Limits on Equilibrium ESG Investing*

Andrei Kirilenko Christoph Loch

September 30, 2024

Please do not quote or circulate without the authors' permission.

ABSTRACT

We derive an equilibrium between stakeholders who care about ESG and investors who only care about profits. We model stakeholder ESG interests as a non-pecuniary variable imbedded into the objective function of a corporate manager. The manager‘s goal is to get funding for a risky project from investors. We assume that the manager knows more about both the project's financial payoff and its non-pecuniary ESG value than the investors. To investors, ESG is just extra noise that pollutes information about the project‘s expected payoff. In equilibrium, investors limit ESG activities rather than support them. The extent of the limit is determined by the degree of information asymmetry - how much less do the investors know than the manager. We extend our baseline model to include ESG externalities, ESG risk ratings, and ESG dividends and subsidies.

*Andrei Kirilenko is with the Cambridge Judge Business School and CEPR and Christoph Loch is with the Cambridge Judge Business School. We humbly thank many colleagues and seminar participants for helpful comments and suggestions. All remaining errors are our own.

I. Introduction

Much has been said about the purpose of a corporation. For decades the prevalent narrative has been that corporations should focus solely on maximizing shareholder profits. More recently, however, corporate managers have publicly committed to engaging with broader stakeholders, including customers, employees, suppliers and local communities, as well as their shareholders. In August 2019 the Business Roundtable, an association of chief executive officers (CEOs) of leading US companies, released a Statement on the Purpose of a Corporation. The statement, which was signed by 181 CEOs, said that corporations "share a fundamental commitment to all of our stakeholders." The ensuing voluntary corporate policies - which have collectively become known as corporate social responsibility (CSR) - gave rise to environmental, social and governance (ESG) considerations. ESG explicitly recognizes and discloses non-pecuniary purposes of a modern corporation. Consequently, ESG allows managers to systematically engage investors into an ongoing shift of the corporate boundary towards stakeholders. For the avoidance of doubt, let us state upfront: in our view, the boundary of a firm is chiefly determined by corporate managers rather than investors. Investors can either cheer or lament the redefined purpose of a corporation, but, unless outright prohibited, they will invest if there are expected returns to be made and/or fees to be collected. Ultimately, it is the task of a corporate manager to find an optimal balance between the firm's stakeholders and its shareholders. The objective of this paper is to describe the fundamentals of how far a corporate manager can extend the boundary of her firm towards the stakeholders. Spoiler alert: it boils down to information asymmetry.

Before we start with a formal model, consider the following illustrative scenario. You work for a global mining company. Your company has been granted a concession to mine copper at a tropical rainforest site. You have been appointed as the manager of this project. This appointment is very important for your career. However, you know that success is not guaranteed. You are told by your boss that the company's corporate strategy is to mitigate possible negative impacts of each project on human rights and the environment. To that end, the company only proceeds with projects that have engaged with relevant stakeholders. These stakeholders are not a part of your company. They are not represented on the company's board. They do not regulate or license the company. However, their actions are deemed to be material for the project's financial payoff. Thus, before the project even starts, your job is to engage with all known stakeholders, including the mayor of the nearby town, the local blogger, local and global activist environmentalists, representatives of the labor union for skilled external contractors, key suppliers, and manufacturers of electric vehicles which would be the customers of the copper mined at this site. You need to make sure that everyone is more or less on board. You are informed that all stakeholders are rational and have non-pecuniary preferences. The mayor wants votes. The local blogger wants clicks and followers. Activist environmentalists want access and respect. The labor union wants safe and clean working conditions. Suppliers want predictability. Customers want sustainability. You are also told that the company does not have sufficient internal resources to finance this project. Thus, once you get the stakeholders engaged, you must present a financial proposal for the project to an outside investor. The investor only cares about her share of the financial payoff from the project; she is neutral on human rights and the environment. The investor is aware of the non-pecuniary aspects of various stakeholders but cannot separate them from the financial performance of the project. To the investor, stakeholder interests are just noise that pollutes the project's payoff. Moreover, she has a less precise knowledge about the project's risky payoff than you. You have the informational advantage but the investor has the funds. Without her money, the project will not proceed. When you present the project's proposal to her, she will bargain for her cash flow rights. You will need to protect your company's interests, including both financial and stakeholder interests. She cannot separate between the financial and non-financial parts of the project in your proposal. However, to get a better understanding of the two parts, she can negotiate with you contractual clauses that limit various stakeholder interests that are valuable to them but not to her. She can specify that you can help the mayor in his campaign but only so much. You can only give one interview per month to the local blogger. Activist environmentalists cannot have unlimited access to the mining site. The labor union can have no more than two air-conditioned resting units for external contractors and so on. If you get the stakeholders to agree, then there must be enough of the financial aspects of the project to share among everyone. To her this would mean that the project is a good investment opportunity and is worth funding. But if you don't get the stakeholders to agree, then the project would not be funded and your boss will replace you. A mutually agreeable outcome, if achieved, fundamentally depends on the degree of information asymmetry - how much more do you know than the investor - and how much of that asymmetry comes from the non-pecuniary interests of various stakeholders.

In this paper we model an equilibrium between stakeholders who care about ESG and shareholders who only care about cash flows from a risky project. We assume that stakeholder ESG interests are a non-pecuniary variable imbedded into the objective function of a corporate manager. The risk averse manager is wealth constrained and needs to obtain financial resources from outside investors in return for a share of the project's cash flows. We assume that information about both the project's financial payoff and its non-pecuniary ESG value to shareholders is asymmetric. By knowing the intricacies of the project, the manager has a more precise signal about the financial payoff than the investors. By directly engaging with the stakeholers, the manager also knows exactly the value of ESG but investors do not. We also assume that there is no separate market for pricing only ESG - markets are incomplete. The manager and competitive investors can, however, negotiate over a normalized variable that describes a set of limiting clauses on non-pecuniary ESG interests. If they agree on these limits, then there is an equilibrium with investment in the project. If they do not, then there isn't. Importantly, we model ESG investing as a process of identifying value and managing risks in commercial opportunities motivated by expected profit. It is not a process in which investors knowingly sacrifice financial returns while supporting a particular set of non-pecuniary values or moral principles (see, e.g., Renneboog *et al.* (2011) and Barber *et al.* (2021) for empirical regularities of such investing).

We ask the following research question: In a model economy with rational agents, can ESG investing be achieved in a competitive equilibrium and, if yes, are there any limits of doing so?

We model two types of rational agents - corporate managers and financial investors. A manager is an agent with the capacity of gaining insight into the value of a project and managing the project's financial and non-financial risks. The manager may be a single individual or a team, and she may manage the project within a new or established, small or large, or innovative or traditional business. We model the manager as a risk averse agent, who manages a single risky commercial project with additional non-pecuniary ESG characteristics that reflect stakeholder preferences. Our assumption about the non-pecuniary nature of ESG is consistent with the literature on CSR. In a comprehensive meta-analysis of this literature, Hong and Shore (2023) find that the majority of CSR studies "generally support the nonpecuniary view." Our assumption about imbedding stakeholder preferences directly into the objective function of the corporate manager is consistent with recent empirical findings. Barzuza *et al.* (2023) show that employees and customers have been increasingly demanding from corporate managers to demonstrate their commitment to ESG. Graham (2022) documents that managers have, in turn, shifted their focus towards better accommodating interests, benefits and preferences of these stakeholders in making more socially responsible corporate decisions. In our model, the manager privately observes information about both financial and ESG characteristics and is wealth-constrained, therefore seeking a financial contribution from an investor.

An investor is an agent with the capacity of investing financial resources into a risky project with the objective of maximizing expected profit net of investment. An investor may be a single individual investing her own financial resources or a fund managing investments on behalf of clients. The investor does not provide advice or governance, does not influence the production technology, and does not own or control any of the project; the investor's liability and involvement are limited to financial resources only. We model investors as competitive, risk-neutral agents who maximize expected profit, conditional on the information available to them. Unlike the manager, investors neither know nor intrinsically value ESG characteristics of the project (this assumption is consistent with the empirical findings of Kruger (2015)). They are neutral - neither pro nor anti-ESG. They only care about sharing in the project's cash flows. To them, a project's non-pecuniary ESG characteristics is just extra noise that pollutes information about the project's financial payoff. However, they treat ESG characteristics as "relevant" in that information about ESG affects their decision to invest (or not) in the project. The assumption that a neutral investor is willing to invest in a green project is consistent with the recent literature on the green consumption-investment tradeoff, e.g., Chen *et al.* (2023) and Sauzet and Zerbib (2024), in which neutral investors allocate some of their capital to green assets.

We use signal extraction modeling in the context of a noisy rational expectation equilibrium framework. In the model, the manager and the investor negotiate over price, cash flow rights, as well as a normalized variable that describes a set of limiting clauses on non-pecuniary ESG characteristics. If they agree on these limits, then there is an equilibrium. If they do not, then there isn't. The main elements of this modeling framework have a long history in economics and finance, including, for example, Diamond and Verecchia (1981) and (1991), who use signal extraction when modeling optimal disclosure and investment decisions, as well as Grossman and Hart (1986) and Hart and Moore (1988), who introduce contractual incompleteness to model optimal contracts. Kirilenko (2001) introduces non-pecuniary characteristics into the modeling framework and derives a solution under a refined equilibrium notion.

We show that under the assumptions of the baseline model, there exists a unique linear equilibrium for ESG investing. In equilibrium, the investor *limits* tolerated ESG activity in order to invest only in projects of higher expected commercial value and the manager accepts the limit in order to get the funding. The key insight from our baseline model is that neutral financial investors will, even if there is no direct cost or incentive conflict, tend to limit ESG activities rather than support them, because these investors cannot directly translate latent non-pecuniary ESG characteristics into profits. This result is consistent with the finding of Geczy *et al.* (2021) who document the presence of limits in investment contracts of funds that invest in projects with a "social-benefit goal alongside the goal of financial performance." Unlike financial returns that, at some point, are actually observed, the value of ESG characteristics remains ultimately latent to investors. Investors can merely adjust allowable ESG activity as an instrument to get clearer signals about expected financial returns. By placing limits on ESG activity though, investors make prices *more* informationally efficient.

In the baseline model, investors view non-pecuniary ESG activities, no matter how beneficial they might be to a broad set of stakeholders, as simply extra noise that pollutes information about the project's financial payoff. To investors, there are no positive or negative externalities associated with ESG. In the extended model, we introduce a positive correlation between ESG and information about the project's financial payoff. We assume that if ESG activities generate net positive impact on the people, society and the environment, these positive externalities are probabilistically linked with a higher financial payoff of the project. Consequently, investors must treat ESG not just as relevant but as "material."

We assume that this link is a positive correlation rather than a certainty because of the lack of societal consensus about what "good" externalities are. Different stakeholders may value different aspects of ESG. Some value emission reduction while others value improvements in diversity and inclusion, but there is not one certain way of linking externalities with profits. Moreover, there is no market or societal mechanism that can aggregate these different views and provide a unified measure of possible benefits. In any case, financial investors need to somehow take into account material ESG externality due to a positive correlation between non-pecuniary ESG benefits and information about the project's financial payoff.

We show that if the signal about the project's financial payoff is positively correlated with its non-pecuniary ESG characteristics, the limit on the material ESG activity becomes *more* rather than less limiting. In other words, the higher is the project's ESG materialilty, the more difficult it is in equilibrium to obtain financial resources from a rational competitive investor. Intuitively, ESG materiality forces investors to *further limit* ESG activities to filter out this additional noise from the signal about the financial payoff. This result is consistent with recent developments in which a number of US states have enacted laws or adopted regulatory actions that prohibit or discourage public entities from considering material but non-pecuniary ESG benefits when investing state resources.

Reducing information asymmetry about a project's ESG materiality could be accomplished via a provision of information by third parties, e.g., ESG ratings. These ESG ratings can, however, be described as *incomplete* insofar as they only rate a correlation between latent ESG attributes and signals about future payoffs. The concept of incomplete ratings, while newly introduced here, is intuitively similar to the concepts of incomplete markets and incomplete contracts. Incomplete markets describe an economic environment in which available securities do not cover all possible future states of the world. Incomplete contracts describe an economic environment in which available contractual clauses do not cover all possible future contingencies. Incomplete ratings describe an economic environment in which available realizations of ratings do not cover all possible future material realizations of the rated latent variable. Incompleteness of ESG ratings could manifest itself in low and unstable correlations among ESG ratings for the same entity offered by different third-party providers, as documented for example by Berg, Kolbel, and Rigobon (2022). Incompleteness of ESG ratings is also consistent with an observation that as some latent ESG characteristics *become* financially material as discussed, for example, in Freiberg, Rogers and Serafeim (2020), more disagreement follows (see, Christensen, Serafeim, and Sikochi (2022)).

Intuitively, having more securities to span additional possible future states of the world makes markets less incomplete and mitigates some allocational failures. Similarly, having a way to vary control over a project/company in situations when available contractual clauses do not cover all possible future contingencies mitigates some contractual failures. Continuing with this intuition, having a more differentiated grid of consistent ESG ratings would quantify idiosyncratic risks of specific ESG attributes and mitigate possible no trade outcomes.

We argue that under incomplete ratings, third-party providers should design robust and transparent processes that make their ratings *consistently disagree* depending on, for example, the number of and the types of companies that they rate. These ESG ratings would disagree across providers but consistently so. Consistent ESG ratings that quantify certain material ESG risks would help investors evaluate and invest in specific ESG characteristics that they care about such as societal benefits of lower CO2 emissions, or greater biodiversity, or reduced inequality. Specific ESG characteristics that matter to the company's stakeholders would then be negotiated between the manager and investors and, if successful, would allow for financing to flow to projects with such ESG characteristics that otherwise may not be financed. If anything, managers should be supported in doing more to reduce information asymmetry about *specific* ESG characteristics that their stakeholders value.

We also investigate a possibility of relaxing limits on ESG investing by allowing the manager to offer the investor an additional deterministic financial incentive that we call an ESG dividend. However, it turns out that neither additional ESG dividend payments between the project's manager and investor nor third party subsidies (as long as they are deterministic) reduce the limits of equilibrium ESG investing. The intuition for this is that a deterministic payout does not reduce the underlying information asymmetry, which is the main source of the investor's problem. This suggests that, for example, public subsidies would not solve this problem. Moreover, public transfers and subsidies are likely to reflect choices and preferences that are not fully shared by all stakeholders. In contrast, ESG ratings, to the extent that they do reduce information asymmetry regarding particular aspects of ESG that matter to specific stakeholders, would lessen ESG limits imposed by the investors.

Our approach can be explained to investment practitioners as follows. If ESG characteristics can be priced directly in the financial market, then investors who have exposure to ESG can diversify away ESG risks by buying and selling assets. That could be done, for example, by constructing and rebalancing a "zero beta" portfolio of assets with respect to a common ESG factor. Asset prices would then reflect these investment decisions and the common ESG factor would be appropriately priced in. However, there is no evidence that supports the existence of a common ESG factor across assets as there may not be a consensus on what valuable ESG characteristics are. What is ESG to some (perfectly rational) investors may not be ESG to others (who are also perfectly rational). In this case, what rational investors can do is to limit the exposure of each company they invest in to its ESG characteristics in a predetermined way. Specifically, investors recognize that managers of a company make corporate decisions that, to some extent, reflect either positive or negative consequences for the wellbeing of the firm's stakeholders, i.e., employees, customers, local communities and the environment. These decisions result in an exposure of the company's asset side to ESG characteristics. The risk of this exposure, along with other risks, is passed on via its liabilities side to the shareholders. By placing limits on what level of ESG characteristics the managers of each company are allowed to have on the asset sides of their companies, investors can keep their ESG exposures within acceptable bounds. That way, exposure to ESG for each company's liability can be priced more efficiently even in the absence of an agreement on a common ESG factor.

In terms of the theoretical literature on ESG investing, Goldstein *et al.* (2022) model ESG as non-financial information partially aggregated into asset prices. At the fundamental level, Goldstein *et al.* (2022) approach ESG investing from an investor's (outsider) point of view, while our model approaches it from the point of view of a manager (insider). This gives rise to several key differences both between our specific modeling approaches and between our results. Namely, Goldstein *et al.* (2022) assume that ESG-specific information directly affects the project's financial payoff, while we - as supported by the empirical findings documented by Hong and Shore (2023) - do not. In contrast, we assume that the ESG benefit of the risky asset has a privately observed value to the manager, but does not produce any direct payoffs/utility for an investor, neither positively nor negatively. At best, an ESG characteristic may have an indirect and stochastic positive effect (modelled as a correlation) with a signal about possible returns from the asset; the investor is risk neutral and considers the ESG only insofar it introduces additional noise about the financial payoff of the project. The investor "instrumentalizes" the noisy ESG value to the manager via a limit on ESG activity as a signal extraction device to gain more information about the financial payoff of the project. This limit is tightened when the ESG activity's positive externality reflects back (via a correlation) to a possibly higher commercial value of the project.

Pastor, Stambaugh, and Taylor (2021) and (2024) model optimal institutional portfolio choice by considering decisions of asset managers/funds that have heterogeneous private tastes for ESG ("green") characteristics when investing in multiple assets. Again, we approach ESG investing from the point of view of a corporate manager whose preferences reflect those of other stakeholders and, in equilibrium, must be balanced against the profit maximization objectives of investors/shareholders. These complimentary points of view can help gain additional insight into the ESG investing decisions. For example, our model could better illustrate recent actions of US states to limit ESG portfolio choices of asset managers despite their private tastes.

There is vibrant empirical literature on ESG investing which we seek to inform with the testable predictions from our model. The non-inclusive and growing list of empirical studies includes Bolton and Kacperczyk (2021), Atta-Darkua *et al.* (2023), Starks, Venkat, and Zhu (2023), Hong and Shore (2023), and Starks (2023) among others.

The rest of the paper is organized as follows. Section II presents the baseline model of equilibrium ESG investing. Section III shows that there exists a unique linear equilibrium for ESG investment and a sufficient condition that describes the permissible extent of information asymmetry that must hold. Section IV presents an extension of the baseline model in which the manager's insight about the financial payoff from a risky project is positively correlated to her privately observed value of the project's ESG benefits and, thus, acquires materiality. Section V introduces the concept of incomplete ratings and derives the equilibrium under incomplete ratings. Section VI presents another extension of the baseline model that allows for an ESG dividend - a side payment that could be paid to either party. Section VII concludes. Proofs are in the Appendix.

II. Baseline Model

We model a two-period economy populated by competitive, risk-neutral investors and risk-averse managers. In the first period, a wealth-constrained manager seeks a financial investment into a commercial project from a profit-maximizing investor. In return for financial investment, the manager offers to the investor $q \geq 0$ shares in the risky cash flows of the project at the price $p > 0$ per share. Shares represent a method that allows the manager and the investor to distribute the project's future cash flows. Shares do not represent a transfer of ownership nor rights to the underlying technology.

The financial payoff is denoted by $\tilde{x}$, where $\tilde{x}$ is a random variable realized in the second period normalized to be in per share terms. We assume that $\tilde{x}$ is distributed normally with mean $\mu > 0$ and variance $\sigma_x^2 > 0$. The manager also receives in period one a private signal about the project's payoff, s, which is not known to the investor. The manager's noisy private signal about x is of the form $s = x + \epsilon$, where ϵ is independently normally distributed with mean zero and variance $\sigma_\epsilon^2 > 0$.

In addition, the manager receives privately observed non-pecuniary ESG attributes valued at v per share. The manager knows v but the investor does not; the investor only knows that v is a normally distributed random variable with mean zero and variance $\sigma_v^2 > 0$, which is uncorrelated with the payoff x (we will later relax this assumption). Unlike x that becomes known to both parties in period two, v remains latent to the investor. The investor does not get to see the realization of v in period two or ever. While this assumption seems strong, it is without the loss of generality. If we assume that neither party gets to know v with certainty, but that the manager has a more precise information about v, the math would be more complicated but the logic would be essentially the same - one party knows more about v than the other.

The investor treats v as "relevant" in that information about v can change the investor's assessment of or decision to invest (or not) in the project. The investor thus demands (since she is the one who holds the purse) that as part of the investment contract, certain conditions accompany the money. The conditions cannot be set directly on v since it is never observed or verified by the investor, but can be set in such a way as to (i) be relative to the share in the project's cash flows, q and (ii) enable the investor to limit attributes valued at v no matter what its realization is going to be.

The limit on ESG attributes, λ, is negotiated as part of the investment contract between the manager, who privately values ESG, and the competitive investor, who only maximizes expected financial profits. We model λ as a scalar, normalized to be between zero and one. Setting λ to zero means that the investor will only invest if the manager agrees to completely forgo all considerations related to the project's non-pecuniary ESG attributes.

More generally, one may think of λ as a set of contractual clauses, investing guidelines, and information disclosure requirements related to relevant ESG activities. Clauses on ESG activities may include, but are not limited to, specific duties of the Impact Committee that place restrictions on strategic and operational actions of the manager, including veto rights and permitted areas for ESG. ESG investing guidelines may include directions on which financially material ESG information—-e.g., contained in ESG ratings—-can be incorporated into the investing process. Some ESG investing guidelines may stipulate investing more in projects with higher ESG ratings, while others may stipulate investing more in projects with lower ESG ratings. Disclosure requirements related to ESG activities include rights to receive ESG-relevant information, rights to use external monitors of ESG impact, and rights to set policies, procedures and standards for internal ESG reporting and external ESG disclosure. Importantly, we do not link λ with compensation, vesting, dilution of ownership, control, or exit.

We define the contract as a triple (p, q, λ) agreed upon before period two. The contract does not include a credibly enforceable commitment not to renegotiate as it is unclear what an enforcement mechanism could possibly be based on. If such a commitment were to exist, then contractual incompleteness would not matter, because the counterparties could credibly agree in the initial period to contribute agreed-upon allocations, regardless of the non-contractibility of some variables.

The formal description of the game, agents, and the definition of equilibrium are as follows. At the start of the game, the manager receives ω, an endowment of a risk-free asset with gross intertemporal return normalized to one and a risky project with a payoff $\tilde{x}$. The manager offers q shares to the investor for the price of p. In the terminal period the manager receives value that consists of three parts: $\omega + (\tilde{x} - p)q + \lambda v q$, where ω is the value of the risk-free asset carried out from the initial period, $(\tilde{x} - p)q$ are the net financial benefits, and $\lambda v q$ are the non-pecuniary ESG benefits of the risky project.

We assume that the manager is risk averse. For simplicity of exposition, we use a well-known negative exponential utility function with a parameter $a > 0$. The manager determines the optimal number of shares q offered to an investor by maximizing the expected utility of the terminal period value conditional on her private knowledge of the signal, s and the value of ESG attributes, v, given the limit on the ESG impact, λ, and the price p offered by the investor, as well as all publicly available information:

$$\max_{q} E\Big\{ -\exp\Big(-a\big(\omega + (\tilde{x} - p)\, q + \lambda v q\big)\Big) | s, v, \lambda \Big\}. \tag{1}$$

Investors are competitive, expected profit maximizers. They lack both technology and experience requisite to managing the risky project. They do not value ESG characteristics, and are attracted solely by the prospects of expected financial return on their investment. Because each investor either already holds or has the ability to hold a diversified portfolio of risky projects, we assume that the investor is risk neutral. Investors do not know s or v. Additional costs or constraints associated with searching for risky investment projects, receiving and processing relevant information, and securing the receipt of cash flows are normalized to zero. The maximum loss for the investor is limited to the amount of investment in the risky project.

The investor offers a price schedule, $p(q)$, given the manager's offer, $q(\lambda)$ for a given λ, and all publicly available information. As investors operate in a competitive environment, they attempt to achieve a competitive target return, which we (without the loss of generality) normalize to zero. Under these assumptions, the solution to the investor's expected profit maximization problem - as long as it exists - must satisfy an expected zero return (getting the investment back) for any q:

$$E\Big\{\tilde{x}|q\Big\} = p(q). \tag{2}$$

All negotiations, followed by the transfer of investment funds, if any, take place in the initial period. In the second period, the payoff, $\tilde{x}$, is realized and observed by both parties, cash flows are distributed, and the project is liquidated. The investor agrees to finance the project only if the expected profit from buying a number of shares, q, for a price, p, is nonnegative.

A strategy for the manager is a rule that specifies an offer, $q(p)$, given the investor's price and the available private, as well as public information. A strategy for the investor is a rule that specifies a price offered to the manager, $p(q)$, and an ESG limit λ, based on the manager's offer and all publicly available information. The distributional and functional assumptions, objective functions of all participants, and the structure of the game are common knowledge.

We define a Bayesian-Nash equilibrium of this game as a strategy function, q, for the manager, a strategy function, $p(q)$, for the investor, and a constant, λ, such that the following conditions hold. First, the agents' strategies yield the best outcomes, conditional on the information available to them. Second, the agents' strategies are mutually consistent in the sense of Bayesian updating. Third, the investor who offers to the manager the least binding limit on ESG, λ, gets to finance the risky project.

III. Equilibrium ESG Investment

We show that there exists a unique linear equilibrium for ESG investment. For the equilibrium to exist, a sufficient condition that describes the permissible extent of information asymmetry must hold. The sufficient condition is expressed via λ. Consider the following proposition.

Proposition 1 (Equilibrium ESG Investment) *If*

$$\lambda^2 > \frac{\sigma_x^2}{\sigma_v^2}\left(\frac{\frac{\sigma_x^2}{\sigma_\epsilon^2}}{\frac{\sigma_x^2}{\sigma_\epsilon^2}+1}\right),$$

then there exists a unique linear equilibrium of the following form:

$$q^\star(\lambda) = \frac{E\big\{\tilde{x}|s\big\} - p + \lambda v}{aVar\big\{\tilde{x}|s\big\} + \phi(\lambda)},$$

$$p^\star(q) = \mu + \phi(\lambda)q,$$

and

$$\lambda^{\star} = inf\{\lambda \in [0,1] | \lambda^2 > \frac{\sigma_x^2}{\sigma_v^2} \left(\frac{\frac{\sigma_x^2}{\sigma_\epsilon^2}}{\frac{\sigma_x^2}{\sigma_\epsilon^2} + 1} \right) \},$$

where $\phi(\lambda) > 0$ *is a constant.*

Proposition 1 states that the equilibrium requires a "normalized" measure of information asymmetry must be smaller than the discounted scale of ESG activity. Moreover, and critically, the equilibrium ESG scale $\lambda^{\star}$ must be limited to the smallest value that satisfies the condition. In other words, the manager must agree to a limited scale of the ESG activities in order to attract investment.

Contractual terms imbedded in equilibrium $\lambda^{\star}$ that limit ESG activities while enabling project financing depend on the variances of signal and noise, and the variance of the private value of ESG. The following corollaries clarify the nature of this dependence.

Corollary 1 ($\lambda^{\star} > 0$)

In equilibrium, $\lambda^{\star}$, is always greater than zero. Investors who offer to disregard all ESG considerations by setting λ to zero do not finance such projects projects in equilibrium.

Corollary 2 (***Higher*** σ_v^2 ***(widepspread "greenwashing") ⇒ smaller*** $\lambda^{\star}$)

Higher variance of v leads to a smaller equilibrium $\lambda^{\star}$. To an investor, v is simply additional noise that makes it harder to extract information about $\tilde{x}$ from q. "Greenwashing" — a widespread provision of misleading information by company managers that their decisions reflect stakeholder ESG preferences when, in fact, they do not — can result in a higher aggregate σ_v^2. If investors suspect widespread greenwashing, they will further limit permissible equilibrium ESG activities.

Corollary 3 ($\sigma_v^2 = 0 \Rightarrow$ ***no trade***)

If $\sigma_v^2 = 0$, an investor can deduce the manager's private signal, s, from the offer schedule, q. This eliminates the informational advantage of the manager, but preserves investor's monopoly over funds. Knowing s, the investor offers such prices that the manager's expected terminal period wealth is equal to her initial cash holdings. This eliminates the manager's incentives to engage and leads to a no trade outcome. Under a no trade outcome, information relevant to the project's payoff is not revealed and risks are not shifted from a risk-averse manager to a risk-neutral investor - a market failure.

Corollary 4 (***Coordinated disclosure of information about*** v ***is not stable***)

As per Corollary 2, a manager would not want to voluntarily disclose v to investors as it would eliminate her informational advantage. However, a manager could be willing to disclose some information about her v via, for example, a third party, e.g., an ESG ratings provider, if it could help reduce σ_v^2. Since σ_v^2 is not specific to a particular manager, but rather reflects the aggregate uncertainty regarding v, a nontrivial proportion of managers would need to coordinate their voluntary information disclosure actions that would together result in a lower σ_v^2. While coordinated disclosure resulting in a lower σ_v^2 would benefit all managers as a group via a higher equilibrium $\lambda^\star$, each individual manager has an incentive not to disclose partial information about her specific v making such a coordinated voluntary disclosure approach suffer from the Prisoner's Dilemma concerns and, thus, not stable, i.e., not possible for some parameter values and knife-edge possible for others. The instability of a coordinated voluntary disclosure of partial information about v would also be consistent with a persistent disagreement among ESG impact ratings and scores for the same company, which has being empirically observed.

IV. Externality, ESG Materiality and Investment

In this section we extend the baseline model to include a core claim of ESG, namely that the taking into account environmental, social and governance effects creates a positive externality, which ultimately financially benefits the project itself. This claim is precisely the thinking behind the slogan "doing well by doing good".

We model it by assuming that the investor, while remaining neutral on ESG, is informed that the manager's private value of the project's ESG benefit v is positively correlated with s, the manager's private signal about the financial payoff on the project. This assumption would mean, for example, that units at a more "green" real estate development are expected to be associated with higher rental/leasing cash flows. However, this externality is not deterministic or guaranteed, but rather, a correlation. In other words, the presence of a positive correlation between s and v is a probabilistic statement. It does not indicate the presence of a structural relationship between $\tilde{x}$ and v.

The presence of this correlation changes how the investor treats non-pecuniary ESG benefits in her investment decision. Specifically, due to a positive correlation between ESG and the project's financial return, the investor must treat ESG benefits as not just as relevant but as "material" as the project's ESG value is associated with a benefit for its commercial value. A "naive" intuition would suggest that a positive commercial payoff association with the ESG activity should increase the project's attractiveness to the investor and lead to a higher tolerance of ESG activities in the project.

Specifically, consider $i = 1, 2, \dots$ managers characterized by different values of v_i, private values of the ESG benefits. Suppose that v_i and $s_i = \tilde{x} + \epsilon_i$ are positively correlated with a correlation coefficient $\rho_i > 0$. We show that such positive correlation leads to a smaller equilibrium $\lambda^\star(\rho_i)$ for all i. Consider the following proposition.

Proposition 2 (ESG Materiality and Equilibrium Investment) *If v_i and s_i are positively correlated with a known correlation coefficient $\rho_i > 0$ for $i = 1, 2, \ldots$, then*
(i) investors set equilibrium $\lambda^\star(\rho_i)$ for each i such that

$$\lambda^\star(\rho_i) = inf\Big\{\lambda \in [0,1] \ \Big| \ \frac{\sigma_v^2}{\sigma_x^2}\left(1 + \frac{\sigma_\epsilon^2}{\sigma_x^2}\right)\lambda^2 + \frac{2}{\sigma_x^2}\rho_i\sigma_v\sigma_\epsilon\lambda - 1 > 0\Big\}$$

and
(ii) for a given σ_v^2, $\lambda^\star(\rho_i) < \lambda^\star$ for all i, where $\lambda^\star$ is defined in Proposition (1).

This proposition offers another key insight from the model in this paper: ESG benefits from a project do *not* easily increase commercial payoffs because of the lack of societal consensus on what "good" externalities are and a lack of market aggregation (and thus measurability) of benefits that do exist. Therefore, the positive externality appears merely as a correlation between ESG value and payoffs (not a certainty), which *increases* the (material) payoff noise for the investor. This additional noise is proportional to $\rho_i\sigma_\epsilon\sigma_v\lambda$, where $\rho_i\sigma_\epsilon\sigma_v$ is project i's material ESG risk and λ is the management of this risk. As a result, for a given σ_v^2, the investor needs to impose *stricter ESG activity limits* in order to be able to elicit a useful signal about the project's commercial payoff. The fact that ESG benefits cannot be translated into tangible payoffs forces rational competitive investors to limit ESG activities more, even though they are not necessarily against them. Intuitively, ESG materiality forces rational investors to further limit ESG activities to get a less noisy signal about the financial payoff of the project.

V. ESG Risk Ratings and Scores

Reducing information asymmetry about the project's ESG materiality could be accomplished via a provision of information by third parties, e.g., ESG ratings or scores. In practice, while there are many providers of ESG ratings, they essentially provide two broad types of ESG ratings or scores - ESG *risk* ratings (or scores) and ESG *impact* scores (or ratings). An ESG impact score aims to measure a company‘s ESG impact. For example, Refinitiv states that its scores "are designed to transparently and objectively measure a company's relative ESG performance, commitment and effectiveness, based on company-reported data.“ Similarly, Moody‘s states that its "data, scores, and assessments can help distinguish meaning from noise to understand companies‘ environmental and social impact.“

In the context of our model, third-party providers of ESG impact scores would offer information to investors about a company‘s v_i and industry-specific σ_v^2. Recall, however, that as per Corollary 4, reducing aggregate σ_v^2 via a coordinated voluntary disclosure of partial information about v_i by managers $i = 1, 2, \ldots$ suffers from the Prisoner's Dilemma concerns which makes it not stable, i.e., not possible for some parameter values and knife-edge possible for others.

In contrast, voluntary disclosure of information by an individual manager about the project's ESG materiality, ρ_i, to the extent that it is beneficial to this manager, does not require any coordination with other managers. This brings us to ESG risk ratings that aim to measure a company‘s exposure to and management of material ESG risks. For example, Bloomberg states that its "ESG scores measure a company's management of financially material ESG issues, [where] financial materiality is defined as the issues that can have a negative or positive impact on a company's financial performance, such as revenue streams, operating costs, cost of capital, asset value and liabilities.“ MSCI states that its ESG risk ratings "aim to measure a company‘s management of financially relevant ESG risks and opportunities.“ S&P states that its "ESG Score measures a company‘s performance on and management of material ESG risks, opportunities, and impacts.“ Sustainalytics states that its ESG risk rating "captures an issuer‘s exposure to material, industry-specific ESG risks and an issuer's management of those risks.“

In the context of our model, third-party providers of ESG risk ratings would offer information to potential investors about $\rho_i\sigma_\epsilon\sigma_v$ for $i = 1, 2, \ldots$, where $\rho_i\sigma_\epsilon\sigma_v$ is project i's material ESG risk. Given the information provided, $\lambda^\star$ is set to "manage“ this risk.

Not that these ESG risk ratings can be described as *incomplete* insofar as they are only rating a correlation between latent ESG attributes and private signals about future payoffs. The concept of incomplete ratings, while newly introduced here, is intuitively similar to the concepts of incomplete markets and incomplete contracts. Incomplete markets describe an economic environment in which available securities do not cover all possible future states of the world. Incomplete contracts describe an economic environment in which available contractual clauses do not cover all possible future contingencies. Incomplete ESG risk ratings describe an economic environment in which available realizations of ratings do not cover all possible future material realizations of the rated latent variable.

Intuitively, having more securities to span additional possible future states of the world makes markets less incomplete and mitigates some allocational failures. Similarly, having a way to vary control over a project/company in situations when available contractual clauses do not cover all possible future contingencies mitigates some contractual failures. Continuing with this intuition, having a more differentiated grid of consistent ESG risk ratings could help quantify idiosyncratic risks of specific ESG attributes and mitigate possible no trade outcomes. Consider the following proposition.

Proposition 3 (Equilibrium Investment With ESG Risk Ratings) *Suppose that manager $i = 1, 2, \ldots$ discloses partial information to a third party - ESG risk ratings provider. Suppose also that this disclosure allows ESG ratings provider to compute a consistent estimator of $\rho_i\sigma_\epsilon\sigma_v$ - a project‘s material ESG risk - denoted by $\hat{\rho}\sigma_\epsilon\sigma_v$, then*

(i) investors compute equilibrium cutoff $\lambda^\star(\hat{\rho})$ such that

$$\lambda^\star(\hat{\rho}) = inf\big\{\lambda \in [0,1] \ \mid \ \frac{\sigma_v^2}{\sigma_x^2}\left(1 + \frac{\sigma_\epsilon^2}{\sigma_x^2}\right)\lambda^2 + \frac{2}{\sigma_x^2}\hat{\rho}\sigma_v\sigma_\epsilon\lambda - 1 > 0\big\} \tag{3}$$

and
(ii) finance projects that satisfy the condition for $\lambda^\star(\hat{\rho})$.

Under the distributional assumptions in the model, a consistent estimator of a project's ESG risk can be constructed as a simple average of the correlation coefficients, i.e., $\hat{\rho} = \sum_{i=1}^{N} \frac{\rho_i}{N}$ for a subset of N projects times the aggregate level of risk $\sigma_\epsilon \sigma_v$. Note that a partial disclosure of information by a number of managers to a third party results in the construction of a consistent estimator of ESG risk. While an investor is risk neutral, having such a consistent estimator is valuable to her in solving the signal extraction problem and achieving an equilibrium outcome.

Consequently, providers of ESG risk ratings should design robust and transparent processes that make their ratings *consistently disagree* depending on, for example, the number of, N and type of industry-specific projects, ρ_i that they rate. These ESG risk ratings would disagree across providers but consistently so. Consistent ESG risk ratings that quantify ESG risks would allow investors to cluster around specific ESG attributes that matter most across the stakeholders as signalled by the manager. As a result, financing would flow to projects with ESG characteristics that otherwise would not be financed.

Note that these ESG risk ratings can be described as *incomplete* insofar as they are only rating a correlation between latent ESG attributes and private signals about future payoffs. The concept of incomplete ratings, while newly introduced here, is intuitively similar to the concepts of incomplete markets and incomplete contracts. Incomplete markets describe an economic environment in which available securities do not cover all possible future states of the world. Incomplete contracts describe an economic environment in which available contractual clauses do not cover all possible future contingencies. Incomplete ESG risk ratings describe an economic environment in which available realizations of ratings do not cover all possible future material realizations of the rated latent variable.

Intuitively, having more securities to span additional possible future states of the world makes markets less incomplete and mitigates some allocational failures. Similarly, having a way to vary control over a project/company in situations when available contractual clauses do not cover all possible future contingencies mitigates some contractual failures. Continuing with this intuition, having a more differentiated grid of consistent ESG risk ratings could help quantify material ESG risks of specific projects/companies and mitigate possible no trade outcomes.

VI. Equilibrium ESG Dividends and Subsidies

Another possible way to avoid market failure in the presence of both financial and non-financial objectives is by providing dividends and subsidies. Consider an extension of the baseline model that allows the manager to offer the investor an ESG dividend denoted by $d \in \mathbb{R}$ (it may be positive or negative, a compensation or a fee) per share over and above the price $p(q)$ for q shares. The party that pays d agrees to a lower total payoff (by the dividend times the number of shares), in effect subsidizing the ESG project by receiving funds from an unspecified external source, e.g., taxpayers or benevolent agents.

In return for accepting d, the investor agrees to adjust $\lambda \in [0,1]$, in the direction requested by the manager. Other than d, all distributional and functional assumptions remain the same as in the baseline case. We show that there exists an equilibrium, which specifies contractual clauses λ and the ESG dividend, d, such that (i) the agents' strategies yield the best outcomes, conditional on the information available to them; the agents' strategies are mutually consistent in the sense of Bayesian updating; (iii) the investor who offers to the manager the best terms on the project's ESG impact, λ, given the proposed ESG dividend, d gets to finance the risky project.

Proposition 4 (Equilibrium ESG Dividend) *There is a pair $(d^{\star\star}, \lambda^{\star\star})$ that supports a unique linear equilibrium of the following form*

$$q((d^{\star\star}, \lambda^{\star\star}), p) = \frac{E\{\tilde{x}|s\} - p - d^{\star\star} + \lambda^{\star\star} v}{aVar\{\tilde{x}|s\} + \phi(\lambda^{\star\star})},$$

$$p((d^{\star\star}, \lambda^{\star\star}), q) = \mu - d^{\star\star} + \phi(\lambda^{\star\star}) q,$$

and

$$\lambda^{\star\star} = inf\{\lambda \in [0,1] | \lambda^2 > \frac{\sigma_x^2}{\sigma_v^2} \left(\frac{\frac{\sigma_x^2}{\sigma_\epsilon^2}}{\frac{\sigma_x^2}{\sigma_\epsilon^2} + 1} \right) \},$$

where $\phi(\lambda) > 0$ is a constant, provided that $\lambda^2 > \frac{\sigma_x^2}{\sigma_v^2} \left(\frac{\frac{\sigma_x^2}{\sigma_\epsilon^2}}{\frac{\sigma_x^2}{\sigma_\epsilon^2}+1} \right)$ for some $\lambda \in [0,1]$.

The tradeoff between equilibrium λ and d is described in the following three corollaries.

Corollary 5 ($d = 0$) *When $d = 0$, Proposition 4 reduces to Proposition 1, the baseline model.*

Corollary 6 ($d \neq 0$) *The presence of a nontrivial ESG dividend $d \neq 0$ makes no difference to the sufficient condition of the existence of equilibrium.*

These results imply that neither additional ESG dividend payments between the project's manager and investor nor third party subsidies — as long as they are deterministic — mitigate the limits imposed on equilibrium ESG activity. The reason is that a deterministic payout does not reduce the underlying information asymmetry, which is the source of the investor's problem. In other words, subsidization of the ESG project by some third party (such as the government or an ESG-minded private investor) will not reduce the problem that competitive rational investors cannot support an emphasis on ESG activities even when they do not bear direct costs from these activities being undertaken. Investors will happily take free or subsidized money, but they will not allow for any impactful ESG activities to take place within the project, which will manifest itself as ESG "impact washing."

VII. Conclusion

We present an equilibrium model of investment in a risky project with non-pecuniary ESG considerations. Investment decisions are made by rational competitive investors who are interested in profits only. The model reflects a view that from the point of view of a neutral financial investor, ESG considerations are idiosyncratic (have zero expected payoff) and are only tenuously connected to financial outcomes because there is no broad societal consensus on what net ESG benefits actually are. Moreover, investors typically cannot internalize or cash out on ESG benefits associated with their investments.

Corporate managers, on the other hand, must — according to empirical evidence, see, e.g., Graham (2022) — increasingly take into account not only the financial interests of their financial investors (shareholders) but also interests, benefits and preferences of other stakeholders such as employees, customers, and the local community. Managers are the ultimate insiders, so we assume that information about ESG characteristics/benefits rests with them rather than outside investors. Managers also receive a signal about the financial payoff of the project, while investors only know the general distribution of the project's payoff. Managers, however, are wealth-constrained and, must therefore trade their informational advantage and ESG interests against financial contributions from investors.

Therefore, though investors are neutral towards ESG in principle (as long as profits are not negatively affected), they have no choice but to act upon the information asymmetries that are ubiquitous in investments, i.e., the fact that managers know more about the project's value than investors. Hence, investors view ESG considerations of the project manager as a source of noise that needs to be reduced in order to elicit a sufficiently predictable signal about the true commercial value of the project. To do so, investors limit ESG activities to the point where only some informationally-acceptable projects get financed. Thus, financial investors with fiduciary responsibilities cannot tolerate the free pursuit of ESG activities even if they are profit neural or tenuously profit enhancing. This prediction of our model — unlike any existing model in which investors are assumed to be to some degree "green" — helps to explain recent decisions undertaken by a number of US states to limit the ability of public pension funds to consider non-pecuniary ESG benefits when making their investment decisions. As a corollary of this result, we show that the more uncertain investors are about ESG, e.g., due to a widepspread "greenwashing", the more they would limit permissible equilibrium ESG activities.

The model enables empirically testable propositions. First, ESG risk ratings and the information on which they are based may offer a market-based conceptual improvement of this conundrum: if they provide a consistent way to evaluate ESG materiality, which we model as a positive correlation (not a certainty) between the project's financial payoff and its non-pecuniary ESG characteristics. These ESG risk ratings can, however, be described as incomplete because the grid of ratings does not cover all possible future material realizations of the rated latent variable. Incompleteness of ESG risk ratings manifests itself in low and unstable correlations among them for the same entity offered by different third-party providers. We argue that under incomplete ratings, third-party providers should design robust and transparent processes that make their risk ratings *consistently disagree* depending on, for example, the number of and the types of companies that they rate. Consistent ESG risk ratings that quantify certain material ESG risks would help investors evaluate and invest in specific ESG characteristics that they care about such as societal benefits of lower CO2 emissions, or greater biodiversity, or reduced inequality. This might represent a large source of value from ESG risk ratings, which can be empirically tested.

Second, subsidization (or "dividends") of ESG projects is predicted to not solve the fundamental information asymmetry problem. If the government gives ESG-related projects money to help attract capital, the project may well get funded, but a key goal of such subsidization, namely to bring about a higher tolerance of the scale of ESG activities in the project, will not be achieved. Again, this prediction can be empirically tested - investors will happily take free or subsidized money, but they will not allow for any impactful ESG activities to take place within the project, which will manifest itself as ESG "impact washing."

One may object to our discussion with an argument that some investors do not operate under perfect competition and actually want to maximize profits rather than achieving zero (or target) expected profits, and they may indeed see profit opportunities from an ESG side activity in a project. Would such investors not value the ESG externality benefit even if it is tenuous? Indeed, making such an assumption would change the behavior of the model presented in this paper. However, we have already discussed the dangers in investing in "idiosyncratic externalities" in the absence of a generally agreed upon monetization of external ESG benefits, which may be seen as drawbacks by different groups of stakeholders. When an investor monetizes an externality, stakeholders may ask, "Which externality is it?". For example, an investor in a soccer stadium (which allegedly supports the development of soccer - and its presumed societal benefits) may at some point be criticized because the stadium was allegedly built on the back of exploited workers from poor countries. Modeling the maximization of ESG-related payoffs would risk coming to incorrect conclusions if the risk of such controversy was not included.

Very few theoretical examinations of ESG investment are available to date with a notable exception of Goldstein *et al.* (2022) which assumes that ESG characteristics provide additional positive financial payoffs with certainty. We think this needs to be justified. If anything, recent empirical evidence, e.g., Hong and Shore (2023), points to the existence of a "greenium", a persistently negative difference (risk-adjusted) in costs of capital for green versus brown firms.

Our proposal is that a key feature of ESG-related projects is their signaling value for investors who are trying to learn more about the true commercial value of a project or asset; it is this signaling value that forces limits on ESG activities. There are some signs that ESG characteristics of assets may acquire an additional signaling value in the future, namely signaling about organizational quality and compliance with future (yet to be considered) regulations and environmental requirements. For example, if one believes that a carbon tax of \$100/ton of CO_2 emissions is inevitable a decade from now, investing in an emissions-reducing project now sends a signal of higher robustness and resilience in the future.

Overall, we believe that developing a framework for including ESG, Socially Responsible or Sustainable considerations into risky investment decisions in competitive market economies is a critically important problem for academic and practitioner finance at this point. Even a limited solution to this problem could unleash large private capital flows into solving critical environmental and societal problems and, in the process, ensuring the long-term legitimacy of established financial actors and market institutions. Understanding the information-related economics of ESG investment is an important requirement in the near future in order to correctly interpret the support that investors can give the transformation of our market economies. Much more theoretical work needs to be done.

References

Atta-Darkua, Vaska, Simon Glossner, Philipp Krueger, and Pedro Matos, 2023, Decarbonizing institutional investor portfolios: Helping to green the planet or just greening your portfolio?, mimeo.

Barber, Brad, Morse, Adair, and Ayako Yasuda, 2021, Impact investing, *Journal of Financial Economics* 139 (1), 162-185.

Barzuza, Michal, Curtis, Quinn, and David Webber, 2023, The Millennial Corporation: Strong Stakeholders, Weak Managers, *Stanford Journal of Law, Business and Finance* 28 (2), 255-306.

Benabou, Roland, and Jean Tirole, 2010, Individual and Corporate Social Responsibility *Economica* 77, 1–19.

Bolton, Patrick, and Marcin Kacperczyk, 2021, Do investors care about carbon risk?, *Journal of Financial Economics* 142, 517–549.

Berg, Florian, Kolbel, Julian, and Roberto Rigobon, 2022, Aggregate Confusion: The Divergence of ESG Ratings, *Review of Finance*, 26 (6), 1315–1344.

Chen, Xuhui, Lorenzo Garlappi, and Ali Lazrak, 2023, Responsible Consumption, Demand Elasticity, and the Green Premium, mimeo.

Christensen, Dane, Serafeim, George, and Anywhere Sikochi, 2022, *Why is Corporate Virtue in the Eye of The Beholder? The Case of ESG Ratings*, The Accounting Review 97 (1), 147–175.

Dai, Xin, Gao, Feng, Lisic, Ling Lei, and Ivy Xiying Zhang, 2023, Corporate Social Performance and the Managerial Labor Market, *Review of Accounting Studies* 28, 307–339.

DeGroot, Morris H., 1970, *Optimal Statistical Decisions* (McGraw-Hill, New York, NY).

Diamond, Douglas W., and Robert E. Verecchia, 1981, Information aggregation in a noisy rational expectations economy, *Journal of Financial Economics* 9, 221-235.

Diamond, Douglas W., and Robert E. Verecchia, 1991, Disclosure, Liquidity, and the Cost of Capital, *Journal of Finance*, 46 (4), 1325-1359.

Hong, Harrison and Edward Shore, 2023, Corporate Social Responsibility, *Annual Review of Financial Economics* 15, 327-350.

Geczy, Christopher, Jeffers, Jessica S., Musto, David K., and Anne M. Tucker, 2021, Contracts with (Social) benefits: The implementation of impact investing, *Journal of Financial Economics* 142 (2), 697-718.

Freiberg, David, Rogers, Jean, and George Serafeim, 2020, How ESG Issues Become Financially Material to Corporations and Their Investors, mimeo.

Goldstein, Itay, Kopytov, Alexandr, Shen, Lin, and Haotian Xiang, 2022, On ESG Investing: Heterogeneous Preferences, Information, and Asset Prices, mimeo.

Graham, John R., 2022, Presidential Address: Corporate Finance and Reality, *Journal of Finance* 77, 1975-2049.

Grossman, Sanford J., and Oliver Hart, 1986, The costs and benefits of ownership, *Journal of Political Economy* 94, 691-719.

Hart, Oliver, and John Moore, 1988, Incomplete contracts and renegotiation, *Econometrica* 56, 755-785.

Kirilenko, Andrei, 2001, Valuation and Control in Venture Finance, *Journal of Finance*, 56 (2), 565-587.

Kruger, Philipp, 2015, Corporate goodness and shareholder wealth, *Journal of Financial Economics*, 115 (2), 304-329.

Pastor, Lubos, Stambaugh, Robert, and Lucian Taylor, 2021, Sustainable investing in equilibrium, *Journal of Financial Economics* 142, 550–571.

Pastor, Lubos, Stambaugh, Robert, and Lucian Taylor, 2024, Green Tilts, mimeo.

Renneboog, Luc, Ter Horst, Jenke, and Chendi Zhang, 2011, Is ethical money financially smart? Nonfinancial attributes and money flows of socially responsible investment funds, *Journal of Financial Intermediation* Volume 20 (4), 562-588.

Sauzet, Maxime, and Olivier Zerbib, 2024, When Green Investors Are Green Consumers, mimeo.

Starks, Laura, Parth Venkat, and Qifei Zhu, 2023, Corporate ESG profiles and investor horizons, mimeo.

Starks, Laura, 2023, Presidential address: Sustainable finance and ESG issues—Value versus values, *Journal of Finance* 78, 1837–1872.

Appendix

Proof of Proposition 1

The functional and distributional assumptions jointly imply that a solution to the manager's expected utility maximization problem (1) must satisfy the following first-order,

$$E\{\tilde{x}|s\} - aVar\{\tilde{x}|s\}q - p'(q)q - p(q) + \lambda v = 0 \tag{4}$$

and second-order optimality conditions,

$$-aVar\{\tilde{x}|s\} - 2p'(q) - p''(q) < 0. \tag{5}$$

Applying the result on conjugate families for samples from a normal distribution (DeGroot (1970)), the manager's posterior distribution of $\tilde{x}$ given the normal prior and the signal s, is also normal with mean

$$E\{\tilde{x}|s\} = \frac{\frac{s}{\sigma_\epsilon^2} + \frac{\mu}{\sigma_x^2}}{\frac{1}{\sigma_\epsilon^2} + \frac{1}{\sigma_x^2}} \tag{6}$$

and variance

$$Var\{\tilde{x}|s\} = \frac{1}{\frac{1}{\sigma_\epsilon^2} + \frac{1}{\sigma_x^2}}. \tag{7}$$

The investor does not observe s and v separately, and therefore cannot distinguish whether the manager's decision to offer q is motivated by the information about the project's expected payoff, $\tilde{x}$, or the value of its ESG impact, v. However, the investor can use the manager's offer schedule to update her information set, construct a posterior distribution of $\tilde{x}$, and then offer the price schedule that solves her expected profit maximization problem using the posterior, if such a solution exists.

Define

$$r(q) \equiv aVar\{\tilde{x}|s\}q + p'(q)q + p(q) \tag{8}$$

and

$$z(q) \equiv \Big(r(q)\Big(\frac{1}{\sigma_\epsilon^2} + \frac{1}{\sigma_x^2}\Big) - \frac{\mu}{\sigma_x^2}\Big)\sigma_\epsilon^2. \tag{9}$$

By construction, as per Equation (4), $r(q) = E\{\tilde{x}|s\} + \lambda v$. Substituting $r(q)$ into $z(q)$ results in

$$z(q) = s + \left(1 + \frac{\sigma_\epsilon^2}{\sigma_x^2}\right)\lambda v = \tilde{x} + \epsilon + \left(1 + \frac{\sigma_\epsilon^2}{\sigma_x^2}\right)\lambda v. \tag{10}$$

That is, $z(q)$ is a descriptive statistic of the risky payoff, $\tilde{x}$, contained in the manager's offer, q, polluted, from the investor's perspective, by two sources of noise, ϵ and v. Note that from the investor's perspective, a higher value of λ increases the payoff-irrelevant noise and, thus, results in a lower informational efficiency of equilibrium prices and in a suboptimality of equilibrium investment.

Since $\tilde{x}$, ϵ and v are independently normally distributed, $z \sim N\left(\tilde{x}, \sigma_z^2\right)$, where $\sigma_z^2 = \sigma_\epsilon^2 + \left(1 + \frac{\sigma_\epsilon^2}{\sigma_x^2}\right)^2 \lambda^2 \sigma_v^2$.

The investor's posterior distribution of $\tilde{x}$ given the normal prior and the offer q, is also normal with mean

$$E\{\tilde{x}|q\} = \frac{\frac{z}{\sigma_z^2} + \frac{\mu}{\sigma_x^2}}{\frac{1}{\sigma_z^2} + \frac{1}{\sigma_x^2}} \tag{11}$$

and variance

$$Var\{\tilde{x}|q\} = \frac{1}{\frac{1}{\sigma_z^2} + \frac{1}{\sigma_x^2}}. \tag{12}$$

Substituting (11) into the condition (2) and then using equations (7), (8) and (9), and collecting terms, results in

$$p'(q)q = \alpha\Big(p(q) - \mu\Big) - \beta q, \tag{13}$$

where

$$\alpha = \frac{\sigma_z^2 - \sigma_\epsilon^2}{\sigma_x^2 + \sigma_\epsilon^2} \tag{14}$$

and

$$\beta = \frac{a\sigma_x^2\sigma_\epsilon^2}{\sigma_x^2 + \sigma_\epsilon^2} \tag{15}$$

are positive constants.

The solution to the first-order differential equation (13) is of the form

$$p(q) = \mu + \frac{\beta}{\alpha - 1} q - C|q|^{\alpha} \quad (16)$$

for $\alpha \neq 1$ and

$$p(q) = \mu + Cq - \beta \ln(|q|) q \quad (17)$$

for $\alpha = 1$, where C is some constant of integration.

There are nine possible combinations of α and C, corresponding to the values of α greater, less than, or equal to 1 and values of C greater, less than, or equal to 0. We can eliminate all but one of these combinations as follows. When $\alpha > 1$ and $C < 0$, the second-order optimality condition (5) is violated. Condition (5) is also violated when $\alpha \leq 1$ and $C \geq 0$. When $\alpha > 1$ and $C > 0$, an investor who offers to finance the project at the same $\alpha > 1$, but with $C = 0$ would be selected over to any investor offering $C > 0$. Finally, $\alpha \leq 1$ and $C < 0$, implies that the investor agrees to pay more per share if she buys more shares. This means that the more shares the manager offers, the higher per share equilibrium price is. This represents an arbitrage opportunity that would come into conflict with the condition (2) and cannot be sustained in equilibrium. This leaves us with a unique combination of α and C that can support an equilibrium, namely, $\alpha > 1$ and $C = 0$. As a result, the equilibrium price rule is unique and linear in q,

$$p^{\star}(q) = \mu + \frac{\beta}{\alpha - 1} q = \mu + \phi(\lambda) q, \quad (18)$$

where

$$\phi(\lambda) = \frac{a\sigma_x^2 \sigma_\epsilon^2}{\left(1 + \frac{\sigma_\epsilon^2}{\sigma_x^2}\right)\left(\left(1 + \frac{\sigma_\epsilon^2}{\sigma_x^2}\right)\lambda^2 \sigma_v^2 - \sigma_x^2\right)} > 0. \quad (19)$$

Thus, a unique equilibrium of the form

$$p^{\star}(q) = \mu + \phi(\lambda) q \quad (20)$$

and

$$q^{\star}(\lambda) = \frac{E\{\tilde{x}|s\} - p + \lambda v}{aVar\{\tilde{x}|s\} + \phi(\lambda)}, \quad (21)$$

exists if $\alpha > 1$, which is equivalent to

$$\lambda^2 > \frac{\sigma_x^2}{\sigma_v^2}\left(\frac{\frac{\sigma_x^2}{\sigma_\epsilon^2}}{\frac{\sigma_x^2}{\sigma_\epsilon^2}+1}\right). \tag{22}$$

Equivalently, this can be expressed as a quadratic inequality

$$\frac{\sigma_v^2}{\sigma_x^2}\left(1+\frac{\sigma_\epsilon^2}{\sigma_x^2}\right)\lambda^2 - 1 > 0. \tag{23}$$

Lastly, in equilibrium, the investor who offers to the manager the best acceptable terms on the project's ESG impact, λ, gets to finance the deal. That is, in equilibrium,

$$\lambda^\star = inf\left\{\lambda \in [0,1] \mid \lambda^2 > \frac{\sigma_x^2}{\sigma_v^2}\left(\frac{\frac{\sigma_x^2}{\sigma_\epsilon^2}}{\frac{\sigma_x^2}{\sigma_\epsilon^2}+1}\right)\right\}. \tag{24}$$

Equivalently,

$$\lambda^\star = inf\left\{\lambda \in [0,1] \mid \frac{\sigma_v^2}{\sigma_x^2}\left(1+\frac{\sigma_\epsilon^2}{\sigma_x^2}\right)\lambda^2 - 1 > 0\right\}. \tag{25}$$

This completes the proof.

Proof of Proposition 2

Consider $i = 1, 2, \ldots$ managers characterized by different values of v_i, private values of non-pecuniary ESG benefits. Suppose that v_i and $s_i = \tilde{x} + \epsilon_i$ are positively correlated with a known and observed by all parties correlation coefficient $\rho_i > 0$.

The investor still does not observe s_i and v_i separately and, thus, still cannot distinguish whether the manager's decision to offer q_i is motivated by the information about the project's expected financial payoff, $\tilde{x}$, or the value of non-pecuniary ESG benefits, v_i. As in the baseline model, the investor can use the manager's offer schedule to update her information set, construct a posterior distribution of $\tilde{x}$, and then offer the price schedule that solves her expected profit maximization problem using the posterior, if such a solution exists.

Recall that in order to do so, the investor forms $z(q)$, which as per Equation (10), is a descriptive statistic of the risky payoff, $\tilde{x}$, contained in the manager's offer, q, polluted by ϵ and v. Recall also that in the baseline case, $\tilde{x}$, ϵ and v were assumed to be independently normally distributed.

Without the loss of generality, we relax the assumption that $\tilde{x}$, ϵ and v are iid and consider the case when the investor is assumed to know ρ_i for each $i = 1, 2, \ldots$. In this case, for each $i = 1, 2, \ldots$, $z_i \sim N\Big(\tilde{x}, \sigma^2_{z_i}\Big)$, where

$$\sigma^2_{z_i} = \sigma^2_\epsilon + \Big(1 + \frac{\sigma^2_\epsilon}{\sigma^2_x}\Big)^2 \lambda^2 \sigma^2_v + 2\rho_i \Big(1 + \frac{\sigma^2_\epsilon}{\sigma^2_x}\Big)\sigma_\epsilon \sigma_v \lambda. \tag{26}$$

The rest of the proof follows the baseline model with z_i instead of z. The necessary and sufficient condition for the existence of unique equilibrium takes the form

$$\frac{\sigma^2_v}{\sigma^2_x}\left(1 + \frac{\sigma^2_\epsilon}{\sigma^2_x}\right)\lambda^2 + 2\rho_i \frac{\sigma_v}{\sigma_x}\frac{\sigma_\epsilon}{\sigma_x}\lambda - 1 > 0. \tag{27}$$

This condition is similar to the baseline condition (23) while accounting for a positive correlation between s and v.

Consequently, for a known ρ_i, investors set equilibrium $\lambda^\star(\rho_i)$ for each i such that

$$\lambda^\star(\rho_i) = inf\Big\{\lambda \in [0,1] \ \Big| \ \frac{\sigma^2_v}{\sigma^2_x}\left(1 + \frac{\sigma^2_\epsilon}{\sigma^2_x}\right)\lambda^2 + 2\rho_i \frac{\sigma_v}{\sigma_x}\frac{\sigma_\epsilon}{\sigma_x}\lambda - 1 > 0\Big\}. \tag{28}$$

This completes the proof of part (i) of the proposition.

To prove part (ii) of the proposition, we need to show that $\lambda^\star$ as defined in (25) is always greater than $\lambda^\star(\rho_i)$ for $i = 1, 2, \ldots$ and $\lambda^\star(\hat{\rho})$.

Define $\bar{\lambda}$ as positive solution of the following quadratic equation:

$$\frac{\sigma^2_v}{\sigma^2_x}\left(1 + \frac{\sigma^2_\epsilon}{\sigma^2_x}\right)\lambda^2 - 1 = 0. \tag{29}$$

Namely,

$$\bar{\lambda} = \frac{\sigma_x}{\sigma_v\sqrt{1 + \frac{\sigma^2_\epsilon}{\sigma^2_x}}}. \tag{30}$$

For each $i = 1, 2, \ldots$, define also $\bar{\lambda}(\rho_i)$ as positive solution of the quadratic equation:

$$\frac{\sigma^2_v}{\sigma^2_x}\left(1 + \frac{\sigma^2_\epsilon}{\sigma^2_x}\right)\lambda^2 + 2\rho_i \frac{\sigma_v}{\sigma_x}\frac{\sigma_\epsilon}{\sigma_x}\lambda - 1 = 0. \tag{31}$$

Plug in $\bar{\lambda}$ into (31) and note that for the equation to hold, $\frac{2\rho_i\sigma_\epsilon}{\sigma_x\sqrt{1+\frac{\sigma^2_\epsilon}{\sigma^2_x}}}$ must be equal to zero for some $i = 1, 2, \ldots$, which is a contradiction, since all of the terms are assumed to be positive.

We have proven that $\bar{\lambda} \neq \bar{\lambda}(\rho_i)$ for $i = 1, 2, \ldots$. Without the loss of generality, define $\kappa(\rho_i) \geq 0$ such that $\bar{\lambda}(\rho_i) = \bar{\lambda}\kappa(\rho_i)$ for $i = 1, 2, \ldots$. We need to show that $\kappa(\rho_i) < 1$ for all $i = 1, 2, \ldots$.

Define, $A = \frac{\sigma_v^2}{\sigma_x^2}\left(1 + \frac{\sigma_\epsilon^2}{\sigma_x^2}\right)$, $B(\rho_i) = 2\rho_i \frac{\sigma_v}{\sigma_x}\frac{\sigma_\epsilon}{\sigma_x}$ and $C = -1$. Note that quadratic equation (31) takes the standard form $A\lambda^2 + B(\rho_i)\lambda + C = 0$. Note also that equation (29) takes the form $A\lambda^2 + C = 0$ with the solution that satisfies $\bar{\lambda}^2 = \frac{1}{A}$.

The positive solution to the quadratic equation (31) is of the form

$$\bar{\lambda}(\rho_i) = \frac{-B(\rho_i) + \sqrt{B^2(\rho_i) + 4A}}{2A}, \tag{32}$$

which after substituting $\bar{\lambda}^2$ for $\frac{1}{A}$ and collecting terms takes the form

$$\bar{\lambda}(\rho_i) = \frac{\bar{\lambda}}{2}\left(\sqrt{B^2(\rho_i)\bar{\lambda}^2 + 4} - B(\rho_i)\bar{\lambda}\right) = \bar{\lambda}\kappa(\rho_i), \tag{33}$$

where

$$\kappa(\rho_i) = \frac{1}{2}\left(\sqrt{B^2(\rho_i)\bar{\lambda}^2 + 4} - B(\rho_i)\bar{\lambda}\right). \tag{34}$$

Note that when $\rho_i = 0$ for some $i = 1, 2, \ldots$, $B(\rho_i) = 0$, $\kappa(\rho_i) = 1$ and $\bar{\lambda}(\rho_i) = \bar{\lambda}$. Suppose that for some $i = 1, 2, \ldots$, $\kappa(\rho_i) > 1$, i.e.,

$$\frac{1}{2}\left(\sqrt{B^2(\rho_i)\bar{\lambda}^2 + 4} - B(\rho_i)\bar{\lambda}\right) > 1, \tag{35}$$

or

$$\sqrt{B^2(\rho_i)\bar{\lambda}^2 + 4} > B(\rho_i)\bar{\lambda} + 2. \tag{36}$$

Taking the square of both sides results in

$$B^2(\rho_i)\bar{\lambda}^2 + 4 > B^2(\rho_i)\bar{\lambda}^2 + 4 + 4B(\rho_i)\bar{\lambda}, \tag{37}$$

which is a contradiction. Thus, for all $i = 1, 2, \ldots$, $\kappa(\rho_i) < 1$ and, consequently, $\bar{\lambda}(\rho_i) < \bar{\lambda}$.

Lastly, since in equilibrium, $\bar{\lambda}$ is minimally close to $\lambda^\star$ and for $i = 1, 2, \ldots$, $\bar{\lambda}(\rho_i)$ is minimally close to $\lambda^\star(\rho_i)$, it must be the case that, for a given σ_v^2, $\lambda^\star(\rho_i) < \lambda^\star$ for all i. This completes the proof.

Proof of Proposition 3

The proof of the proposition follows exactly the same steps as in Proposition 2 but under the assumption that while each $\rho_i > 0$ for $i = 1, 2, \ldots$ is not observable to investors, they receive from a third party a consistent estimator of each $\rho_i > 0$ for $i = 1, 2, \ldots$ denoted by $\hat{\rho} > 0$. Under the distributional assumptions in the model, a consistent estimator of a project‘s ESG risk can be constructed as a simple average of the correlation coefficients, i.e., $\hat{\rho} = \sum_{i=1}^{N} \frac{\rho_i}{N}$ for a subset of N projects times the aggregate level of risk $\sigma_\epsilon \sigma_v$.

While investors still do not observe s_i and v_i for any project separately, they can use a specific manager's offer schedule to update their information sets, construct a posterior distribution of $\tilde{x}$, and then offer the price schedule that solves the expected profit maximization problem using the posterior, if such a solution exists.

Consequently, if each $\rho_i > 0$ is not known, but investors receive a consistent estimator $\hat{\rho} > 0$, then they set cutoff $\lambda^\star(\hat{\rho})$ as follows

$$\lambda^\star(\hat{\rho}) = inf\Big\{\lambda \in [0, 1] \ \mid \ \frac{\sigma_v^2}{\sigma_x^2}\left(1 + \frac{\sigma_\epsilon^2}{\sigma_x^2}\right)\lambda^2 + 2\hat{\rho}\frac{\sigma_v}{\sigma_x}\frac{\sigma_\epsilon}{\sigma_x}\lambda - 1 > 0\Big\} \tag{38}$$

and only finance projects that satisfy the condition for $\lambda^\star(\hat{\rho})$. This completes the proof of the proposition.

Proof of Proposition 4

Denote by $d \in \mathbb{R}$ an ESG dividend expressed in per share terms. The objective function of the manager takes the following form

$$\max_q E\Big\{ - \exp\Big(- a\big(W_0 + (\tilde{x} - d - p)\, q + \lambda v q\big)\Big)|s, v, \lambda, d\Big\}, \tag{39}$$

where the price p per share is net of the ESG dividend.

Similarly, the solution to the investor's expected profit maximization problem adjusted for d must satisfy

$$E\Big\{\tilde{x}|q\Big\} = p(q) + d. \tag{40}$$

Note that while d is not restricted to be nonnegative, it is expressed to represent an additional payment from the to manager to the investor.

A solution to the manager's expected utility maximization problem (39) must satisfy the following first-order,

$$E\big\{\tilde{x}|s\big\} - aVar\big\{\tilde{x}|s\big\}q - p'(q)q - p(q) - d + \lambda v = 0 \tag{41}$$

and second-order optimality conditions,

$$-aVar\{\tilde{x}|s\} - 2p'(q) - p''(q) < 0. \tag{42}$$

We proceed to define

$$r(q) \equiv aVar\{\tilde{x}|s\}q + p'(q)q + p(q) \tag{43}$$

and

$$z_d(q) \equiv \Big(r(q)\Big(\frac{1}{\sigma_\epsilon^2} + \frac{1}{\sigma_x^2}\Big) - \frac{\mu}{\sigma_x^2}\Big)\sigma_\epsilon^2 + \Big(1 + \frac{\sigma_\epsilon^2}{\sigma_x^2}\Big)d. \tag{44}$$

By construction, as per Equation (41), $r(q) = E\{\tilde{x}|s\} + \lambda v - d$. Substituting $r(q)$ into $z_d(q)$ results in

$$z_d(q) = s + \Big(1 + \frac{\sigma_\epsilon^2}{\sigma_x^2}\Big)\lambda v = \tilde{x} + \epsilon + \Big(1 + \frac{\sigma_\epsilon^2}{\sigma_x^2}\Big)\lambda v. \tag{45}$$

That is, $z_d(q)$ is a descriptive statistic of the risky payoff, $\tilde{x}$, contained in the manager's offer, q, polluted, from the investor's perspective, by two sources of noise, ϵ and v.

As in the baseline case, since $\tilde{x}$, ϵ and v are independently normally distributed, $z_d \sim N\Big(\tilde{x}, \sigma_z^2\Big)$, where $\sigma_z^2 = \sigma_\epsilon^2 + \Big(1 + \frac{\sigma_\epsilon^2}{\sigma_x^2}\Big)^2 \lambda^2 \sigma_v^2$.

The investor's posterior distribution of $\tilde{x}$ given the normal prior and the offer q, is also normal with mean

$$E\{\tilde{x}|q\} = \frac{\frac{z_d}{\sigma_z^2} + \frac{\mu}{\sigma_x^2}}{\frac{1}{\sigma_z^2} + \frac{1}{\sigma_x^2}} \tag{46}$$

and variance

$$Var\{\tilde{x}|q\} = \frac{1}{\frac{1}{\sigma_z^2} + \frac{1}{\sigma_x^2}}. \tag{47}$$

Substituting (46) into the condition (40) and then using equations (7), (43) and (44), and collecting terms, results in

$$p'(q)q = \alpha\Big(p(q) - (\mu - d)\Big) - \beta q, \tag{48}$$

where α and β are defined in Equations (14) and (15), respectively - same as in the baseline case.

The rest of the proof follows the baseline case in that the equilibrium price rule adjusted for d is unique and linear in q,

$$p^{\star\star}(d, \lambda, q) = \mu - d + \phi(\lambda)q, \tag{49}$$

where $\phi(\lambda)$ is defined in (19) and

$$q^{\star\star}(d, \lambda, p) = \frac{E\{\tilde{x}|s\} - p - d + \lambda v}{aVar\{\tilde{x}|s\} + \phi(\lambda)}, \tag{50}$$

exists if $\alpha > 1$, which is again equivalent to

$$\lambda^2 > \frac{\sigma_x^2}{\sigma_v^2}\left(\frac{\frac{\sigma_x^2}{\sigma_\epsilon^2}}{\frac{\sigma_x^2}{\sigma_\epsilon^2} + 1}\right). \tag{51}$$

This completes the proof.