Model-free ESG Frontier

Kai Li and Bo Sun*

February 13, 2026

Abstract

This paper develops a novel, model-free framework for ESG-integrated asset pricing, grounded in the classical Law of One Price (LOOP) and a new principle proposed in this study—the Law of One ESG-adjusted Price (LOOE). LOOE extends LOOP into the ESG dimension, enabling consistent valuation of sustainability attributes across portfolios, including long-short strategies. Under mild conditions on asset structure and investor heterogeneity, LOOP and LOOE arise endogenously in competitive equilibrium. Together, these assumptions support the construction of an ESG-adjusted stochastic discount factor (ESDF), which jointly prices financial and ESG risks without relying on explicit utility functions. Using the SDF–ESDF pair, we derive a generalized four-fund separation theorem and a two-beta representation linking asset returns to market and ESG exposures. Empirically, the two-beta model enables flexible fund selection along the ESG-efficient frontier and substantially improves pricing performance relative to both the CAPM and the ESG factor model of Pástor et al. (2021). By introducing LOOE and extending classical pricing theory, this paper provides a flexible and empirically robust foundation for ESG investing.

JEL Codes: G11, G12, G23, M14, Q5

Keywords: Portfolio choice, ESG, Socially responsible investing, Impact investing, Sustainable investing, CSR, Carbon, Governance

*Kai Li (kaili825@gmail.com), Peking University HSBC Business School and Tsinghua University PBC School of Finance; and Bo Sun (harevysun@gmail.com), Peking University HSBC Business School. We are grateful to Hengjie Ai, Zhi Da, Bing Han, Xiaoji Lin, Thomas Sargent, and Chaoyi Zhao for their helpful comments and suggestions. We also thank the audiences at the 2025 MRS International Risk Conference and the HEC–HKUST Sustainable Finance Webinar for valuable feedback. We acknowledge financial support from the Theme-based Research Scheme T31-603/21-N of the Hong Kong Research Grants Council (RGC). The usual disclaimer applies.

1 Introduction

The explosive growth of ESG investing—now surpassing tens of trillions of dollars in assets—poses a fundamental challenge to classical asset pricing theory. Traditional models are built to price financial risk and return, yet fail to accommodate non-pecuniary preferences that increasingly shape investor behavior and market outcomes.

Recent theoretical advances (e.g., Pedersen et al. (2021) and Pástor et al. (2021)) incorporate ESG preferences into asset pricing using utility-based equilibrium models. While elegant, these frameworks impose restrictive assumptions—homogeneous ESG preferences, specific utility functions, and structural equilibrium—which limit empirical flexibility and fail to generalize to heterogeneous investor environments or long-short portfolios.

To address these limitations, we propose a model-free ESG pricing framework built on two minimal but powerful principles: the Law of One Price (LOOP) and a novel contribution of this paper, the Law of One ESG-adjusted Price (LOOE). These assumptions allow us to derive both a traditional stochastic discount factor (SDF) and an ESG-adjusted SDF (ESDF), enabling arbitrage-free valuation across financial and sustainability dimensions—without assuming any specific investor utility.

Building on this foundation, we derive a four-fund separation theorem and a two-beta representation that simplify portfolio construction and empirical implementation in ESG-integrated markets. Our model accommodates heterogeneous preferences, long-short strategies, and generalizes prior utility-based approaches as special cases.

Specifically, we introduce the ESG-adjusted price as a linear functional that parallels the usual pricing functional but scales asset prices by their ESG scores. As shown in Section 2, this construct preserves linearity and remains well-defined even for portfolios with zero market value, thereby overcoming the limitations of standard value-weighted ESG scores when evaluating long–short or market-neutral strategies. We also introduce an alternative no-arbitrage condition, the Law of One ESG-adjusted Price (LOOE), which requires payoff-equivalent portfolios to carry the same ESG-adjusted price. LOOP and LOOE jointly provide the analytical foundation of our unified pricing framework.

Importantly, these principles are not imposed ad hoc. Section 2 shows that each of several economically transparent assumptions—non-redundancy of basis assets, a competitive equilibrium with both ESG-indifferent and ESG-sensitive investors, or consistency of ESG-score replication—is by itself sufficient for LOOP and LOOE to hold (Proposition 3.2). Appendix A.1 illustrates this equilibrium logic in a parsimonious general-equilibrium model with heterogeneous investors, showing that the mere coexistence of ESG-indifferent and

ESG-sensitive investors guarantees LOOP and LOOE regardless of their relative proportions. Building on the same setup, Appendix A.2 further demonstrates that changes in the share of ESG-sensitive investors (or the strength of their ESG preferences) instead affect the magnitude of the green premium, amplifying it as ESG preferences become more prevalent or more intense.

The core intuition underlying our approach is that LOOP provides the basis for constructing the stochastic discount factor (SDF) m^*, while LOOE underpins the ESG-adjusted stochastic discount factor (ESDF) g^*. The SDF captures traditional sources of systematic risk and preserves consistency with classical asset pricing, whereas the ESDF tilts state prices toward outcomes that are unfavorable in ESG terms, thereby incorporating the non-pecuniary tastes of ESG-conscious investors. Together, m^* and g^* deliver a unified, SDF-based description of asset prices in which financial and sustainability risks are jointly accounted for. In the subsections that follow, we formalize the ESG-adjusted risk-neutral probability measure induced by g^* and provide a state-level interpretation based on Arrow–Debreu securities, which clarifies how g^* reweights ESG-critical and ESG-benign states relative to the benchmark SDF.

While LOOP allows us to construct the SDF through projection methods, we further extend the framework by introducing the ESG mean–variance frontier and four-fund separation for a more efficient construction of the SDF. By drawing an analogy to the Hansen and Richard (1987) framework, we show that an ESG-efficient mean–variance frontier naturally emerges. This ESG mean–variance frontier, with its associated four-fund separation and two-beta representation, simplifies the SDF construction. In particular, it is sufficient to focus on two linearly independent portfolios on the ESG mean–variance frontier to span the relevant SDF space. These two portfolios, which represent optimal combinations of financial and ESG characteristics, reduce the dimensionality of the SDF representation, offering both theoretical parsimony and empirical tractability relative to using the full asset space.

To provide more details, our four-fund separation theorem provides a decomposition of the asset space into four orthogonal components: the basic return, the ESG-adjusted return, the excess return, and a residual component with zero expected return. This theorem generalizes and unifies several existing "four-fund separation" results, encompassing both traditional financial and ESG-adjusted perspectives within a single framework. By characterizing the asset space in this way, the theorem offers a versatile foundation for constructing portfolios that balance risk, return, and sustainability considerations, enabling a coherent integration of ESG factors into financial decision-making.

Further extending this framework, we derive the two-beta representation theorem, which

allows any asset return to be expressed as a linear combination of two representative returns. One typical case is that one corresponding to the traditional market index and the other corresponding to ESG exposure. This theorem provides a simplified and effective approach for asset pricing that integrates both financial and ESG dimensions, offering an empirical framework for analyzing how these factors influence asset pricing.

We empirically validate the two-beta representation using U.S. equity data and demonstrate that it significantly improves pricing performance relative to both the traditional CAPM and the two-factor model proposed by Pástor et al. (2021) (PST) in ESG-sorted portfolios.

To conduct the analysis, we construct testing portfolios sorted by market capitalization and ESG beta, where ESG beta is defined as the regression coefficient of an asset's excess return on the MSCI USA ESG Screened Index. The results indicate that the CAPM fails to account for cross-sectional variation in realized returns, and although the PST model offers some improvement by introducing an ESG-sorted factor, it still leaves substantial pricing errors. By contrast, augmenting the CAPM with the maximum Sharpe ratio portfolio (MSP) beta yields markedly better explanatory power. Moreover, Generalized Method of Moments (GMM) estimations confirm that the proposed two-beta model consistently outperforms both CAPM and PST, as evidenced by smaller pricing errors and superior model fit across multiple test portfolios.

Our framework surpasses prior approaches in several dimensions. Unlike Pedersen et al. (2021) and Pástor et al. (2021), we do not impose explicit investor utility functions, allowing for heterogeneous ESG preferences without constraining portfolio selection or equilibrium conditions. Moreover, we introduce a broader and more flexible fund selection mechanism for the four-fund separation, providing both sufficient and necessary conditions for selecting representative funds. This greater flexibility enables a better decomposition of ESG-adjusted returns, enhancing empirical pricing accuracy. Additionally, our model extends separation theorems to all portfolios, including long-short portfolios, addressing a critical limitation in previous ESG frameworks. Together, our model offers a well-defined and practical approach for empirical implementation, addressing key challenges that previous models have faced in incorporating ESG factors into asset pricing.

Our work contributes to several strands of literature. First, we build upon the foundational work of Hansen and Richard (1987), who investigated the role of conditioning information in dynamic asset pricing models. They demonstrated that incorporating conditioning information is crucial for deriving testable restrictions and understanding the mean-variance frontier of asset returns. Our research extends this framework by integrating ESG factors as

additional conditioning information, thereby enhancing the asset pricing models to account for ESG considerations.

Our research also contributes to the literature on ESG integration in asset pricing models. While traditional models primarily focus on financial metrics, recent studies have emphasized the importance of incorporating ESG factors to better capture risks and opportunities associated with sustainability. Recent studies have begun to address this gap by incorporating ESG factors into asset pricing frameworks. For instance, Pedersen et al. (2021) and Pástor et al. (2021) have developed multifactor models that include ESG attributes, highlighting the impact of investor preferences for sustainability on asset prices.Avramov et al. (2024) develops a dynamic equilibrium model that integrates ESG demand and supply . Our research builds upon these advancements by introducing a more generalized four-fund separation theorem that integrates ESG factors, providing a comprehensive framework for understanding how ESG considerations influence asset pricing.

Besides, the relationship between ESG factors and stock performance has been extensively studied, yielding mixed results. Some research indicates that firms with strong ESG profiles may experience lower returns due to investor preferences for sustainable investments, which can drive up prices and compress future yields. For instance, Hong and Kacperczyk (2009) found that "sin" stocks—companies involved in alcohol, tobacco, and gaming—often outperform their peers, suggesting that social norms may lead to undervaluation of these firms. Baker et al. (2022) and Zerbib (2019) report that green bonds are issued at a premium and offer lower yields compared to traditional bonds, driven by investors' environmental preferences. Hsu et al. (2023) propose and model a new systematic risk related to environmental policy uncertainty can leads to a premium for high carbon emission firms. Chava (2014) and El Ghoul et al. (2011) further suggest that greener firms face a lower implied cost of capital, consistent with the idea that ESG preferences reduce required returns for green firms. Conversely, other studies suggest that robust ESG practices can enhance financial performance. Edmans (2011) demonstrated that companies with high employee satisfaction, a component of the social aspect of ESG, tend to achieve superior long-term returns. Similarly, Gompers et al. (2003) found that firms with strong governance structures often outperform those with weaker governance.Kempf and Osthoff (2007) found that stocks with high ESG ratings outperformed their lower-rated counterparts during the 1992–2004 period, likely due to strengthened ESG preferences among investors. Our research contributes to this discourse by introducing the concept of "ESG beta," which measures a stock's sensitivity to ESG factors. We provide empirical evidence that this ESG beta is a significant determinant of stock performance, offering a nuanced perspective that reconciles the varying findings in existing literature.

Finally, we engage with the economic theories of discrimination, notably taste-based discrimination as introduced by Becker (2010), Fama and French (2007) and statistical discrimination as discussed by Phelps (1972). Becker's model posits that employers' prejudices lead to discriminatory practices, while Phelps' theory suggests that discrimination arises from employers using group averages to make decisions under uncertainty. Our study draws parallels between these theories and the financial markets, examining how ESG factors may influence investment decisions and potentially lead to discriminatory practices in portfolio management.

The remainder of the paper is organized as follows. Section 2 discusses the theoretical foundation of the ESG-adjusted pricing framework, including the definition and implications of ESG-adjusted price. Section 3 presents the no arbitrage condition and establishes LOOP and LOOE as the core assumptions of our model. Section 4 develops the ESG-adjusted Stochastic Discount Factor and discusses its use in pricing both financial and ESG-adjusted returns. Section 5 explores the ESG-efficient frontier, presenting the four-fund separation and showing how these results extend classical portfolio theory to incorporate ESG factors. Section 6 compares our framework with existing literature, highlighting key improvements. Section 7 introduces the two-beta representation theorem, demonstrating its empirical relevance. Section 8 presents the empirical tests, validating our model's performance. Finally, Section 9 concludes the paper.

2 Assets with ESG

Consider a financial market with N basis assets, where each asset $n \in \{1, \ldots, N\}$ delivers a state-contingent payoff $x_{n,s}$ in state $s = 1, 2, \ldots, S$. Let $X_{N\times S} \in \mathbb{R}^{N\times S}$ denote the matrix of asset payoffs, where each row $x_n \in \mathbb{R}^S$ represents the payoffs of basis asset n across all states. The vector of asset prices is $P \in \mathbb{R}^N$, with p_n denoting the price of asset n.

Let $\theta \in \mathbb{R}^N$ denote a portfolio vector, where θ_n is the holding (in units) of basis asset n. Throughout, we adopt the convention that subscripts refer to vector components (e.g., θ_n is the n-th entry of θ), while superscripts index distinct portfolios (e.g., θ^i and θ^j denote two different portfolios). For each basis asset n, let $\theta^{\text{basis},n} \in \mathbb{R}^N$ denote the portfolio vector with a 1 in the n-th position and 0 elsewhere. That is,

$$\theta_k^{\text{basis},n} = \begin{cases} 1 & \text{if } k = n, \\ 0 & \text{otherwise.} \end{cases}$$

We define the *pricing function* $q : \mathbb{R}^N \to \mathbb{R}$ to represent the market value of any portfolio. This function satisfies two economically natural properties: linearity in portfolio weights and consistency with observed basis asset prices. Formally, we can restate these two properties as the defining features of the pricing functional.

Definition 2.1 (Pricing Functional)**.** *The pricing function $q : \mathbb{R}^N \to \mathbb{R}$ satisfies:*

1. ***Linearity:*** *$q(w_1\theta^i + w_2\theta^j) = w_1 q(\theta^i) + w_2 q(\theta^j)$ for all portfolios $\theta^i, \theta^j \in \mathbb{R}^N$ and scalars $w_1, w_2 \in \mathbb{R}$;*
2. ***Basis consistency:*** *$q(\theta^{basis,n}) = p_n$ for all $n = 1, \ldots, N$.*

Under the definition, the pricing function admits a closed-form representation: for any portfolio $\theta \in \mathbb{R}^N$, its market value is given by

$$q(\theta) = \theta' P = \sum_{n=1}^{N} \theta_n p_n,$$

where $P \in \mathbb{R}^N$ is the vector of basis asset prices. This characterization follows directly from linearity and the requirement that each basis asset be priced correctly. It provides the foundation upon which we introduce ESG considerations in portfolio valuation.

With the growing prominence of sustainable investing, portfolio choice increasingly incorporates ESG attributes. Let $E \in \mathbb{R}^N$ denote the vector of ESG scores, where e_n is the ESG score assigned to basis asset $n \in \{1, \ldots, N\}$ by a rating agency or derived from a specified methodology. These scores are intended to summarize each asset's non-financial impacts across ESG dimensions. To evaluate the ESG characteristics of portfolios, we use $s^e(\cdot)$ to denote the *ESG score function*, with $s^e(\theta)$ representing the ESG score assigned to portfolio $\theta \in \mathbb{R}^N$.

In practice, a common approach—adopted widely in both academic studies and industry practice—is to define $s^e(\theta)$ as the value-weighted average ESG score of the portfolio's constituents. That is, for any portfolios $\theta^i, \theta^j \in \mathbb{R}^N$ and scalars $w_1, w_2 \in \mathbb{R}$ such that $q(w_1\theta^i + w_2\theta^j) \neq 0$, the ESG score function satisfies the property:

$$s^e(w_1\theta^i + w_2\theta^j) = \frac{w_1 q(\theta^i) s^e(\theta^i) + w_2 q(\theta^j) s^e(\theta^j)}{w_1 q(\theta^i) + w_2 q(\theta^j)}.$$

This value-weighted formulation ensures that the ESG attributes of a portfolio are aggregated proportionally to the capital allocated across assets.

However, it suffers from a critical shortcoming: for any long-short or dollar-neutral portfolio with $q(\theta) = 0$, the ESG score is undefined because of the zero denominator. This

issue is not merely technical; it precludes ESG evaluation for a broad class of economically meaningful strategies – including market-neutral, factor-based, and arbitrage portfolios – as emphasized by Pedersen et al. (2021).

To address this limitation, we introduce a second valuation operator that mirrors the standard pricing function, but instead of being consistent with observed market prices it adheres to a natural ESG-weighted valuation rule. Specifically, this functional evaluates each asset not by its price p_n alone, but by the product $p_n \cdot e_n$, which intuitively captures both financial value and ESG quality. We formalize this by defining:

Definition 2.2 (ESG-adjusted Pricing Functional)**.** *The ESG-adjusted pricing function* $q^e : \mathbb{R}^N \to \mathbb{R}$ *satisfies:*

1. ***Linearity:*** $q^e(w_1\theta^i + w_2\theta^j) = w_1 q^e(\theta^i) + w_2 q^e(\theta^j)$ *for all portfolios* $\theta^i, \theta^j \in \mathbb{R}^N$ *and scalars* $w_1, w_2 \in \mathbb{R}$*;*
2. ***Basis consistency:*** $q^e(\theta^{basis,n}) = p_n^e \equiv p_n e_n$ *for all* $n = 1, \ldots, N$*.*

That is, just as the pricing function $q(\cdot)$ maps portfolios to market value by linear extension from individual prices p_n, the ESG-adjusted pricing function $q^e(\cdot)$ maps portfolios to a scalar value by linearly extending from ESG-weighted basis asset values $p_n^e \equiv p_n e_n$. This yields the following closed-form representation:

$$q^e(\theta) = \sum_{n=1}^{N} \theta_n p_n^e = \theta' P^e,$$

where $P^e = (p_1^e, \ldots, p_N^e)'$ is the vector of ESG-adjusted basis prices.

Moreover, for any portfolio with positive market value, this functional satisfies the identity:

$$q^e(\theta) = q(\theta) \cdot s^e(\theta), \quad \text{if } q(\theta) \neq 0.$$

This establishes a one-to-one correspondence between the ESG score function and the ESG-adjusted pricing function over the domain $\{\theta : q(\theta) \neq 0\}$, while also highlighting that $q^e(\cdot)$ remains well defined for all θ, including long-short and market-neutral portfolios where $q(\theta) = 0$ but ESG content remains economically meaningful.

In economic terms, the ESG-adjusted price provides a coherent and continuous measure of a portfolio's net ESG exposure. For portfolios with long exposure to ESG-positive assets, $q^e(\theta) > 0$ indicates a net allocation to sustainable capital. For portfolios short ESG assets or long ESG-adverse positions, $q^e(\theta) < 0$, capturing the extent and direction of ESG risk

assumed. Crucially, this construction preserves the linear structure of portfolio aggregation and the heterogeneity of ESG signals across assets.

Having established these extended pricing concepts, we now introduce key asset subspaces that are foundational to the model-free ESG pricing framework. Let $\mathcal{Y}$ denote the set of all attainable payoffs,

$$\mathcal{Y} = \{\theta' X : \theta \in \mathbb{R}^N\}.$$

Define the *return space* as the set of payoffs that can be constructed at unit cost,

$$\mathcal{R} = \{\theta' X : q(\theta) = 1,\ \theta \in \mathbb{R}^N\}.$$

The *excess return space* is the set of payoffs with zero net market value,

$$\mathcal{R}^e = \{\theta' X : q(\theta) = 0,\ \theta \in \mathbb{R}^N\},$$

and the *excess ESG-adjusted return space* is the set of payoffs with zero net ESG-adjusted price,

$$\mathcal{R}^{ee} = \{\theta' X : q^e(\theta) = 0,\ \theta \in \mathbb{R}^N\}.$$

These spaces serve as the mathematical foundation for decomposing payoffs and returns along both financial and ESG dimensions, providing the essential building blocks for the analysis of pricing, portfolio construction, and factor representation in the presence of both pecuniary and non-pecuniary considerations.

3 No Arbitrage Condition

Definition 2.1 and 2.2 provide necessary linearity constraints for pricing and ESG-adjusted pricing of portfolios. However, these axioms alone do not guarantee economic consistency. In particular, they do not restrict the relative valuation of portfolios that deliver identical payoffs, leaving open the possibility of arbitrage. To establish a robust and internally coherent pricing framework—both in financial and ESG dimensions—we introduce a set of foundational no-arbitrage conditions.

The most fundamental principle is the *Law of One Price* (LOOP), which asserts that any two portfolios delivering the same state-contingent payoffs must be assigned the same market price.

Assumption 3.1 (Law of One Price (LOOP))**.** *All portfolios with the same payoff must have*

the same price. That is,

$$\text{if } (\theta^i)'X = (\theta^j)'X, \text{ then } q(\theta^i) = q(\theta^j)$$

for any portfolios $\theta^i, \theta^j \in \mathbb{R}^N$.

LOOP is a foundational axiom of financial economics, underlying classic results from arbitrage pricing theory (Ross, 1976) to the construction of risk-neutral valuation. It ensures that arbitrage opportunities—where investors could generate costless profits by exploiting price discrepancies between portfolios with identical payoffs—cannot persist in equilibrium. Moreover, LOOP implies that pricing is ultimately a function of payoffs, not merely of portfolio weights or structure.

As ESG investing becomes increasingly prominent, it is natural to seek a parallel consistency condition for the non-financial attributes of portfolios. In both academic literature and investment practice, ESG scores are used to summarize the sustainability profile of a portfolio. To ensure logical coherence in ESG attribution, we introduce the *Law of One ESG Score*, which requires that portfolios with the same economic exposure receive the same ESG evaluation.

Assumption 3.2 (Law of One ESG Score)**.** *All portfolios with the same payoff must have the same ESG score. That is,*

$$\text{if } (\theta^i)'X = (\theta^j)'X, \text{ then } s^e(\theta^i) = s^e(\theta^j)$$

for any portfolios $\theta^i, \theta^j \in \mathbb{R}^N$.

The rationale for the Law of One ESG Score closely parallels that of LOOP. When two portfolios deliver identical cash-flow distributions across all states, any ESG scoring methodology that reflects true economic exposure—rather than superficial asset labels or construction artifacts—should assign them the same ESG score. This requirement is consistent with industry practice: leading ESG data providers and institutional investors attribute ESG characteristics as functions of portfolio exposure rather than based on nominal composition.

While the Law of One ESG Score offers a meaningful consistency condition for ESG scores, it is often more practical to work with the ESG-adjusted pricing function $q^e(\cdot)$, which aggregates financial and ESG information in a way that mirrors the standard pricing function $q(\cdot)$. To this end, we introduce a parallel no-arbitrage condition—the Law of One ESG-adjusted Price (LOOE)—which ensures that portfolios with the same payoff receive the same ESG-adjusted valuation.

Assumption 3.3 (Law of One ESG-adjusted Price (LOOE)). *All portfolios with the same payoff have the same ESG-adjusted price. That is,*

$$\text{if } (\theta^i)'X = (\theta^j)'X, \text{ then } q^e(\theta^i) = q^e(\theta^j)$$

for any two portfolios $\theta^i, \theta^j \in \mathbb{R}^N$.

When combined with LOOP, this condition is in fact equivalent to requiring consistency of ESG scores across payoff-equivalent portfolios, as formalized below:

Proposition 3.1 (Equivalence of LOOE and the Law of One ESG Score under LOOP). *Suppose Assumption 3.1 holds. Then, Assumption 3.3 is equivalent to the Assumption 3.2.*

This proposition implies that under the no-arbitrage principle of LOOP, ensuring consistency in ESG-adjusted prices is mathematically tantamount to requiring consistency in ESG scores across payoff-equivalent portfolios. Henceforth, we use only LOOP and LOOE to refer to our no-arbitrage assumptions, and we do not separately invoke the Law of One ESG Score principle in the remainder of the paper.

Taken together, LOOP and the LOOE impose consistency across financial and non-financial dimensions of asset valuation. While conceptually distinct, both are jointly necessary to rule out arbitrage. LOOP ensures pricing consistency across portfolios with identical payoffs, whereas LOOE enforces ESG attribution consistency. Violations of either principle open the door to arbitrage—either in price or in sustainability rating—for portfolios with the same economic substance. Hence, both are foundational for a coherent pricing framework, with robustness to shifts in investor preferences.[1]

We now characterize three sufficient conditions—related to asset structure and investor composition—under which LOOP and LOOE are guaranteed to hold.

Proposition 3.2 (Sufficient Conditions for LOOP and LOOE). *LOOP and LOOE hold simultaneously if any of the following conditions are satisfied:*

1. ***Non-redundancy:*** *The set of basis assets has full rank. That is, asset payoffs are linearly independent, so any attainable payoff corresponds to a unique portfolio. In this case,* $(\theta^i)'X = (\theta^j)'X$ *implies* $\theta^i = \theta^j$, *and hence* $q(\theta^i) = q(\theta^j)$ *and* $s^e(\theta^i) = s^e(\theta^j)$. *This*

[1] If investors care about only a specific dimension of ESG—say, the environmental (E) component—the same logic applies. The relevant no-arbitrage condition would then require consistency in ESG-adjusted prices with respect to the environmental score only. The LOOE principle remains valid as long as it reflects the specific ESG segment valued by investors.

condition is implicitly assumed in models such as Pedersen et al. (2021) and Pástor et al. (2021), which assume invertible covariance matrices.

2. ***Competitive equilibrium with heterogeneous investors:*** *The market includes both ESG-indifferent and ESG-sensitive investors.*[2] *If two portfolios with the same payoff differ in price, ESG-indifferent investors arbitrage the price gap, enforcing LOOP. If they differ in ESG score, ESG-sensitive investors strictly prefer the higher-scoring portfolio, leading to excess demand and ESG score convergence. These two investor types are often referred to as "value investors" and "values investors," respectively, in the terminology of Starks (2023). Thus, any misalignment violates optimality and cannot persist in equilibrium.*
3. ***ESG replication consistency:*** *If a linear combination of basis assets replicates the payoff of a portfolio, then it must also replicate its ESG score. This ensures that any replicating combination of basis assets inherits the ESG score of the payoff it replicates.*

These conditions clarify when LOOP and LOOE are upheld either mechanically (under non-redundancy), as equilibrium outcomes disciplined by heterogeneous investor incentives, or by construction through ESG assignment rules. Appendix A.1 provides a simple competitive-equilibrium model with ESG-indifferent and ESG-sensitive investors that implements condition (ii) and shows how their coexistence enforces both price and ESG-adjusted price consistency. Appendix A.2 embeds a green and a brown asset into a similar general-equilibrium environment and derives a closed-form greenium: stronger ESG preferences or a larger share of ESG-sensitive investors raise the relative price of the green asset, while LOOP and LOOE continue to hold.

When none of these conditions is satisfied, violations may arise—though often due to frictions or hidden risk exposures rather than genuine arbitrage. For example, two portfolios may appear payoff-equivalent under standard metrics but differ under tail risks or conditional states. Bonds from the same issuer but with different maturities or covenants may differ in default exposure; a firm selectively defaulting on "brown" bonds before "green" ones implies their state-contingent payoffs are not truly identical. In such cases, neither LOOP nor LOOE applies.

In sum, LOOP and LOOE are robust principles, holding either under structural assumptions or as equilibrium outcomes in the presence of heterogeneous investors. Apparent violations often reflect market frictions or misidentified payoff equivalence rather than true

[2] As long as both ESG-indifferent and ESG-sensitive investors are present in the market, the exact proportion of each type does not affect the validity of LOOP and LOOE. As demonstrated in Appendix A.1 (and extended in the greenium setting of Appendix A.2), the coexistence of heterogeneous preferences—not their specific distribution—is the key equilibrium driver.

breaches of no-arbitrage logic.

3.1 ESG-adjusted risk-neutral probability

In the classical setting, absence of arbitrage implies the existence of a risk-neutral probability measure $\mathbb{Q}$ under which all traded assets earn the risk-free rate in expectation, $E^{\mathbb{Q}}[R_i] = R_f$. The change of measure from the physical probabilities $\{\pi_s\}$ to $\mathbb{Q}$ is implemented by the SDF m^* via

$$\tilde{\pi}_s \equiv \frac{\pi_s m_s^*}{E[m^*]}, \qquad E^{\mathbb{Q}}[R_i] = \sum_s \tilde{\pi}_s R_{i,s} = R_f,$$

so states with low consumption and high marginal utility receive higher $\tilde{\pi}_s$.

In the ESG-adjusted setting, the pricing equation for an asset i with payoff vector $X_i = (X_{i,1}, \ldots, X_{i,S})$, price p_i, and ESG score e_i is

$$q_i^e \equiv e_i p_i = E[g^* X_i] = \sum_s \pi_s g_s^* X_{i,s},$$

where g^* is the ESG-adjusted SDF. Define the ESG-adjusted pseudo-risk-neutral probabilities by

$$\tilde{\pi}_s^e \equiv \frac{\pi_s g_s^*}{E[g^*]}.$$

Then

$$e_i p_i = E[g^* X_i] = E[g^*]\, E^{\tilde{\pi}^e}[X_i].$$

Using the gross return $R_{i,s} \equiv X_{i,s}/p_i$ and dividing by $e_i p_i$ yields

$$1 = E[g^*]\, E^{\tilde{\pi}^e}\left[\frac{R_i}{e_i}\right].$$

For the risk-free asset with payoff $X_f = R_f$, price $p_f = 1$, and ESG score e_f, the ESG-adjusted price is $q_f^e = e_f = E[g^* R_f]$, which implies $E[g^*] = e_f/R_f$. Substituting this normalization gives

$$E^{\tilde{\pi}^e}\left[\frac{R_i}{e_i}\right] = \frac{R_f}{e_f}.$$

This identity is the ESG analogue of the classical relation $E^{\mathbb{Q}}[R_i] = R_f$. Under the ESG-adjusted measure $\tilde{\pi}^e$, all assets earn the same "ESG-adjusted risk-free rate" once returns are scaled by their ESG scores. The relevant object of valuation is the return per unit ESG, R_i/e_i, rather than R_i alone.

This representation also admits an aggregate, state-level interpretation. Let $w_{M,i}$ denote

the weight of asset i in the market (or aggregate) portfolio, with state-contingent returns $R_{i,s}$ and ESG scores $e_i > 0$. A natural state-level indicator of aggregate "return per unit ESG" is

$$\Psi_s \equiv \sum_i w_{M,i} \frac{R_{i,s}}{e_i},$$

a portfolio-weighted average of asset-level return–ESG ratios in state s. States with low Ψ_s deliver little financial compensation per unit of ESG exposure at the market level, whereas states with high Ψ_s deliver relatively high compensation per unit ESG. To avoid confusion with the traditional terminology of financially good and bad states, we say that a state s is *ESG-critical* if Ψ_s is low relative to a benchmark (for example, the ESG-adjusted risk-free ratio R_f/e_f or the cross-state average of $\{\Psi_s\}_s$), and *ESG-benign* if Ψ_s is high relative to that benchmark.

The indicator Ψ_s can be decomposed into a purely financial component and an ESG-related cross-sectional component. Define the market return in state s by

$$R_{M,s} \equiv \sum_i w_{M,i} R_{i,s},$$

and the portfolio-weighted cross-sectional average of inverse ESG scores by

$$\bar{\phi} \equiv \sum_i w_{M,i} \frac{1}{e_i}.$$

Let $E_M[\cdot]$ and $\mathrm{cov}_M(\cdot,\cdot)$ denote, respectively, the cross-sectional mean and covariance under weights $\{w_{M,i}\}_i$. Then

$$\Psi_s = \sum_i w_{M,i} R_{i,s} \frac{1}{e_i} = E_M\big[R_{i,s}\big]\, E_M\Big[\frac{1}{e_i}\Big] + \mathrm{cov}_M\Big(R_{i,s}, \frac{1}{e_i}\Big),$$

that is,

$$\Psi_s = R_{M,s}\, \bar{\phi} + \mathrm{cov}_M\Big(R_{i,s}, \frac{1}{e_i}\Big).$$

Because $w_{M,i}$ and e_i are fixed at time 0, variation in Ψ_s across states is driven entirely by the pattern of returns $\{R_{i,s}\}_i$. A low value of Ψ_s (an ESG-critical state) can arise through two channels: (i) a low market return $R_{M,s}$, which coincides with the usual notion of a financially bad state, and/or (ii) a negative cross-sectional covariance $\mathrm{cov}_M(R_{i,s}, 1/e_i)$, which means that high-ESG assets (low $1/e_i$) deliver relatively high returns while low-ESG (brown) assets underperform in that state. The first channel captures the traditional risk dimension; the second captures an ESG-specific rotation in the cross-section, in which green assets

outperform brown assets precisely in states that are ESG-critical (for instance, following a severe climate-related disaster or an unexpected tightening of environmental regulation).

In this sense, the ESG-adjusted probabilities $\tilde{\pi}_s^e$ tilt the pricing weight toward states that are simultaneously financially severe (low $R_{M,s}$) and/or characterized by strong relative performance of green assets (negative $\mathrm{cov}_M(R_{i,s}, 1/e_i)$), and tilt away from states in which the market delivers high return per unit ESG. The ESG-adjusted SDF g^* therefore incorporates both financial risk and ESG preferences: it increases the valuation weight of states in which aggregate payoffs are low and in which green assets outperform brown assets, and decreases the weight of states that are financially favorable and ESG-benign.

3.2 Arrow–Debreu securities and state-level ESG scores

The state-level structure of g^* can be further clarified by considering Arrow–Debreu securities. Let $\{X^{(j)}\}_{j=1}^S$ denote a complete set of Arrow–Debreu payoffs, where the j-th Arrow–Debreu claim pays

$$X_s^{(j)} = \mathbf{1}\{s = j\},$$

i.e., one unit of consumption only in state j. Denote its price by p_j, its ESG score by e_j, and the components of m^* and g^* by m_j^* and g_j^*, respectively. In the classical setting,

$$p_j = E[m^* X^{(j)}] = \pi_j m_j^*,$$

so a high Arrow–Debreu price p_j identifies a financially bad state (high marginal utility), while a low p_j identifies a financially good state.

Under ESG valuation, the ESG-adjusted price of $X^{(j)}$ is

$$q_j^e = E[g^* X^{(j)}] = \pi_j g_j^*.$$

The LOOE relation $q_j^e = e_j p_j$ implies

$$\pi_j g_j^* = e_j \pi_j m_j^* \quad \Rightarrow \quad g_j^* = e_j m_j^*,$$

and hence

$$e_j = \frac{q_j^e}{p_j} = \frac{g_j^*}{m_j^*}.$$

Thus, in the Arrow–Debreu basis, the ESG score of the j-th Arrow–Debreu security is exactly the multiplicative distortion between the ESG-adjusted SDF and the traditional SDF in state

j.

This characterization separates the financial and ESG dimensions at the state level. The Arrow–Debreu prices $\{p_j\}$ describe how financially severe each state is: a high p_j corresponds to a financially bad state, a low p_j to a financially good state. The ESG scores $\{e_j\}$ describe how ESG considerations reweight those states relative to the risk-free benchmark. Using e_f as reference, we call a state j *ESG-critical* if $e_j > e_f$ and *ESG-benign* if $e_j < e_f$. In an ESG-critical state, the ESG-adjusted SDF g^* places more weight on payoffs than m^* would, so an extra unit of payoff is particularly valuable from an ESG perspective; in an ESG-benign state, the opposite holds.

These notions are consistent with observed return patterns for green and brown assets. Suppose there is a green asset G and a brown asset B with ESG scores $e_G > e_B$, and state-contingent payoffs $\{X_{G,s}\}_s$ and $\{X_{B,s}\}_s$. By linearity of the ESG-adjusted price functional with respect to Arrow–Debreu claims,

$$q^e_G = \sum_s w_{G,s} q^e_s = \sum_s w_{G,s} p_s e_s, \qquad q^e_B = \sum_s w_{B,s} q^e_s = \sum_s w_{B,s} p_s e_s,$$

where the weights $w_{i,s}$ reflect the contribution of state s to the present value of X_i. Consider a state j in which brown assets earn relatively high returns and green assets relatively low returns. Then $w_{B,j} \gg w_{G,j}$, so consistency with $e_G > e_B$ requires e_j to be relatively low; otherwise the strong exposure of B to state j would drive q^e_B up more than q^e_G. Such a state is ESG-benign. Conversely, in a state k in which green assets perform relatively well and brown assets relatively poorly, we must have e_k relatively high to maintain $e_G > e_B$; this state is ESG-critical.

In summary, the ESG score of an Arrow–Debreu security provides a state-level classification that is directly analogous to, but conceptually distinct from, the financial classification based on state prices. A state is ESG-critical if its Arrow–Debreu claim carries a high ESG score $e_j > e_f$ (or, equivalently, the aggregate return-per-ESG indicator Ψ_j is relatively low), and ESG-benign if $e_j < e_f$ (or Ψ_j is relatively high). Through the relation

$$\tilde{\pi}^e_s \propto \pi_s g^*_s = \pi_s m^*_s e_s,$$

the ESG-adjusted risk-neutral probabilities tilt mass toward ESG-critical states and away from ESG-benign states, just as classical risk-neutral probabilities tilt toward financially bad (high-p_s) states. Taken together, Arrow–Debreu prices $\{p_s\}$ and state-level ESG scores $\{e_s\}$ yield a two-dimensional description of each state in terms of financial risk and sustainability relevance.

4 ESG-adjusted Stochastic Discount Factor

Having established that a market governed by both LOOP and LOOE admits a unified and arbitrage-free valuation of both financial and ESG-adjusted prices, it is natural to inquire about the existence and structure of the pricing kernel that implements this logic. The fundamental insight of modern asset pricing is that the absence of arbitrage, formalized by LOOP, guarantees the existence of a stochastic discount factor (SDF) that consistently links prices and payoffs across all securities. Extending this logic, LOOE implies an analogous pricing kernel, called ESG-adjusted stochastic discount factor, which serves as the kernel for ESG-adjusted valuation. In this section, we formally establish the existence and uniqueness of these discount factors and discuss their economic interpretation.

If LOOP holds universally, it implies that a unique SDF exists, which equitably prices all payoffs, ensuring that market prices reflect expected discounted cash flows without allowing for arbitrage Cochrane (2009). The following proposition formalizes this connection between LOOP and the existence of a unique SDF that consistently prices all payoffs within a given space, demonstrating that LOOP inherently implies a stochastic discount factor.

Proposition 4.1 (Existence of SDF under LOOP)**.** *If LOOP holds, there exists a unique payoff $m^* \in \mathcal{Y}$, such that the price of any payoff is given by*

$$q(y) = E\left(m^* y\right)$$

for all $y \in \mathcal{Y}$.

The proof of Proposition 4.1 can be found in Hansen and Richard (1987). This result suggests that under LOOP, m^* serves as a unique pricing kernel, aligning payoffs and prices in a consistent manner through the expected value operator. The existence of such an SDF implies that all securities in the market are priced to prevent arbitrage, as any misalignment between prices and payoffs would violate LOOP. Thus, this proposition provides a formal link between LOOP and the structure of asset pricing, wherein the SDF encapsulates the risk-adjusted discounting necessary to ensure pricing consistency across all payoffs in the given space.

Building on this foundation, we extend the logic to ESG-adjusted valuation. The presence of LOOE guarantees the existence of a unique ESG-adjusted stochastic discount factor, which integrates both financial and ESG considerations into a single pricing kernel. The following proposition formalizes this claim.

Proposition 4.2 (Existence of ESG-Adjusted SDF under LOOE)**.** *If LOOE holds, there exists a unique payoff $g^* \in \mathcal{Y}$, such that the ESG-adjusted price of any payoff is given by*

$$q^e(y) = E(g^* y)$$

for all $y \in \mathcal{Y}$.

Proof. Let y_0 be a payoff with nonzero ESG-adjusted price and let z_0 be its conditional projection onto $\mathcal{R}^{ee}$. Then $q^e(y_0) = q^e(y_0 - z_0)$ since $q^e(z_0)$ is zero. Define the benchmark return r^{e*} by

$$r^{e*} = \frac{y_0 - z_0}{q^e(y_0)}.$$

Notice that r^{e*} is conditionally orthogonal to $\mathcal{R}^{ee}$ and $q^e(r^{e*}) = 1$. Given any y in $\mathcal{Y}$, $y - q^e(y)r^{e*}$ is in $\mathcal{R}^{ee}$ since

$$q^e(y - q^e(y)r^{e*}) = q^e(y) - q^e(y)q^e(r^{e*}) = 0.$$

Hence, $E[r^{e*}(y - q^e(y)r^{e*})] = 0$ or

$$E(r^{e*}y) = q^e(y)E\left[(r^{e*})^2\right].$$

We digress to show by contradiction that $\Pr\{E[(r^{e*})^2] = 0\} = 0$. Suppose $E[(r^{e*})^2] = 0$ on a set A where $\Pr(A) > 0$. Hence on $A, r^{e*} = 0$. Define a zero payoff

$$z = \begin{cases} r^{e*} & \text{on } A, \\ 0 & \text{otherwise.} \end{cases}$$

By Definition 2.2, $q^e(z) = 0$, but $q^e(z) = q^e(r^{e*}) = 1$ on A, establishing the contradiction.

Returning to the main proof, divide the previous equation by $E[(r^{e*})^2]$ and define

$$g^* = \frac{r^{e*}}{E[(r^{e*})^2]}$$

to obtain

$$q^e(y) = E(g^* y).$$

The element g^* is clearly unique since if g' is any element of $\mathcal{Y}$ for which $q^e(y) = E(g'y)$ for all y, we have $E(g^{*2}) = E(g^* g') = E(g'^2)$ so that $E[(g^* - g')^2] = 0$.

Since $E(g^{*2}) = \frac{1}{E[(r^{e*})^2]}$, it follows that $\Pr\{E(g^{*2}) > 0\} = 1$. □

The existence and uniqueness of g^* under LOOE establish g^* as the ESG-adjusted stochastic discount factor (ESDF), or ESG-sensitive pricing kernel, in this framework. Just as the standard SDF m^* encapsulates all risk adjustments necessary for financial pricing under LOOP, the ESDF g^* incorporates both risk and ESG considerations, assigning appropriate weights to states of the world in the evaluation of ESG-adjusted prices.

Economically, the stochastic discount factor m^* serves as the lens through which risk-averse investors evaluate uncertain future payoffs: it places higher weight on states that are more unpleasant or less likely, ensuring that asset prices reflect not only expected returns but also risk exposures. The SDF thus ensures that assets providing payoffs in adverse conditions are valued more highly, in accordance with standard asset pricing theory.

The ESG-adjusted stochastic discount factor g^* plays an analogous but enriched role. It adjusts the valuation of future payoffs not only for risk but also for ESG characteristics, placing additional weight on states deemed "ESG-unpleasant"—that is, states where sustainability factors impose higher costs or lower benefits. In practice, g^* can be interpreted as capturing both the market's risk aversion and its ESG preferences. For example, if certain future scenarios are characterized by heightened regulatory scrutiny or negative ESG outcomes, the ESDF will discount payoffs in those states more heavily, much as the traditional SDF discounts outcomes associated with high risk or low marginal utility.

While both the SDF and ESDF serve as pricing kernels mapping future payoffs to present value, the key distinction lies in their domain of adjustment. The SDF (m^*) reflects pure financial risk, whereas the ESDF (g^*) internalizes both risk and the evolution of ESG attributes across states and over time. Mathematically, while m^* and g^* may take similar forms as state-dependent weights in expected value operators, g^* generally embeds a richer structure, as it accounts for the joint dynamics of risk factors and ESG considerations. In this sense, the ESDF extends the classical framework by systematically pricing sustainability, ensuring that both pecuniary and non-pecuniary dimensions are harmonized in the valuation process. The ESDF thereby provides an essential foundation for integrating ESG into asset pricing models, enabling consistent, arbitrage-free valuation across the expanding frontier of sustainable finance.

5 The ESG Efficient Frontier

Building on the framework established by the concept of an ESG-adjusted stochastic discount factor, we now explore the implications of ESG considerations within a portfolio optimization setting. As investors increasingly incorporate ESG factors into their decision-making,

traditional mean-variance optimization approaches must adapt to account for these additional dimensions of value. Specifically, we consider an optimization problem where investors seek to minimize portfolio risk while achieving both a target expected return and a desired level of ESG performance. This introduces the concept of an ESG-efficient frontier, which extends the traditional efficient frontier by embedding ESG criteria alongside financial risk and return.

Problem 5.1 (Mean Variance Optimization Conditional on ESG Score)**.** *Under the ESG scheme, investors face the following problem*

$$\begin{aligned} &\min \sigma^2(r),\ r \in \mathcal{R} \\ &s.t.\ E(r) = \mu \\ &\quad ESG(r) = e \end{aligned}$$

This setup incorporates the traditional objectives of risk and return while conditioning on a specific ESG score. The introduction of ESG constraints implies that investors are not solely focused on financial performance but are also guided by non-financial considerations that align with their sustainability preferences. The result is an ESG-efficient portfolio that optimally balances return, risk, and ESG impact.

To further formalize the relationship between ESG performance and asset pricing, we propose an equivalent formulation of the ESG-efficient optimization problem. The following proposition states that the optimization problem above is equivalent to a similar mean-variance optimization where the ESG score constraint is replaced by an ESG-adjusted price condition, $q^e(r) = e$.

Proposition 5.1 (Equivalent Formulation of ESG-Constrained Mean-Variance Problem)**.** *Problem 5.1 is equivalent to:*

$$\begin{aligned} &\min \sigma^2(r),\ r \in \mathcal{R} \\ &s.t.\ E(r) = \mu \\ &\quad q^e(r) = e \end{aligned}$$

and the ESG Mean Variance Frontier is equivalent to:

$$\{r \in R : r = argmin\ \sigma^2(r)\ s.t.\ E(r) = \mu\ and\ q^e(r) = e\ for\ some\ \mu\ and\ e\}$$

This equivalence suggests that, under the ESG pricing scheme, targeting a specific ESG

score is analogous to selecting a payoff with a specific ESG-adjusted price. Thus, by constraining the portfolio's ESG-adjusted price rather than its score, investors effectively achieve the same risk-adjusted ESG outcome.

Building on our previous definition of the ESG-adjusted mean-variance optimization problem, we now turn to a fundamental result in portfolio theory: the four-fund separation theorem. This theorem provides a structural representation of the asset space $\mathcal{R}$ by decomposing it into distinct, orthogonal components that align with both traditional financial objectives and ESG considerations. By constructing portfolios across these components, investors can achieve a balanced allocation that integrates risk, return, and ESG performance in a coherent manner.

Given the pricing function is nontrivial, we assume that $\Pr\{q(m^*) = 0\} = 0$. This assumption allows us to define a well-behaved return in the payoff space $\mathcal{Y}$ as:

$$r^* = \frac{m^*}{q(m^*)}$$

Under this setup, the payoff space $\mathcal{Y}$ can be expressed as:

$$\mathcal{Y} = \{r : r = r^* + z \text{ for some } z \in \mathcal{R}^e\}$$

where $\mathcal{R}^e$ represents the excess return space, containing payoffs with a price of zero. This decomposition allows us to separate $\mathcal{Y}$ into components aligned with the basic return r^* and excess returns within $\mathcal{R}^e$.

The space $\mathcal{R}^e$ is a conditionally complete linear subspace within $\mathcal{Y}$, which means that it spans all payoffs with zero price. Furthermore, we use the conditional expectation operator $E(\cdot)$ to define a continuous linear functional on $\mathcal{R}^e$, a tool that will enable further decomposition of the asset space.

Next, to incorporate ESG considerations into this framework, we impose the following assumption:

Assumption 5.1. *There exists a payoff r_g in $\mathcal{R}^e$ for which $q^e(r_g) \neq 0$.*

With Assumption 5.1, Proposition 4.2 ensures the existence of a unique payoff $r_g^* \in \mathcal{R}^e$, such that the ESG-adjusted price $q^e(z)$ for any $z \in \mathcal{R}^e$ can be represented as $q^e(z) = E\left(zr_g^*\right)$. This result introduces r_g^* as a pivotal component in $\mathcal{R}^e$, defining one conditional dimension within the space, with all other dimensions orthogonal to r_g^* having a fixed ESG score. This leads to the following lemma:

Lemma 5.1. *$\mathcal{R}^e$ can be represented as $\mathcal{R}^e = \{z : z = wr_g^* + k$ for some w in I and some k in $K\}$ where K is defined as $K = \mathcal{R}^e \cap \mathcal{R}^{ee} = \{z$ in $\mathcal{R}^e : q^e(z) = 0\}$.*

In this representation, the space $\mathcal{R}^e$ is decomposed into a component aligned with the ESG-adjusted return r_g^* and a residual space K, where all payoffs have a zero ESG-adjusted price.

Combining Lemma 5.1 with our initial decomposition of $\mathcal{Y}$ yields the following representation of $\mathcal{R}$:

$$\mathcal{R} = \left\{r : r = r^* + wr_g^* + k \text{ for some } w \in I \text{ and } k \in K\right\}$$

To complete the decomposition, we introduce an additional assumption for analyzing the structure of K :

Assumption 5.2. *There exists a payoff k_0 in K for which $E(k_0) \neq 0$.*

Under this assumption, we can apply Theorem 4.2 to derive the following lemma:

Lemma 5.2. *K can be represented as $K = \{k : k = vr_e^* + n$ for some v in I and some n in $N\}$ where N is defined as $N = \{k$ in $K : E(k) = 0\} = \{z$ in $\mathcal{Y} : q(z) = 0, q^e(z) = 0, E(z) = 0\}$.*

By combining Lemma 5.1 and Lemma 5.2, we arrive at a comprehensive representation of the asset space, leading to the four-fund separation theorem.

Theorem 5.1 (Four-Fund Separation Theorem)**.** *$\mathcal{R}$ can be represented as*

$$\mathcal{R} = \{r : r = r^* + wr_g^* + vr_e^* + n \text{ for some } w, v \in \mathbb{R} \text{ and } n\}$$

where

- *r^*, r_g^*, r_e^*, n are pairwise orthogonal;*
- *$r_g^*, r_e^*, n \in \mathcal{R}^e$;*
- *$r_e^*, n \in \mathcal{R}^{ee}$;*
- *n provides zero expected return.*

This theorem establishes that any return $r \in \mathcal{R}$ can be constructed as a linear combination of four distinct components: the basic return r^*, the ESG-adjusted return r_g^*, the excess return r_e^*, and the residual component n with zero expected return. Each of these components represents an independent source of value, and their orthogonality ensures that they are uncorrelated with one another. This four-fund decomposition offers investors a robust

framework for portfolio construction, enabling them to allocate resources across financial and ESG dimensions according to their specific objectives, while respecting the intrinsic structure of the asset space.

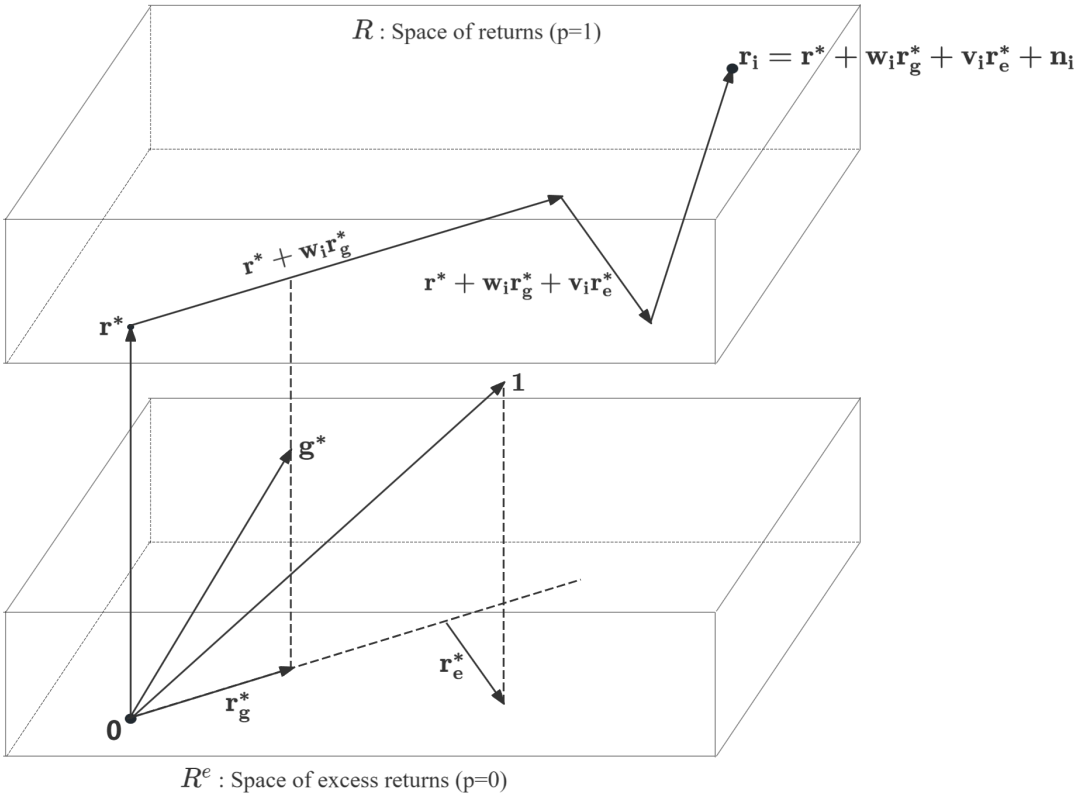


Figure 1: Orthogonal Decomposition and Four-Fund Separation. This figure illustrates the decomposition of returns, where r_i can be decomposed into four orthogonal parts: r^*,$w_i r_g^*$,$v_i r_e^*$ and n_i. r^* is the return that can act as the stochastic discount factor, and is orthogonal to $\mathcal{R}$ and $\mathcal{R}^e$ (because it is a 3D graph, the orthogonality is drawn approximately). r_g^* is the projection of g^* to $\mathcal{R}^e$ space. r_e^* is the projection of unit vector $\mathbf{1}$ to $\mathcal{R}^{ee}$ space, and n_i is the remaining part.

Following the four-fund separation theorem, which offers a structural decomposition of the asset space $\mathcal{R}$ into four orthogonal components, we turn our attention to several propositions that refine this framework within the context of the ESG-efficient frontier. These propositions address the unique structure of ESG-efficient portfolios and reveal conditions under which portfolios that meet both financial and ESG objectives can be effectively constructed with reduced dimensionality. Specifically, we explore a "three fund" separation result that applies uniquely to the ESG-efficient frontier, demonstrating how this streamlined structure aids in constructing portfolios with integrated ESG considerations.

The first proposition establishes that the ESG Efficient Frontier, representing portfolios that optimally balance risk, return, and ESG performance, can be simplified by using three fundamental components. Formally, the ESG Efficient Frontier can be represented as follows.

Proposition 5.2 (Three-fund Separation for ESG Frontier)**.** *The ESG Efficient Frontier can be represented as*

$$\textit{ESG-EF} = \{r : r = r^* + wr_g^* + vr_e^* \textit{ for some } w, v \in \mathbb{R}\}$$

where r^*, r_g^*, r_e^* *are defined as previous theorem.*

This result implies that, unlike the general asset space $\mathcal{R}$, which requires four components to represent all possible payoffs, only three distinct funds are needed to span the ESG Efficient Frontier: the basic return r^*, the ESG-adjusted return r_g^*, and the excess return r_e^*. Notably, the residual component n (which has zero expected return) is unnecessary for constructing ESG-efficient portfolios, as only the components contributing to financial or ESG-adjusted pricing are relevant for efficiency on the frontier. This three-fund structure therefore reduces the complexity of portfolio construction on the ESG frontier, enabling investors to achieve any point on the ESG Efficient Frontier by allocating resources across just these three components.

Proposition 5.3 (Convexity of the ESG-Efficient Frontier)**.** *Any combination of ESG efficient portfolios is an ESG efficient portfolios.*

Building upon this structure, the second proposition asserts that any combination of ESG-efficient portfolios remains an ESG-efficient portfolio. This result highlights the convex nature of the ESG Efficient Frontier: any weighted average of two or more portfolios on this frontier will also be ESG efficient. Consequently, this property provides a high degree of flexibility for investors, allowing them to construct diversified portfolios within the ESG efficient space while preserving optimality. By blending multiple ESG-efficient portfolios, investors can tailor their portfolio choices to align with specific ESG and financial objectives, assured that such combinations will remain on the frontier.

The third proposition further emphasizes the robustness of the ESG Efficient Frontier by demonstrating that it can be spanned by any three linearly independent ESG-efficient portfolios.

Proposition 5.4 (Three-Portfolio Spanning of the ESG-Efficient Frontier)**.** *ESG efficient frontier can be spanned by any three ESG efficient portfolio as long as they cannot construct each other.*

This spanning property implies that a well-chosen set of three ESG-efficient portfolios allows investors to replicate any desired portfolio on the ESG Efficient Frontier, simplifying portfolio construction within the multi-dimensional ESG framework. By selecting a diverse trio of portfolios, investors can effectively explore the full range of efficient options on the frontier without the need for additional components.

Together, these propositions reveal important insights into the structure and flexibility of the ESG Efficient Frontier. The three fund separation result shows that any ESG-efficient portfolio can be constructed from three primary components, while the convexity property allows for flexibility in portfolio design through combinations of ESG-efficient portfolios. Finally, the spanning property ensures that a carefully selected trio of ESG-efficient portfolios can represent the entire frontier, greatly simplifying the practicalities of multi-asset portfolio management with ESG integration. These properties underscore the coherence and utility of the ESG-efficient framework, allowing investors to achieve their desired balance of risk, return, and ESG considerations with a streamlined and robust approach.

5.1 Geometric Slices of the ESG-Efficient Frontier

To illustrate the geometry of the ESG-efficient frontier, we use simulated data rather than empirical returns. We generate 10 risky assets and a risk-free asset with mean and covariance parameters loosely calibrated to U.S. equities. ESG scores for the risky assets are drawn uniformly from $[-5, 5]$, while the risk-free asset is assigned an ESG score of zero. Using these inputs, we construct the ESG-efficient frontier and plot two-dimensional slices along the expected return and ESG dimensions.

Figure 2 reports slices of the frontier at fixed expected return levels of 5, 10, and 15. Each curve traces the minimum variance attainable across ESG levels for a given return target. The slices are hyperbolic with parallel asymptotes, reflecting the underlying mean–variance structure. As the target return increases, both the minimum achievable standard deviation and the ESG level at which it is attained shift upward, indicating that higher-return portfolios on the ESG-efficient frontier tend to involve higher ESG exposure.

Figure 3 provides the complementary view by slicing the frontier at fixed ESG levels of -10, -5, 0, 5, and 10. The ESG $= 0$ slice is approximately linear, corresponding to the classical mean–variance frontier without ESG tilts, whereas non-zero ESG targets yield hyperbolic slices. The ESG $= 5$ slice and its mirror at ESG $= -5$ form a symmetric pair; as the ESG level becomes more positive, both the minimum variance and the associated expected return increase. Taken together, these slices highlight that the ESG-efficient frontier

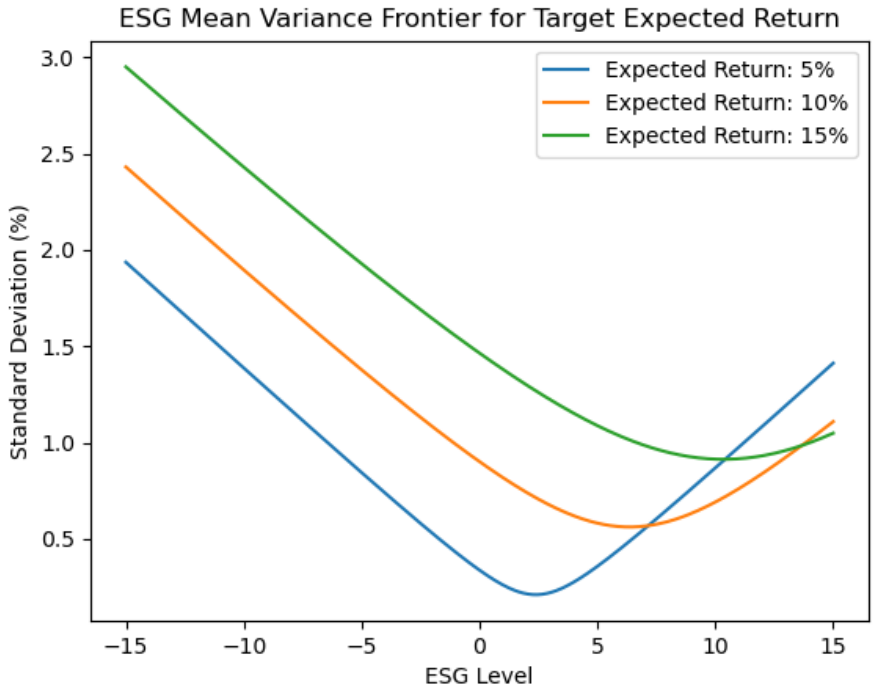


Figure 2: Simulated slices of the ESG-efficient frontier at fixed expected return levels (5, 10, and 15). The frontier is constructed from 10 simulated risky assets and a risk-free asset with risk and return parameters calibrated to U.S. equities. Risky-asset ESG scores are drawn from $[-5, 5]$ and the risk-free ESG score is set to zero. Each curve plots the minimum variance attainable across ESG levels for a given return target.

is a three-dimensional surface in return–risk–ESG space and that tightening ESG objectives bends the standard mean–variance frontier in a systematic and tractable manner.

5.2 Tangency Portfolio and Maximum Sharpe Ratio Portfolio

In traditional frameworks such as the Capital Asset Pricing Model (CAPM), the tangency portfolio formed by combining the risk-free asset with the market portfolio is also the portfolio that achieves the maximum Sharpe ratio. This equivalence arises because the introduction of a risk-free asset transforms the efficient frontier of risky assets from a hyperbola into a straight line, where the tangency point represents both the optimal portfolio and the maximum Sharpe ratio portfolio. However, in our ESG-integrated framework, this relationship no longer holds due to the distinct structure of the ESG-efficient frontier.

In our model, the ESG-efficient frontier retains a hyperbolic shape even after introducing the risk-free asset. This is because ESG considerations add a new dimension to the efficient

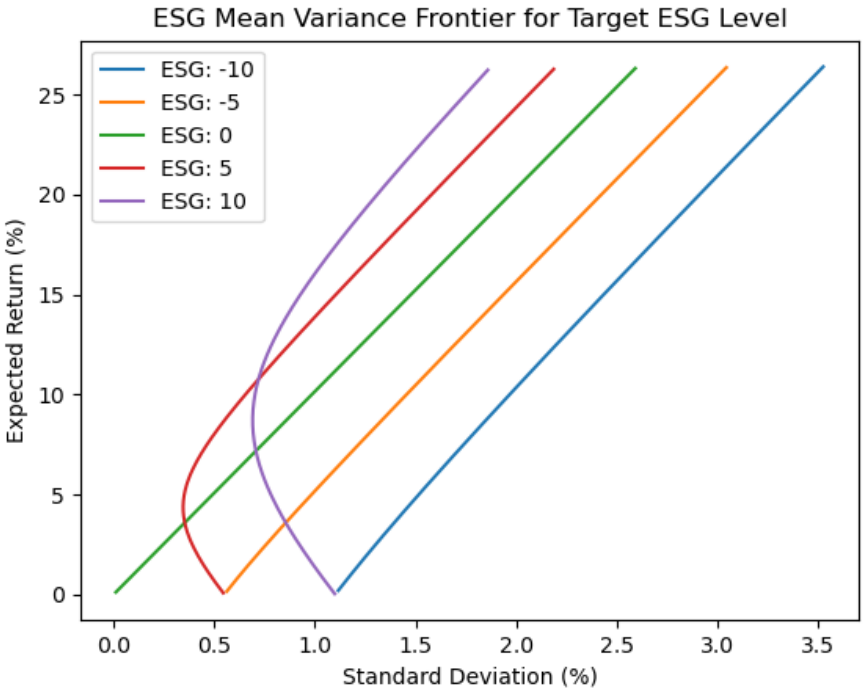


Figure 3: Simulated slices of the ESG-efficient frontier at fixed ESG levels (-10, -5, 0, 5, and 10), using the same 10 simulated risky assets and risk-free asset as in Figure 2. The ESG $= 0$ slice is approximately linear, while non-zero ESG targets generate hyperbolic mean–variance tradeoffs with parallel asymptotes.

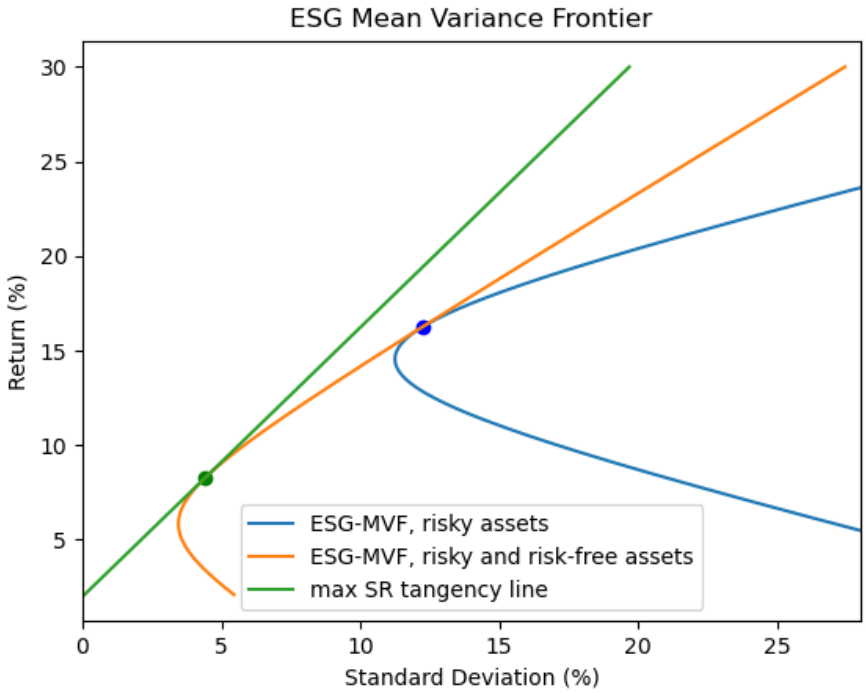


Figure 4: Simulated ESG Mean Variance Frontier, With and Without Risk-free Asset. This plot illustrates the ESG Mean Variance frontier slice constructed from 30 simulated risky assets and a risk-free asset, where ESG level is selected to be 5. The hyperbolic frontier formed by risky assets is shown in blue, alongside the extended frontier incorporating the risk-free asset (in yellow) with their intersect the tangency portfolio (the blue point). The max SR tangency line is depicted in green, and its intersection with the extended frontier identifies the maximum Sharpe ratio (SR) portfolio.

frontier, creating a "bigger" hyperbola that incorporates both financial and ESG attributes. The tangency portfolio, defined as the point of intersection between the hyperbolas formed by the efficient frontiers of risky assets (with and without the risk-free asset), reflects the portfolio that balances financial and ESG objectives as determined by the stochastic discount factor and ESG-adjusted discount factor.

In contrast, the maximum Sharpe ratio (SR) portfolio is the portfolio where the tangency line representing combinations of the risk-free asset and the tangency portfolio intersects with the outer hyperbolic frontier formed by risky and risk-free assets. While the tangency portfolio optimizes the trade-off between financial returns, risk, and ESG attributes, the maximum SR portfolio exclusively focuses on maximizing the financial risk-adjusted return. These two portfolios may not coincide in our framework due to the additional constraints and preferences introduced by ESG considerations.

This distinction highlights the unique flexibility and broader applicability of our model. Unlike CAPM, where the inclusion of a risk-free asset simplifies the efficient frontier into a linear relationship, our model preserves the hyperbolic structure, enabling a more nuanced representation of investor preferences. Specifically, we have the following propositions.

Proposition 5.5 (Tangency Portfolio Spanning of the ESG-Efficient Frontier)**.** *The ESG-efficient frontier with a risk-free asset can be spanned by any two distinct tangency portfolios together with the risk-free asset.*

Proof. Let T_1 and T_2 be two distinct tangency portfolios, corresponding to different ESG levels, and let F denote the risk-free asset. It suffices to show that T_1, T_2, and F are linearly independent.

Suppose, to the contrary, that they are linearly dependent. Then T_1, T_2, and F lie on the same line in portfolio space, so there exist weights such that

$$T_2 = \alpha T_1 + (1 - \alpha)F$$

for some scalar α. The line through F and T_1 is the capital allocation line associated with T_1, and by definition it is tangent to the ESG-efficient frontier at T_1. Since T_2 also lies on this line and is itself an efficient portfolio, the same line would be tangent to the frontier at both T_1 and T_2. This contradicts the strict convexity of the efficient frontier, which implies that a line through the risk-free asset can be tangent to the frontier at most at one point.

He nce T_1, T_2, and F must be linearly independent. Because the risky part of the ESG-efficient frontier is a two-dimensional affine set, any efficient risky portfolio can be written as a linear combination of T_1 and T_2, and any portfolio on the full ESG-efficient frontier (including positions in the risk-free asset) can be generated by T_1, T_2, and F. □

The previous result shows that any two distinct tangency portfolios, together with the risk-free asset, span the entire ESG-efficient frontier. We now identify an economically distinguished tangency portfolio: the aggregate market portfolio that arises in competitive general equilibrium.

Proposition 5.6 (Efficiency of the Market Portfolio)**.** *Under competitive general equilibrium, the market-value-weighted portfolio of risky assets is a tangency portfolio on the ESG-efficient frontier.*

To see the intuition, note first that the market portfolio is, by definition, the value-weighted aggregate of all investors' risky holdings. In equilibrium, net supply of the risk-free

asset is zero, so the aggregate portfolio of risky assets coincides with the market portfolio and lies in the risky-asset space. Because each investor chooses an optimal combination of the risk-free asset and risky portfolios given prices, the aggregate holdings must lie on the ESG-efficient frontier: if the market portfolio were inefficient at its ESG level, all investors could strictly improve by shifting toward some more efficient portfolio, contradicting optimality and market clearing.

A tangency portfolio is defined as the point where the ESG-efficient frontier meets the line through the risk-free asset that delivers the highest Sharpe ratio for a given ESG level. Suppose the market portfolio at its ESG level were not such a tangency portfolio. Then there would exist another portfolio with the same ESG exposure but a higher Sharpe ratio along the line with the risk-free asset. Every investor could then tilt from the market portfolio toward this alternative, improving their risk–return tradeoff without changing ESG exposure. In equilibrium, such a strictly dominated aggregate allocation cannot persist. Hence the market portfolio must lie at the tangency point, and therefore is a tangency portfolio on the ESG-efficient frontier.

6 Two Special Cases

We now show that our general ESG portfolio framework reproduces the known fund-separation results of Pedersen et al. (2021) and Pástor et al. (2021) as special cases. Under the assumptions of Pedersen et al. (2021), our model yields the familiar four-fund separation result. Under the assumptions of Pástor et al. (2021), it yields the three-fund separation. In each case we recover their main portfolios and the corresponding efficient frontier, and then note how our broader framework subsumes these results while allowing additional flexibility.

6.1 ESG-Efficient Portfolios Without Risk-Free Asset

Following Pedersen et al. (2021), consider a market of n risky assets with payoff vector x (mean μ, covariance Σ) and assume a risk-free asset is excluded from the payoff space. Let each risky asset have unit price (so the price vector is $\mathbf{1}$) and ESG score e_i, collected in the vector e. Assume Σ is invertible and that $\mathbf{1}, \mu, e$ are linearly independent. In this setting one finds three linearly independent portfolios on the ESG-efficient frontier: the minimum-variance portfolio, the (traditional) tangency portfolio, and the ESG-tangency portfolio. These are given by solving projections of $\mathbf{1}$, μ, and e onto the payoff space. The following proposition formalizes this:

Proposition 6.1 (Four-Fund Separation Without Risk-Free Asset)**.** *If we treat payoff of basis risky assets as a vector of random variables x with mean of μ and covariance matrix of Σ. Then, given that basis risky assets has price vector of $\mathbf{1}$ and ESG score vector of e, the following three portfolios:*

$$\begin{aligned}\theta_{mv} &= \Sigma^{-1}\mathbf{1}\\ \theta_{\mu} &= \Sigma^{-1}\mu\\ \theta_{esg} &= \Sigma^{-1}e\end{aligned}$$

are ESG-efficient portfolios.

Proof. By Riesz representation theorem, we can express m^* and g^* as:

$$m^* = \mathbf{1}'E(xx')^{-1}x = \mathbf{1}'(\mu\mu' + \Sigma)^{-1}x \tag{1}$$
$$g^* = e'E(xx')^{-1}x = e'(\mu\mu' + \Sigma)^{-1}x \tag{2}$$

Then, by the Sherman-Morrison formula, since Σ is invertible,

$$(\Sigma + \mu\mu')^{-1} = \Sigma^{-1} - \frac{\Sigma^{-1}\mu\mu'\Sigma^{-1}}{1 + \mu'\Sigma^{-1}\mu}. \tag{3}$$

Apply Equation (3) to (1) and (2) we have:

$$m^* = \mathbf{1}'(\Sigma^{-1} - \frac{\Sigma^{-1}\mu\mu'\Sigma^{-1}}{1 + \mu'\Sigma^{-1}\mu})x \tag{4}$$
$$g^* = e'(\Sigma^{-1} - \frac{\Sigma^{-1}\mu\mu'\Sigma^{-1}}{1 + \mu'\Sigma^{-1}\mu})x \tag{5}$$

Notice that portfolio m^* and g^* lie on the ESG-efficient frontier. Since risk-free asset is not included in the payoff space $\mathcal{Y}$, we know that the projection of $\mathbf{1}$ onto $\mathcal{Y}$ is ESG-efficient. We have:

$$\theta_{m^*} = \Sigma^{-1}\mathbf{1} - \frac{\mu'\Sigma^{-1}\mathbf{1}}{1 + \mu'\Sigma^{-1}\mu} \cdot \Sigma^{-1}\mu \tag{6}$$
$$\theta_{g^*} = \Sigma^{-1}e - \frac{\mu'\Sigma^{-1}e}{1 + \mu'\Sigma^{-1}\mu} \cdot \Sigma^{-1}\mu \tag{7}$$
$$\theta_{\text{proj}(\mathbf{1})} = E(x)E(xx')^{-1} = \frac{1}{1 + \mu'\Sigma^{-1}\mu} \cdot \Sigma^{-1}\mu \tag{8}$$

By Equation (6),(7),(8), we can represent $\theta_{\text{mv}}, \theta_{\mu}, \theta_{\text{esg}}$ as:

$$\theta_{\text{mv}} = \theta_{m^*} + (\mu'\Sigma^{-1}\mathbf{1})\theta_{\text{proj}(\mathbf{1})} \tag{9}$$

$$\theta_{\mu} = (1 + \mu'\Sigma^{-1}\mu)\theta_{\text{proj}(\mathbf{1})} \tag{10}$$

$$\theta_{\text{esg}} = \theta_{g^*} + (\mu'\Sigma^{-1}e)\theta_{\text{proj}(\mathbf{1})} \tag{11}$$

By proposition 5.3, $\theta_{\text{mv}}, \theta_{\mu}, \theta_{\text{esg}}$ are also ESG-efficient. □

Because Σ^{-1} is nonsingular and $\mathbf{1}, \mu, e$ are assumed independent, the three portfolios $\theta_{\text{mv}}, \theta_{\mu}, \theta_{\text{esg}}$ are linearly independent. By Proposition 5.4, the risky ESG-efficient frontier is spanned by these three portfolios. Adding the risk-free asset yields exactly the four-fund separation theorem of Pedersen et al. (2021): any optimal portfolio is a combination of the risk-free asset and these three risky funds. Thus our framework recovers the Pedersen et al. result as a special case. Moreover, we allow greater generality. In Pedersen's model each investor is assumed to have a specific CARA utility that adds a fixed-weight ESG term. By contrast, our approach does not require a particular utility form. Each investor i may have a utility that is increasing in expected wealth and in the portfolio ESG score. This flexibility means our results hold under weaker, heterogeneous preference assumptions. For example, an ESG preference need not enter the utility as a fixed additive term. The conclusion – that the ESG frontier is spanned by $\theta_{\text{mv}}, \theta_{\mu}, \theta_{\text{esg}}$ – remains valid regardless of the particular utility specification.

6.2 ESG-Efficient Portfolios Including Risk-Free Asset

Next we consider the setting of Pástor et al. (2021), where a risk-free asset of return r_f is included in the investment opportunity set. Again assume risky asset returns are normal with mean μ, covariance Σ, and constant risk aversion, as in Pastor's model. Under these assumptions one finds that only two risky portfolios (plus the risk-free) are needed to span the ESG-efficient frontier. In particular, define the traditional tangency portfolio on the excess-return space and the ESG portfolio by:

$$\theta_{\text{tan}} = \Sigma^{-1}(\mu - r_f\mathbf{1}),$$
$$\theta_{\text{esg}} = \Sigma^{-1}e.$$

We claim that θ_{tan}, θ_{esg}, and the risk-free asset lie on the ESG-efficient frontier.

Proposition 6.2 (Three-Fund Separation With Risk-Free Asset)**.** *With a risk-free asset of return r_f (price 1, ESG score 0), the two risky portfolios θ_{tan}, θ_{esg} together with the risk-free asset are ESG-efficient. Hence any optimal portfolio is a combination of these two risky funds plus the risk-free asset.*

Proof. First, following the proof of proposition 6.1, if we exclude risk free asset in our payoff space, $\theta_{\text{mv}}, \theta_{\mu}$ are ESG-efficient. Now let's consider the case with risk free asset. Because risk free asset is the minimum variance asset, it lies on the new ESG efficient frontier. $\Sigma^{-1}\mu, \Sigma^{-1}\mathbf{1}$ does not necessarily lie on the new ESG efficient frontier, but at least one of their combination do. Then, m^* can be expressed as

$$m^* = a_0 r_f + a_1 \mathbf{1}'\Sigma^{-1}x + a_2 \mu'\Sigma^{-1}x \tag{12}$$

where a_0, a_1, a_2 are scalar coefficient. This comes from the fact that m^* lies on the new ESG efficient frontier. By pricing equation:

$$E[m^*(x - r_f)] = 0 \tag{13}$$

comparing coefficient of μ and $\mathbf{1}$ leads to:

$$a_0 r_f + (\mu'\Sigma^{-1}\mathbf{1})a_1 + (\mu'\Sigma^{-1}\mu)a_2 + a_2 = 0 \tag{14}$$

$$-a_0 r_f^2 - (\mu'\Sigma^{-1}\mathbf{1})a_1 r_f - (\mu'\Sigma^{-1}\mu)a_2 r_f + a_1 = 0 \tag{15}$$

Multiply equation 14 by r_f and add to equation 15 we have:

$$a_1 = -r_f a_2 \tag{16}$$

which implies that $\theta_{\text{tan}} = \Sigma^{-1}(\mu - r_f\mathbf{1})$ lies on the new ESG efficient frontier. Lastly, since risk free asset has an ESG score of zero, we have:

$$E[g^*(x - r_f)] = 1$$

Following similar derivation as proposition 6.1,

$$\theta_{g^*} = \Sigma^{-1}e - \frac{(\mu - r_f\mathbf{1})'\Sigma^{-1}e}{1 + (\mu - r_f\mathbf{1})'\Sigma^{-1}(\mu - r_f\mathbf{1})} \cdot \Sigma^{-1}(\mu - r_f\mathbf{1}) \tag{17}$$

lies on the ESG efficient frontier. Since $\theta_{\text{tan}} = \Sigma^{-1}(\mu - r_f\mathbf{1})$ is already on the frontier, we finish the proof by asserting that $\theta_{\text{esg}} = \Sigma^{-1}e$ is on the ESG efficient frontier. □

Because Σ^{-1} is invertible and we assume $\mu - r_f\mathbf{1}$ and e are not collinear, the two portfolios $\theta_{\tan}$ and θ_{esg} together with risk-free asset are linearly independent. Proposition 5.4 then implies that the ESG efficient frontier is spanned by $\theta_{\tan}$, θ_{esg}, and the risk-free asset, recovering the three-fund separation of Pástor et al. (2021). Any optimal portfolio can be formed by combining the risk-free asset with the tangency and ESG portfolios above.

Again, our general framework encompasses this result under weaker assumptions. For example, we do not require all investors to have the same risk aversion or an additive ESG preference in their utility. In fact, any two linearly independent ESG-efficient portfolios (plus the risk-free asset) would suffice to span the frontier. This flexibility can be used to choose portfolios that are most orthogonal and reduce multicollinearity in empirical implementations. Moreover, we allow arbitrary long-short positions. Even if an investor's desired portfolio involves short positions, it can still be decomposed into these three funds (or two funds if the portfolio itself is efficient).

In summary, our model recovers the classical four-fund separation of Pedersen et al. and the three-fund separation of Pastor et al. as special cases. The key advantage of our approach is that it does so without imposing strict utility or market assumptions. The only requirement is that some investors prefer high return to low return, and some investors prefer high ESG to low ESG. Furthermore, our results allow a wide choice of spanning portfolios on the frontier, any combination of which yields valid fund separations. For instance, one could select any three linearly independent portfolios on the ESG-efficient frontier (in the risk-only case) or any two in the case with a risk-free asset. This broadens the practical applicability of the fund separation theorems. Finally, our framework also identifies edge cases: for example, if the ESG scores e lie in the span of $\mathbf{1}, \mu$, then ESG considerations do not add an extra dimension beyond expected return, and the four-fund case automatically reduces to a three-fund case. Such scenarios are naturally captured by our general model.

7 Two-Beta Representation

Building on the structural decomposition provided by the four-fund separation theorem and the properties of the ESG Efficient Frontier, we now present the two-beta representation theorem. This theorem offers a compact, factor-based approach to representing the stochastic discount factor (SDF) and ESG-adjusted stochastic discount factor (ESDF) for any asset in the return space $\mathcal{R}$. By demonstrating that both discount factors can be expressed as linear combinations of two independent returns, this result introduces a streamlined representation of asset pricing that aligns with both traditional financial objectives and ESG considerations.

The two-beta representation theorem is grounded in the idea that for any two linearly independent returns r_1 and r_2, both the SDF and ESDF can be effectively described through a two-factor structure. This implies that any return's expected value within $\mathcal{R}$ can be captured by these two factors, making it possible to model risk and return with a reduced number of parameters. This result not only simplifies the pricing process but also provides a practical framework for ESG-integrated portfolios, allowing investors to quantify the influence of both financial and ESG factors on expected returns in a systematic way.

Theorem 7.1 (Two-Beta Representation Theorem)**.** *For any two linearly independent returns $r_1, r_2 \in \mathcal{R}$ such that:*

1. $c_0 + c_1 r_1 + c_2 r_2 = \mathbf{0}$ *implies* $c_0 = c_1 = c_2 = 0$,
2. *each return* r_j $(j = 1, 2)$ *admits the decomposition*

$$r_j = r^* + w_j r_g^* + v_j r_e^*.$$

There exist constants $a_0, a_1, a_2, b_0, b_1, b_2$ such that:

$$m = a_0 + a_1 r_1 + a_2 r_2 \quad \text{is an SDF for all } r_i \in \mathcal{R}, \tag{18}$$

$$g = b_0 + b_1 r_1 + b_2 r_2 \quad \text{is an ESDF for all } r_i \in \mathcal{R}, \tag{19}$$

Furthermore, we have the following two-factor model.

$$\mathbb{E}(r_i) = \gamma + \beta_{i,1} \left[\mathbb{E}(r_1) - \gamma\right] + \beta_{i,2} \left[\mathbb{E}(r_2) - \gamma\right], \quad \text{for all } r_i \in \mathcal{R} \tag{20}$$

Proof. We first establish (18). Since the risk-free rate r_f lies on the ESG-efficient frontier, the four-fund separation result (Proposition 5.2) implies

$$r_f = r^* + w_0 r_g^* + v_0 r_e^*. \tag{21}$$

Because the SDF $m^* \in \mathcal{R}$, it admits the decomposition

$$m^* = r^* + w^* r_g^* + v^* r_e^* + n^*, \tag{22}$$

where n^* is orthogonal to all ESG-efficient portfolios by construction. By the SDF pricing condition and orthogonality, $\mathbb{E}[m^* n^*] = 0$ and hence $\mathbb{E}[(n^*)^2] = 0$, implying $n^* = 0$ almost surely. Thus

$$m^* = r^* + w^* r_g^* + v^* r_e^*. \tag{23}$$

Using the decompositions

$$
\begin{aligned}
r_f &= r^* + w_0 r_g^* + v_0 r_e^*, \\
r_1 &= r^* + w_1 r_g^* + v_1 r_e^*, \\
r_2 &= r^* + w_2 r_g^* + v_2 r_e^*,
\end{aligned}
$$

any linear combination of (r_f, r_1, r_2) can be written as

$$
a_0 r_f + a_1 r_1 + a_2 r_2 = \big(a_0 + a_1 + a_2\big) r^* + \big(a_0 w_0 + a_1 w_1 + a_2 w_2\big) r_g^* + \big(a_0 v_0 + a_1 v_1 + a_2 v_2\big) r_e^*. \tag{24}
$$

Matching (24) to (23), we obtain a linear system in (a_0, a_1, a_2). A sufficient condition for a solution to exist is that the 3×3 coefficient matrix

$$
\begin{pmatrix}
1 & 1 & 1 \\
w_0 & w_1 & w_2 \\
v_0 & v_1 & v_2
\end{pmatrix} \tag{25}
$$

is invertible. This holds precisely when r_f, r_1, r_2 are linearly independent in $\mathcal{R}$, or equivalently, when $c_0 + c_1 r_1 + c_2 r_2 = \mathbf{0}$ implies $c_0 = c_1 = c_2 = 0$. Since r_f is constant, we can absorb $a_0 r_f$ into a constant term and write

$$
m^* = a_0 + a_1 r_1 + a_2 r_2,
$$

which proves (18). The same argument applies to the ESG-adjusted discount factor g^*, yielding (19).

We now prove the two-factor representation (20). From $m^* = a_0 + a_1 r_1 + a_2 r_2$, the SDF pricing condition for any $r_i \in \mathcal{R}$ gives

$$
1 = \mathbb{E}(m^* r_i) = \mathbb{E}(m^*)\mathbb{E}(r_i) + \operatorname{cov}(m^*, r_i), \tag{26}
$$

so

$$
\mathbb{E}(r_i) = \frac{1}{\mathbb{E}(m^*)} - \frac{\operatorname{cov}(m^*, r_i)}{\mathbb{E}(m^*)}. \tag{27}
$$

Define $\gamma := 1/\mathbb{E}(m^*)$. Using $m^* = a_0 + a_1 r_1 + a_2 r_2$ and linearity of covariance,

$$
\operatorname{cov}(m^*, r_i) = a_1 \operatorname{cov}(r_1, r_i) + a_2 \operatorname{cov}(r_2, r_i). \tag{28}
$$

Let $f := (r_1, r_2)'$ and denote

$$k_i := \begin{pmatrix} \operatorname{cov}(r_1, r_i) \\ \operatorname{cov}(r_2, r_i) \end{pmatrix}, \qquad \Sigma := \begin{pmatrix} \operatorname{var}(r_1) & \operatorname{cov}(r_1, r_2) \\ \operatorname{cov}(r_1, r_2) & \operatorname{var}(r_2) \end{pmatrix}.$$

Condition (i) implies that r_1 and r_2 are linearly independent, so Σ is invertible. Define the beta vector

$$\beta_i := \Sigma^{-1} k_i = \begin{pmatrix} \beta_{i,1} \\ \beta_{i,2} \end{pmatrix},$$

and the 2×1 vector

$$\lambda := -\frac{1}{\mathbb{E}(m^*)} \Sigma \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}.$$

Then (27) can be written compactly as

$$\begin{aligned} \mathbb{E}(r_i) - \gamma &= -\frac{1}{\mathbb{E}(m^*)} \operatorname{cov}(m^*, r_i) = -\frac{1}{\mathbb{E}(m^*)} \begin{pmatrix} a_1 & a_2 \end{pmatrix} k_i \\ &= \beta_i' \lambda = \beta_{i,1} \lambda_1 + \beta_{i,2} \lambda_2. \end{aligned} \tag{29}$$

Applying (29) to r_1 and r_2 themselves yields

$$\begin{pmatrix} \mathbb{E}(r_1) - \gamma \\ \mathbb{E}(r_2) - \gamma \end{pmatrix} = \begin{pmatrix} \lambda_1 \\ \lambda_2 \end{pmatrix},$$

so $\lambda = \big(\mathbb{E}(r_1) - \gamma,\, \mathbb{E}(r_2) - \gamma\big)'$. Substituting this back into (29) gives

$$\mathbb{E}(r_i) = \gamma + \beta_{i,1}\big[\mathbb{E}(r_1) - \gamma\big] + \beta_{i,2}\big[\mathbb{E}(r_2) - \gamma\big], \tag{30}$$

which is exactly (20). This completes the proof. □

This theorem has several important interpretations and implications. First, it establishes that any asset's price can be determined by referencing just two representative returns, along with the risk-free rate as a baseline. By providing a two-beta structure, the theorem enables investors to understand how each return in $\mathcal{R}$ is influenced by these two fundamental factors, capturing both the financial and ESG-adjusted perspectives of asset valuation. This two-factor approach aligns with the spirit of classical multifactor models, yet it is distinct in explicitly incorporating ESG-adjusted valuation through the ESDF.

Importantly, our framework encompasses the single-beta representation of Hansen and Richard (1987) as a special case. In the absence of ESG considerations, if the market portfo-

lio lies on the traditional efficient frontier, it can serve as a one-beta benchmark for pricing assets via the SDF. However, in ESG-integrated environments, the market portfolio may not lie on the ESG-efficient frontier due to non-pecuniary constraints or heterogeneous investor preferences. This complicates the construction of a single-beta pricing model, as identifying a representative portfolio that simultaneously captures both financial and ESG dimensions becomes empirically challenging. In contrast, the two-beta representation permits the use of any two linearly independent returns—such as the market portfolio and a second well-identified ESG-tilted portfolio—which are both empirically estimable and theoretically sufficient to span the relevant payoff space. This flexibility enhances the tractability and testability of the model.

A further strength of the two-beta representation is its dual application to both the SDF and the ESDF. Equations (18) and (19) demonstrate that both pricing kernels can be represented using the same pair of factors. This parallel structure allows for consistent pricing of assets across both financial and ESG dimensions, and facilitates direct inference on how each return contributes to expected value under both perspectives. In practice, this enables empirical researchers and portfolio managers to analyze whether price deviations are driven by traditional risk-return tradeoffs or by ESG-adjusted valuation considerations, using a unified linear framework.

Finally, the two-beta representation reduces the dimensionality of the asset pricing problem without loss of generality. Because any return in $\mathcal{R}$ can be projected onto the space spanned by r_1 and r_2, both the expected return and its pricing kernel (SDF or ESDF) can be analyzed using only two coefficients per asset. This not only simplifies estimation and inference but also offers a transparent way to integrate sustainability metrics into return-based models, bridging the gap between theoretical asset pricing and ESG-aware portfolio construction.

8 Empirical Evidence

This section provides an empirical validation of our theoretical framework by evaluating the performance of the two-factor model implied by the two-beta representation theorem. We test whether asset returns can be better explained by their exposures to the market portfolio and the maximum Sharpe ratio portfolio (MSP), compared to two widely used benchmarks: the single-factor CAPM and the two-factor model of Pástor et al. (2021) (PST), which augments CAPM with an ESG-sorted factor.

Rather than presenting individual empirical exercises in isolation, we organize this section

around model comparison. Specifically, we assess the predictive and pricing performance of our proposed two-factor model against CAPM and PST, using a series of complementary empirical approaches. These include expected return regressions (via Fama-MacBeth methodology), graphical analysis of realized versus fitted returns, and generalized method of moments (GMM) estimation of the stochastic discount factor.

Our objective is to demonstrate that the two-factor model derived from the ESG-efficient frontier offers superior empirical fit and pricing consistency. The evidence presented in this section aims to validate the theoretical claims established earlier in the paper and to highlight the practical relevance of incorporating ESG-adjusted pricing in asset valuation.

8.1 Data and Methodology

This subsection outlines the empirical setup that supports the validation of our two-factor model. We begin by describing the data sources and sample construction. We then detail the formation of testing portfolios and the estimation procedures for the two factor portfolios used in our framework. Finally, we introduce the benchmark models against which our proposed specification is compared. These components collectively provide the empirical foundation for the comparative analysis presented in the following section.

8.1.1 Data Sources

Our empirical analysis uses monthly return data from July 2012 to September 2024, comprising 147 months. Equity return data are sourced from the Center for Research in Security Prices (CRSP). We include all common stocks incorporated in the United States and listed on the NYSE, AMEX, or NASDAQ exchanges, with a CRSP share code of 10 or 11. The CRSP universe is widely regarded as a comprehensive and representative dataset for U.S. equity markets, and serves as the primary input for constructing test portfolios and estimating factor exposures.

To measure ESG-related exposure, we use the MSCI USA ESG Screened Index. This index excludes companies involved in controversial weapons, tobacco, thermal coal, and those that fail to meet specific ESG standards, and thus serves as a practical benchmark for ESG-conscious investing. We use the monthly return series of this index to estimate ESG beta exposures. Unlike firm-level ESG scores—which often vary significantly across providers and may embed judgmental adjustments—the use of a return-based ESG index allows us to sidestep subjectivity and retain consistency across the sample.

8.1.2 Construction of Testing Portfolios

To assess the pricing performance of various asset pricing models, we construct several sets of test portfolios sorted on ESG exposure and traditional firm characteristics.

Our primary test assets are 100 double-sorted portfolios. At the end of each June from 2012 onward, we first sort firms into deciles based on market capitalization. Within each size decile, we estimate an ESG beta for each firm using 24 to 60 months of past returns. ESG beta is defined as the slope coefficient from a time-series regression of excess firm returns on the MSCI USA ESG Screened Index. These betas reflect the firm's exposure to ESG-related return variation. Firms within each size bucket are then sorted into deciles based on their estimated ESG beta, resulting in a 10-by-10 grid of 100 equally weighted portfolios. Portfolio returns are tracked over the following 12 months.

To complement this primary specification, we also construct 5 and 10 portfolios sorted solely on ESG beta. These single-sorted portfolios provide a simpler cross-sectional structure and help isolate ESG effects without size stratification.

As robustness checks, we include two additional test asset sets. The first comprises 21 mixed portfolios, including five ESG beta-sorted portfolios, six size-momentum portfolios, and five industry portfolios, following the structure used in Hsu et al. (2023). The second includes 25 portfolios formed from a 5-by-5 double sort on firm size and operating profitability. These portfolios allow us to test the generalizability of our model across traditional risk dimensions.

8.1.3 Estimation of Factor Portfolios

The two factors used in our empirical specification are the CRSP market index and the maximum Sharpe ratio portfolio (MSP) at ESG level equal to 1.

The selection of these two portfolios is guided by theoretical, empirical, and practical considerations. First, from a theoretical perspective, both the market portfolio and the MSP are tangency portfolios, as established in our earlier results. Tangency portfolios, along with the risk-free asset, span the ESG-efficient frontier. This ensures that these two portfolios satisfy the spanning conditions required by the two-beta representation theorem, allowing us to consistently project returns onto the frontier and capture the joint variation in financial and ESG risk.

Second, these portfolios are implementable in practice. The CRSP value-weighted index is a widely accepted and transparent proxy for the market portfolio. For the MSP, we adopt

estimation methods from Ao et al. (2019), Ledoit and Wolf (2017), Markowitz (1952), and Jensen et al. (2024), which are specifically designed to perform well in high-dimensional environments. These methods are well-established in the empirical asset pricing literature and allow us to construct stable and efficient estimates of the maximum Sharpe ratio portfolio.

Third, and importantly, using these two portfolios avoids direct reliance on firm-level ESG ratings. ESG scores from commercial providers often exhibit substantial disagreement and are susceptible to methodological opacity, as emphasized in Avramov et al. (2022) and Berg et al. (2022). Our approach avoids these challenges by constructing an endogenous ESG factor through return-based methods, ensuring robustness to subjective ESG classifications and enhancing empirical credibility.

We estimate the MSP using the procedure in Ao et al. (2019), which involves the following steps:

1. Let $R = [R_1, R_2, \ldots, R_T]'$ denote the matrix of returns for N assets over T months. Compute the sample mean vector $\hat{\mu}$ and the sample covariance matrix $\hat{\Sigma}$:

$$\hat{\mu} = \frac{1}{T}\sum_{t=1}^{T} R_t, \quad \hat{\Sigma} = \frac{1}{T-1}\sum_{t=1}^{T}(R_t - \hat{\mu})(R_t - \hat{\mu})'.$$

2. Compute the adjusted Sharpe ratio:

$$\hat{\theta} = \frac{(T-N-2)\hat{\theta}_s - N}{T}, \quad \text{where } \hat{\theta}_s = \hat{\mu}'\hat{\Sigma}^{-1}\hat{\mu}.$$

3. Define the corrected target return:

$$\hat{r}_c = \frac{\sigma(1+\hat{\theta})}{\sqrt{\hat{\theta}}},$$

   where σ is a chosen risk level.

4. Estimate portfolio weights w via constrained regression:

$$\min_{w} \frac{1}{T}\sum_{t=1}^{T}(\hat{r}_c - w'R_t)^2, \quad \text{s.t. } \|w\|_1 \leq \lambda,\ w'e = 1,$$

   where λ controls the degree of sparsity.

5. Use cross-validation to select the optimal λ, and solve for final portfolio weights $\hat{w}$.

8.1.4 Benchmark Models

To evaluate the performance of our two-factor model, we compare it against two benchmark specifications.

The first is the Capital Asset Pricing Model (CAPM), which explains expected excess returns using only the market factor, proxied by the CRSP value-weighted index.

The second benchmark is the ESG-augmented two-factor model developed by Pástor et al. (2021), which we denote as PST model. In this model, a second factor is added to capture ESG preferences. This ESG factor is constructed as a long-short portfolio sorted by firm-level ESG scores, with weights proportional to ESG ratings. While the PST model incorporates sustainability preferences, it relies on exogenous ESG scores that are potentially noisy and methodologically opaque.

Together, these two benchmarks represent traditional and ESG-aware approaches to pricing. Our goal is to demonstrate that our model—rooted in the ESG-efficient frontier and constructed without reliance on noisy ESG labels—achieves superior explanatory power across a range of empirical tests.

8.2 Results and Model Comparison

In this subsection, we compare the empirical performance of our proposed two-factor model against two benchmark specifications: the CAPM and the PST model. To conduct a comprehensive empirical assessment, we employ three classes of tests: expected return predictability, Fama-MacBeth regressions, and generalized method of moments (GMM) estimation of the stochastic discount factor.

The following three figures visualize the fitted expected returns against the realized average returns for a set of test portfolios under each model. Fitted values are calculated using the estimated parameters from each specification, and realized returns are computed as the time-series averages of monthly returns over the sample period. Ideally, if the model accurately prices the cross-section of returns, points should lie along the 45° line. Figure 5 displays the results for the CAPM model, which uses only the market beta as an explanatory variable. Figure 6 corresponds to the PST specification, which augments the CAPM with a long-short ESG factor based on score-based portfolio weights. Figure 7 presents the results for our two-factor model that includes both the market beta and the beta with respect to the maximum Sharpe ratio portfolio.

In addition to visual diagnostics, we implement Fama-MacBeth regressions to evaluate

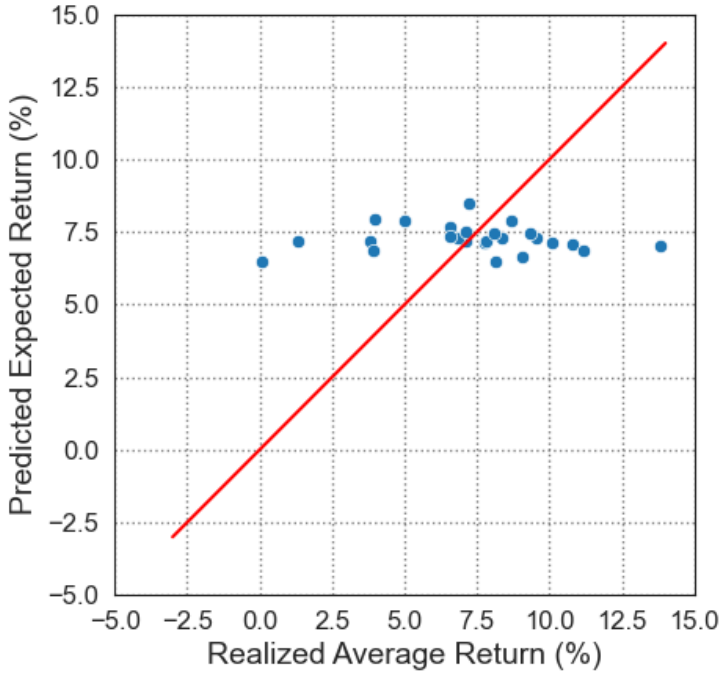


Figure 5: Fitted expected returns versus realized average returns, CAPM Model. In this scatter plot, each point corresponds to a portfolio. The horizontal axis represents the realized average return, while the vertical axis displays the fitted expected return from the regression model:

$$\mathrm{E}[R_{i,t}] = c_0 + c_1 \beta^i_{\mathrm{CAPM}},$$

where β^i_{CAPM} is the market beta obtained from an OLS regression of portfolio i's returns on the CRSP value-weighted market index. The 45° line serves as a benchmark for perfect model fit.

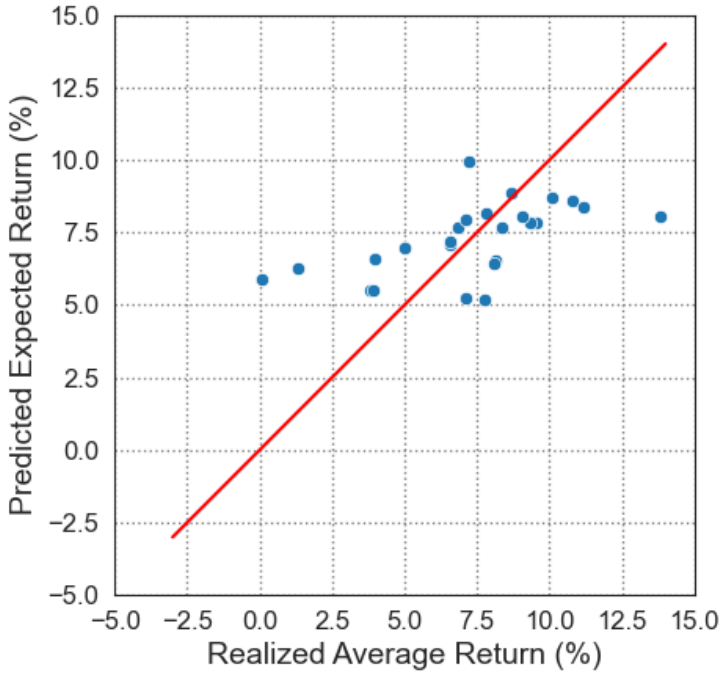


Figure 6: Fitted expected returns versus realized average returns, CAPM Model with PST Factor. Each point corresponds to a portfolio. Fitted returns are obtained from the two-factor model:

$$\mathrm{E}[R_{i,t}] = c_0 + c_1\beta^i_{\mathrm{CAPM}} + c_2\beta^i_{\mathrm{PST}},$$

where β^i_{PST} is the loading on the ESG factor constructed according to Pástor et al. (2021). All coefficients are estimated using time-series regressions. The 45° line provides a benchmark for explanatory accuracy.

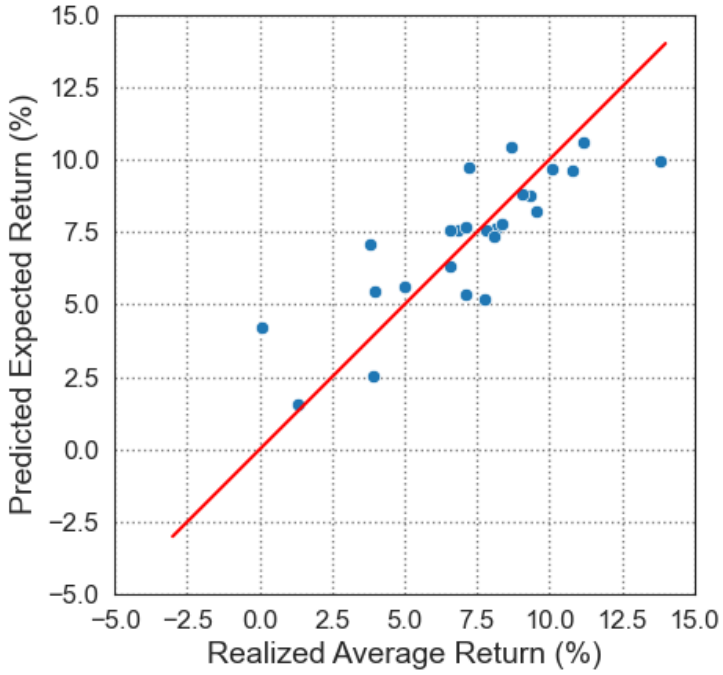


Figure 7: Fitted expected returns versus realized average returns, CAPM Model with MSP Factor. Each point represents a test portfolio. Fitted values are obtained from the model:

$$\mathrm{E}[R_{i,t}] = c_0 + c_1\beta^i_{\mathrm{CAPM}} + c_2\beta^i_{\mathrm{MSP}},$$

where β^i_{MSP} is the beta with respect to the maximum Sharpe ratio portfolio. Rolling-window OLS regressions are used to estimate the factor loadings. The 45° line represents the ideal prediction benchmark.

model performance across portfolios sorted by ESG beta. Table 1 presents the full results. Panel A shows the average excess returns for five ESG-sorted portfolios and the high-minus-low (H–L) spread. Panel B reports regression estimates from the CAPM model, including alphas and market betas. Panel C includes the additional ESG factor from the PST model. Panel D provides results for our model using the market and MSP factors. All betas are estimated using rolling-window regressions with 24 to 60 months of data.

To evaluate the prices of these risk factors more formally, we employ the methodology outlined by Cochrane (2009, pp. 236–239). We define the stochastic discount factor (SDF) as:

$$\text{SDF}_t = 1 - \lambda \times \text{MKT}_t - \lambda_e \times \text{MSP}_t. \tag{31}$$

The moment condition is given by:

$$\text{E}[R_i^e] = -\text{cov}(\text{SDF}_t, R_i^e), \tag{32}$$

where R_i^e denotes excess returns. We estimate this equation using the Generalized Method of Moments (GMM), replacing conditional moments with unconditional moments and using Newey-West standard errors with three lags.

We apply this procedure first to a set of 21 mixed portfolios comprising five ESG-sorted portfolios, six size-momentum portfolios, and five industry portfolios, as in Hsu et al. (2023). Table 2 summarizes the estimated prices of risk under five model specifications: (1) MSP only, (2) CAPM only, (3) CAPM + MSP, (4) PST only, and (5) CAPM + PST. Panel A reports the estimated prices of risk, t-statistics, p-values, and measures of model fit (SSQE, MAPE, J-statistic). Panels B and C report GMM-implied betas and alphas under specifications (2) and (3), respectively.

As a robustness check, we repeat the GMM analysis using 25 portfolios sorted by firm size and operating profitability (a 5×5 grid). Table 3 displays the corresponding results across the same five model specifications. This exercise allows us to examine whether the explanatory advantages of our two-factor model extend to traditional risk-based portfolios.

8.2.1 Comparison with CAPM

From a theoretical perspective, the CAPM posits that the market portfolio alone prices all assets, which implies that a single beta with respect to the market suffices to explain expected returns. However, in our framework, which incorporates ESG preferences, the market

Table 1: Asset Pricing Factor Tests. Panel A reports the average excess returns and associated t-statistics for five portfolios sorted by ESG beta, from the lowest (L) to the highest (H), as well as the high-minus-low (H–L) portfolio. Panel B shows the results from the CAPM model, including alphas and market betas (β_{MKT}). Panel C presents the CAPM augmented with the PST factor, showing alphas, market betas, and PST betas (β_{PST}). Panel D reports the results for the CAPM augmented with the MSP factor, showing alphas, market betas, and MSP betas (β_{MSP}). This comparison evaluates the relative explanatory power of the CAPM, the PST model, and our proposed model in capturing the cross-sectional variation in returns across ESG beta-sorted portfolios.

Panel A: Excess Return						
	L	2	3	4	H	$H-L$
$E[R_i - R_f]$	7.65	9.67	9.47	9.84	9.61	1.97
t-stat	0.84	1.43	1.58	1.83	1.76	0.37
Panel B: CAPM						
	L	2	3	4	H	$H-L$
α	-11.46	-5.86	-4.76	-3.42	-3.39	8.07
t-stat	-1.93	-1.61	-1.61	-1.53	-1.27	1.63
β_{MKT}	1.52	1.24	1.13	1.06	1.04	-0.49
t-stat	14.00	18.55	20.92	25.74	21.23	-5.38
Panel C: CAPM + PST						
	L	2	3	4	H	$H-L$
α	-4.77	-1.05	-1.14	-1.20	-5.18	-0.41
t-stat	-0.71	-0.29	-0.41	-0.57	-1.83	-0.06
β_{MKT}	1.34	1.12	1.05	1.03	1.05	-0.29
t-stat	10.17	16.15	19.20	24.80	18.93	-2.38
β_{PST}	-12.03	-8.43	-6.61	-4.27	-1.18	10.85
t-stat	-3.95	-5.25	-5.23	-4.46	-0.92	3.83
Panel D: CAPM + MSP						
	L	2	3	4	H	$H-L$
α	-2.08	0.01	-0.12	-0.46	-1.99	0.09
t-stat	-0.50	0.01	-0.06	-0.25	-0.75	0.03
β_{MKT}	1.53	1.24	1.14	1.06	1.04	-0.49
t-stat	20.53	27.84	30.32	32.49	21.83	-8.19
β_{MSP}	-12.43	-7.79	-6.14	-3.93	-1.86	10.57
t-stat	-12.12	-12.68	-11.89	-8.77	-2.84	12.73

Table 2: Estimating the Market Price of Risk: 21 Mixed Portfolios. This table presents GMM estimates of the stochastic discount factor using 21 test portfolios comprising five ESG beta-sorted portfolios, six size-momentum portfolios, and five industry portfolios. MSP represents the return on the maximum Sharpe ratio portfolio, serving as an ESG risk proxy. Panel A reports the estimated prices of risk from five specifications: (1) MSP only, (2) MKT only, (3) MKT + MSP, (4) PST only, and (5) MKT + PST. All SDFs are normalized such that $E[m] = 1$. t-statistics and p-values are based on Newey-West standard errors with three lags. Measures of fit include the sum of squared pricing errors (SSQE), mean absolute pricing errors (MAPE), and the J-statistic for overidentifying restrictions. Panel B reports GMM-implied market betas and alphas for the five ESG beta-sorted portfolios under the CAPM specification (2). Panel C reports market and MSP betas and alphas under the CAPM + MSP specification (3).

Panel A: Price of risk					
	(1)	(2)	(3)	(4)	(5)
MKT		0.19	0.21		0.34
t-stat		2.76	2.47		2.82
p		0.01	0.01		0.01
MSP	0.12		0.33		
t-stat	0.67		2.80		
GMSR p	0.50		0.00		
PST				0.19	0.32
t-stat				0.12	2.28
PST p				0.90	0.02
SSQE(%)	15.62	5.07	0.81	11.79	9.04
MAPE(%)	91.38	47.33	17.31	78.00	65.66
J-Test	17.92	20.86	14.58	14.45	15.34
p	0.27	0.14	0.41	0.49	0.36

Panel B: SDF (MKT) in Panel A (2)						
	L	2	3	4	H	H-L
β^i_{MKT}	6.72	5.54	4.97	4.73	4.89	-1.84
t-stat	13.69	14.67	16.83	15.85	17.53	-5.73
α^i	0.59	0.82	0.99	0.93	0.73	0.14
t-stat	1.14	2.88	4.00	4.81	3.13	0.40

Panel C: SDF (MKT+GMSR) in Panel A (3)						
	L	2	3	4	H	H-L
β^i_{MKT}	6.72	5.54	4.97	4.73	4.89	-1.84
t-stat	15.06	14.91	17.04	15.87	18.08	-6.25
β^i_{MSP}	1.44	0.79	0.54	0.40	0.37	-1.07
t-stat	1.77	2.34	2.30	2.37	1.35	-1.88
α^i	0.59	0.82	0.99	0.93	0.73	0.14
t-stat	1.20	4.98	4.13	4.92	3.23	0.41

Table 3: Estimating the Market Price of Risk: 25 Size and Operating Profitability Portfolios. This table presents GMM estimates of the stochastic discount factor (SDF), specified as SDF $= 1 - \lambda \times$ MKT $- \lambda_e \times$ MSP(PST), using 5×5 portfolios sorted on firm size and operating profitability. MSP represents the return on the maximum Sharpe ratio portfolio, serving as an ESG risk proxy. Panel A reports estimated prices of risk for five model specifications: (1) MSP only, (2) CAPM (market-only), (3) CAPM + MSP, (4) PST only, and (5) CAPM + PST. All models are normalized such that $E[m] = 1$. t-statistics and p-values are computed using Newey-West standard errors with three lags. Model fit is assessed via the sum of squared pricing errors (SSQE), mean absolute pricing errors (MAPE), and the J-statistic for overidentifying restrictions.

Panel A: Price of risk					
	(1)	(2)	(3)	(4)	(5)
MKT		0.24	0.25		0.24
t-stat		3.54	3.15		2.57
p		0.00	0.00		0.01
MSP	0.45		0.08		
t-stat	1.39		2.03		
MSP p	0.16		0.04		
PST				-0.48	0.35
t-stat				-3.32	2.75
PST p				0.00	0.01
SSQE(%)	25.30	4.21	1.19	11.08	1.45
MAPE(%)	98.37	33.20	16.65	56.82	20.00
J-Test	14.75	24.60	24.73	21.57	18.66
p	0.93	0.43	0.36	0.61	0.72

portfolio generally fails to span the ESG-efficient frontier. While the single-beta representation theorem remains valid, the market portfolio no longer serves as the representative tangency portfolio. This marks a key departure from traditional mean-variance analysis. The ESG dimension introduces an additional axis of investor preference that the market portfolio alone cannot capture. Although a single tangency portfolio that incorporates ESG efficiently does exist in theory, it is challenging to identify empirically. By contrast, our two-factor model selects two empirically estimable portfolios—the CRSP market index and the maximum Sharpe ratio portfolio—to span the ESG-efficient frontier. Thus, our model is empirically more suitable for environments where ESG factors influence investor behavior, and it should outperform the CAPM in empirical applications.

Empirical results are consistent with this theoretical expectation. Figure 5 and Figure 7 illustrate the fitted expected returns versus realized average returns under the CAPM and our model, respectively. In the CAPM plot, fitted values are nearly uniform and largely fail to match the dispersion in realized returns. The regression line is close to horizontal, suggesting that the market beta alone offers limited explanatory power across the cross section. In contrast, our model produces fitted returns that vary more meaningfully across portfolios, and the points cluster much more closely around the 45° line. This improved alignment highlights the added explanatory value of the MSP factor.

Table 1 further supports these findings through a Fama-MacBeth regression analysis. Panel B, which reports CAPM estimates, shows that the high-minus-low (H–L) portfolio generates an alpha of 8.07% with a t-statistic of 1.63, indicating that CAPM fails to explain return variation related to ESG beta. In contrast, Panel D shows that the alpha under our model drops to 0.09% and is statistically insignificant (t = 0.03). Additionally, the MSP betas in Panel D display a smooth and monotonic pattern across ESG-sorted portfolios, ranging from –12.43 to –1.86, with all t-statistics significant and consistent with the sorting dimension. These results confirm that our model captures ESG-related cross-sectional variation that the CAPM misses.

We further validate model performance using GMM estimation of the stochastic discount factor. Table 2 reports results for 21 mixed portfolios. Under the CAPM (Specification 2), the mean absolute pricing error (MAPE) is 47.33%, and the sum of squared pricing errors (SSQE) is 5.07%. When the MSP factor is added (Specification 3), MAPE drops to 17.31% and SSQE falls to 0.81%, indicating substantial improvement in pricing accuracy. The J-statistic also falls from 20.86 to 14.58, and remains statistically insignificant in both cases, confirming model validity. In terms of factor prices, the market price of risk remains significant (p = 0.01) across both specifications, while the MSP factor becomes statistically significant only in

the augmented model (p = 0.00), indicating that it captures pricing information orthogonal to the market.

Table 3 corroborates these findings using 25 portfolios sorted on size and operating profitability. Under the CAPM, MAPE is 33.20% and SSQE is 4.21%. Adding the MSP factor reduces MAPE to 16.65% and SSQE to 1.19%. These improvements confirm that our two-factor model better prices both ESG-oriented and traditional risk-based portfolios.

In summary, our two-factor model outperforms the CAPM from both theoretical and empirical perspectives. While the CAPM is unable to account for ESG-related variation in asset returns, our model explicitly incorporates ESG considerations via the MSP factor. This enables superior cross-sectional fit, lower pricing errors, and a more accurate representation of investor preferences in the presence of sustainability concerns.

8.2.2 Comparison with PST

The two-factor model of Pástor et al. (2021) represents a major step forward in ESG asset pricing by introducing an explicit ESG factor alongside the market portfolio. However, as illustrated in Section 6, the PST framework relies on a set of strong theoretical assumptions that limit its applicability in more general settings. First, their derivation depends on a representative agent with a specific expected utility function and a constant coefficient of relative risk aversion γ. Moreover, they assume that the market portfolio is ESG-neutral in several ways—not only that its value-weighted ESG score is zero, but also that its equal-weighted and beta-weighted ESG scores vanish. These assumptions are necessary for the ESG factor to emerge cleanly from their model, but they are difficult to justify empirically and unlikely to hold in practice.

To reconnect with the theoretical foundation, the structure developed in this paper does not depend on a specific utility function or representative agent, and imposes no assumptions regarding the ESG neutrality of the market portfolio. Instead, the framework is derived from fundamental asset pricing principles and the law of one ESG-adjusted price. The two-beta representation holds as long as two linearly independent tangency portfolios span the ESG-efficient frontier. As a result, the model is more general and empirically tractable than the PST specification.

A second key distinction concerns estimation. The PST factor is constructed from ESG score rankings, which introduces considerable measurement noise. ESG ratings are subject to inter-provider disagreement and lack consistency, as emphasized in Berg et al. (2022) and Avramov et al. (2022). Sorting portfolios based on noisy ESG scores may attenuate factor

loadings and inflate pricing errors. In contrast, the present approach avoids reliance on ESG ratings. The MSP is estimated using high-dimensional shrinkage techniques as in Ao et al. (2019), which yield stable factor portfolios and improved out-of-sample performance. This estimation procedure reduces noise and enhances empirical reliability.

Empirical comparisons confirm the advantages of the proposed model. Figure 6 shows that the PST model improves upon CAPM in fitted-vs.-realized returns: the scatter becomes more upward sloping and the regression line more responsive to variation. However, Figure 7 reveals that the two-beta model achieves markedly better alignment, with fitted returns closely clustering around the 45° line. This pattern indicates improved explanatory power in capturing cross-sectional return variation.

Table 1 presents a detailed Fama-MacBeth regression analysis. In Panel C, the PST model reduces the high-minus-low alpha to -0.41%, compared to 8.07% under CAPM. Nonetheless, the alpha remains statistically insignificant, and the estimated PST betas show less monotonicity and lower significance. In contrast, Panel D shows that the proposed two-factor model nearly eliminates pricing errors (alpha $= 0.09\%$, $t = 0.03$), and MSP betas vary smoothly across ESG-sorted portfolios, with all t-statistics strongly significant. These findings underscore the improved explanatory performance of the two-beta specification.

The GMM estimation results reinforce these conclusions. In Table 2, the PST-only specification (4) produces substantial pricing errors, with MAPE $= 78.00\%$, SSQE $= 11.79\%$, and an insignificant ESG risk price ($\lambda_e = 0.19$, $p = 0.90$). The PST+market specification (5) offers marginal improvements (MAPE $= 65.66\%$, SSQE $= 9.04\%$) but still underperforms the two-beta model (Specification 3), which achieves MAPE $= 17.31\%$ and SSQE $= 0.81\%$. All models satisfy overidentification tests, but the error metrics clearly favor the ESG-integrated specification. Similar patterns hold in Table 3, where the PST-only model generates a significant negative ESG risk price ($\lambda_e = -0.48$, $p < 0.01$) and elevated pricing errors (MAPE $=$ 56.82%, SSQE $= 11.08\%$). Even in the PST+market case, pricing errors remain higher than those from the two-beta model (MAPE $= 20.00\%$, SSQE $= 1.45\%$).

In sum, our proposed two-beta model outperforms the PST benchmark both theoretically and empirically. Theoretically, it relaxes restrictive assumptions on utility form and market neutrality while maintaining a valid pricing relation grounded in no-arbitrage. Empirically, it delivers improved pricing accuracy, greater explanatory power, and lower estimation noise. These findings affirm the robustness and relevance of the ESG-adjusted two-beta representation.

9 Conclusions

This paper presents a comprehensive, model-free framework for integrating Environmental, Social, and Governance (ESG) factors into asset pricing and portfolio management, providing a significant advancement over traditional models. By introducing the concept of ESG-adjusted prices, we address the limitations of conventional ESG scores, offering a consistent and scalable measure that applies across all types of portfolios, including long-short structures. Anchored in the principles of the LOOP and its ESG counterpart, LOOE, our framework establishes a unified no-arbitrage foundation that ensures consistency in pricing both financial and ESG dimensions.

Our key contributions include the four-fund separation theorem with minimal assumptions, which generalizes existing results by decomposing the asset space into orthogonal components that reflect both financial returns and ESG considerations. The two-beta representation theorem further simplifies the modeling of asset returns, linking them to two fundamental risk factors: the traditional market beta and the ESG beta. These theoretical advancements provide a more flexible and empirically applicable approach to asset pricing, enabling a clearer understanding of how ESG factors influence risk and return.

Empirically, we validate these theoretical contributions by demonstrating the superior explanatory power of the two-beta representation. In particular, we show that incorporating the beta on the maximum Sharpe ratio (MSP) portfolio significantly improves the performance of standard asset pricing models. Compared to both the traditional CAPM and the ESG-augmented model of Pástor et al. (2021) (PST), our specification achieves lower pricing errors and stronger cross-sectional fit. These findings underscore the growing relevance of ESG-related factors in asset pricing and offer a flexible, empirically grounded framework for portfolio construction that integrates financial and sustainability objectives.

While our framework makes significant strides in advancing ESG integration in asset pricing, it also opens several important avenues for future research. One key area is extending the analysis to dynamic settings, where the evolution of ESG preferences over time could influence asset prices and market equilibria. Additionally, investigating the implications of ESG factors across different global markets under varying regulatory, political, and cultural contexts could provide further insights into the role of sustainability in global finance. Finally, further empirical studies could explore sector-specific ESG dynamics and their interaction with broader macroeconomic conditions, which would enhance the practical applicability of our framework. Our work sets the foundation for a deeper exploration of ESG's role in shaping sustainable finance and asset pricing. It also provides tools that both academics and

practitioners can use to better understand and integrate ESG factors into their investment decisions.

References

Ao, M., Yingying, L., and Zheng, X. (2019). Approaching mean-variance efficiency for large portfolios. *The Review of Financial Studies*, 32(7):2890–2919.

Avramov, D., Cheng, S., Lioui, A., and Tarelli, A. (2022). Sustainable investing with esg rating uncertainty. *Journal of financial economics*, 145(2):642–664.

Avramov, D., Lioui, A., Liu, Y., and Tarelli, A. (2024). Dynamic esg equilibrium. *Management Science*.

Baker, M., Bergstresser, D., Serafeim, G., and Wurgler, J. (2022). The pricing and ownership of us green bonds. *Annual review of financial economics*, 14(1):415–437.

Becker, G. S. (2010). *The economics of discrimination*. University of Chicago press.

Berg, F., Koelbel, J. F., and Rigobon, R. (2022). Aggregate confusion: The divergence of esg ratings. *Review of Finance*, 26(6):1315–1344.

Chava, S. (2014). Environmental externalities and cost of capital. *Management science*, 60(9):2223–2247.

Cochrane, J. H. (2009). *Asset pricing: Revised edition*. Princeton university press.

Edmans, A. (2011). Does the stock market fully value intangibles? employee satisfaction and equity prices. *Journal of Financial economics*, 101(3):621–640.

El Ghoul, S., Guedhami, O., Kwok, C. C., and Mishra, D. R. (2011). Does corporate social responsibility affect the cost of capital? *Journal of banking & finance*, 35(9):2388–2406.

Fama, E. F. and French, K. R. (2007). Disagreement, tastes, and asset prices. *Journal of financial economics*, 83(3):667–689.

Gompers, P., Ishii, J., and Metrick, A. (2003). Corporate governance and equity prices. *The quarterly journal of economics*, 118(1):107–156.

Hansen, L. P. and Richard, S. F. (1987). The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica: Journal of the Econometric Society*, pages 587–613.

Hong, H. and Kacperczyk, M. (2009). The price of sin: The effects of social norms on markets. *Journal of Financial Economics*, 93(1):15–36.

Hsu, P.-H., Li, K., and Tsou, C.-Y. (2023). The pollution premium. *The Journal of Finance*, 78(3):1343–1392.

Jensen, T. I., Kelly, B. T., Malamud, S., and Pedersen, L. H. (2024). Machine learning and the implementable efficient frontier. *Swiss Finance Institute Research Paper*, (22-63).

Kempf, A. and Osthoff, P. (2007). The effect of socially responsible investing on portfolio performance. *European financial management*, 13(5):908–922.

Ledoit, O. and Wolf, M. (2017). Nonlinear shrinkage of the covariance matrix for portfolio selection: Markowitz meets goldilocks. *The Review of Financial Studies*, 30(12):4349–4388.

Markowitz, H. (1952). Modern portfolio theory. *Journal of Finance*, 7(11):77–91.

Pástor, L., Stambaugh, R. F., and Taylor, L. A. (2021). Sustainable investing in equilibrium. *Journal of financial economics*, 142(2):550–571.

Pedersen, L. H., Fitzgibbons, S., and Pomorski, L. (2021). Responsible investing: The esg-efficient frontier. *Journal of financial economics*, 142(2):572–597.

Phelps, E. S. (1972). The statistical theory of racism and sexism. *The american economic review*, 62(4):659–661.

Ross, S. A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, 13(3):341–360.

Starks, L. T. (2023). Presidential address: Sustainable finance and esg issues—value versus values. *The Journal of Finance*, 78(4):1837–1872.

Zerbib, O. D. (2019). The effect of pro-environmental preferences on bond prices: Evidence from green bonds. *Journal of banking & finance*, 98:39–60.

10 Appendix A.1: Illustrating LOOP and LOOE

This appendix examines how the classical Law of One Price (LOOP) and our proposed Law of One ESG-adjusted Price (LOOE) can jointly hold in a market equilibrium where a continuum of investors may differ in their ESG sensitivity.

A potential challenge to LOOE arises because, unlike market prices, ESG ratings are set externally (by rating agencies) and do not adjust through trading activity. If two assets had the same payoff and the same price but different ESG scores, classical arbitrage would not directly eliminate the ESG discrepancy since one cannot "trade" ESG scores. An ESG-sensitive investor would strictly prefer the higher-rated asset at the same price, whereas an ESG-indifferent investor would be indifferent between them. Absent a price adjustment mechanism for ESG, one might argue that the Law of One Price could hold (since prices are equal) while the Law of One ESG would fail (ESG scores differ). However, we argue that such a situation cannot persist in equilibrium when ESG-sensitive investors are present. If a higher-ESG asset offered the same return for the same price, ESG-minded investors would flock to it, creating excess demand that would drive its price up relative to the lower-ESG asset. Conversely, if the higher-ESG asset became too expensive, ESG-indifferent investors would favor the cheaper, lower-ESG asset for the same payoff, exerting downward pressure on the green asset's price. Through these opposing forces, the market finds an equilibrium where both LOOP and LOOE hold together: identical-payoff assets will either trade at the same price and have the same ESG profile, or any difference in ESG is accompanied by a consistent price premium.

We illustrate this intuition with a simple two-period, two-asset toy model featuring a continuum of investors who may differ in their ESG preferences. In this model, we consider three scenarios that differ in the distribution of investor types. Scenario 1 features a fully ESG-indifferent investor base. Scenario 2 assumes all investors are ESG-sensitive, but potentially to different degrees. Scenario 3 introduces heterogeneity: a fraction $p \in [0, 1]$ of investors are ESG-indifferent and the remaining fraction $1-p$ are ESG-sensitive. In each case, we examine whether a market equilibrium exists and whether the no-arbitrage conditions (LOOP and LOOE) hold.

The analysis shows that when all investors are indifferent to ESG (Scenario 1), the assets are priced equally (LOOP holds), and no ESG premium arises (LOOE is violated, since ESG is simply ignored). When all investors value ESG (Scenario 2), the higher-rated asset commands a strictly higher price, introducing an ESG-driven price premium consistent with LOOE. Even though classical LOOP in its strict form is violated, no arbitrage exists because

investors consider the assets as fundamentally different due to ESG. Finally, when a mix of investor types is present (Scenario 3), we demonstrate that no stable equilibrium can exist with identical-payoff assets trading at the same price but different ESG ratings. In fact, the attempt to accommodate both investor types leads to a pricing contradiction: ESG-indifferent investors reject any ESG premium, while ESG-sensitive investors demand one. The resulting excess supply or demand prevents market clearing, underscoring that in any true equilibrium LOOP and LOOE must effectively align.

10.1 Model Setup

Consider a two-period model ($t = 0, 1$) with uncertainty at $t = 1$. For simplicity, suppose there are two equally likely states at time 1 (denote $s = 1, 2$ with $\Pr(s) = \pi_s$). There are two assets, Green (G) and Brown (B), each in unit supply. Both assets have identical financial payoffs across states: if $X(s)$ denotes the payoff of one unit of either asset in state s, then

$$X^G(s) = X^B(s) = X(s) \text{ for all state } s.$$

However, the assets differ in their ESG characteristics. Let e^G and e^B be their ESG scores, with $e^G > e^B$. We normalize $e^B = 0$ (no ESG value for the Brown asset) and take $e^G > 0$ (e.g. $e^G = 1$ for one unit of "ESG benefit" per Green asset held). The assets' payoffs X^G, X^B, ESG scores e^G, e^B and the probabilities of each state are common knowledge.

There is a continuum of investors indexed by $i \in [0, 1]$. Each has log utility for consumption and may derive additional utility from holding the Green asset. Specifically, investor i's utility is

$$U_i(\theta_i^G, \theta_i^B) = \mathbb{E}[\log C_{i,1}] + \gamma_i e^G \theta_i^G$$

where $C_{i,1}$ is consumption of investor i at time 1, θ_i^G and θ_i^B are the quantities of G and B held by i at time 0, and $\gamma_i \geq 0$ measures investor i's strength of preference for ESG (the extra utility per unit of Green held). An ESG-indifferent (traditional) investor has $\gamma_i = 0$, valuing only financial returns, while an ESG-sensitive investor has $\gamma_i > 0$ and gains non-pecuniary "warm-glow" utility from holding Green. There is no extra term for holding Brown since $e^B = 0$.

Let p^G and p^B denote the time-0 market prices of the Green and Brown assets, respectively. Each investor faces a budget constraint at $t = 0$:

$$\theta_i^G p^G + \theta_i^B p^B \leq W_i$$

where W_i is investor i's initial wealth. We assume all investors are endowed symmetrically with one unit of total wealth: each is endowed with one unit of the market portfolio consisting of one unit of each asset, i.e. initial endowment $(w_i^G, w_i^B) = (1, 1)$. Then, $W_i = p^G + p^B$. Similarly, for each possible state s at $t = 1$, investors have consumption constraint:

$$C_{i,1}(s) \leq \theta_i^G X^G(s) + \theta_i^B X^B(s)$$

meaning consumption in state s is the payoff of the portfolio held. Given $X^G(s) = X^B(s)$, this simplifies to $C_{i,1}(s) = (\theta_i^G + \theta_i^B)X(s)$. Thus, $\theta_i^G + \theta_i^B$ represents investor i's total units of the payoff (the sum of green and brown holdings), which, together with the state payoff $X(s)$, determines consumption. There is no additional income or trading at $t = 1$, so the portfolio payoff is entirely consumed.

Let $p \in [0, 1]$ denote the proportion of ESG-indifferent investors in the economy, so that a fraction $1 - p$ are ESG-sensitive. Formally, $\gamma_i = 0$ for all $i \in [0, p)$ and $\gamma_i > 0$ for all $i \in [p, 1]$.

Lastly, market-clearing requires that aggregate demand equals aggregate supply:

$$\int_0^1 \theta_i^G \, di = 1, \quad \int_0^1 \theta_i^B \, di = 1,$$

since the total supply of each asset is 1 unit.

We now analyze three scenarios for the investor base and determine the equilibrium outcomes in each case, paying special attention to whether LOOP and LOOE are satisfied. For analytical tractability, we normalize the Brown asset's price to 1 (i.e. treat $p^B = 1$ as the numeraire) throughout the discussion. This means all other prices and wealth can be interpreted in units of the Brown asset.

10.2 Scenario 1: All Investors ESG-Indifferent

In this scenario, all investors are purely return-driven with no taste for ESG ($\gamma_i = 0$ for all $i \in [0, 1]$). Each investor cares only about maximizing expected log consumption. Given that Green and Brown have identical payoffs, an ESG-indifferent investor perceives the two assets as perfect substitutes. We expect that in equilibrium the two assets should have the same price (no price premium for green), as any deviation would create an arbitrage opportunity for these investors.

To verify this, we set up the Lagrange function for investor i with a Lagrange multiplier

λ_i on the budget constraint:

$$\mathcal{L}_i(\theta_i^G, \theta_i^B, \lambda_i) = \mathbb{E}_s\{\log[(\theta_i^G + \theta_i^B)X(s)]\} + \lambda_i \left(p^G + 1 - \theta_i^G p^G - \theta_i^B\right)$$

Since $X^G = X^B = X$, the investor's period-1 consumption is $C_i(s) = (\theta_i^G + \theta_i^B)X_s$. Let $Z_i \equiv \theta_i^G + \theta_i^B$ denote the total units of payoff held by investor i. The utility becomes $\mathbb{E}[\log(Z_i X_s)] = \mathbb{E}[\log Z_i + \log X_s] = \log Z_i + \mathbb{E}[\log X_s]$ (since Z_i is non-random and $\mathbb{E}[\log X_s]$ is fixed). Thus, apart from an additive constant, the investor simply wants to maximize $Z_i = \theta_i^G + \theta_i^B$ – i.e. hold as much total payoff as possible.

The first-order conditions (FOCs) for interior optimum holdings (θ_i^G, θ_i^B) are:

$$\frac{\partial \mathcal{L}_i}{\theta_i^G} = \frac{1}{\theta_i^G + \theta_i^B} - \lambda_i p^G = 0$$
$$\frac{\partial \mathcal{L}_i}{\theta_i^B} = \frac{1}{\theta_i^G + \theta_i^B} - \lambda_i = 0$$

These conditions simplify to $\frac{1}{Z_i} = \lambda_i p^G$ and $\frac{1}{Z_i} = \lambda_i$. Setting them equal yields a simple relation:

$$\lambda_i(p^G - 1) = 0 \tag{33}$$

Equation 33 is the key optimality condition for an ESG-indifferent investor. It says that either $\lambda_i = 0$ or $p^G - 1 = 0$. In economic terms, if the green asset's price p^G differs from the brown's price 1, an indifferent investor would unboundedly shift their portfolio toward the relatively cheaper source of payoff until the budget constraint is no longer binding (which would correspond to $\lambda_i = 0$, a corner solution). In a well-behaved equilibrium with interior solutions for investors (i.e. $\lambda_i > 0$, meaning each investor fully uses their budget), it must be that $p^G - 1 = 0$. Therefore, in equilibrium we have $p^G = 1$.

In words, the green asset must trade at the same price as the brown asset. The Law of One Price holds, as expected: two assets with identical cash flows have identical prices when no one values the ESG aspect. The market-clearing conditions can indeed be satisfied with $p^G = p^B = 1$ (for instance, each investor could simply hold their initial endowment of 0.5 unit of G and 0.5 unit of B; any redistribution between G and B is immaterial to them, since payoffs are the same and prices are equal).

It is important to note that in this scenario LOOE does not hold – nor would we expect it to, given the investor base. LOOE would require that a difference in ESG characteristics be matched by a price difference, but here $e^G > e^B$ while $p^G = p^B$. The green asset carries a

higher ESG score but no price premium. This reflects the intuition that if no investors care about ESG, then ESG traits have no effect on demand or price. There is no "greenium" in equilibrium. Thus, Scenario 1 yields an equilibrium where LOOP holds (trivial no-arbitrage pricing based on payoffs), but the condition implied by LOOE is violated – which is perfectly consistent because the preconditions for LOOE (the presence of ESG-sensitive demand) are absent.

10.3 Scenario 2: All Investors ESG-Sensitive

We now consider the opposite extreme: all investors derive utility from holding the green asset, in addition to the asset payoffs. Suppose each investor $i \in [0, 1]$ has $\gamma_i > 0$, with potentially different magnitudes across investors (reflecting different intensities of ESG preference). In this scenario, every investor is willing to pay a premium for the green asset because it provides extra non-financial utility. We expect to find an equilibrium where $p^G > p^B$, i.e. the green asset is overpriced relative to brown from a pure payoff perspective – yet no arbitrage exists because investors do not view the two assets as identical goods (green provides an extra "service"). This will illustrate the LOOE principle: the price difference compensates for the ESG difference, so that an equilibrium can occur without arbitrage opportunities.

Once again we derive the optimality conditions for a representative investor i (now with $\gamma_i > 0$). The Lagrangian FOCs become:

$$\frac{\partial \mathcal{L}_i}{\theta_i^G} = \frac{1}{\theta_i^G + \theta_i^B} + \gamma_i e^G - \lambda_i p^G = 0$$
$$\frac{\partial \mathcal{L}_i}{\theta_i^B} = \frac{1}{\theta_i^G + \theta_i^B} - \lambda_i = 0$$

where the only change from Scenario 1 is the additional $+\gamma_i e^G$ term in the first equation, reflecting the marginal utility kicker for holding an extra unit of G. Using the second equation, we again have $\lambda_i = 1/(\theta_i^G + \theta_i^B) = 1/Z_i$. Substituting this into the first equation and rearranging, we obtain the counterpart to equation 33 for an ESG-sensitive investor:

$$\frac{1}{\theta_i^G + \theta_i^B}(p^G - 1) = \lambda_i(p^G - 1) = \gamma_i e^G \tag{34}$$

Equation 34 states that, in optimum, the price gap $p^G - 1$ is proportional to the investor's ESG preference γ_i. If p^G were equal to 1, the left-hand side would be zero, but the right-hand side $\gamma_i e^G$ is positive – a contradiction – meaning an interior optimum requires $p^G \neq 1$. In fact, since $\lambda_i > 0$ (the budget will be fully spent in equilibrium), equation 34 implies

$p^G - 1 > 0$. Thus, $p^G > 1$: the green asset must trade at a higher price than the brown asset. This confirms that an ESG price premium emerges when all investors value ESG. No investor perceives an arbitrage in this situation, because the higher price is offset by the extra utility $\gamma_i e^G$ gained from holding green.

We can go further to solve for the equilibrium magnitude of the green price premium. In equilibrium, prices adjust such that asset markets clear given investors' demands. The total demand across investors for payoff units must equal the total supply of 2 units (one from G, one from B), and individual demands satisfy

$$\theta_i^G + \theta_i^B = \frac{p^G - 1}{\gamma_i e^G}.$$

Integrating over all investors $i \in [0, 1]$ yields the market-clearing condition for payoff units:

$$\int_0^1 \left(\theta_i^G + \theta_i^B\right) di = \int_0^1 \frac{p^G - 1}{\gamma_i e^G} \, di = 2.$$

Solving for $p^G - 1$ gives:

$$p^G - 1 = 2e^G \cdot \left(\int_0^1 \frac{1}{\gamma_i} \, di\right)^{-1}. \tag{35}$$

Equation 35 is the equilibrium greenium formula in the continuum-investor setting. It shows that the price premium is increasing in the ESG score e^G and decreasing in the average "inverse" ESG sensitivity $\int_0^1 \frac{1}{\gamma_i} \, di$, which is the harmonic mean of the γ_i's across the population. Intuitively, a higher average willingness to pay for ESG (higher γ_i's) results in a larger green premium.

In Scenario 2, we have found a well-defined equilibrium in which the green asset is overpriced relative to its payoff from a traditional viewpoint, yet correctly priced once ESG utility is considered. The Law of One ESG-adjusted Price holds: the difference in ESG score e^G leads to a proportional difference in price, preventing any utility-enhancing arbitrage. Traditional LOOP in the narrow sense (identical payoffs, identical prices) is violated here, but importantly no investor can exploit this as an arbitrage. Investors do not view the assets as identical: the green asset provides an extra "consumption dividend" in utility terms. Thus, no trade can yield a risk-free profit given the preferences – the higher price of the green asset is justified by its extra non-financial payoff. This scenario mirrors real-world situations in which all market participants price in ESG attributes, leading to a sustainability premium or "greenium" in asset prices.

10.4 Scenario 3: Mixed Population of ESG-Indifferent and ESG-Sensitive Investors

Finally, we turn to the mixed case where some investors are ESG-indifferent and others are ESG-sensitive. Suppose a fraction $p \in (0,1)$ of the investor continuum has $\gamma = 0$ (cares only about financial return), and the remaining $1-p$ fraction has $\gamma > 0$ (cares about ESG). This scenario is of particular interest because it tests whether the market can satisfy both the classical no-arbitrage condition (driven by ESG-indifferent investors) and the ESG-based pricing condition (driven by ESG-sensitive investors) simultaneously. We will see that a tension arises: each type of investor pushes the price in a different direction, and in this setting, no single price p^G can clear the market while satisfying both groups' optimality conditions. In other words, under these assumptions, a competitive equilibrium with both types of investors does not exist. This outcome highlights that if ESG-indifferent and ESG-sensitive investors cohabit the market, the combination of LOOP and LOOE can only hold under more nuanced conditions (e.g., with additional frictions or if one type of investor effectively drops out).

To analyze this scenario, we utilize the key optimality conditions derived earlier. For ESG-indifferent investors (fraction p), the condition is given by equation 33:

$$\lambda(p^G - 1) = 0.$$

For ESG-sensitive investors (fraction $1-p$), the condition is given by equation 34:

$$\lambda(p^G - 1) = \gamma e^G > 0.$$

Let us consider what these imply. There are two cases to examine:

- **Case (i): $p^G = 1$.** If the green asset were priced equally to the brown ($p^G - 1 = 0$), then ESG-indifferent investors' condition is satisfied for any λ (trivially, $0 = 0$). These investors would be content that there is no apparent arbitrage between G and B. However, ESG-sensitive investors' condition would require $0 = \gamma e^G$, which is impossible since $\gamma e^G > 0$. In utility terms, if $p^G = 1$, ESG-sensitive investors see the green asset as a strictly better deal than brown (same price, same payoff, plus extra utility), and would want to allocate all their wealth to G. Since the total supply of G is fixed at one unit, the resulting excess demand from the $1-p$ fraction of investors would drive up the price. Hence, $p^G = 1$ cannot be an equilibrium.
- **Case (ii): $p^G > 1$.** In this situation, the green asset carries a price premium. ESG-

sensitive investors' condition can now be satisfied, since a $p^G > 1$ can solve $\lambda(p^G - 1) = \gamma e^G$ for some $\lambda > 0$. However, ESG-indifferent investors now perceive G and B as payoff-identical, and strictly prefer the cheaper one — B. They will not voluntarily hold G at $p^G > 1$, and would instead shift wealth toward B, even shorting G to increase their B holdings (since short-selling is allowed).

The problem is that while ESG-indifferent investors want to offload G, ESG-sensitive investors are not willing to absorb the full supply (including both the endowment and shorted amounts) unless p^G is lowered. At the prevailing $p^G > 1$, aggregate demand from ESG-sensitive investors falls short of aggregate supply (including short sales) from ESG-indifferent investors. Hence, the market does not clear at any $p^G > 1$ either.

The two cases above demonstrate that there is no single price p^G that can satisfy both types of investors simultaneously in our model. At $p^G = 1$, ESG-sensitive investors have unfilled excess demand for G; at any $p^G > 1$, ESG-indifferent investors supply more green than ESG-sensitive investors are willing to hold. The fundamental issue is that the marginal rate of substitution between the green and brown asset differs between the two groups: ESG-sensitive investors always value green more (in utility terms) than brown, whereas ESG-indifferent investors value them purely by price-payoff ratio. There is no price that equalizes their marginal valuations across the entire continuum. In a competitive equilibrium, prices must adjust so that aggregate demand meets supply, with all investors optimizing at those prices. Here, aggregate demand for green cannot equal supply in a smooth way: either the price is too low (demand exceeds supply), or too high (supply exceeds demand).

Crucially, this breakdown occurs regardless of the relative population share $p \in (0, 1)$. Even a small fraction of ESG-sensitive investors creates excess demand at equal prices, while even a small fraction of ESG-indifferent investors resists any ESG premium. Thus, the mere coexistence of heterogeneous preferences prevents equilibrium.

The implication is that a market in which LOOP and LOOE cannot jointly hold—specifically, where two assets deliver identical payoffs yet differ in ESG score—cannot reach an equilibrium if both ESG-indifferent and ESG-sensitive investors are present. In our model, the coexistence of investor types under such conditions leads to incompatible optimality conditions: ESG-sensitive investors demand a green premium, while ESG-indifferent investors reject it. As a result, no price can simultaneously satisfy both groups' optimal allocations and clear the market. This failure of equilibrium underscores the necessity of LOOP and LOOE holding jointly: when both investor types are active, the only consistent market outcome is one in which assets with identical payoffs command identical prices (LOOP), and ESG-adjusted prices reflect ESG differences (LOOE). Hence, the assumptions of LOOP and

LOOE are not only economically meaningful but are also equilibrium requirements in markets with heterogeneous investor preferences.

In summary, the toy model demonstrates that when the market reaches equilibrium, it aligns with both LOOP and LOOE in the sense that identical-payoff assets will not trade at the same price unless they are also identical in ESG characteristics. If nobody cares about ESG (Scenario 1), the assets trade at the same price and effectively have no meaningful difference aside from an unused ESG label. If everybody cares about ESG (Scenario 2), the asset with superior ESG commands a premium exactly sufficient to equilibrate the extra demand, thereby upholding a generalized "one price" principle when utility is properly accounted for. But if some care and some do not (Scenario 3), a uniform price cannot reconcile the two perspectives – indicating that in any true equilibrium, such a stark investor dichotomy must be resolved either by price segmentation or by one class of investors exiting the field. Thus, our analysis supports the idea that LOOE and LOOP hold together in equilibrium: the market will not allow a persistent situation of two identical payoff assets with equal prices but different ESG scores. Any divergence in ESG will be accompanied by a price divergence in equilibrium, whereas without any ESG-sensitive investors, prices converge despite ESG differences.

11 Appendix A.2: Greenium in General Equilibrium

This appendix develops a parsimonious general-equilibrium model with a green and a brown asset to characterize how ESG preferences generate a greenium in competitive equilibrium. In the presence of both ESG-indifferent and ESG-sensitive investors, we solve for equilibrium portfolio holdings and prices with LOOP and LOOE automatically satisfied. In particular, we derive a set of equation system for the equilibrium price p^G of the green asset and the implied greenium—the price differential between the green and brown asset. The magnitude of this greenium depends monotonically on the ESG preference parameter γ, the proportion p of ESG-sensitive investors, and the ESG score e^G of the green asset, providing a simple structural interpretation of how changes in ESG preferences are transmitted into equilibrium prices.

11.1 Model Setup

We consider a stylized two-period general equilibrium model with two risky assets: a green asset G and a brown asset B. The economy has two possible states of the world at date

$t = 1$: state 1 occurs with probability $\frac{1}{3}$ and state 2 with probability $\frac{2}{3}$.

The asset payoffs are given by:

- Green asset: $x^G = (2, 2)$ in states 1 and 2.
- Brown asset: $x^B = (1, 4)$ in states 1 and 2.

So the payoff matrix is:

$$X = \begin{bmatrix} 2 & 1 \\ 2 & 4 \end{bmatrix},$$

There is a continuum of measure one of investors, each endowed with $(1, 1)$ units of green and brown assets, respectively. An investor i chooses allocations (θ_i^G, θ_i^B) subject to a budget constraint at $t = 0$ and consumption determined by asset payoffs at $t = 1$.

Each investor has adjusted CRRA utility:

$$U_i = \mathbb{E}\left[-\frac{1}{C_i}\right] + \gamma_i e^G \theta_i^G,$$

where C_i is consumption at $t = 1$, $\gamma_i \geq 0$ represents investor i's ESG preference, and $e^G > 0$ is the ESG score of the green asset. The brown asset is normalized to have ESG score zero.

Investors fall into two types:

- ESG-sensitive investors: proportion p, each with $\gamma_i = \gamma > 0$.
- ESG-indifferent investors: proportion $1 - p$, each with $\gamma_i = 0$.

At time $t = 0$, asset prices are denoted p^G and p^B. Each investor chooses allocations (θ_i^G, θ_i^B) to maximize U_i subject to:

$$\begin{aligned} \text{(Budget constraint)} \quad & p^G \theta_i^G + p^B \theta_i^B \leq p^G + p^B; \\ \text{(Consumption)} \quad & C_{i,s} = x_s^G \theta_i^G + x_s^B \theta_i^B, \quad \text{for } s = 1, 2. \end{aligned}$$

Market-clearing requires:

$$\int_0^1 \theta_i^G \, di = 1, \quad \int_0^1 \theta_i^B \, di = 1.$$

In the following subsections, we derive the equilibrium SDF m^*, the ESG-adjusted SDF g^*, and the corresponding risk-neutral probabilities implied by each pricing kernel.

11.2 Traditional SDF and Risk-Neutral Probability (All ESG-Indifferent Investors)

We consider the benchmark case where all investors are ESG-indifferent: that is, $\gamma_i = 0$ for all i. We normalize the Brown asset's price to $p^B = 1$, and solve for the equilibrium price p^G of the Green asset.

So in this case, each investor maximizes expected utility:

$$U_i = \mathbb{E}\left[-\frac{1}{C_i}\right] = -\sum_{s=1}^{2} \pi_s \frac{1}{C_{i,s}},$$

subject to the budget constraint:

$$p^G \theta_i^G + p^B \theta_i^B \leq p^G + p^B,$$

and consumption in state s:

$$C_{i,s} = x_s^G \theta_i^G + x_s^B \theta_i^B.$$

First-Order Conditions. Let λ_i be the Lagrange multiplier on the budget constraint. The Lagrangian is:

$$\mathcal{L}_i = -\sum_{s=1}^{2} \pi_s \frac{1}{C_{i,s}} + \lambda_i \left(p^G + p^B - p^G \theta_i^G - p^B \theta_i^B\right).$$

Taking derivatives with respect to θ_i^G and θ_i^B, we get:

$$\frac{\partial \mathcal{L}_i}{\partial \theta_i^G} = \sum_{s=1}^{2} \pi_s \frac{x_s^G}{(C_{i,s})^2} - \lambda_i p^G = 0,$$
$$\frac{\partial \mathcal{L}_i}{\partial \theta_i^B} = \sum_{s=1}^{2} \pi_s \frac{x_s^B}{(C_{i,s})^2} - \lambda_i p^B = 0.$$

No-Trade Equilibrium. Given identical endowments and preferences, the no-trade allocation $(\theta_i^G, \theta_i^B) = (1, 1)$ is a candidate equilibrium. With payoffs:

$$x^G = (2, 2), \quad x^B = (1, 4),$$

the state-contingent consumption is:

$$C_1 = 2 + 1 = 3, \quad C_2 = 2 + 4 = 6.$$

Using $\pi_1 = \frac{1}{3},\ \pi_2 = \frac{2}{3}$, we evaluate the first-order conditions at $(1, 1)$:

$$\lambda_i p^G = \frac{1}{3} \cdot \frac{2}{3^2} + \frac{2}{3} \cdot \frac{2}{6^2} = \frac{2}{27} + \frac{4}{108} = \frac{2}{27} + \frac{1}{27} = \frac{3}{27} = \frac{1}{9},$$

$$\lambda_i p^B = \frac{1}{3} \cdot \frac{1}{3^2} + \frac{2}{3} \cdot \frac{4}{6^2} = \frac{1}{27} + \frac{8}{108} = \frac{1}{27} + \frac{2}{27} = \frac{3}{27} = \frac{1}{9}.$$

Hence,

$$\lambda_i p^G = \lambda_i p^B \quad \Rightarrow \quad p^G = p^B = 1.$$

Conclusion. In this homogeneous, ESG-indifferent setting, no trade is optimal and consistent with asset prices $p^G = p^B = 1$.

11.3 ESG-Adjusted SDF and System of Equations

We now consider the case where the investor population consists of a mix of ESG-indifferent and ESG-sensitive investors. Specifically, a proportion $p \in (0, 1)$ of investors are ESG-sensitive with a common ESG preference parameter $\gamma > 0$, while the remaining $1 - p$ are ESG-indifferent with $\gamma = 0$. All investors continue to have identical endowments and utility takes the form:

$$U_i = \mathbb{E}\left[-\frac{1}{C_i}\right] + \gamma_i e^G \theta_i^G,$$

where $C_i = \sum_{s=1}^{2} \pi_s C_{i,s}$, and consumption in each state is given by $C_{i,s} = x_s^G \theta_i^G + x_s^B \theta_i^B$. As before, we normalize the Brown asset's price to $p^B = 1$, and attempt to determine p^G along with investor allocations.

Lagrangian Functions. Let θ_1^G, θ_1^B be the allocations of the ESG-indifferent investor and θ_2^G, θ_2^B those of the ESG-sensitive investor. Their respective multipliers are λ_1 and λ_2. The Lagrangian functions are:

(1) ESG-indifferent investor ($\gamma_1 = 0$):

$$\mathcal{L}_1 = -\sum_{s=1}^{2} \pi_s \frac{1}{C_{1,s}} + \lambda_1 \left(p^G + 1 - p^G \theta_1^G - \theta_1^B\right).$$

(2) ESG-sensitive investor ($\gamma_2 = \gamma > 0$):

$$\mathcal{L}_2 = -\sum_{s=1}^{2} \pi_s \frac{1}{C_{2,s}} + \gamma e^G \theta_2^G + \lambda_2 \left(p^G + 1 - p^G \theta_2^G - \theta_2^B\right).$$

First-Order Conditions. For the ESG-indifferent investor:

$$\frac{\partial \mathcal{L}_1}{\partial \theta_1^G} = \sum_{s=1}^{2} \pi_s \frac{x_s^G}{(C_{1,s})^2} - \lambda_1 p^G = 0,$$
$$\frac{\partial \mathcal{L}_1}{\partial \theta_1^B} = \sum_{s=1}^{2} \pi_s \frac{x_s^B}{(C_{1,s})^2} - \lambda_1 = 0.$$

For the ESG-sensitive investor:

$$\frac{\partial \mathcal{L}_2}{\partial \theta_2^G} = \sum_{s=1}^{2} \pi_s \frac{x_s^G}{(C_{2,s})^2} + \gamma e^G - \lambda_2 p^G = 0,$$
$$\frac{\partial \mathcal{L}_2}{\partial \theta_2^B} = \sum_{s=1}^{2} \pi_s \frac{x_s^B}{(C_{2,s})^2} - \lambda_2 = 0.$$

Each investor also satisfies their budget constraint with equality (due to strict convexity of preferences):

$$p^G \theta_1^G + \theta_1^B = p^G + 1,$$
$$p^G \theta_2^G + \theta_2^B = p^G + 1.$$

Market Clearing. Let total market demand equal total endowment (which is 1 unit of each asset). Then:

$$(1-p)\theta_1^G + p\theta_2^G = 1, \quad (1-p)\theta_1^B + p\theta_2^B = 1.$$

Since one of the budget equations is redundant with these, we can summarize the system as follows.

System of Equations. Letting:

$$C_{i,s} = x_s^G \theta_i^G + x_s^B \theta_i^B, \quad \text{with } x^G = (2,2),\ x^B = (1,4),\ \pi = \left(\tfrac{1}{3}, \tfrac{2}{3}\right),$$

the system of 7 equations in 7 unknowns:

$$(\theta_1^G, \theta_1^B, \theta_2^G, \theta_2^B, \lambda_1, \lambda_2, p^G)$$

is given by:

$$\begin{aligned}
&(1)\ \frac{1}{3}\cdot\frac{2}{(C_{1,1})^2}+\frac{2}{3}\cdot\frac{2}{(C_{1,2})^2}=\lambda_1 p^G,\\
&(2)\ \frac{1}{3}\cdot\frac{1}{(C_{1,1})^2}+\frac{2}{3}\cdot\frac{4}{(C_{1,2})^2}=\lambda_1,\\
&(3)\ \frac{1}{3}\cdot\frac{2}{(C_{2,1})^2}+\frac{2}{3}\cdot\frac{2}{(C_{2,2})^2}+\gamma e^G=\lambda_2 p^G,\\
&(4)\ \frac{1}{3}\cdot\frac{1}{(C_{2,1})^2}+\frac{2}{3}\cdot\frac{4}{(C_{2,2})^2}=\lambda_2,\\
&(5)\ p^G\theta_1^G+\theta_1^B=p^G+1,\\
&(6)\ p^G\theta_2^G+\theta_2^B=p^G+1,\\
&(7)\ (1-p)\theta_1^G+p\theta_2^G=1.
\end{aligned}$$

11.4 Numerical Mapping from ESG Preference to Green Asset Price

Given the nonlinear nature of the equilibrium system, it is difficult to solve analytically for the Green asset price p^G as a function of the ESG preference intensity γe^G. However, we can invert this mapping numerically. That is, we treat p^G as a given parameter and determine the implied γe^G consistent with equilibrium. The procedure proceeds as follows, assuming $p = \frac{1}{2}$ and normalizing $p^B = 1$:

1. **Fix a value of** $p^G \in [1.0, 1.5]$.
2. **Compute the ratio of investor 1's state-contingent consumption:** From first-order conditions (1) and (2), compute the ratio $C_{1,1}/C_{1,2}$. This ratio alone suffices, since only proportions matter at this stage.
3. **Recover the ratio of allocations** θ_1^G/θ_1^B**:** Given:

   $$C_{1,1} = 2\theta_1^G + \theta_1^B, \quad C_{1,2} = 2\theta_1^G + 4\theta_1^B,$$

   use the known ratio $C_{1,1}/C_{1,2}$ to solve for θ_1^G/θ_1^B.
4. **Solve allocations using market clearing and budget constraints:** Use conditions

(5)–(7), along with the known ratio θ_1^G/θ_1^B, to solve for:

$$(\theta_1^G, \theta_1^B,\ \theta_2^G, \theta_2^B).$$

5. **Compute investor 2's state-contingent consumption:**

$$C_{2,1} = 2\theta_2^G + \theta_2^B, \quad C_{2,2} = 2\theta_2^G + 4\theta_2^B.$$

6. **Backsolve the implied ESG preference intensity γe^G:** From investor 2's first-order conditions (3) and (4), isolate and compute:

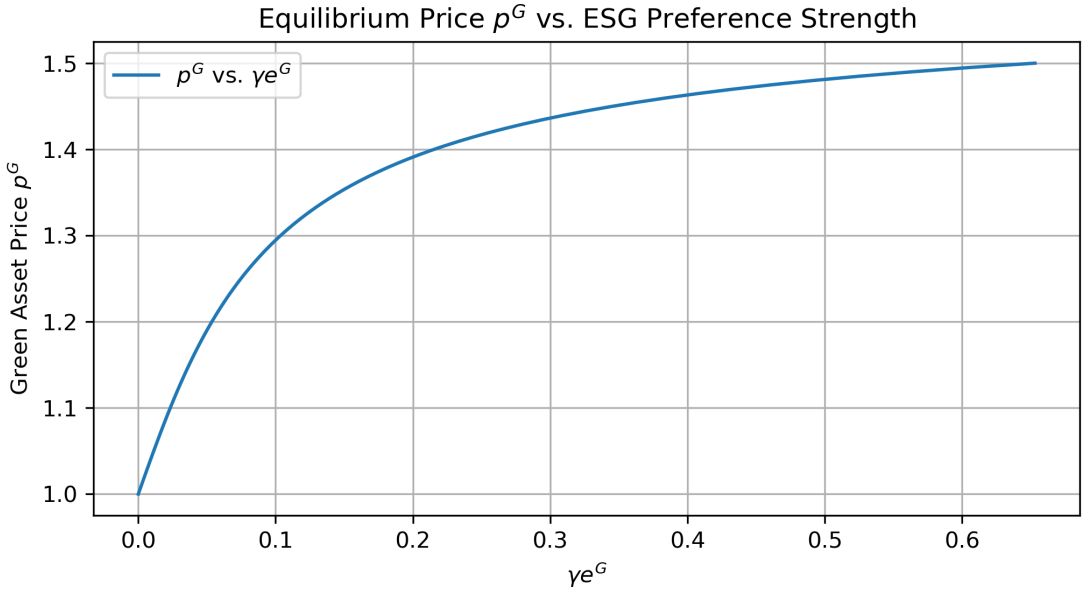
$$\gamma e^G = \left(\frac{1}{3} \cdot \frac{2}{(C_{2,1})^2} + \frac{2}{3} \cdot \frac{2}{(C_{2,2})^2}\right) - p^G \cdot \left(\frac{1}{3} \cdot \frac{1}{(C_{2,1})^2} + \frac{2}{3} \cdot \frac{4}{(C_{2,2})^2}\right).$$

By varying p^G across a range of plausible values, we numerically generate the functional relationship:

$$p^G \mapsto \gamma e^G,$$

which characterizes how equilibrium pricing responds to ESG demand.

Figure 1: Mapping from ESG Preference Strength γe^G to Green Asset Price p^G

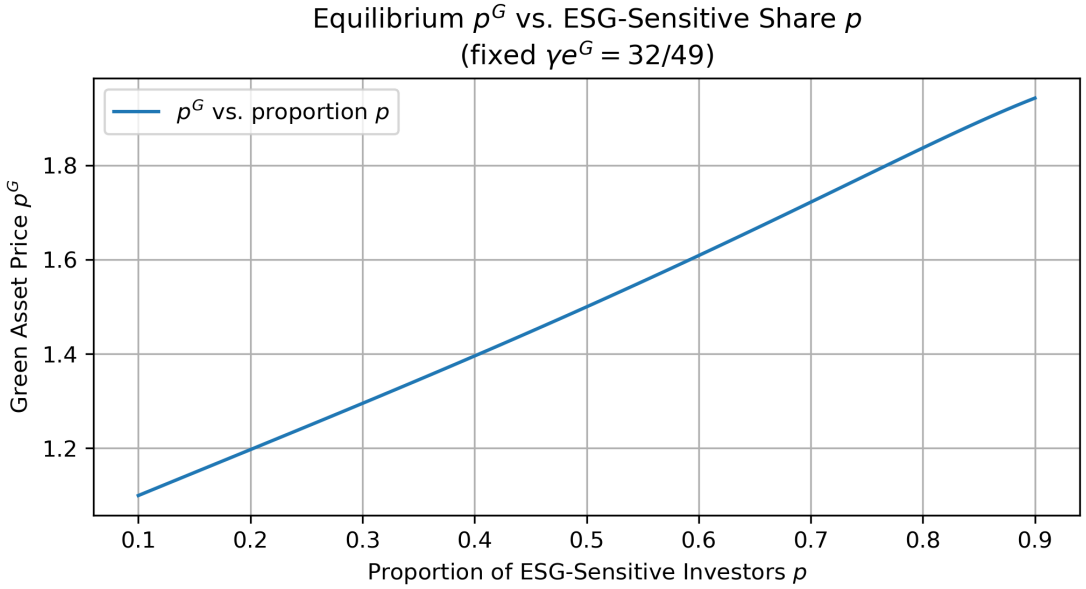


Interpretation. The figure shows how the equilibrium price of the green asset, p^G, responds to increasing ESG preference strength, measured by the product γe^G. Several insights emerge:

- *Monotonicity.* The price p^G increases monotonically with γe^G. As investors care more intensely about sustainability, they are willing to allocate more wealth toward the green asset, pushing its price upward in equilibrium.
- *Diminishing Marginal Effect.* The slope of the curve is strictly decreasing: the marginal impact of γe^G on p^G diminishes as ESG preferences strengthen. Intuitively, once the price premium becomes large, further increases in ESG preference deliver less marginal utility per unit of green investment, so equilibrium prices adjust more slowly.
- *Importance of Fundamental Differences.* Unlike many stylized settings, the Green and Brown assets in this model do *not* have identical payoff structures. The Brown asset pays $(1, 4)$, while the Green pays $(2, 2)$. The green asset is therefore less risky but also less exposed to upside. This asymmetry means that the observed price premium is not just a function of ESG demand, but interacts with fundamental return profiles. In this context, the ESG-adjusted SDF must jointly capture heterogeneity in risk and sustainability preferences.

We further examine how the equilibrium green price p^G varies with the proportion of ESG-sensitive investors in the market. Fixing the ESG preference intensity at $\gamma e^G = \frac{32}{49}$, we numerically solve for the value of p^G that clears markets for varying values of $p \in [0.1, 0.9]$.

Figure 2: Mapping from ESG-Sensitive Share p to Green Asset Price p^G (fixed $\gamma e^G = \frac{32}{49}$)



Interpretation. As the share of ESG-sensitive investors rises, the Green asset becomes more highly demanded in equilibrium, leading to a strictly increasing price p^G. In contrast to the diminishing marginal effects seen with preference strength, the relationship between p^G and p is nearly linear, reflecting the direct weight placed on ESG-sensitive valuations in aggregate demand.

To conclude the model, not only the intensity of ESG motives, but also the breadth of investor participation, plays a central role in shaping ESG-adjusted pricing and the structure of the stochastic discount factor.

12 Appendix B: Proofs

Proof of proposition 3.1. □