A Different Approach to Agency Theory and Implications For ESG

*Jonathan Bonham and Amoray Riggs-Cragun**

ABSTRACT

In conventional agency theory, the agent is modeled as exerting unobservable "effort" that influences the distribution over outcomes the principal cares about. Recent papers instead allow the agent to choose the entire distribution, an assumption that better describes the extensive and flexible control that CEOs have over firm outcomes. Under this assumption, the optimal contract rewards the agent directly for outcomes the principal cares about, rather than for what those outcomes reveal about the agent's effort. This article briefly summarizes this new agency model and discusses its implications for contracting on ESG activities.

CONTENTS

INTRODUCTION

The relationship between a firm's owners and its CEOs is well described as a principal–agent relationship—the owners (the principal) hire a CEO (the agent) to act on their behalf.[1] The CEO may prefer to take

* Jonathan Bonham, University of Illinois-Chicago, UIC Business; Amoray Riggs-Cragun, University of Chicago, Booth School of Business. This article was prepared for the 14th annual Berle Symposium,

actions that are not in the best interest of the firm owners, and firm owners cannot perfectly observe the CEO's actions.

Because conflicting preferences over the agent's action are central to the agency problem, the way that the agent's action is modeled should be somewhat descriptive of the agency setting being studied. Following the seminal work of Bengt Holmström in 1979, the agent's action is typically modeled as a level of effort that serves as a parameter in the distribution over the outcome of interest.[2] For example, the agent's effort might simply shift the mean of the distribution, such that higher effort results in higher expected outcomes.[3]

However, the job of a CEO is complex and multifaceted, involving a wide range of actions and decisions that span across domains, such as strategic planning, financial management, human resources, budgeting, product development, risk management, and corporate social responsibility.[4] The classic effort model is not descriptive of this rich action space, else the enormous library of books on how to successfully run a business could be replaced by the maxim, "try hard." The role of a CEO is perhaps better captured by the *generalized distribution* approach taken by recent studies, where the agent can influence the *entire joint distribution* over firm outcomes.[5]

"Developing a 21st Century Corporate Governance Model," held at the University of Southern California Gould School of Law on March 24–25, 2023. The authors can be contacted at jbonham@uic.edu and amoray@chicagobooth.edu.

1. *See, e.g.*, Michael C. Jensen & Kevin J. Murphy, *Performance Pay and Top-Management Incentives*, 98 J. POL. ECON. 225, 225–226 (1990).

2. *See generally* Bengt Holmström, *Pay for Performance and Beyond*, 107 AM. ECON. REV. 1753 (2017).

3. *See generally* Bengt Holmstrom & Paul Milgrom, *Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design*, 7 J. L. ECON. & ORG. 24 (1991).

4. *See generally* Marwan Altarawneh, Rohami Shafie & Rokiah Ishak, *CEO Characteristics: A Literature Review and Future Directions*, 19 ACAD. STRATEGIC MGMT. J. 1 (2020).

5. *See*, *e.g*., Benjamin Hébert, *Moral Hazard and the Optimality of Debt*, 85 REV. ECON. STUD. 2214, 2221 (2018); Jonathan Bonham, *Shaping Incentives Through Measurement and Contracts* 1 (Working Paper, 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3691705 [https://perma.cc/3ACC-7DQR]; Jonathan Bonham & Amoray Riggs-Cragun, *Contracting on Aggregated Accounting Estimates* 1 (Chi. Booth, Rsch. Paper No. 22–04, 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3962557 [https://perma.cc/C8N8-KE9W]; Jonathan Bonham & Amoray Riggs-Cragun, *Motivating ESG Activities Through Contracts, Taxes and Disclosure Regulation* 9 (Chi. Booth, Rsch. Paper No. 22–05, 2022), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4016659 [https://perma.cc/AHK3-UHUN]; Jonathan Bonham & Amoray Riggs-Cragun, *Contracting on Information About Value* 1 (Chi. Booth, Rsch. Paper No. 22–03, 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3892838 [https://perma.cc/GY3J-4EJC]; George Georgiadis, Doron Ravid & Balázs Szentes, *Flexible Moral Hazard Problems* 1–2 (Working Paper, 2023), https://www.kellogg.northwestern.edu/faculty/georgiadis/FlexibleMoralHazard.pdf [https://perma.cc/K6KD-RN68]; Daniel F. Garrett, George Georgiadis, Alex Smolin & Balázs Szentes, *Optimal Technology Design*, 209 J. ECON. THEORY 1, 2 (2023).

The way the agent's action is modeled significantly changes the optimal contract that the principal offers to the agent. Under the classic effort model, the optimal contract rewards the agent for outcomes indicating a higher likelihood that the agent took the desired level of effort; that is, the agent is rewarded for how hard they appear to be working.[6] Applied to the CEO setting, this suggests that executive compensation contracts should be based on metrics like hours spent in the office, emails sent, or words typed—all of which could be easily tracked on a company issued laptop. However, in reality, executive compensation packages do not include those types of "effort-informative" metrics.[7] Instead, they tend to be based on signals that are informative about *what shareholders actually care about*, like accounting profits, stock returns, and more recently, ESG metrics.[8] The generalized distribution model predicts exactly that: CEO compensation should depend not on signals of the CEO's *input*, but on the *outcome(s)* that shareholders care about.[9] For example, if shareholders care about a firm's carbon footprint, shareholders should reward the CEO not for how many meetings the CEO held about ESG or how much company money the CEO threw at ESG issues, but instead for measures of the firm's actual carbon emissions.

The remainder of this article proceeds as follows. Section I provides an overview of the agency problem modeled under the classic effort approach and under the generalized distribution approach, and discusses the differences in the fundamental insights that emerge from each approach. Section II discusses some implications of these main insights in the context of motivating a CEO to take ESG-oriented actions. Section III concludes.

I. REVISITING THE AGENCY PROBLEM

To formalize ideas, consider a firm owned by shareholders who care about a vector of random variables, $\mathbf{x}$. The vector may include financial outcomes, such as true economic profits, and ESG outcomes, such as the firm's carbon footprint or the happiness of its employees. Shareholders care about these various outcomes according to some *shareholder welfare* function, $\boldsymbol{\Omega}(\mathbf{x})$. Shareholders are represented by a principal (e.g., the board of directors), who contracts with a risk-averse CEO (the agent) to take ac-

6. Holmström, *supra* note 2, at 75.

7. David De Angelis & Yaniv Grinstein, *Performance Terms in CEO Compensation Contracts*, 19 REV. FIN. 619, 620 (2015).

8. *See* Shira Cohen, Igor Kadach, Gaizka Ormazabal & Stefan Reichelstein, *Executive Compensation Tied to ESG Performance: International Evidence*, 61 J. ACCT. RSCH. 805, 805 (2023).

9. *See* Bonham & Riggs-Cragun, *Contracting on Information About Value*, *supra* note 5, at 1–13.

tions that improve shareholder welfare. Because the principal cannot observe the agent's actions, and because the agent's preferences for actions differ from shareholders', the principal must elicit effort from the agent using a performance-based inducement. The principal cannot directly observe the realizations of the outcomes **x** that the shareholders intrinsically value. Instead, they observe **y**, a set of verifiable and contractible signals, such as accounting profits, factory emissions readings, and employee turnover. Before the agent acts, the principal makes a one-shot offer of an incentive contract, *s*, which pays the agent $s(\mathbf{y})$ when the realized outcome is **y**. If the agent does not accept the contract, their payoff from outside options is $\overline{U}$.

The following two subsections present how this setting is analyzed under the classic "effort" model and under the generalized distribution model.

A. The Classic Approach: Actions Modeled as "Effort"

Under the classic formulation, the agent's action is represented by a scalar, $a \in A \subseteq \mathbb{R}$, which is interpreted as effort.[10] This action serves as a parameter in the joint distribution over **x** and **y**, denoted $f(\mathbf{x},\mathbf{y};a)$.[11] The agent suffers disutility of effort in the form of an increasing and convex function, $c(a)$.[12] The agent's utility is additively separable in compensation and effort, and can be written as $U(s(\mathbf{y})) - c(a)$.[13]

The principal's aim is to propose a contract and action pair (s,a) that maximizes their expected utility subject to two constraints.[14] First, the proposed pair must make the agent's expected utility at least as great as $\overline{U}$; this *individual rationality* (IR) constraint ensures that the agent is willing to accept the contract.[15] Second, the proposed pair must be *incentive compatible* (IC); that is, given the proposed scheme $s(\mathbf{y})$, the agent chooses the proposed *a* voluntarily.[16] Holmström simplified the incentive compatibility constraint by assuming that the first-order approach is valid; in other words, that the agent's expected utility is globally concave in *a* given

$$\begin{aligned} \max_{s,a} \quad & \int (\Omega(\mathbf{x}) - s(\mathbf{y})) f(\mathbf{x},\mathbf{y};a) d\,\mathbf{x}\,d\,\mathbf{y} \\ \text{s.t.} \quad & \int [U(s(\mathbf{y})) - c(a)] f(\mathbf{x},\mathbf{y};a) d\,\mathbf{x}\,d\,\mathbf{y} \geq \bar{U} \\ \text{and} \quad & \int U(s(\mathbf{y})) f_a(\mathbf{x},\mathbf{y};a) d\,\mathbf{x}\,d\,\mathbf{y} = c'(a) \end{aligned} \tag{1}$$

10. Holmström, *supra* note 2, at 75–87.
11. *Id.*
12. *Id.*
13. *Id.*
14. *Id.*
15. *Id.*
16. *Id.*

$s(\mathbf{y})$.[17] Given this assumption, the first-order condition from the agent's problem is sufficient for $s(\mathbf{y})$ to be incentive compatible.[18] With these classic ingredients in place, the principal's program is given as follows, where $f_a(\mathbf{x},\mathbf{y},a)$ denotes the derivative of $f(\mathbf{x},\mathbf{y},a)$ with respect to a.[19]

Letting λ and μ be the multipliers on the IR and IC constraints, pointwise optimization of the Lagrangian with respect to $s(\cdot)$ produces the following iconic characterization of the optimal contract.[20]

$$\frac{1}{U'(s(\mathbf{y}))} \quad = \quad \lambda + \mu \cdot \frac{f_a(\mathbf{y};a)}{f(\mathbf{y};a)} \tag{2}$$

It follows from the above characterization that the shape of the optimal incentive scheme depends on the shape of the *likelihood ratio*, $\frac{f_a(\mathbf{y},a)}{f(\mathbf{y},a)}$, which gives an indication of the likelihood that the agent exerted the desired level of effort given the realization $\mathbf{y}$.[21] In other words, the optimal contract behaves as if the principal were making inferences about the agent's action.[22] An important implication of this result is that performance measures are useful for contracting to the extent that they are incrementally informative about the agent's action.[23] This result is termed the *informativeness principle*, and it is the "main predictive content" of the classic agency model.[24] Applied to the shareholder–CEO setting, Holmström's informativeness principle predicts that CEOs should be rewarded for how hard they appear to be working.[25]

B. A Different Approach: Actions Modeled as Distributions

Now assume that rather than choosing a distribution from the restricted parametric set $\{f(\mathbf{x},\mathbf{y};a)|a \in \mathrm{A}\}$, the agent can choose *any* distribution $f(\mathbf{x},\mathbf{y})$ from the set of all joint distributions over $\mathbf{x}$ and $\mathbf{y}$. To model the cost of choosing a distribution, let the agent's cost $C(f)$ be given by the *Kullback–Leibler divergence* of $f(\mathbf{x},\mathbf{y})$ from some reference distribution $g(\mathbf{x},\mathbf{y})$[26]:

$$C(f) = D\left(f||g\right) \quad \equiv \quad \int f(\mathbf{x},\mathbf{y}) \ln\left(\frac{f(\mathbf{x},\mathbf{y})}{g(\mathbf{x},\mathbf{y})}\right) d(\mathbf{x},\mathbf{y}). \tag{3}$$

17. *See id.*
18. *See id.*
19. *See id.*
20. *Id.*
21. *Id.*
22. *See id.*
23. *See id.*
24. *See* Oliver Hart & Bengt Holmstrom, *The Theory of Contracts*, *in* ADVANCES IN ECONOMIC THEORY: FIFTH WORLD CONGRESS 83 (Truman Fassett Bewley ed., 1987).
25. *See id.*
26. Benjamin Hébert, *Moral Hazard and The Optimality of Debt*, 85 REV. ECON. STUD. 2214, 2215 (2018); Bonham & Riggs-Cragun, *Contracting on Aggregated Accounting Estimates*, *supra* note

The Kullback–Leibler divergence (or *relative entropy*) is a measure of the dissimilarity between two distributions.[27] Here, the divergence captures the cost incurred by the agent when they take actions to implement some proposed distribution, f, rather than their preferred distribution, g. The agent incurs zero cost from implementing their preferred distribution, i.e., when they set $f = g$. The larger the divergence of f from g, the larger the agent's disutility from implementing f.

The principal must choose a contract-action pair that maximizes shareholder welfare, $\mathbf{\Omega(x)}$, minus the expected cost of compensating the agent. The principal's program is constructed as follows.[28]

$$\begin{aligned} \max_{f,s} \quad & \int\int(\Omega(\mathbf{x}) - s(\mathbf{y}))f(\mathbf{x},\mathbf{y})d\,\mathbf{x}\,d\,\mathbf{y} \\ \text{s.t.} \quad & \int\int s(\mathbf{y})f(\mathbf{x},\mathbf{y})d\,\mathbf{x}\,d\,\mathbf{y} - \int\int f(\mathbf{x},\mathbf{y})\ln\left(\frac{f(\mathbf{x},\mathbf{y})}{g(\mathbf{x},\mathbf{y})}\right)d\,\mathbf{x}\,d\,\mathbf{y} \geq 0 \\ & f(\mathbf{x},\mathbf{y}) = g(\mathbf{x},\mathbf{y})e^{s(\mathbf{y})-\nu-1} \text{ for all } (\mathbf{x},\mathbf{y}) \\ & \nu = \ln\left(\int\int g(\mathbf{x},\mathbf{y})e^{s(\mathbf{y})-1}d\,\mathbf{x}\,d\,\mathbf{y}\right) \end{aligned} \tag{4}$$

Pointwise Lagrangian optimization yields the following characterization of the optimal contract, where λ and η are constants.[29]

$$\frac{1}{U'(s(\mathbf{y}))} = \mathbb{E}[\Omega(\mathbf{x})|\mathbf{y}] - s(\mathbf{y}) + \lambda - \eta, \tag{5}$$

Notice that $\mathbf{y}$ enters this expression only through the contract $s(\mathbf{y})$ and the conditional expectation $\mathrm{E}[\mathbf{\Omega(x)}|\mathbf{y}]$. Then, for the expression to be maintained for all realizations $\mathbf{y}$, the contract $s(\mathbf{y})$ will vary with some signal y_i if and only if $\mathrm{E}[\mathbf{\Omega(x)}|\mathbf{y}]$ varies with y_i, that is, if y_i contains incremental information about shareholder welfare, $\mathbf{\Omega(x)}$.

In other words, *a signal is useful for contracting if and only if it is informative about shareholder welfare*. This contrasts with the classic model, where performance measures are valuable when they are informative about the agent's action.[30] Implications of this result for ESG are explored next.

II. INCENTIVIZING ESG ACTIVITIES

The generalized distribution approach to agency theory can be used to gain insights into how executive compensation contracts can be used to improve a firm's ESG outcomes. Several such insights are summarized

---

5, at 13; Bonham & Riggs-Cragun, *Motivating ESG Activities through Contracts, Taxes and Disclosure Regulation*, *supra* note 5, at 9.

27. *See generally* SHUN-ICHI AMARI, INFORMATION GEOMETRY AND ITS APPLICATIONS (2016).

28. *See* Bonham & Riggs-Cragun, *Motivating ESG Activities Through Contracts, Taxes and Disclosure Regulation*, *supra* note 5, at 9–10.

29. *See id.* at 12–13.

30. *Compare id.* at 8–18, *with* Holmström, *supra* note 2, at 75–87.

below, some of which are developed in the recent working paper Bonham and Riggs-Cragun (2022b) (BR).[31]

A. Why Firms Might Contract on ESG

Equation (5) suggests three potential reasons that a given ESG measure might be included in a CEO's compensation contract. First, it might be a measure of an ESG outcome that shareholders intrinsically value. For example, if shareholders intrinsically care about the company's racial diversity, then the compensation contract might include a measure of the firm's workforce diversity. Second, the measure could be informative about a financial outcome that shareholders value. For example, shareholders might not *intrinsically* care about racial diversity but may believe that a lack of racial diversity will lead to negative public scrutiny, boycotts, or reduced demand; in this case, the compensation contract may include a measure of diversity because it is informative about future profits. Finally, the measure could be subject to taxes, subsidies, or regulatory oversight. For example, if a firm operates in a jurisdiction with a carbon tax regime, we might expect the compensation contract to include measures of the firm's carbon emissions.

B. The Importance of Good Measurement

BR conceptualize the strength of an ESG measurement system by how difficult it is for the manager to greenwash, or more specifically, by the correlation between the ESG measure and the true ESG outcome in the agent's cost-minimizing distribution g.[32] Unless the measurement system is perfect, the manager always engages in some greenwashing in equilibrium, making the reported ESG measure look better than the true ESG outcome. The more reliable the ESG measurement system, the larger the improvement in true ESG outcomes that can be induced by the contract. Importantly, when greenwashing is costless to the manager, *no* real improvements in ESG can be induced by contracting on ESG measures.

C. Incentivizing Target-Level ESG Performance

In some cases, shareholders may prefer an interior level of an ESG outcome. Consider gender diversity as an example. A workforce with 0% women is not gender diverse, but neither is a workforce with 100% women. Shareholders who value gender diversity may prefer some interior target percentage of women employed by their firm (e.g., 50%). BR show

---

31. *See generally* Bonham & Riggs-Cragun, *Motivating ESG Activities Through Contracts, Taxes and Disclosure Regulation*, *supra* note 5.

32. *Id.* at 9.

that shareholders could incentivize such a target with a contract that is hump-shaped in the ESG measure, imposing increasingly harsh penalties for deviations from a contractual target.[33] In response to this contract, the agent moves the expected ESG outcome closer to the target and reduces the variance of the ESG outcome. The worse the quality of the contractible ESG measure, the higher the contractual target must be to induce the desired target in the actual ESG outcome.

D. Motivating Green Innovation with ESG-Performance-Based Shares

Shareholders may have complementary preferences over financial and ESG performance. For example, shareholders may prefer lower firm emissions, but not if that means low profits; likewise, shareholders may want higher profits, but not if that means profits earned with very high emissions. Shareholders with these preferences would like the agent to engage in "green innovation"; that is, they want the agent to lower emissions in ways that are actually profitable for the firm. BR show that this can be achieved by giving the agent ESG-performance-based stock awards, where the size of the stock award is increasing in ESG performance.[34]

CONCLUSION

A disconnect has long existed between the predictions of classic parametric agency theory and executive compensation practice. As this article suggests, the generalized distribution approach is perhaps a better fit for studying executive pay. With shareholders becoming increasingly concerned with ESG issues, the generalized approach may provide a useful theoretical foundation for developing ESG-oriented compensation practices.

33. *Id.* at 3–4.
34. *Id.*