous and depends on the CEO's performance and the board's belief about the CEO's type. Thus, unlike Adams and Ferreira (2007), our model has implications for CEO turnover and compensation.

A key innovation of our paper is to model the relationship between a CEO and a board as a knowledge hierarchy, as in Garicano (2000) and Garicano and Rossi-Hansberg (2006). Following the original contribution of Ide and Talamàs (2025), we introduce AI within a knowledge-hierarchy setup. As in Nikolowa (2015), agents learn about their knowledge over time. As in Fuchs, Garicano, and Rayo (2015), there is asymmetric information about

[7] To the best of our knowledge, the advisory role of boards is first mentioned in the economics and finance literature by Fama and Jensen (1983), who argue that directors "*provide an important support function to the top managers in dealing with specialized decision problems*" (p. 314).

[8] See also Malenko (2024) for a comprehensive review of the literature on information flows in corporate governance.

agents' knowledge.

Our analysis shows that AI adoption by CEOs and boards can have unintended consequences for corporate governance due to the interactions between the board's monitoring and advising roles. For evidence that boards perform an important advisory role (which at times conflicts with its monitoring role), see Coles, Daniel, and Naveen (2008), Linck, Netter, and Yang (2008), Faleye, Hoitash, and Hoitash (2011), Field, Lowry, and Mkrtchyan (2013), Harford and Schonlau (2013), Adams, Akyol, and Verwijmeren (2018), de Haas, Ferreira, and Kirchmaier (2021), and Ewens and Malenko (2024).

There are a few theoretical papers on AI in organizations and governance. An early contribution is Athey et al. (2020), who introduces an AI agent to Aghion and Tirole's model of authority. Zhong (2025) develops a framework to study the interactions between humans and technologies in a multi-layer decision-making problem. The model implies that the optimal human-AI authority allocation depends on the tradeoff between correcting errors versus introducing new ones, with AI typically best as the final decision maker. Huang and Ouyang (2025) develop a cheap-talk model of AI-based advising in which AI agents are unbiased (unlike human agents) but face difficulties in codifying soft information. More closely related to our work is Chen and Han (2026). They develop a model in which biased CEOs may have access to AI technology. Because AI increases the granularity and precision of hard information available to CEOs, there is more scope for biased decisions, making agency problems more severe. Firms can mitigate some of these negative effects by changing contracts, but these fixes come at a cost. They conclude that AI adoption may reduce profits and welfare.

Finally, our paper connects to recent empirical work documenting AI's unintended consequences on worker behavior. Cowgill, Hernandez-Lagos, and Wright (2024) show that generative AI reduces screening accuracy by 4-9% as less-qualified individuals utilize AI to mimic expert communications, thereby fundamentally altering how information flows between applicants and evaluators. Chen et al. (2025) show that AI adoption leads to lower worker motivation. These empirical findings underscore the importance of AI

as a source of information cost. Our model shows that firms can partially mitigate these adverse effects of AI adoption by adjusting their governance structures.

We present our model setup in Section 2. Section 3 shows the impact of AI on board learning. Section 4 solves the model in a pre-AI economy. Our main results are then presented in Section 5. Section 6 extends the model to allow for AI adoption by the board. Section 7 concludes. All omitted proofs are in the Appendix.

2 Model Setup

In this section, we present our model setup. We begin by presenting a simple model of knowledge hierarchies in Subsection 2.1. Following Ide and Talamàs (2025), we work with Garicano's (2000) knowledge hierarchy framework, augmented by AI. In Subsection 2.2, we apply the model to the relationship between a CEO and a board.

2.1 AI in a Knowledge Hierarchy

A human agent operates a technology represented by a continuum of problems. Each problem has *difficulty* x, which is uniformly distributed over $[0, 1]$. The agent is endowed with *knowledge* indexed by $z \geq 0$. Time is continuous but divided into periods of length 1, which we call one "unit of time." As in Garicano (2000), at the beginning of a period, the agent randomly draws a problem of unknown type x. If $x \leq z$, an agent who spends $l \leq 1$ units of time on the problem can solve it with probability l. We set the value of each solved problem to 1. If $x > z$, the agent cannot solve the problem regardless of the time she spends on it, in which case the output is zero.

To assist with problem-solving, the agent also has access to an AI model. Formally, we represent the capabilities of the AI model by a pair of parameters $(\alpha, \kappa) \in [1, \overline{\alpha}] \times [0, 1]$. The AI model has two capabilities. First, it serves as a tool to help the agent solve problems. Second, it can solve some problems autonomously.

If the agent uses AI as a tool, she can solve any problem that has difficulty $x \leq \alpha z$. That is, parameter α measures how much AI enhances the agent's ability to solve problems. We call α AI's *skill-augmenting capability*. Conceptually, AI's skill-augmenting capability is no different from other technological innovations—such as computers—that help agents perform complex tasks and solve more problems.

A key difference between AI models and other skill-augmenting technologies is that AI models can autonomously solve certain problems with minimal human input. Let κ denote AI's knowledge as an autonomous problem solver. That is, AI models can also be used as a substitute for human agents, in which case they solve problems of difficulty $x \leq \kappa$. We call κ AI's *problem-solving capability*. It is this capability that arguably differentiates AI models from other innovations. While α represents AI's ability to complement an agent's knowledge, κ represents AI's ability to be a substitute for human decision-making. Because AI can be a substitute for humans in some tasks, we may alternatively think of AI as a hybrid between an oracle and an agent (Jiang (2025)). Henceforth, we also refer to AI models as *AI agents*. Despite this language, in the model, we leave all implementation decisions to humans, thus AI agentic autonomy is limited to the provision of advice.

Firms organize production such that agents with spare time capacity can offer "help" to other agents. Specifically, if, at the beginning of a period, an agent draws a problem she cannot solve, she can ask for the help of another agent (as in Garicano (2000), Garicano and Rossi-Hansberg (2006), and Ide and Talamàs (2025), among others). We call agents who specialize in providing help *advisors*, while agents who draw problems are called *workers*. We assume that a firm can hire at most one worker and one advisor per period. That is, firms are structured as two-layer knowledge hierarchies. If an advisor with knowledge $z_a \geq x$ spends l units of time assisting a worker, they solve the problem with probability l. If $z_a < x$, the problem cannot be solved.

Under complete information about knowledge levels and problem difficulty, a worker with access to an AI agent with capabilities (α, κ) would proceed as follows. If she draws a problem with difficulty $x \leq \tilde{z} := \max\{\alpha z, \kappa\}$, she can solve it without involving the

advisor. If $x > \tilde{z}$, the worker asks the advisor for help if and only if $x \leq z_a$. The net effect of AI is identical to that of a technology that makes the worker's "effective knowledge" equal to $\tilde{z}$. In other words, whether AI is used as a tool to help the worker or as a substitute for the worker does not matter.

Things are different when information is incomplete. Suppose that workers are uncertain about their types, that is, they know only the underlying distribution of knowledge, $z \sim F(\cdot)$. Additionally, as in Nikolowa (2015), suppose that workers cannot observe the difficulty of a problem; a worker learns only whether she can solve the problem immediately after drawing it. Thus, even if AI's problem-solving capability κ is known, the worker still doesn't know whether AI can solve a problem, because the problem's difficulty is unknown. We also assume that AI agents do not know whether they can solve a problem. That is, an AI agent always offers a solution when asked, even when it cannot solve the problem. In other words, AI models make mistakes.[9]

Suppose it takes at least one unit of time for the outcome of the decision to be realized, and waiting is costly. In that case, the worker may have to decide whether to solve the problem herself (with the help of AI) or to ask *either* AI to solve it *or* the advisor for help. Thus, in such cases, the optimal decision protocol requires the worker first to assess whether she can solve the problem independently, with the help of AI. If the worker cannot solve the problem, she decides whether to assign it to AI or ask the advisor for help.

[9] An AI agent may know it makes mistakes, and may even know the probability that it cannot solve a problem. All we need is for the AI agent not to know with certainty whether it can or cannot solve a particular problem. It is certainly the case that, with the right prompt, an LLM may admit not to know the answer to a particular question, especially for a fully objective question, such as a math problem. However, for questions about management and strategy, the human agent is likely to remain uncertain about the quality of an AI-generated answer.

2.2 CEOs and Boards

Here, we apply and adapt the framework just described to the relationship between a CEO and a board. We consider a firm owned by one or multiple shareholders, who are residual claimants of the firm's cash flow. Shareholders have no specialized human capital and, thus, cannot become CEOs or directors. All agents—CEOs, directors and shareholders—are risk-neutral, have zero discount rates, have zero outside utility, and are protected by limited liability.

There are two periods. The firm must hire one CEO and one director in each period. We set the board size to one. In the language of Subsection 2.1, the CEO plays the role of a worker, and the director serves as an advisor. Each CEO is endowed with knowledge $z \sim U[0,1]$. That is, both x and z are uniformly distributed in the population. No one observes z. Knowledge is persistent. The CEO has access to an AI agent with capabilities (α, κ). Directors (i.e., advisors) have knowledge $z_a = 1$. This assumption is stronger than necessary; we only need the board's expected knowledge to be sufficiently high.

Within a period, the timing of actions is as follows.

(i) The shareholders sign contracts with a CEO with (unobservable) knowledge $z \sim U[0,1]$ and a board of (observable) *ex ante* type $\pi \in [0,1]$. The CEO enjoys a private benefit $B > 0$.

(ii) The CEO decides whether to exert effort at a private cost c. If the CEO exerts effort, she finds a "deal" (i.e., a problem) with (unknown) difficulty $x \sim U[0,1]$. The CEO then learns her "*ex post* type," that is, whether $x \leq \alpha z$ (type g) or $x > \alpha z$ (type b). After learning her type, she chooses one of three actions: $m \in \{m_a, m_b, m_c\}$, where m_a denotes asking the AI agent to solve the problem autonomously, m_b denotes disclosing the deal to the board, and m_c denotes solving the problem by herself.

(iii) The state of the world, $\omega \in \{\omega^{tough}, \omega^{weak}\}$, is realized, where π (the board's *ex ante* type) denotes the probability of $\omega = \omega^{tough}$. The board then decides whether to fire

or retain its incumbent CEO.

(iv) Output $y_t \in \{0, 1\}$ is realized. The CEO earns a bonus $w_t \geq 0$ if $y_t = 1$.

Figure 1 provides a summary of the timing.

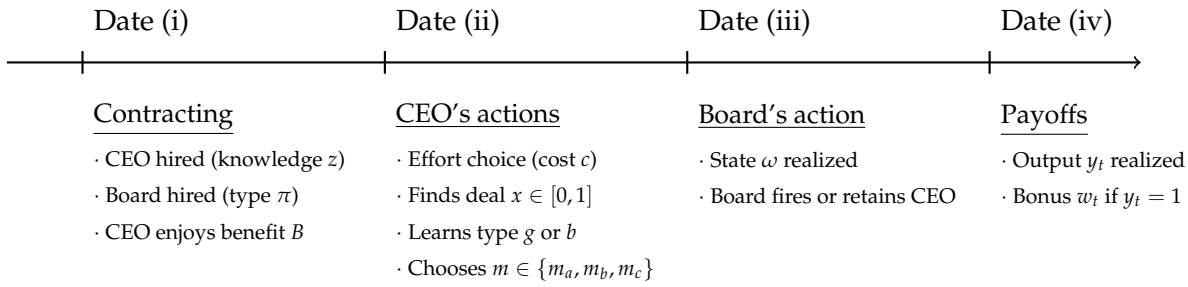


Figure 1: Timing of the game within a period.

We now provide a more detailed description of the game. At Date (i), shareholders choose a board with *ex ante* type $\pi \in [0, 1]$, denoting its "independence" (i.e., its degree of alignment with shareholders' interests). Formally, we think of π as the result of an action that stochastically reduces the board's private cost of firing a CEO: with probability π, the board is "tough" (state of the world ω^{tough}) and faces no cost of firing the CEO, while with probability $1 - \pi$, the board is "weak" (state of the world ω^{weak}) and the cost of firing the CEO is prohibitively high. The state of the world, ω, is realized at Date (iii). Note that shareholders can commit not to fire a CEO by choosing $\pi = 0$. In contrast, the board cannot commit to always firing a CEO; even if $\pi = 1$, firing a CEO must still be incentive-compatible for the board. Thus, board monitoring is a real option to replace the CEO in state ω^{tough}.

Our interpretation is that π is determined by the board's structure and composition. A less independent board faces a greater cost of monitoring the CEO, and is thus less likely to replace an underperforming CEO.[10] As in Adams and Ferreira (2007), we call π

[10]For simplicity, we assume that shareholders face no cost of choosing π; we could easily add a direct cost of independence without changing the nature of the problem.

the board's *monitoring intensity*. The interpretation is that CEOs regularly take actions to entrench themselves, such as co-opting board members. This entrenchment is a byproduct of the collaboration between the CEO and the board. By working closely with the CEO, the board becomes CEO-friendly. A friendlier board has lower monitoring intensity.

In Date (i) of period $t = 1$, the firm offers a take-it-or-leave-it contract to a CEO with unknown knowledge $z \sim U[0,1]$. In Date (i) of period $t = 2$, the firm offers a take-it-or-leave-it contract to either the incumbent CEO (i.e., the same CEO as in $t = 1$) or a new CEO (again with knowledge $z \sim U[0,1]$). At the beginning of each period, an employed CEO enjoys a nonpecuniary private benefit of $B > 0$. CEOs have no external wealth. Firms can offer CEOs bonuses contingent on output $y_t \in \{0,1\}$ at the end of each period. Without loss of generality, we assume CEOs receive positive bonuses only when the output is high (i.e., $y_t = 1$). Let w_t denote the bonus at the end of period t conditional on $y_t = 1$, $t \in \{1,2\}$. We assume that the firm can commit to one-period contracts only. Thus, at $t = 1$, the firm cannot commit to bonuses that depend on the CEO's employment status at $t = 2$. That is, at each new period, the firm and the CEO must negotiate a new contract. We also assume that the CEO can "credibly commit" to consuming any compensation received at the end of a period, for example, by borrowing against that amount. This assumption implies that the firm cannot recoup w_1 as part of the negotiation of the $t = 2$ contract, thus avoiding the introduction of lifetime contracts through the back door.

In reality, CEO contracts are either fixed-term or "at will" (i.e., with no explicit termination date; see Gonzalez-Uribe and Groen-Xu (2017)). Fixed-term contracts usually last no longer than five years, but renewal is common. Most companies opt for at-will contracts, which do not provide protection against dismissal and are, in effect, a series of short-duration contracts that are continuously renegotiated. We can interpret our one-period contracts either as fixed-term contracts with the possibility of renewal or as at-will contracts.[11]

[11]Our assumption of short or fixed-term contracts is realistic, but it is also with a loss of generality. At the cost of increased complexity, we could introduce alternative frictions that deliver similar qualitative results. However, the model would lose much of its simplicity and tractability.

At the beginning of Date (ii), if the CEO exerts effort at a fixed personal cost of $c > 0$, she identifies a "deal" of unknown difficulty $x \sim U[0,1]$. A deal is a problem that needs to be solved. The CEO does not observe x. One interpretation is that c is the CEO's personal cost of "deal sourcing."

Once the CEO identifies a deal, she still needs to spend time assessing it before deciding whether to invest. Immediately after identifying a deal, the CEO learns whether she can assess the deal by herself (with the help of AI as a skill-augmenting technology). That is, the CEO learns that she is one of two *ex post* types: $x \leq \alpha z$ (type g) or $x > \alpha z$ (type b). After learning her type, she chooses one of three actions: $m \in \{m_a, m_b, m_c\}$. The board can observe $\hat{m} \in \{m_b, \overline{m}_b\}$ where $\overline{m}_b := \{m_a, m_c\}$. To allow for mixed strategies, let $P_t(\tau)$ denote the period-t probability distribution over m (i.e., the CEO's behavior strategy) conditional on the CEO's knowledge of her *ex post* type, $\tau \in \{b, g\}$. $P_t(\tau)$ determines $p_t(\tau) \in [0,1]$, which is the probability that a CEO of *ex post* type τ in period t asks the board for advice, i.e., the probability of $m = m_b$. If the CEO does not ask the board for help (i.e., $m \in \{m_a, m_c\}$), the board is unaware of the deal and, thus, cannot act.[12]

The board's time and attention are constrained. If the CEO asks the board for help, the board reduces the time and attention it allocates to routine "compliance tasks." Specifically, diverting the board's attention away from the compliance tasks destroys one unit of output with probability $\eta \in (0,1)$. That is, if the board spends one unit of time helping the CEO assess a deal, the expected output is $1 - \eta$. Thus, if the CEO can perform the assessment herself (i.e., if her *ex post* type is g), it is best not to involve the board. Henceforth, we refer to $1 - \eta$ as the *board's problem-solving capability*.[13]

At Date (iii) of $t = 1$, the board decides whether to replace the CEO. If the board is "weak" (state of the world ω^{weak}), the CEO is retained. If the board is "tough" (state of the world ω^{tough}), the CEO is fired with probability $\rho(\hat{m}) \in [0,1]$, where $\hat{m} \in \{m_c, \overline{m}_c\}$ and $\rho(.)$

[12] Even if the board knows that the CEO has identified a deal, it is unaware of x unless the CEO shares this information with the board.

[13] Parameter η plays the role of the "cost of help" in standard knowledge hierarchy models (e.g., Ide and Talamàs (2025)).

is optimally chosen by the board to maximize the second period expected profit. CEOs do not want to be replaced because they would lose the nonpecuniary private benefit of $B > 0$ in $t = 2$.

An important timing assumption is that the decision to fire or retain the CEO happens at Date (iii), which is before cash flows are realized (Date (iv)). The interpretation is that corporate policies may have long-term consequences, and decisions to retain a CEO may occur before the full impact of these actions is known. Formally, all we need is that the CEO's decision to ask the board for help at Date (ii) has incremental informativeness for the board's decision to fire or retain the CEO. Our timing assumptions retain this property in the simplest possible manner.

3 AI and Board Learning

In this section, we begin our analysis by focusing on the board's learning problem. In period $t = 1$, the CEO learns more about her own knowledge by trying to solve a problem at Date (ii). If the CEO learns she cannot solve a problem (i.e., her *ex post* type is b), she decides whether to seek help from the AI agent ($m = m_a$) or the board ($m = m_b$). To focus on the more interesting case, we make the following assumption:

Assumption 1. *The board has greater problem-solving capability than the AI agent:* $1 - \eta > \kappa$.

Assumption 1 implies that a CEO who identifies a deal she cannot assess faces a trade-off. On the one hand, if she asks the board to help, she will receive better advice than that of the AI agent. On the other hand, the board may learn that she was unable to assess the deal. In that case, the board may decide to fire the CEO on Date (iii). If Assumption 1 does not hold, no such trade-off exists, and the CEO does not ask the board for help. In that case, the CEO is never fired, regardless of the monitoring intensity. This illustrates one way in which AI can be detrimental to corporate governance: if an AI agent is better than the board at assessing deals and solving problems, the board will not receive any

information about the CEO's ability and will be unable to monitor her effectively. In contrast, under Assumption 1, the board's advice remains valuable, and the results are more nuanced.

The unconditional probability that a randomly chosen CEO is of *ex post* type g is

$$\Pr(g) = \Pr(x \leq \alpha z) = \int_0^1 \int_0^1 \mathbb{1}_{x \leq \alpha z} \, dx \, dz. \tag{1}$$

For a given z, the CEO solves all problems with $x \leq \alpha z$. Thus,

$$\Pr(g) = \int_0^1 \min(\alpha z, 1) \, dz = \int_0^{1/\alpha} \alpha z \, dz + \int_{1/\alpha}^1 1 \, dz = 1 - \frac{1}{2\alpha}. \tag{2}$$

As expected, if the AI agent has greater skill-augmenting capabilities, the probability that the CEO solves the problem is higher. In contrast, and by definition, the AI agent's problem-solving capability does not affect the CEO's ability to solve the problem independently.

Let S_2 denote the event in which a retained CEO can solve the problem in $t = 2$. Let $\Pr(S_2 \mid \tau)$ denote the probability of S_2 conditional on the CEO being "good" or "bad," $\tau \in \{g, b\}$. We have the following result.

Lemma 1 (Learning). *If the CEO from $t = 1$ is retained in $t = 2$, the probability that she can solve a newly drawn problem is $\Pr(S_2|g) = \frac{2(3\alpha-2)}{3(2\alpha-1)}$ if she is good and $\Pr(S_2|b) = \frac{1}{3}$ if she is bad.*

Note that the AI agent's skill-augmenting capability makes the "good" CEO better, but has no effect on the expected output of a "bad" CEO in period 2. The latter result arises because when AI is more skill-augmenting (larger α), the CEO who fails to solve the problem is, in expectation, of even lower type. This negative selection effect exactly cancels out the positive skill-augmenting effect in period 2.

When facing a bad CEO, if the board is "tough," it has the option to replace the CEO with a new CEO of unknown type. A newly-hired CEO has an average type of $E[z] = \frac{1}{2}$; call this type $\tau = n$. Note that, in $t = 2$, there are no career concerns, and thus CEOs

always ask the board for help if they cannot solve a problem. With a newly-hired CEO, the expected output is $E[y_2 \mid n] = (1 - \frac{1}{2\alpha}) + \frac{1}{2\alpha}(1 - \eta) = 1 - \frac{\eta}{2\alpha}$. If a bad incumbent CEO is retained, the expected output is $E[y_2 \mid b] = 1 - \frac{2\eta}{3}$. We define the expected output gain from replacing a bad CEO with a CEO of average type as

$$M(\alpha, \eta) := 1 - \frac{\eta}{2\alpha} - 1 + \frac{2\eta}{3} = \eta\left(\frac{2}{3} - \frac{1}{2\alpha}\right) > 0. \tag{3}$$

We can interpret $M(\alpha, \eta)$ as the value of the option to replace a bad CEO, or the *(marginal) benefit of monitoring*. Note that this value increases with the AI agent's skill-augmenting capability, α, and with the cost of board advising, η. By contrast, AI's problem-solving capability does not affect the value of board learning.

If a CEO of type τ is employed in $t = 2$, the shareholders must offer her a contract with contingent bonus $w_{2\tau}$, where the subscript τ indicates that the shareholders know that the CEO is of type τ. Unless bonus $w_{2\tau}$ is sufficiently high, the CEO will not exert effort in period $t = 2$. To focus on the more interesting case, we make the following assumption.

Assumption 2. *Exerting effort is always efficient:* $1 - \frac{2}{3}\eta \geq c$.

Assumption 2 implies that it is efficient for a "bad" CEO to exert effort in $t = 2$. That is, shareholders find it optimal to induce all (known) CEO types to exert effort in $t = 2$. This assumption also implies that an equilibrium without board learning does not exist, as we show next.

Lemma 2 (No Uninformative Equilibrium). *There is no equilibrium with positive effort provision in both periods in which* $p_t(b) = p_t(g)$ *for* $t \in \{1, 2\}$.

Recall that $p_t(\tau)$ is the probability that a CEO with *ex post* type $\tau \in \{b, g\}$ asks the board for advice (i.e., $m = m_b$). Lemma 2 implies that CEOs of different *ex post* types must communicate with the board with different probabilities. It then follows that there is no full pooling of types in equilibrium; with some positive probability, the board must learn something about τ from the equilibrium play. In equilibrium, there could be either full

separation of types or partial separation, in which case both types may choose the same action, but with different probabilities.

4 The Pre-AI Boardroom

In this section, we consider the case of a "pre-AI economy," where $\alpha = 1$ and $\kappa = 0$. This case serves as a benchmark. Solving this simpler version reveals the mechanisms more clearly and helps build intuition. We extend the benchmark model to the general case in the next section.

At the beginning of the first period, shareholders choose a pair (w_1, π). For each (w_1, π), players play a sequential game of incomplete information. The solution concept is Perfect Bayesian Equilibrium.

To solve for an equilibrium, note first that because $\kappa = 0$, at Date (ii) of either period action m_c weakly dominates m_a, and thus only two choices are relevant: $m \in \{m_b, m_c\}$. Because there are no career concerns in $t = 2$, the CEO's decision whether to ask the board for help in this period is trivial. If the CEO knows how to solve the problem, she will choose m_c because it delivers w_2 with probability 1, while choosing m_b delivers w_2 with probability $1 - \eta$. Similarly, if the CEO does not know how to solve the problem, it is strictly better to ask for help and gain w_2 with probability $1 - \eta$.

In the first period, at Date (iii) the board updates its beliefs about the CEO after observing $\hat{m} = m \in \{m_b, m_c\}$. Lemma 2 implies that, if effort is provided in the first period (which is the only equilibrium type we need to consider), both nodes at this date (m_b, or "ask for help," and m_c, or "do not ask for help") must be reached with positive probability in equilibrium. Thus, the board's beliefs about the CEO's type are fully pinned down by Bayesian updating, and there are no off-equilibrium beliefs to be determined. Because the game is sequential and beliefs are uniquely determined, for each (w_1, π) there is a (generically) unique equilibrium of the continuation game that follows. Since shareholders move first, they select their preferred equilibria by choosing (w_1, π). We refer to the

shareholders' preferred equilibrium as the *optimal equilibrium*.[14]

We have the following result.

Lemma 3 (Good CEOs Don't Ask for Help). *In equilibrium,* $p_1(g) = 0$.

Intuitively, because ex post "good" CEOs are better at solving problems than the board, not asking the board for help must be seen as a positive signal by the board. By contrast, ex post "bad" CEOs face a trade-off: while asking for help increases the chances of solving the problem, it also increases the probability that the CEO is fired. Thus, only type-b CEOs can choose mixed strategies in equilibrium: while $p_1(g) = 0$, we must have $p_1(b) > 0$ (due to Lemma 2). It follows immediately from Lemma 3 that after observing m_b, the board knows that the CEO must be of the bad type. Thus, a tough board should always replace a CEO who asks for help: $\rho(m_b) = 1$.

To streamline the exposition, we will henceforth restrict the analysis to pure-strategy equilibria. As we will formally show in the proof of Proposition 1, restricting the analysis to pure strategies is without loss of generality because an optimal equilibrium must be in pure strategies.

Under pure strategies, Lemmas 2 and 3 jointly imply $p_1(b) = 1$, that is, an *ex post* bad CEO always asks the board for help. That is, in equilibrium, bad and good CEOs must make different choices at Date (ii), and these choices fully reveal the CEO's type. Accordingly, we solve the model under the assumption that shareholders expect a bad CEO to request the board's help at $t = 1$. Then, we show that this belief is consistent with equilibrium play, that is, $p_1(b) = 1$.

Solving the model requires checking six incentive compatibility constraints. Working backwards, we begin at Date (ii) in $t = 2$ (note that Date (iii) is redundant in $t = 2$ because the CEO cannot stay for another period). At this point, the CEO has no career concerns and thus does not need to be incentivized to ask the board for help when she is unable to

[14] Generic uniqueness allows for a measure zero of pairs (w_1, π) to induce multiple equilibria in the continuation game. We will show in the proof of Proposition 1 that this multiplicity does not create difficulties for finding the optimal equilibrium, which is always unique.

assess a deal. However, the CEO still needs to be incentivized to exert effort c to identify a deal.

Suppose first that the CEO from $t = 1$ is retained in $t = 2$. If the CEO is of type b, from Lemma 1 we know that $E[z \mid b] = \Pr(S_2 \mid b) = \frac{1}{3}$. Let w_{2b} denote the bonus offered to a CEO of type b conditional on $y_2 = 1$. The probability that the output is $y_2 = 1$ if the CEO is of type b is $\frac{1}{3} + \frac{2}{3}(1 - \eta) = 1 - \frac{2}{3}\eta$. A bad CEO's "effort incentive constraint" at Date (ii) of $t = 2$ is

$$\Big(1 - \frac{2}{3}\eta\Big) w_{2b} - c + B \geq B. \tag{4}$$

The left-hand side of (4) is the bad CEO's payoff in $t = 2$ if she exerts effort. By exerting effort at a personal cost c, the CEO produces $y_2 = 1$ with probability $1 - \frac{2}{3}\eta$, in which case she earns w_{2b}, on top of her private benefit, B. If she does not exert effort, she collects her private benefit, which is the right-hand side of (4). At Date (i) of $t = 2$, if shareholders know that the CEO is of type b, they will offer the lowest possible bonus such that (4) holds, implying $w^*_{2b} = \frac{3c}{3-2\eta}$.

Similarly, if a retained CEO is of type g, from Lemma 1 we know that $E[z \mid g] = \Pr(S_2 \mid g) = \frac{2}{3}$, and the effort incentive constraint is

$$\Big(1 - \frac{\eta}{3}\Big) w_{2g} - c + B \geq B, \tag{5}$$

implying $w^*_{2g} = \frac{3c}{3-\eta}$.

In the case of a newly appointed CEO, which we denote by type n, we have that $E[z \mid n] = \Pr(S_2 \mid n) = \frac{1}{2}$, and the effort incentive constraint is

$$\Big(1 - \frac{\eta}{2}\Big) w_{2n} - c + B \geq B, \tag{6}$$

implying $w^*_{2n} = \frac{2c}{2-\eta}$.

By definition, a "weak" board will never replace the CEO. At Date (iii) of $t = 1$, because the marginal benefit of monitoring is positive ($M(1, \eta) > 0$), a "tough" board will always

replace a CEO of type b, but will not replace a CEO of type g. Intuitively, there is no reason for a board to retain an *ex post* bad CEO because an outside replacement is expected to be more qualified. Similarly, an *ex post* good CEO is always retained because she is expected to be more qualified than outsiders.[15] That is, in equilibrium we must have $\rho(m_b) = 1$ and $\rho(m_c) = 0$.

Consider now Date (ii) of $t = 1$. Recall that w_1 is the bonus conditional on $y_1 = 1$. First, we need to make sure that a "good" CEO does not want to pretend she is "bad":

$$w_1 + B \geq (1-\eta)w_1 + (1-\pi)\Big[\Big(1-\frac{\eta}{3}\Big)w_{2b}^* - c + B\Big]. \tag{7}$$

The left-hand side of (7) is the (incremental) payoff a good CEO expects by solving the problem herself, without involving the board. In that case, the CEO guarantees w_1 and knows that she will be retained in $t = 2$. She will then enjoy her private benefit B and exert effort c in return for a second-period bonus of w_{2g}. In that case, her $t = 2$ utility is $\big(1-\frac{\eta}{3}\big)w_{2g}^* - c + B = B$. The right-hand side of (7) is the payoff a good CEO expects by involving the board. If a good CEO asks the board for help, the board solves the problem with probability $1-\eta$ and, with probability $1-\pi$, it retains the CEO (whom the board now thinks is bad) and offers her a second-period bonus of w_{2b}^*.[16]

Second, still on Date (ii), a bad CEO must be incentivized to ask the board for help:

$$(1-\eta)w_1 + (1-\pi)B \geq B. \tag{8}$$

Condition (8) is the "(asking for) help incentive constraint." If a CEO who cannot assess

[15] An implicit assumption here is that only the firm (or its board) observes m and the decision of whether to retain the CEO, implying that an *ex post* good CEO does not have better outside opportunities. Allowing CEO retention decisions to signal quality to competing employers (as in Waldman (1984), Greenwald (1986), Li (2013), and Ferreira and Nikolowa (2023), among others) complicates the analysis but does not add significant insights.

[16] We assume that when a good CEO asks the board for help, the assessment of the deal is delegated to the board. The board succeeds at assessing the deal with probability $1-\eta$; the CEO's knowledge here is irrelevant.

a deal asks the board for help, the board will solve the problem with probability $1-\eta$, in which case the CEO receives w_1. By revealing her type, the CEO is now retained with probability $1-\pi$, in which case she enjoys B.[17] If the CEO chooses not to ask for help, she will not solve the problem, but will be retained with probability 1, because $\rho(m_c)=0$. If retained, the CEO is guaranteed to enjoy her private benefit B at $t=2$. Because the board thinks the CEO is good, it will offer her a bonus of w_{2g}^*. The CEO will choose not to exert effort because her IC in (4) is violated at w_{2g}^*.

The help IC constraint can also be written as

$$\pi \leq \frac{(1-\eta)w_1}{B}. \tag{9}$$

Intuitively, for the CEO to share information with the board, the board must be sufficiently "friendly," that is, its monitoring intensity π should not be too high. The right-hand side of (9) is the ratio between the CEO's benefit from the board's help and the cost of sharing information with the board. A "bad" CEO benefits from asking the board for help because the board can solve a problem with probability $1-\eta$, in which case the CEO receives w_1. The cost of sharing information is that, if the board is "tough," the CEO will be fired and lose her private benefit, B. This logic is reminiscent of Adams and Ferreira's (2007) model.

The sixth and final constraint is the first-period "effort incentive compatibility" constraint:

$$\underbrace{B-c}_{t=1 \text{ utility}} + \frac{1}{2}\underbrace{(w_1+B)}_{\text{g's utility}} + \frac{1}{2}\underbrace{((1-\eta)w_1+(1-\pi)B)}_{\text{b's utility}} \geq 2B. \tag{10}$$

The left-hand side of (10) is the CEO's lifetime expected payoff if she exerts effort at $t=1$ and behaves optimally throughout. By exerting effort, the CEO identifies a deal and eventually learns her type (g or b), each with equal probabilities. The right-hand side of (10) is the CEO's payoff if she chooses instead not to exert effort at $t=1$. In that case, the first-period output is $y_1=0$, implying that she does not earn w_1. In addition, she does

[17] The CEO is also promised bonus w_{2b}^* in $t=2$ in exchange for effort c, but because the effort constraint in (4) binds, the CEO's exactly breaks even.

not learn her type; she believes her knowledge is average ($E[z] = \frac{1}{2}$). Because she does not ask the board for help, she is retained in $t = 2$ and offered a bonus of w_{2g}^*. Since w_{2g}^* violates the second-period IC constraint for type n (see (6)), she does not exert effort at $t = 2$. Thus, her lifetime expected utility is $2B$.

The shareholders' problem at Date (i) of $t = 1$ is to choose a first-period bonus $w_1 \geq 0$ and a board monitoring intensity (i.e., independence) $\pi \in [0, 1]$ to maximize its expected profit, subject to constraints (7), (8) and (10), taking w_{2b}^*, w_{2g}^*, and w_{2n}^* as given.[18] This is a linear programming problem. To solve it, we proceed by initially assuming that (7) is slack. After solving for the optimal values, we then check (and confirm) that (7) is indeed slack at these values. With a slight abuse of terminology, we call a pair (w_1, π) a *governance structure*.[19]

Consider first the case in which $B \leq (1 - \eta)2c$. In this case, it can be readily checked that, for all $\pi \in [0, 1]$, the effort IC constraint (10) implies the help IC constraint (8). Thus, the optimal board monitoring intensity π^* must be a corner solution that binds constraint (10) only. From (11), we see that the marginal benefit of monitoring a bad CEO is $M(1, \eta) = \frac{\eta}{6}$. Thus, the firm wants to monitor more intensively as long as the marginal benefit exceeds what the bad CEO expects to lose from monitoring, which is B. If $M(1, \eta) > B$, the optimal board monitoring intensity is $\pi^* = 1$; if $M(1, \eta) < B$, the optimal board structure is $\pi^* = 0$. Thus, although a friendly board might be optimal in this case, the governance structure is determined not by the CEO's incentives to communicate but rather by the incentives for effort provision, unlike Adams and Ferreira (2007).

The interesting case is, of course, when the effort IC constraint (10) does not imply the help IC constraint (8). To make sure that this happens for any value of η, from now on, we make the following assumption.

Assumption 3. *The effort IC constraint does not imply the help IC constraint: $B \geq 2c$.*

[18] In some cases, it might be optimal to promise more than the full first-period output to the CEO ($w_1 > 1$), if that is necessary to induce CEOs to reveal their types.

[19] Strictly speaking, the governance structure should also include the second-period bonuses: (w_{2b}, w_{2g}, w_{2n}).

Figure 1 illustrates constraints (8) and (10) on the (w_1, π) plane, for a set of parameters that satisfy Assumption 3. Note that the help IC constraint passes through the origin and has a smaller slope than the effort IC constraint. The incentive-compatible region is the gray area to the right of both boundaries.

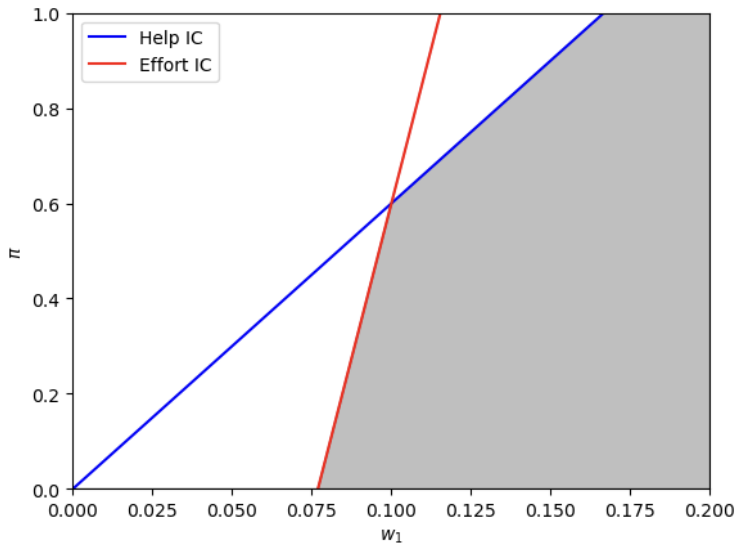


Figure 2: Incentive-compatible governance structures

Given the optimal $t = 2$ bonuses, and assuming that constraints (8) and (10) bind, the firm's expected profit at the beginning of $t = 1$ is

$$\begin{aligned} V(w_1, \pi) &= \frac{1}{2}\Big(\underbrace{2 - \frac{\eta}{3} - c - w_1}_{\text{Profit from good CEO}} \Big) + \frac{1}{2}\Big(\underbrace{2 - \frac{5}{3}\eta - c - (1-\eta)w_1 + \pi\frac{\eta}{6}}_{\text{Profit from bad CEO}} \Big) \\ &= 2 - \eta - c - (1 - \frac{\eta}{2})w_1 + \frac{\pi}{2}M(1, \eta). \end{aligned} \tag{11}$$

From (11) we can write the *isoprofit* for a given level of profit V as

$$\pi(w_1) = \frac{1}{M(1,\eta)}[(2-\eta)w_1 + 2(\eta + c) - 4 + 2V]. \tag{12}$$

The profit increases as we move towards the Northwest in Figure 2. Shareholders must

thus choose a governance structure on the highest isoprofit that intersects the incentive-compatible region.

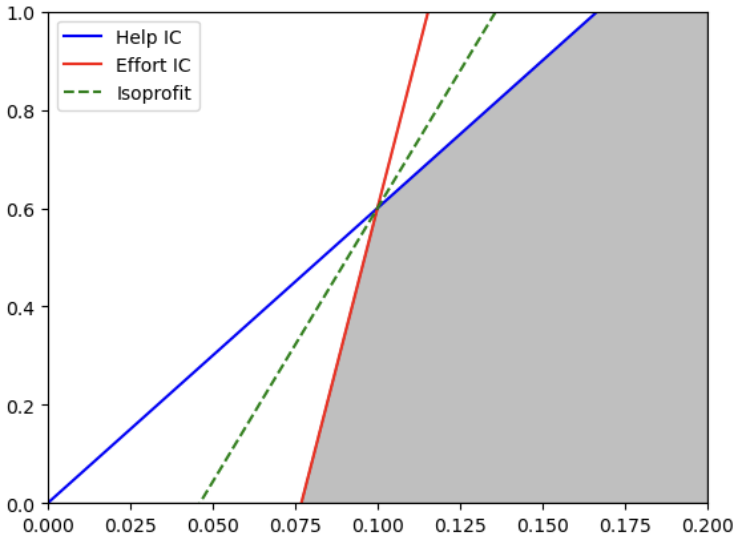


Figure 3: Optimal governance structure when the solution is interior

Figure 3 shows the unique governance structure (w_1^*, π^*) when the solution is interior. This case happens when the isoprofit's slope is between the help IC's slope and the effort IC's slope. This condition can be written as

$$\frac{M(1,\eta)}{B} \in \left(1, \frac{2-\eta}{1-\eta}\right). \tag{13}$$

To understand (13), notice that if the benefit of monitoring $M(1,\eta)$ is less than the cost of monitoring B, then we must be at a corner solution where the optimal monitoring is zero ($\pi^* = 0$). If, instead, $M(1,\eta) > B$, an interior solution obtains unless the benefit of monitoring is "too high" (i.e., $M(1,\eta) > \frac{2-\eta}{1-\eta}B$), in which case the optimal solution requires $\pi^* = 1$.

The following proposition formally characterizes the optimal governance structure.

Proposition 1. *The (generically unique)*[20] *optimal governance structure is*

$$(w_1^*, \pi^*) = \begin{cases} (\frac{2c}{2-\eta}, 0) & \text{if } M(1,\eta) < B \\ (2c, \frac{(1-\eta)2c}{B}) & \text{if } \frac{M(1,\eta)}{B} \in (1, \frac{2-\eta}{1-\eta}) \\ (\frac{B}{1-\eta}, 1) & \text{if } M(1,\eta) > (\frac{2-\eta}{1-\eta})B \end{cases} \tag{14}$$

Proposition 1 shows that a CEO-friendly board ($\pi^* < 1$) is optimal unless the benefit of monitoring is sufficiently high ($M(1,\eta) > \frac{2-\eta}{1-\eta}B$). Under an interior solution, (w_1^*, π^*) solves (8) and (10):

$$w_1^* = 2c \text{ and } \pi^* = \frac{(1-\eta)2c}{B}. \tag{15}$$

Replacing (w_1^*, π^*) in (7) confirms that this constraint is slack (see the Appendix for the proof). The optimal board "friendliness" is $1 - \pi^* > 0$.

5 AI Adoption by the CEO

In this section, we modify the benchmark model to allow the CEO to have access to an AI agent. We start with the simpler case in which the AI agent has problem-solving capabilities ($\kappa > 0$) but no skill-augmenting capabilities ($\alpha = 1$) (Subsection 5.1), before considering the case in which AI has both capabilities (Subsection 5.2). Since the analysis follows essentially the same steps as in the benchmark case, for brevity, we present only the steps necessary to understand the intuition behind the results.

[20] In the knife-edge cases of $M(1,\eta) = B$ or $M(1,\eta) = (\frac{2-\eta}{1-\eta})B$, there are multiple values of π that maximize profit.

5.1 CEO's Adoption of AI with Problem-Solving Capabilities

Let $\kappa > 0$ and $\alpha = 1$. The only difference from the benchmark model is that the help IC constraint is now

$$(1-\eta)w_1 + (1-\pi)B \geq \kappa w_1 + B. \tag{16}$$

Intuitively, a CEO who does not know how to assess a deal can now ask the AI agent for help, rather than consulting the board.[21] That is, a bad CEO who chooses not to ask the board for help must choose m_a as it strictly dominates m_c. This option delivers expected payoff $\kappa w_1 + B$, making the help IC constraint harder to satisfy. In turn, shareholders now need a friendlier board to induce the CEO to communicate truthfully with the board.

Figure 4 illustrates the introduction of an AI agent with problem-solving capabilities. The help IC constraint becomes flatter due to $\kappa > 0$, while the effort IC remains the same. The intersection now implies a lower monitoring intensity and a lower first-period bonus in an interior solution. Solving (10) and (16) simultaneously leads to

$$w_{1AI}^* = \frac{2c}{1+\kappa} < w_1^*. \tag{17}$$

and

$$\pi_{AI}^* = \frac{(1-\eta-\kappa)2c}{B(1+\kappa)} < \pi^*. \tag{18}$$

That is, when AI has problem-solving capabilities, in an interior solution, firms must choose a lower monitoring intensity. Intuitively, by offering an imperfect alternative to the board's advice, AI increases the CEO's payoff from not seeking help from the board. This makes the CEO less likely to reveal information that can lead to her dismissal. To encourage the CEO to communicate with the board, the board must become even friendlier. As a consequence, CEOs become more entrenched and are replaced less often. More CEO

[21] As previously discussed, the AI agent always offers a solution, which may not be correct. Thus, the CEO does not know whether the AI solution is correct, implying that the sequential strategy of first asking the AI agent and then the board leads to the same expected outcome as bypassing the AI agent and asking the board directly.

entrenchment is inefficient because CEOs are kept even when a better replacement is available.

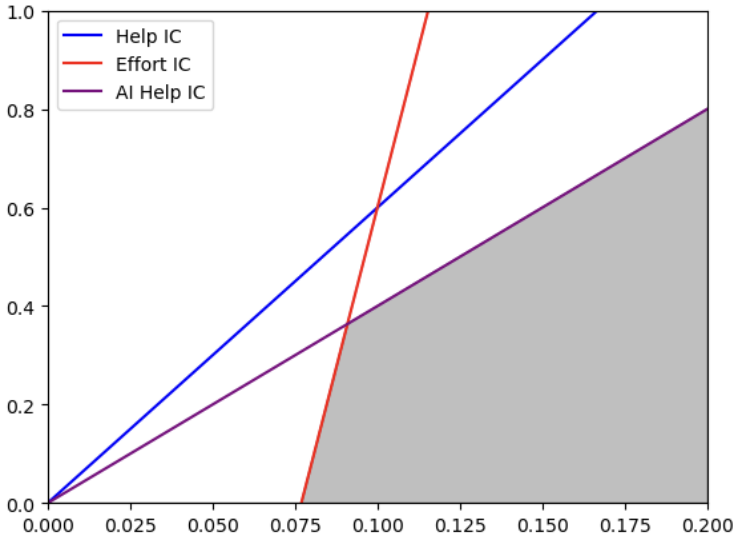


Figure 4: Incentive-compatible governance structures under AI ($\kappa > 0$ and $\alpha = 1$)

The next proposition formally states these results.

Proposition 2 (Governance and AI's problem-solving capability). *Let $\kappa \in [0, 1-\eta]$ and $\alpha = 1$. The optimal governance structure is interior and given by (17) and (18) if and only if*

$$\frac{M(1,\eta)}{B} \in \left(1, \frac{2-\eta}{1-\eta-\kappa}\right). \tag{19}$$

Under an interior solution, an increase in κ decreases the first-period bonus, the board's monitoring intensity, and the firm's profit. The CEO's expected lifetime utility is independent of κ.

Proposition 2 implies that improvements in AI's problem-solving capabilities can be inefficient: an increase in κ reduces firm value without affecting the CEO's utility. The firm is worse off with a larger κ because it needs to monitor less intensively to encourage the CEO to communicate with the board. Lower monitoring intensity increases CEO en-

trenchment, which is inefficient. Because the CEO's lifetime effort IC constraint is binding, the entrenchment cost is fully paid by the firm.

Because the first-period bonus is lower under $\kappa > 0$, while the second-period bonuses are the same, incentive pay falls with κ. We can write the first-period CEO's expected lifetime compensation as

$$E(w) = (1 - \frac{\eta}{2})w^*_{1AI} + \frac{1}{2}(1 - \frac{\eta}{3})w^*_{2g} + \frac{1}{2}(1 - \pi^*_{AI})(1 - \frac{2\eta}{3})w^*_{2b}. \tag{20}$$

The next corollary shows that expected total compensation also falls with κ.

Corollary 1. *Expected total compensation is decreasing in κ.*

Thus, an increase in κ reduces both the average bonus and expected total compensation. That is, while on the one hand, the CEO benefits from increased job stability, on the other hand, her compensation falls with improvements in AI's problem-solving capabilities. Because the first-period CEO's utility is $2B$ regardless of κ, the CEO is indifferent to the value of κ.

5.2 CEO's Adoption of AI with Problem-Solving and Skill-Enhancing Capabilities

We now extend the model to consider the case in which the AI agent possesses skill-augmenting capabilities (i.e., $\alpha \in [1, \overline{\alpha}]$), in addition to its problem-solving capabilities.

We note that the help IC constraint in (16) does not change. The second-period bonus for a bad CEO, w^*_{2b}, also does not change. The second-period bonus for a good CEO becomes

$$w^*_{2gAI} = \frac{3c}{3 - \frac{\eta}{2\alpha - 1}} < w^*_{2g} \tag{21}$$

and the second-period bonus for a newly-appointed CEO becomes

$$w^*_{2nAI} = \frac{2c}{2 - \frac{\eta}{\alpha}} < w^*_{2n}. \tag{22}$$

Intuitively, because the CEO is more competent with the help of AI, incentivizing her to exert effort is easier, requiring smaller bonuses overall.

The first-period effort incentive compatibility constraint becomes

$$B - c + (1 - \frac{1}{2\alpha})(w_1 + B) + \frac{1}{2\alpha}\left((1-\eta)w_1 + (1-\pi)B\right) \geq 2B. \tag{23}$$

Note that a larger α relaxes the effort IC because it increases the probability that the CEO can solve the problem by herself (State g).

The firm's expected profit at the beginning of $t = 1$ is

$$V(w_1, \pi) = \frac{2\alpha - 1}{2\alpha}\big(\underbrace{2 - \frac{\eta}{3(2\alpha - 1)} - c - w_1}_{\text{Profit from good CEO}}\big) + \frac{1}{2\alpha}\big(\underbrace{2 - \frac{5}{3}\eta - c - (1-\eta)w_1 + \pi M(\alpha, \eta)}_{\text{Profit from bad CEO}}\big). \tag{24}$$

Note that—all else constant—a larger α increases the firm's profit in both states (g and b), and also increases the probability that the CEO is perceived as "good." Because the profit is larger for type g than b for any governance structure (w_1, π), we conclude that α is a positive profit shifter.

The next proposition shows how α affects governance.

Proposition 3 (Governance and AI's skill-augmenting capability). *Let $\kappa \in [0, 1 - \eta]$ and $\alpha \geq 1$. The optimal governance structure is interior and given by*

$$w^*_{1AI} = \frac{2c\alpha}{2\alpha - 1 + \kappa} < w^*_1. \tag{25}$$

and

$$\pi^*_{AI} = \frac{(1-\eta-\kappa)2c\alpha}{B(2\alpha-1+\kappa)} < \pi^*, \tag{26}$$

if and only if

$$\frac{M(\alpha,\eta)}{B} \in \Big(1, \frac{2\alpha-\eta}{1-\eta-\kappa}\Big). \tag{27}$$

Under an interior solution, an increase in α decreases the first-period bonus and the board's monitoring intensity, and increases the firm's profit. The CEO's expected lifetime utility is independent of κ.

Proposition 3 shows that the monitoring intensity also decreases with the AI agent's skill-augmenting capability, α. An increase in α makes the CEO more productive, thereby increasing the benefit from exerting effort and thus loosening the CEO's effort IC constraint. This allows the firm to incentivize effort provision with a lower first-period bonus, w_1. But a lower first-period bonus makes the help IC constraint harder to meet, which then forces the firm to decrease board independence, π. Overall, as in the case of an increase in κ, both board independence and CEO compensation fall when α increases, as the following corollary shows.

Corollary 2. *Expected total compensation decreases with α.*

Despite the equilibrium displaying lower board independence, the firm is better off with a higher α because the increase in CEO productivity, coupled with the fall in compensation (in both periods), more than compensates for the (inefficient) lower CEO turnover rate. This result is robust and not driven by specific assumptions on parameters. This occurs because an increase in α both expands the incentive-compatibility region and shifts the profit function upward. Intuitively, α expands the technological frontier without (directly) affecting the CEO's incentives to share information. Thus, increasing α must be welfare increasing. Because the effort IC constraint binds, the CEO's utility is unaffected by α, and the firm captures all of the welfare increase in the form of profits.

To summarize, our analysis indicates that improvements in AI (both κ and α) lead to

a decrease in board independence and CEO turnover. AI improvements also reduce CEO incentive compensation, regardless of which capability is enhanced (κ or α). The welfare implications of AI adoption and improvements are ambiguous, as they depend on the relative strength of each AI capability (κ or α) and how these capabilities are affected by technological advances. If CEOs use AI primarily as a substitute for the board's advice, total welfare is reduced.

Form an empirical perspective, we note that these predictions relate the adoption of AI by the CEO, all else held constant. If a shock increases AI usage by the firm as a whole, e.g., knowledge workers now use AI more intensively, our predictions do not necessarily hold, because the resulting improvements in firm productivity may indirectly affect the determination of the optimal governance structures. Thus, empirical tests of our predictions must identify shocks to CEOs' use of AI that are orthogonal to other shocks affecting firm productivity and governance.

6 AI Adoption the Board

So far, we have assumed that AI adoption only benefits the CEO. But what if board members can also use AI when performing their tasks? In this section, we extend the model to consider how AI affects the board's performance in its advisory role (Subsection 6.1) and monitoring role (Subsection 6.2).

6.1 AI and the Board's Advisory Role

Suppose the board has access to an AI agent with problem-solving capability κ. If the CEO asks the board for advice, the board can either spend one unit of time assessing the deal x, in which case output is $y_t = 1$ with probability $1 - \eta$, or ask the AI agent for an assessment, in which case output is $y_t = 1$ with probability κ. Assumption 1 ($1 - \eta > \kappa$) implies that the board always prefers to assess the deal rather than delegate this task to

the AI agent. Thus, AI's problem-solving capabilities are irrelevant to the board.[22]

Things are more interesting when the board can benefit from AI's skill-augmenting capabilities. Although we have assumed that the board's knowledge is $z_a = 1$, AI can still reduce the cost of board advice by reducing η. For example, AI may reduce the time the board needs to spend on routine compliance tasks. Alternatively, we may assume that the board's knowledge is $z_a < 1$, and that the AI agent's skill-augmenting capability enables the board to solve harder problems *and* perform more routine compliance tasks per unit of time.

For simplicity, assume that the board's problem-solving capability is $\beta(1-\eta)$, where $\beta \in [1, \overline{\beta}]$ is the AI's *board skill-augmenting capability*. To find an interior equilibrium, we replace $1-\eta$ with $\beta(1-\eta)$ in the help IC constraint (16) and the effort IC constraint (23). The optimal first-period bonus is unchanged and given by (25). The optimal monitoring intensity becomes

$$\pi^*_{AI} = \frac{(\beta(1-\eta)-\kappa)2c\alpha}{B(2\alpha-1+\kappa)}. \tag{28}$$

It is immediate that increasing β increases the optimal monitoring intensity. Intuitively, an increase in β makes the board a better advisor, providing further incentives for the CEO to ask for help. Thus, all else constant, shareholders can choose a more independent board while still meeting the CEO's help IC constraint. We conclude that AI adoption by the board can mitigate some of the adverse effects of AI adoption by CEOs.

Because advances in AI are likely to shift (α, κ, β) in the same direction (but perhaps by different magnitudes), the comparative static exercises we have performed thus far have some apparent limitations. We are interested in shocks that improve (α, κ, β) simultaneously. Let a denote a latent variable measuring AI's overall capabilities. The interpretation is that a higher a increases each of AI's specific capabilities. The effect of changes in a on the monitoring intensity π^*_{AI} depends on the particular correlation structure among the

[22] If we drop Assumption 1 and let $\kappa \geq 1-\eta$, the immediate conclusion is that asking the board for advice is dominated by asking the AI agent. In this case, the CEO never communicates with the board and remains fully entrenched, independent of π.

elements in $(\alpha(a), \kappa(a), \beta(a))$. Define the *relative knowledge advantage* the board has over AI as an advisor as

$$R(a) = \beta(a)(1-\eta) - \kappa(a). \tag{29}$$

Assumption 1 and $\beta(a) \in [1, \overline{\beta}]$ imply $R(a) > 0$. We say that a shock $\epsilon_a > 0$ to the latent variable a is *biased against the board* if $R(a+\epsilon_a) - R(a) < 0$, *biased towards the board* if $R(a+\epsilon_a) - R(a) > 0$, and *neutral* if $R(a+\epsilon_a) - R(a) = 0$. Assuming differentiability, as $\epsilon_a \to 0$, the direction of the bias is determined by the sign of $R'(a)$.

We then have

$$\frac{\partial \pi^*_{AI}}{\partial a} = \frac{2c}{B(2\alpha(a) - 1 + \kappa(a))}\Big(R'(a)\alpha(a) - R(a)\frac{\alpha'(a)(1-\kappa(a)) + \kappa'(a)\alpha(a)}{2\alpha(a) - 1 + \kappa(a)}\Big). \tag{30}$$

Note that since $\alpha'(a)(1-\kappa(a)) + \kappa'(a)\alpha(a) > 0$, an AI improvement can increase the board's monitoring intensity only if it is *sufficiently* biased towards the board, that is, if

$$R'(a) > R(a)\frac{\alpha'(a)(1-\kappa(a)) + \kappa'(a)\alpha(a)}{\alpha(a)(2\alpha(a) - 1 + \kappa(a))}. \tag{31}$$

To help with interpretation, we rearrange (31) and simplify notation to get

$$\varepsilon_{Ra} \geq \frac{|\varepsilon_{w_1\alpha}|}{1-\kappa}\Big((1-\kappa)\varepsilon_{\alpha a} + \kappa\varepsilon_{\kappa a}\Big), \tag{32}$$

where ε_{yx} denotes the elasticity of variable y with respect to parameter x. Because the right-hand side of (32) is positive, for π^*_{AI} to increase with advances in AI, we need the board's relative knowledge advantage to be sufficiently responsive to AI technology shocks (i.e., high ε_{Ra}). This condition is harder to obtain when (i) the first-period bonus is very sensitive to AI's skill-augmenting capability (i.e., high $|\varepsilon_{w_1\alpha}|$), (ii) AI's skill-augmenting capability is very sensitive to AI technology shocks (i.e., high $\varepsilon_{\alpha a}$), and (iii) AI's problem-solving capability is very sensitive to AI technology shocks (i.e., high $\varepsilon_{\kappa a}$).

A natural benchmark is to consider changes in a that are *unbiased*, that is, changes in a

such that $\varepsilon_{Ra} = 0$. In that case, we have

$$\frac{\partial \pi^*_{AI}}{\partial a} = -2cR(a)\frac{\alpha'(a)(1-\kappa(a)) + \kappa'(a)\alpha(a)}{B(2\alpha(a) - 1 + \kappa(a))^2} < 0. \tag{33}$$

We thus conclude that a technological change that improves AI capabilities in an unbiased manner will lead to a decrease in the board's monitoring intensity. To put it differently, an AI improvement can increase the board's monitoring intensity only if it is sufficiently biased towards the board, that is, if the board's relative advantage over AI as an advisor increases significantly.

To understand the intuition, note that an increase in κ affects the optimal monitoring intensity through two routes. First, it reduces incentives to ask for help for a given first-period bonus. Second, it reduces the optimal first-period bonus, further reducing incentives to ask for help. In contrast, an increase in β increases incentives to ask for help for a given bonus, but does not affect the optimal bonus (note that the optimal bonus (25) is independent of η as well). Therefore, a shock of similar magnitude to both κ and $\beta(1-\eta)$ would not change incentives to ask for help for a given first-period bonus, but would still lower the bonus, tightening the help IC constraint. The net effect on the monitoring intensity is thus negative. In addition, if the same shock also increases α, the first-period bonus is further reduced, further contributing to the tightening of the help IC constraint.

6.2 AI and the Board's Monitoring Role

We now consider the case in which AI can assist the board in performing its monitoring task. Suppose that, with the help of AI, the board may independently learn some of the information the CEO possesses. For example, with the use of AI, the board may find it easier to process market information about the firm.[23] We maintain the assumption that the CEO must exert effort to identify a deal. That is, the board cannot identify a deal

[23]See Ferreira, Ferreira, and Raposo (2011) for evidence that the board's monitoring function is helped by information embedded in stock prices.

unknown to the CEO. For concreteness, suppose that at some point before Date (iii), with probability $\lambda \in [0,1]$ the board learns both about the existence of a deal (with unknown difficulty x) and whether $z \geq x$ or $z < x$, that is, whether the CEO can assess the deal.[24] At first blush, it may seem that a bad CEO would now be more willing to ask the board for help, because the board may learn that the CEO is bad even without the CEO asking for help. However, this reasoning is incomplete because when the board learns about a problem the CEO cannot solve, the board can solve it with probability $\beta(1-\eta)$, which benefits the CEO through her first-period bonus w_1.

The help IC constraint is

$$\beta(1-\eta)w_1 + (1-\pi)B \geq (1-\lambda)(\kappa w_1 + B) + \lambda\left(\beta(1-\eta)w_1 + (1-\pi)B\right), \tag{34}$$

which implies

$$\beta(1-\eta)w_1 + (1-\pi)B \geq \kappa w_1 + B. \tag{35}$$

That is, the board's ability to independently obtain the CEO's information has no impact on the help IC constraint. In an equilibrium in which the help IC constraint is met, the CEO always shares information with the board, thus λ also does not affect the effort IC constraints or the isoprofits, implying that the equilibrium is unaffected by λ.

The intuition behind this result stems from the fact that the help IC must bind in the optimal contract. When this constraint binds, the CEO receives the same expected payoff from seeking the board's help as from withholding information. Consequently, the CEO is indifferent between full disclosure and probabilistic disclosure, i.e., revealing information with any probability.

We can interpret this probability as the likelihood that the board discovers the problem through independent inspection. Crucially, because the CEO's payoff remains unchanged whether the board learns through voluntary disclosure or independent discovery, the help

[24] It is less reasonable to assume that the board can learn about the CEO's ability to assess a particular deal without learning also what the deal is.

IC constraint continues to hold even when the board can inspect. The binding nature of the help IC ensures that the optimal contract is independent of the board's inspection capability as long as it induces the CEO to disclose the information voluntarily.

It's worth noting that if the board could obtain verifiable evidence of the CEO's information withholding and could write contracts contingent on such evidence, this would make it easier for the CEO to communicate the information. However, such contractibility would violate a fundamental premise of our model: that the problems CEOs face are inherently non-contractible due to their complexity and the difficulty of ex-ante specification. That is, these are the types of strategic decisions and opportunities that cannot be fully anticipated or described in formal contracts.

Although AI may be helpful to the board in other ways, our conclusion here is that, within the logic of the model, the board's ability to obtain information independently of the CEO does not affect the CEO's incentives to exert effort or communicate with the board.

7 Discussion and conclusions

This paper develops a comprehensive framework for understanding how AI transforms the relationship between CEOs and corporate boards. Our model captures two distinct technological capabilities of AI: its problem-solving capability, which allows AI to substitute for human judgment, and its skill-augmenting capability, which enhances human decision-making abilities. We show that both capabilities reduce board monitoring intensity and CEO compensation, but through different mechanisms.

While AI's problem-solving capability unambiguously destroys value by making it more difficult for the board to identify and replace bad CEOs, its skill-augmenting capability can enhance firm value despite reducing board monitoring. When boards also adopt AI, some negative effects are mitigated, but only if technological advances are sufficiently biased toward enhancing board capabilities relative to those of the CEO.

Our analysis reveals a fundamental tension in AI adoption: technology that enhances CEOs' capabilities also reduces their dependence on board advice, disrupting traditional information flows that underpin effective governance. Firms respond by reducing board independence to restore communication incentives, leading to increased CEO entrenchment without corresponding benefits to CEOs themselves. This result adds a cautionary note to the conventional wisdom that technological advancement improves organizational efficiency. Instead, we show that AI adoption necessitates a reconfiguration of governance structures, adjusting both the intensity of monitoring and compensation packages.

The framework we develop extends beyond the CEO-board relationship to any setting where principals need agents to both exert effort and communicate truthfully. Consider a sales manager overseeing field agents who must both pursue leads (effort) and report difficulties in dealing with potential customers (communication). If AI tools enable agents to handle customer interactions better, they may stop sharing crucial information with managers. Similarly, procurement officers need to both identify suppliers (effort) and evaluate them. Better AI might cause the procurement officer to bypass marginal suppliers to the manager, causing operational risks. In each case, our model predicts that organizations must reduce monitoring intensity and adjust compensation to maintain information flows.

A broad theoretical insight from our analysis is that AI can paradoxically increase transaction costs by tightening incentive compatibility constraints. This occurs because AI affects not only the equilibrium payoff of the agent, but also his *off-equilibrium payoff*: the value of not asking for help in our setting. Organizations that fail to recognize this and adapt their governance structures accordingly may find that AI adoption leads to worse outcomes than the pre-AI era. Our framework, therefore, highlights an essential challenge for organizational design in the AI era: managing the incentive costs associated with AI adoption.

Appendix: Proofs

Proof of Lemma 1: Using the law of total probability, we have

$$\Pr(S_2|g) = \int_0^1 \Pr(S_2|z,g) \cdot f(z|g)\, dz.$$

Since problems are independent given z, then

$$\Pr(S_2|z,g) = \Pr(S_2|z) = \min(\alpha z, 1).$$

From Bayes' theorem, the posterior distribution of z given g is

$$f(z|g) = \frac{\Pr(g|z) \cdot f(z)}{\Pr(g)} = \frac{\min(\alpha z, 1)}{1 - \frac{1}{2\alpha}}.$$

Therefore

$$\Pr(S_2|g) = \frac{1}{1 - \frac{1}{2\alpha}} \int_0^1 [\min(\alpha z, 1)]^2\, dz = \frac{1}{1 - \frac{1}{2\alpha}} \int_0^{1/\alpha} (\alpha z)^2\, dz + \int_{1/\alpha}^1 1\, dz = \frac{2(3\alpha - 2)}{3(2\alpha - 1)}.$$

Similarly, we have

$$\Pr(S_2|b) = \int_0^1 \Pr(S_2|z,b) \cdot f(z|b)\, dz.$$

Since problems are independent given z, we have

$$\Pr(S_2|z,b) = \Pr(S_2|z) = \min(\alpha z, 1).$$

If $z > 1/\alpha$, the CEO can solve any problem. Therefore, $f(z|b) = 0$ for $z > 1/\alpha$. For $z \leq 1/\alpha$, we have $f(z|b) = 2\alpha(1 - \alpha z)$. Then

$$\Pr(S_2|b) = \int_0^{1/\alpha} \alpha z \cdot 2\alpha(1 - \alpha z)\, dz = 2\alpha^2 \left[\frac{z^2}{2} - \frac{\alpha z^3}{3}\right]_0^{1/\alpha} = \frac{1}{3}.$$

□

Proof of Lemma 2: Suppose the board is tough. At Date (iii) of $t = 1$, it must then decide whether to retain or fire the CEO. Suppose $p_1(\tau)$ is uninformative about τ, i.e., $p_1(\tau) = p$ is a constant. Thus, the board expects the incumbent CEO to be good with probability $1 - \frac{1}{2\alpha}$ regardless of whether the board observes m_b or $\{m_a, m_c\}$. Note that, in period 2, a CEO has no career concerns and thus does not need to be incentivized to ask the board for help when she cannot assess a deal. However, the CEO still needs to be incentivized to exert effort c to identify a deal. If a CEO of type τ is to exert effort, her second-period *effort incentive constraint* must hold:

$$\Big(\Pr(S_2 \mid \tau) + (1 - \Pr(S_2 \mid \tau))(1 - \eta)\Big) w_2 - c + B \geq B. \tag{A.1}$$

Suppose first that the board fires the CEO and hires a replacement. We denote the type of the replacement by n. We have $\Pr(S_2 \mid n) = 1 - \frac{1}{2\alpha}$. The board will make the new CEO's incentive constraint bind, which implies a second-period expected profit of

$$v_{2n} = 1 - \frac{\eta}{2\alpha} - c. \tag{A.2}$$

Suppose instead that the firm retains the incumbent CEO, and offers a wage that is sufficient to incentivize the bad CEO, $w_{2b} = \frac{3c}{3-2\eta}$ (here we consider the case in which the CEO exerts effort in the first period, so she privately learns about her type). Under this contract, both CEO types, b and g, will exert effort in period 2. Let v_{2gb} denote the second-period profit when the firm retains an incumbent of unknown type and both types are induced to exert effort. After algebra, we get

$$v_{2gb} = \frac{(3 - 2\eta - 3c)(2\alpha - \eta)}{2\alpha(3 - 2\eta)}$$

and thus

$$v_{2n} - v_{2gb} = \frac{c\eta(4\alpha - 3)}{2\alpha(3 - 2\eta)} > 0,$$

implying that the board strictly prefers hiring a new CEO to retaining a CEO of unknown type with a contract that induces both types to exert effort.

Alternatively, the board may choose w_2g such that it incentivizes only the good CEO. In this case, we have expected profit

$$v_{2g} = \frac{2\alpha - 1}{2\alpha}(1 - \frac{\eta}{3(2\alpha - 1)} - c) = \frac{1}{2\alpha}(1 - \frac{\eta}{3} - c),$$

and thus

$$v_{2n} - v_{2g} \geq 1 - \frac{\eta}{2} - c - \frac{1}{2}(1 - \frac{\eta}{3} - c) = \frac{1}{2}(1 - \frac{2\eta}{3} - c) > 0,$$

which follows from Assumption 2. Thus, if the board does not learn anything, it always prefers to replace the CEO. But if a bad CEO knows she will be replaced, she will ask for advice, because her expected payoff will be $(1 - \eta)w_1 > 0$ (because w_1 must be positive if the CEO is to exert effort in period 1, knowing she will be replaced in period 2). This implies that a bad CEO always shares information with the board, i.e., $p_1(b) = 1$. By contrast, a good CEO prefers not to ask the board for help, because she can solve the problem with probability 1, while the board solves it only with probability $1 - \eta$, implying $p_1(g) = 0$, contradicting $p_1(b) = p_1(g) = p$. We conclude that $p_1(g) \neq p_1(b)$.

Because there are no career concerns in $t = 2$, trivially $p_2(b) = 1$ and $p_2(g) = 0$. □

Proof of Lemma 3. Conditional on knowing that her own ex post type is g at $t = 1$, a CEO who chooses m_c expects to receive w_1 with probability 1. By choosing m_b instead, the CEO expects instead $(1 - \eta)w_1$. Thus, a good CEO would only choose w_b with strictly positive probability if asking the board for help increases her probability of retention, i.e., if $\rho(m_b) < \rho(m_c)$. But if $\rho(m_b) < \rho(m_c)$, action $m = m_b$ is strictly dominant for a CEO of type b, implying $p_1(b) = 1$. Lemma 2 then implies $p_1(g) < 1$. This means that when board observes m_b it must revise down its expectation of the CEO's knowledge, as the CEO is

more likely to be bad than good. But then choosing $\rho(m_b) < \rho(m_c)$ is not rational for shareholders, because a replacement CEO is equally likely to be good or bad, and the proof of Lemma 2 shows that replacement always dominates retention when the incumbent is equally likely to be good or bad. We conclude that we must have $p_1(g) = 0$. □

Proof of Proposition 1: **Case 1:** $M(1,\eta) < B$. In this case, the isoprofit is steeper than the effort IC constraint. As argued in the text, in this case, the optimal π must be zero. This implies that the help IC constraint is slack; the "bad" CEO always asks for help. The optimal first-period bonus is determined by the effort IC constraint alone:

$$B - c + \frac{1}{2}(w_1^* + B) + \frac{1}{2}\left((1-\eta)w_1^* + B\right) = 2B, \tag{A.3}$$

which implies $w_1^* = \frac{2c}{2-\eta}$. Replacing it in (11) yields $V(w_1^*, 0) = 2 - \eta - 2c$, which is positive given Assumption 2. Because violating the effort IC constraint leads to zero profit, choosing $w_1^* = \frac{2c}{2-\eta}$ and $\pi^* = 0$ must be the unique optimal solution for this set of parameters.

Case 2: $\frac{M(1,\eta)}{B} \in (1, \frac{2-\eta}{1-\eta})$. As illustrated in Figure 3, the tangency between the highest isoprofit and the IC region happens when both (8) and (10) bind. Solving for (w_1^*, π^*) yields:

$$w_1^* = 2c \text{ and } \pi^* = \frac{(1-\eta)2c}{B}. \tag{A.4}$$

To confirm that constraint (7) is slack, we replace (w_1^*, π^*) in (7) to get slack

$$\frac{c\big(6B - 5B\eta + 2c\eta(1-\eta)\big)}{B(3-2\eta)} > 0 \quad \text{for all } \eta \in (0,1),\ B > 0,\ c > 0.$$

Thus, the IC constraint in (7) is slack at the optimal solution.

To confirm that (A.4) is the unique optimal governance structure in this case, we also need to consider values of (w_1, π) that violate (8) under pure strategies. If any equilibrium also violates the first-period effort IC constraint, its overall profit is bounded above by v_{2n},

which is less than the profit under the optimal pure-strategy equilibrium. Thus, we do not need to consider these "shirking equilibria," as they are strictly dominated.

If (8) is violated, there is no pure-strategy (without shirking) equilibrium: with $\rho(m_c) = 0$ the bad type strictly prefers m_c, leading to pooling (ruled out by Lemma 2). With $\rho(m_c) = 1$, m_b is a dominant action for type b, so we must have $p_1(g) < 1$. Then, m_c would reveal type g and $\rho(m_c) = 1$ cannot be sequentially rational. Thus, for such values, equilibria with first-period effort must involve mixed strategies. We now show that such mixed-strategy equilibria always deliver lower profits than some pure-strategy equilibria associated with governance structure (w_1, π). That is, mixed-strategy equilibria are never optimal equilibria.

Suppose (w_1, π) is such that only a strictly mixed-strategy equilibrium exists where first-period effort constraint is satisfied. Let $p_1(b) = p \in (0,1)$ denote the equilibrium probability that a CEO of type b asks the board for help. Let $\rho(m_c) = \rho \in (0,1)$ be the equilibrium probability that a CEO who does not ask for help is replaced conditional on the board being tough ($\omega = \omega^{tough}$). The bad CEO must be indifferent between m_b and m_c:

$$(1-\eta)w_1 + (1-\pi)B = \pi(1-\rho)B + (1-\pi)B, \tag{A.5}$$

which implies $\pi = \frac{(1-\eta)w_1}{(1-\rho)B}$. We now consider two cases.

Suppose first that $w_1 \geq 2c$. If a strict mixed-strategy equilibrium exists for governance structure (w_1, π), a tough board must be indifferent between firing or retaining a CEO who did not ask for help, implying the expected profit from retaining the incumbent must equal $v_{2n} = 1 - \frac{\eta}{2} - c$, which is the expected second-period profit when a new CEO is hired. The weak board after m_c always retains the incumbent, facing the exact same posterior and offering the same contract. So its profit from retention is also v_{2n}. Thus, the unconditional expected second-period profit must be

$$v_{2n} - \frac{p}{2}(1-\pi)M(1,\eta).$$

We get this expression by noting that profit after m_c must be v_{2n}, while after m_b profit is v_{2n} if the CEO is replaced and v_{2b} if retained. The unconditional probability of retaining a bad CEO is $\frac{p}{2}(1-\pi)$, in which case the firm "loses" the marginal benefit of monitoring, $M(1,\eta)$.

By contrast, under (w_1^*, π^*), the optimal pure-strategy equilibrium implies an expected period 2 profit of

$$\frac{1}{2}v_g + \frac{1}{2}(\pi^* v_{2n} + (1-\pi^*)v_{2b}) = \frac{1}{2}(v_{2n} + M(1,\eta) + \pi^* v_{2n} + (1-\pi^*)(v_{2n} - M(1,\eta))$$

$$= v_{2n} + \frac{\pi^*}{2}M(1,\eta).$$

We conclude that the second-period profit is strictly larger under the pure-strategy equilibrium for the optimal governance structure, (w_1^*, π^*). Now we need to compare first-period profits. For a mixed-strategy equilibrium under (w_1, π), we have the first-period profit

$$\frac{1}{2}(1-w_1) + \frac{1}{2}p(1-\eta)(1-w_1),$$

while for the optimal pure strategy equilibrium we have

$$\frac{1}{2}(1-w_1^*) + \frac{1}{2}(1-\eta)(1-w_1^*).$$

Since $w_1^* = 2c$ and $w_1 \geq 2c$, it follows that the first-period profit is also larger under the pure-strategy equilibrium for the optimal governance structure. We conclude that, if $w_1 \geq 2c$, any mixed-strategy equilibria is dominated by the optimal pure-strategy equilibrium under (w_1^*, π^*).

Suppose now that $w_1 < 2c$. To prove this result, we first find an upper bound to the profit under (w_1, π) and then show that this upper bound is strictly smaller than the profit under the pure-strategy equilibrium induced by (w_1^*, π^*).

Define the *first-best surplus*, S^{FB}, as the total surplus to all players if the game is played

cooperatively: (i) the CEO exerts effort, (ii) a bad CEO always asks for help, and (iii) the board always replaces a bad CEO and retains a good CEO. S^{FB} denotes social surplus—output net of effort costs plus CEOs' private benefits. Since the firm employs one CEO each period and the employed CEO enjoys B, at least $2B$ of social surplus is non-extractable by shareholders. Let $\overline{V}^{FB} := S^{FB} - 2B$ denote the *extractable first-best surplus*: the maximum surplus shareholders can possibly appropriate. The logic is that an employed CEO can get at least B in that period by not exerting effort, so at least $2B$ of the surplus must go to CEOs.

In a pure-strategy equilibrium where both (8) and (10) bind, the $t = 1$ expected utility of an employed CEO is $2B$, which is the minimum that a CEO must get to exert effort at $t = 1$. Because (w_1^*, π^*) implies a pure strategy equilibrium where both (8) and (10) bind, the profit is $V(w_1^*, \pi^*) = \overline{V}^{FB} - \frac{1}{2}(1 - \pi^*)M(1, \eta)$. The intuition is that the only source of inefficiency under a pure-strategy optimal equilibrium is that with probability $\frac{1}{2}(1 - \pi^*)$ a bad CEO is retained, which implies a surplus loss of $M(1, \eta)$. In addition, in this equilibrium shareholders capture all of the extractable surplus

Define $\overline{V}^m(w_1, \pi)$ as the extractable surplus under a mixed-strategy equilibrium associated with (w_1, π). By definition, the extractable surplus is an upper bound for the profit under this equilibrium: $\overline{V}^m(w_1, \pi) \geq V(w_1, \pi)$.

We have

$$\overline{V}^m(w_1, \pi) = \overline{V}^{FB} - \frac{1 - \pi(p + (1 - p)\rho)}{2} M(1, \eta) - \frac{\pi\rho}{2} M(1, \eta) - \frac{1 - p}{2}(1 - \eta).$$

We derive this expression by subtracting from the extractable first-best surplus the expected losses from inefficient actions. There are three sources of inefficiency: (i) sometimes a bad CEO is retained (the second term on the right-hand side), (ii) sometimes a good CEO is replaced (the third term on the right-hand side), and (iii) sometimes a bad CEO does

not ask for help (the fourth term on the right-hand side). This expression simplifies to

$$\overline{V}^m(w_1, \pi) = \overline{V}^{FB} - \frac{1 - p\pi(1-\rho)}{2} M(1, \eta) - \frac{1-p}{2}(1-\eta).$$

Now, use the fact that $\pi = \frac{(1-\eta)w_1}{(1-\rho)B}$ and $w_1 < 2c = w_1^*$ to conclude that $\pi(1-\rho) < \pi^*$, so that

$$\frac{1 - p\pi(1-\rho)}{2} > \frac{1-\pi^*}{2},$$

implying $V(w_1^*, \pi^*) > \overline{V}^m \geq V(w_1, \pi)$. We conclude that the pure-strategy equilibrium under (w_1^*, π^*) has strictly higher profit than any mixed-strategy equilibrium.

As a final point, note that at (w_1^*, π^*), mixed strategy equilibria may also exist, with $p \in (0, 1)$ and $\rho = 0$. However, any such equilibrium trivially has a lower profit than the pure-strategy equilibrium. Furthermore, such equilibria are not robust in the following sense. If under (w_1^*, π^*) the CEO is expected to play a strict mixed strategy, then shareholders can choose instead contract $(w_1^* + \varepsilon, \pi^*)$ with $\varepsilon > 0$ infinitesimally small, which breaks the bad CEO's indifference between asking or not asking for help, leading to a unique equilibrium in pure strategies where the help IC constraint is (slightly) slack. The profit in this equilibrium is arbitrarily close to $V(w_1^*, \pi^*)$. We thus conclude that such mixed-strategy equilibria cannot be optimal and would never be chosen by shareholders.

Case 3: $M(1, \eta) > \frac{2-\eta}{1-\eta} B$. In this case, the isoprofit is flatter than the help IC constraint and, as argued in the text, the optimal board structure must be $\pi^* = 1$. The effort IC constraint is not binding, thus the help IC constraint determines the optimal first-period bonus: $w_1^* = \frac{B}{1-\eta}$. Replacing it in (11) yields $V(\frac{B}{1-\eta}, 1) = 2 - \eta - c - \frac{1-\frac{\eta}{2}}{1-\eta} B + \frac{\eta}{12}$. Because $M(1, \eta) > \frac{2-\eta}{1-\eta} B \Rightarrow \frac{\eta}{6} > \frac{2-\eta}{1-\eta} B \Rightarrow \frac{\eta}{12} > \frac{1-\frac{\eta}{2}}{1-\eta} B$, we have $V(\frac{B}{1-\eta}, 1) > 0$. □

Proof of Proposition 2: The effort IC constraint is independent of κ, and its slope is $\frac{2-\eta}{B}$. The isoprofit is also independent of κ, and its slope is $\frac{2-\eta}{M(1,\eta)}$. As we can see from Figure 3, an interior solution requires $M(1, \eta) > B$.

The help IC constraint's slope is now $\frac{1-\eta-\kappa}{B}$. The slope must be smaller than that of the isoprofit, implying

$$\frac{M(1,\eta)}{B} < \frac{2-\eta}{1-\eta-\kappa}.$$

All other steps are similar to those in the proof of Proposition 1 and are omitted for brevity. From (17) we have

$$\frac{\partial w_{1AI}^*}{\partial \kappa} = -\frac{2c}{(1+\kappa)^2} < 0,$$

and from (18) we have

$$\frac{\partial \pi_{AI}^*}{\partial \kappa} = \frac{2c(\eta-2)}{B(1+\kappa)^2} < 0.$$

From (11) we have

$$\frac{\partial V_{AI}^*}{\partial \kappa} = (1-\frac{\eta}{2})\frac{2c}{(1+\kappa)^2} + \frac{c(\eta-2)M(1,\eta)}{B(1+\kappa)^2} = \frac{c(2-\eta)}{(1+\kappa)^2}\Big(1-\frac{M(1,\eta)}{B}\Big) < 0,$$

which follows because $M(1,\eta) > B$ in an interior solution. Because the effort IC constraint in (10) binds, the CEO's lifetime utility is $2B$. □

Proof of Corollary 1: Because the first-period effort IC constraint binds, we must have

$$E(w) + B - c + \frac{1}{2}(B-c) + \frac{1}{2}(1-\pi_{AI}^*)(B-c) = 2B,$$

which implies

$$E(w) = 2c + \frac{\pi_{AI}^*}{2}(B-c).$$

Since $B > c$ and π_{AI}^* is decreasing in κ, then $E(w)$ is decreasing in κ. □

Proof of Proposition 3: The effort IC constraint can be written as

$$\pi \leq \frac{(2\alpha-\eta)w_1 - 2\alpha c}{B}. \tag{A.6}$$

The isoprofit now becomes

$$\pi(w_1) = \frac{1}{M(\alpha,\eta)}[(2\alpha-\eta)w_1 + 2(\eta+c\alpha) - 4\alpha + 2\alpha V]. \tag{A.7}$$

As we can see from Figure 3, an interior solution requires the effort IC constraint to be steeper than the isoprofit, which implies $M(\alpha,\eta) > B$.

The help IC constraint is independent of α, and its slope is $\frac{1-\eta-\kappa}{B}$. The slope must be smaller than that of the isoprofit, implying

$$\frac{M(\alpha,\eta)}{B} < \frac{2\alpha-\eta}{1-\eta-\kappa}.$$

All other steps are similar to those in the proof of Proposition 1 and are omitted for brevity. From (25) we have

$$\frac{\partial w_{1AI}^*}{\partial \alpha} = \frac{2c(\kappa-1)}{(2\alpha-1+\kappa)^2} < 0,$$

and from (26) we have

$$\frac{\partial \pi_{AI}^*}{\partial \alpha} = \frac{(1-\eta-\kappa)2c(\kappa-1)}{B(2\alpha-1+\kappa)^2} < 0.$$

It can be shown that $\frac{\partial V_{AI}^*}{\partial \alpha} > 0$ by brute force, but that route is long and not illuminating. A simpler and more intuitive (and elegant) proof is to note that increasing α (i) relaxes the effort IC constraint, (ii) does not affect the help IC constraint, and (iii) (strictly) increases profit for any given (w_1, π) (the arguments are presented in text, and follow directly from inspection of these expressions). It thus follows from (i)+(ii)+(iii) that the firm must do strictly better when maximizing profit under a larger α.

Because the effort IC constraint in (23) binds, the CEO's lifetime utility is $2B$. □

Proof of Corollary 2: Because the first-period effort IC constraint binds, we must have

$$E(w) + B - c + (1 - \frac{1}{2\alpha})(B-c) + \frac{1}{2\alpha}(1-\pi_{AI}^*)(B-c) = 2B,$$

which implies

$$E(w) = 2c + \frac{\pi^*_{AI}}{2\alpha}(B-c) = 2c + \frac{(1-\eta-\kappa)c}{B(2\alpha-1+\kappa)}(B-c),$$

which is decreasing in α. □

References

ADAMS, R. B., A. C. AKYOL, AND P. VERWIJMEREN (2018): "Director Skill Sets," *Journal of Financial Economics*, 130, 641–662.

ADAMS, R. B. AND D. FERREIRA (2007): "A Theory of Friendly Boards," *Journal of Finance*, 62, 217–250.

AGHION, P. AND J. TIROLE (1997): "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105, 1–29.

ATHEY, S. C., K. A. BRYAN, AND J. S. GANS (2020): "The Allocation of Decision Authority to Human and Artificial Intelligence," *AEA Papers and Proceedings*, 110, 80–84.

BABINA, T., A. FEDYK, A. HE, AND J. HODSON (2025): "Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition," in *Technology, Productivity, and Economic Growth*, ed. by S. Basu, L. Eldridge, J. Haltiwanger, and E. Strassner, University of Chicago Press, vol. 83 of *Studies in Income and Wealth*.

BALDENIUS, T., N. MELUMAD, AND X. MENG (2014): "Board Composition and CEO Power," *Journal of Financial Economics*, 112, 53–68.

BALDENIUS, T., X. MENG, AND L. QIU (2019): "Biased Boards," *The Accounting Review*, 94, 1–27.

BARANCHUK, N. AND P. H. DYBVIG (2009): "Consensus in Diverse Corporate Boards," *The Review of Financial Studies*, 22, 715–747.

BURKART, M., D. GROMB, AND F. PANUNZI (1997): "Large Shareholders, Monitoring, and the Value of the Firm," *The Quarterly Journal of Economics*, 112, 693–728.

BURKART, M., S. MIGLIETTA, AND C. OSTERGAARD (2023): "Why Do Boards Exist? Governance Design in the Absence of Corporate Law," *The Review of Financial Studies*, 36, 1788–1836.

CHAKRABORTY, A. AND B. YILMAZ (2017): "Authority, Consensus, and Governance," *The Review of Financial Studies*, 30, 4267–4316.

CHEMMANUR, T. J. AND V. FEDASEYEU (2018): "A Theory of Corporate Boards and Forced CEO Turnover," *Management Science*, 64, 4798–4817.

CHEN, J. J. AND B. Y. HAN (2026): "When Corporate AI Adoption Backfires," Working Paper.

CHEN, Y., J. GONG, J. LI, AND Z. ZHAO (2025): "Better Technology, Worse Motivation: GenAI's Mediocrity Trap," Available at SSRN: `https://ssrn.com/abstract=5208163`.

COLES, J. L., N. D. DANIEL, AND L. NAVEEN (2008): "Boards: Does One Size Fit All?" *Journal of Financial Economics*, 87, 329–356.

COWGILL, B., P. HERNANDEZ-LAGOS, AND N. L. WRIGHT (2024): "Does AI cheapen talk? Theory and evidence from global entrepreneurship and hiring," *Theory and Evidence From Global Entrepreneurship and Hiring (July 26, 2024). Columbia Business School Research Paper*.

CSASZAR, F. A., H. KETKAR, AND H. KIM (2024): "Artificial Intelligence and Strategic Decision-Making: Evidence from Entrepreneurs and Investors," *Strategy Science*, 9, 322–345.

DE HAAS, R., D. FERREIRA, AND T. KIRCHMAIER (2021): "The Inner Workings of the Board: Evidence from Emerging Markets," *Emerging Markets Review*, 48, 100777.

DELL'ACQUA, F., E. MCFOWLAND III, E. R. MOLLICK, H. LIFSHITZ-ASSAF, K. KELLOGG, S. RAJENDRAN, L. KRAYER, F. CANDELON, AND K. R. LAKHANI (2023): "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," Tech. Rep. Working Paper No. 24-

013, Harvard Business School Technology & Operations Mgt. Unit, the Wharton School Research Paper. Available at SSRN: https://ssrn.com/abstract=4573321.

DONALDSON, J. R., N. MALENKO, AND G. PIACENTINO (2020): "Deadlock on the Board," *The Review of Financial Studies*, 33, 4445–4488.

DRYMIOTES, G. AND K. SIVARAMAKRISHNAN (2021): "Strategic Director Appointments," *Journal of Accounting Research*, 59, 1303–1347.

EWENS, M. AND N. MALENKO (2024): "Board Dynamics over the Startup Life Cycle," *Journal of Finance*, forthcoming.

FALEYE, O., R. HOITASH, AND U. HOITASH (2011): "The Costs of Intense Board Monitoring," *Journal of Financial Economics*, 101, 160–181.

FAMA, E. F. AND M. C. JENSEN (1983): "Separation of Ownership and Control," *The Journal of Law and Economics*, 26, 301–325.

FERREIRA, D., M. A. FERREIRA, AND B. MARIANO (2018): "Creditor Control Rights and Board Independence," *Journal of Finance*, 73, 2385–2422.

FERREIRA, D., M. A. FERREIRA, AND C. RAPOSO (2011): "Board Structure and Price Informativeness," *Journal of Financial Economics*, 99, 523–545.

FERREIRA, D. AND R. NIKOLOWA (2023): "Talent Discovery and Poaching Under Asymmetric Information," *Economic Journal*, 133, 201–234.

FIELD, L., M. LOWRY, AND A. MKRTCHYAN (2013): "Are Busy Boards Detrimental?" *Journal of Financial Economics*, 109, 63–82.

FUCHS, W., L. GARICANO, AND L. RAYO (2015): "Optimal Contracting and the Organization of Knowledge," *The Review of Economic Studies*, 82, 632–658.

GARICANO, L. (2000): "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 108, 874–904.

GARICANO, L. AND E. ROSSI-HANSBERG (2006): "Organization and Inequality in a Knowledge Economy," *The Quarterly Journal of Economics*, 121, 1383–1435.

GONZALEZ-URIBE, J. AND M. GROEN-XU (2017): "CEO Contract Horizon and Innovation," Available at SSRN: https://ssrn.com/abstract=2633763.

GREENWALD, B. C. (1986): "Adverse Selection in the Labour Market," *Review of Economic Studies*, 53, 325–347.

HARFORD, J. AND R. J. SCHONLAU (2013): "Does the Director Labor Market Offer Ex Post Settling-up for CEOs? The Case of Acquisitions," *Journal of Financial Economics*, 110, 18–36.

HARRIS, M. AND A. RAVIV (2008): "A Theory of Board Control and Size," *The Review of Financial Studies*, 21, 1797–1832.

HUANG, J. AND S. OUYANG (2025): "Soft Information, Hard Decisions: AI Advising," Working Paper.

IDE, E. AND E. TALAMÀS (2025): "Artificial Intelligence in the Knowledge Economy," *Journal of Political Economy*, 133, 3762–3800.

INDERST, R. AND H. M. MUELLER (2010): "CEO Replacement Under Private Information," *The Review of Financial Studies*, 23, 2935–2969.

JIANG, W. (2025): "Corporate Finance and Governance with Artificial Intelligence: Old and New," in *Artificial Intelligence in Finance*, ed. by T. Foucault, L. Gambacorta, W. Jiang, and X. Vives, CEPR Press, chap. 4.

JIANG, X. AND V. LAUX (2023): "What Role Do Boards Play in Companies with Visionary CEOs?" *Journal of Accounting Research*, 61, 1757–1796.

KENNY, G., M. KOWALKIEWICZ, AND K. OOSTHUIZEN (2025): "How CEOs Are Using Gen AI for Strategic Planning," *Harvard Business Review Digital Article*, online.

KORINEK, A. (2025): "AI Agents for Economic Research," NBER Working Paper 34202, National Bureau of Economic Research.

LARCKER, D. F., A. SERU, B. TAYAN, AND L. YOLER (2025): "The Artificially Intelligent Boardroom," Tech. Rep. 1, Stanford Graduate School of Business.

LEVIT, D. (2012): "Expertise, Structure, and Reputation of Corporate Boards," SSRN Working Paper.

——— (2020): "Words Speak Louder Without Actions," *Journal of Finance*, 75, 91–131.

LEVIT, D. AND N. MALENKO (2016): "The Labor Market for Directors and Externalities in Corporate Governance," *Journal of Finance*, 71, 775–808.

LI, J. (2013): "Job Mobility, Wage Dispersion, and Technological Change: An Asymmetric Information Perspective," *European Economic Review*, 60, 105–126.

LINCK, J. S., J. M. NETTER, AND T. YANG (2008): "The Determinants of Board Structure," *Journal of Financial Economics*, 87, 308–328.

MALENKO, N. (2014): "Communication and Decision-Making in Corporate Boards," *The Review of Financial Studies*, 27, 1486–1532.

——— (2024): "Information Flows, Organizational Structure, and Corporate Governance," in *Handbook of Corporate Finance*, ed. by D. Denis, Edward Elgar Publishing, chap. 14, 511–547.

NIKOLOWA, R. (2015): "Career Dynamics and Span of Control," *Economics Letters*, 128, 6–8.

RAHEJA, C. G. (2005): "Determinants of Board Size and Composition: A Theory of Corporate Boards," *Journal of Financial and Quantitative Analysis*, 40, 283–306.

SHEKSHNIA, S. AND V. YAKUBOVICH (2025): "How Pioneering Boards Are Using AI," *Harvard Business Review*, 103, 56–64.

SONG, F. AND A. V. THAKOR (2006): "Information Control, Career Concerns, and Corporate Governance," *Journal of Finance*, 61, 1845–1896.

WALDMAN, M. (1984): "Job Assignments, Signalling, and Efficiency," *RAND Journal of Economics*, 15, 255–267.

WARTHER, V. A. (1998): "Board Effectiveness and Board Dissent: A Model of the Board's Relationship to Management and Shareholders," *Journal of Corporate Finance*, 4, 53–70.

ZHONG, H. (2025): "Optimal Integration: Human, Machine, and Generative AI," *Management Science*, forthcoming.