Finanziato dall'Unione europea NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

Finanziato nell'ambito del Piano Nazionale di Ripresa e Resilienza PNRR. Missione 4, Componente 2, Investimento 1.3 Creazione di "Partenariati estesi alle università, ai centri di ricerca, alle aziende per il finanziamento di progetti di ricerca di base"

GRINS FOUNDATION

GRINS – Growing Resilient, INclusive and Sustainable

"9. Economic and financial sustainability of systems and territories"

Codice Identificativo: **PE00000018**

Finanziato nell'ambito del Piano Nazionale di Ripresa e Resilienza PNRR Missione 4 – Componente 2

SPOKE 4

D4.1.3 – Policy briefs and best practices for SME's access to sustainable finance instruments

October 2025

Finanziato dall'Unione europea NextGenerationEU | Ministero dell'Università e della Ricerca | Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA | UNIVERSITÀ DEGLI STUDI DI BERGAMO | GRINS FOUNDATION

Toward Principled Metrics for ESG-Aware Risk and Reward

Based on the paper: *"An Axiomatic Risk-Reward Framework for Sustainable Investing" (2025)*, G. Torri, R. Giacometti, D. Dentcheva, S. T. Rachev, W. B. Lindquist. Preprint available at https://arxiv.org/html/2309.05866v3.

---

1. Context & Problem Statement

The rise of sustainable / ESG investing

Over the past decade, demand for investment strategies incorporating environmental, social, and governance (ESG) criteria has surged. Investors increasingly expect portfolios to align with sustainability goals, not just financial returns. Market practices are heterogeneous: some strategies simply screen out "bad" firms, others factor in ESG scores in a pre-selection step, and still others attempt to jointly optimize sustainability and financial performance.

Joint optimization over sustainability and financial performance is typically done by either integrating the ESG scores of the assets in the objective function or introducing constraints to the optimization. The variability over time of the ESG scores is often disregarded, due to difficulties in forecasting future evolutions of such variables, and due to the lack of high frequency data (ESG scores are typically issued with annual frequency).

The ESG scores are thus used alongside the financial returns but are not integrated in the measurement of risk and reward, limiting the power of existing frameworks to account for the tradeoffs between returns and ESG performance. We argue for a bivariate **principled**, **axiomatic** extension of classical risk / reward measures to an ESG-aware setting. The problem is becoming increasingly relevant in light of the appearance of high frequency ESG data, such as the Factset Truvalue Scores, that are updated daily.

Key gap: how to measure ESG-inclusive risk and reward

- Traditional risk measures (e.g. Value at Risk, expected shortfall) focus solely on financial returns, ignoring sustainability fluctuations or ESG score volatility.
- ESG scores are often treated deterministically or exogenously (i.e. fixed), rather than as stochastic processes evolving over time.
- Investors with strong ESG preferences (beyond pure financial optimization) require risk/reward metrics that internalize ESG considerations, rather than treating them as side constraints.

Thus, the central question: **Can we define coherent, axiomatic risk and reward measures that integrate both monetary returns and ESG outcomes, in a way that respects investor preferences (tradeoffs)?**

We point out that the main goal of this work is not to use ESG analysis to improve financial performance, but we aim instead to develop a framework that integrates investors' preferences towards sustainable investment per se, for ethical reasons or to satisfy some specific institutional mandate.

---

2. Conceptual Framework & Definitions

Bivariate modeling: returns + ESG

Each asset (or portfolio) is modelled through a bivariate random vector

$$(X, Y)$$

where:

- X= monetary component (e.g. cumulative log-return over a horizon)
- Y= ESG or sustainability component (e.g. cumulative ESG score over the same horizon)

Thus, the asset's outcome is not just "money gained or lost," but also "how sustainable" the performance was, as captured by the ESG score.

Because $X$and Ymay be correlated (e.g. a firm's ESG changes may correlate with its profits or risk exposures), a bivariate framework allows capturing nuanced dependencies between financial and nonfinancial outcomes.

A possible interpretation of the sustainability component is to consider it as the cumulative sustainability over a period of time, computed as the integral over time of a stochastic process esg_t, proxied by the ESG score issued by a data provider. This implies that sustainability can be treated as a (positive or negative) externality associated to a company.

Investor preference parameter (λ)

To allow different weights or tradeoffs between ESG and financial concerns, the framework introduces a scalar parameter λ (lambda). That parameter encodes the relative weight an investor places on ESG versus monetary risks or rewards. In effect, λ captures the investor's ESG sensitivity or orientation.

This makes the resulting risk and reward measures flexible and personalized: as λ varies, one recovers purely financial ($\lambda = 0$) or purely ESG ($\lambda = 1$) extremes.

---

3. ESG-Coherent Risk Measures

Focusing on risk, we extend the notion of coherent risk measures (Artzner et al., 1999) to the bivariate ESG + monetary setting. Recall that in the univariate framework, a risk measure $\rho(r_T)$, where r_T is the random variable representing the returns of an asset over time T, is defined coherent if it satisfies four axioms: subadditivity, positive homogeneity, translation invariance, and monotonicity.

ESG-coherent risk measures must satisfy analogous, "multivariate" axioms (denoted SUB-M, PH-M, MO-M, LH-M) for a functional

$$\boldsymbol{\rho}_\lambda(X_T, Y_T)$$

that maps from bivariate outcomes to a scalar "risk level."

Let X_T, Y_T denote the random return and ESG of an asset over the period T, and let $\lambda \in [0,1]$, the axioms that define ESG-coherent risk measures are:

- **SUB-M (Subadditivity)**: if $X_{1,T}, X_{2,T} \in \mathcal{X}_2$, then $\boldsymbol{\rho}_\lambda(X_{1,T} + X_{2,T}) \leq \boldsymbol{\rho}_\lambda(X_{1,T}) + \boldsymbol{\rho}_\lambda(X_{2,T})$; That is the risk of a combined position should not exceed the sum of risks (captures diversification across assets and across ESG + monetary dimensions).
- **PH-M (Positive Homogeneity)**: if $\beta \in \mathbb{R}_+$ and $X_T \in \mathcal{X}_2$, then $\boldsymbol{\rho}_\lambda(\beta X_T) = \beta \boldsymbol{\rho}_\lambda(X_T)$; Scaling both monetary and ESG components by a positive constant, scales risk accordingly.

- **MO-M (Monotonicity)**: if $X_{1,T}, X_{2,T} \in \mathcal{X}_2$ and $(r_{1,T} \leq r_{2,T} \wedge ESG_{1,T} \leq ESG_{2,T})$ a.s., then $\boldsymbol{\rho}_\lambda(X_{1,T}) \geq \boldsymbol{\rho}_\lambda(X_{2,T})$; If one position is component-wise "worse or equal" than another, its risk should not be lower.
- **LH-M (Lambda Homogeneity)**: if $a = [a_1, a_2]' \in \mathbb{R}^2$ is deterministic, then $\boldsymbol{\rho}_\lambda(a) = -((1-\lambda)a_1 + \lambda a_2)$. This property determines how a "safe" ESG + monetary position (deterministic) contributes to risk, given λ preferences.

The framework also establishes a **translation invariance extension (TI-M)** in this bivariate context, implied by the four axioms SUB-M, PH-M, MO-M, LH-M.

An important consequence is **diversification** not just across assets, but **between monetary risk and ESG risk**: the total ESG risk is bounded by the sum of pure monetary risk and pure ESG risk.

We underline that the **parameter $\boldsymbol{\lambda}$** reflects the **individual trade-off** between the financial and ESG components for each investor, and that $\lambda = 0$, $\lambda = 1$ represent the extreme cases in which the investor considers only one of the two dimensions. In such cases, the ESG-coherent risk measure collapses to the univariate case of a coherent risk measure computed on either the returns, or the ESG.

A **dual representation** of ESG-coherent risk measures is also derived, analogous to classical risk duality (i.e., as supremum of expectations over a dual "risk envelope").

ESG-coherent risk measures can be obtained also by **extending existing univariate risk measure**. As an example, the following is an extensions of the Average Value at Risk (AVaR):

$$ESG\text{-}AVaR_{\lambda,\tau}(X,Y) := AVaR_\tau\big((1-\lambda)X + \lambda Y\big),$$

where $X$and Y are the period return and ESG, respectively, and $AVaR_\tau(X)$ is the univariate average value at risk (also known as CVaR or expected shortfall). It can be shown that it satisfies the axioms, and in case of $\lambda = 0$ it reduces to the AVaR computed on returns only.

---

4. ESG Reward Measures and Reward-Risk Ratios

Parallel to risk, the authors define **ESG reward measures** that assign a "positive payoff value" to bivariate outcomes (X, Y). These reward measures must align with the idea that higher monetary return and higher ESG outcomes are preferable. An example is the ESG-return $ESG\text{-}\mu_\lambda(X,Y)$ defined as the linear combination if the expectations of X and Y:

$$ESG\text{-}\mu_\lambda(X,Y) := \big((1-\lambda)\mathbb{E}(X) + \lambda\mathbb{E}(Y)\big).$$

From those, one can define **ESG reward-risk ratios (RRRs)**:

$$\text{ESG-RRR} = \frac{\text{ESG reward}}{\text{ESG risk}}.$$

These serve as performance metrics that balance reward and risk in a unified ESG-aware way. The authors generalize classical risk–reward ratios (Sharpe ratio, Sortino ratio, etc.) into this ESG context by replacing numerator and denominator with their bivariate equivalents.

Several variants of ESG reward-risk ratios are discussed in the paper, depending on how one defines the reward and risk components in practice (e.g. using excess returns, thresholded ESG scores, etc.).

An example is the ESG-STAR ratio:

$$ESG\text{-}STARR_{\alpha,\lambda}(X,Y) := \frac{ESG\text{-}\mu_\lambda(X,Y)}{ESG\text{-}AVaR_{\lambda,\alpha}(X,Y)}.$$

5. Empirical Illustration

To showcase the utility of the framework, we construct three simple portfolios using the constituents of the **Dow Jones Industrial Average (DJIA) index**, selecting all the assets in the index, the 10 with the highest average ESG score, and the ones with the lowest ESG score, respectively. Stock prices and ESG scores are retrieved from Factset. In particular, the Factset Truvalue Insight Scores, issued daily for all the stocks in the index are used. The ESG score for each portfolio is computed by averaging the scores of the components weighted by portfolio weight. The empirical work provides a proof-of-concept: the methodology is implementable with real data and yields actionable insights regarding ESG-aware portfolio selection.

We consider a four-years period from January 1st 2020 to December 31st 2023. Figure 1 shows the evolution of the portfolio wealth, assuming an initial investment of one dollar, and the average ESG score for the portfolios. We see that the evolution of the three portfolios is similar, leading to similar final wealth. More in details, the portfolio that includes all the assets leading to a slightly higher final wealth, followed by the "Top ESG" and the "Bottom ESG". As expected, the ESG scores are highest for the "Top ESG" and lowest for the "Bottom ESG" by construction.

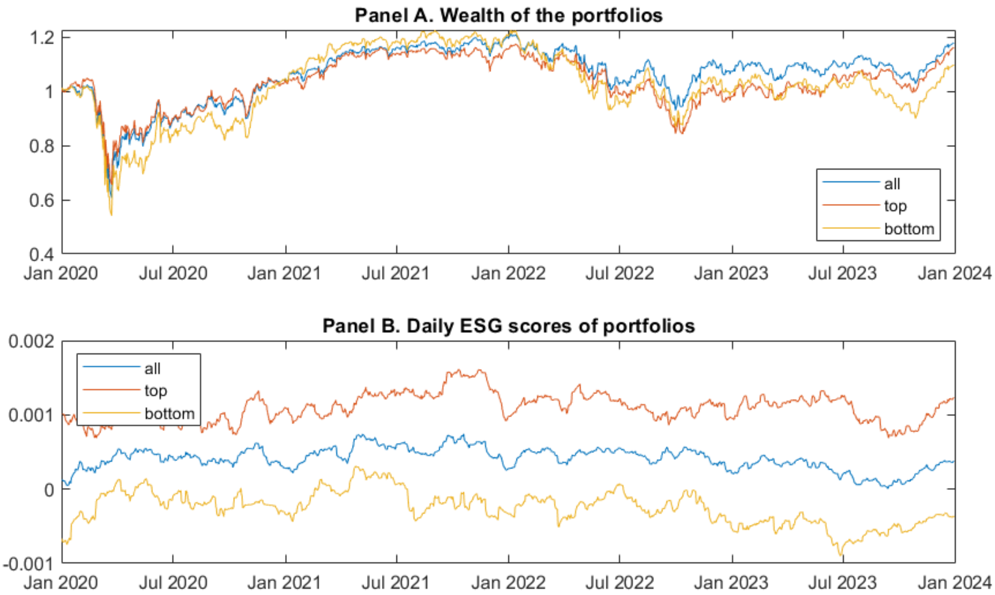

Figure 1- Portfolio Wealth and ESG score.

Figure 2 compares the $ESG\text{-}AVaR$, $ESG\text{-}\mu$, and $ESG\text{-}STAR$ ratio for three different levels of lambda: $\lambda = 0$ (left column, no ESG concerns), $\lambda = 0.25$ (center column, intermediate ESG orientation), and $\lambda = 0.5$ (right column, strong ESG orientation).

For a traditional investor with no ESG concerns, the "all" portfolio that includes all the assets in the index is the most sensitive choice, since it has the lowest risk and the highest return. The "top" portfolio follows, being characterized by slightly smaller risk and lower returns (reflected by a lower ESG-STAR ratio compared to the "all". The "bottom" portfolio ranks worst, showing higher risk and lower return. For investors with ESG concerns (i.e. $\lambda > 0$), the ranking of the portfolio changes: the "top" portfolio has lower ESG risk and higher ESG return compared to the "all" portfolio, and its advantage grows with λ (compare the central and right plots). This means that the simple portfolio screening based on ESG scores leads to a slight worsening of the performance of the portfolio when evaluated using traditional measures, but it leads to relevant improvements when evaluated using the proposed ESG-coherent measures. Finally, the "bottom" portfolio becomes strongly uncompetitive for ESG-oriented investors.

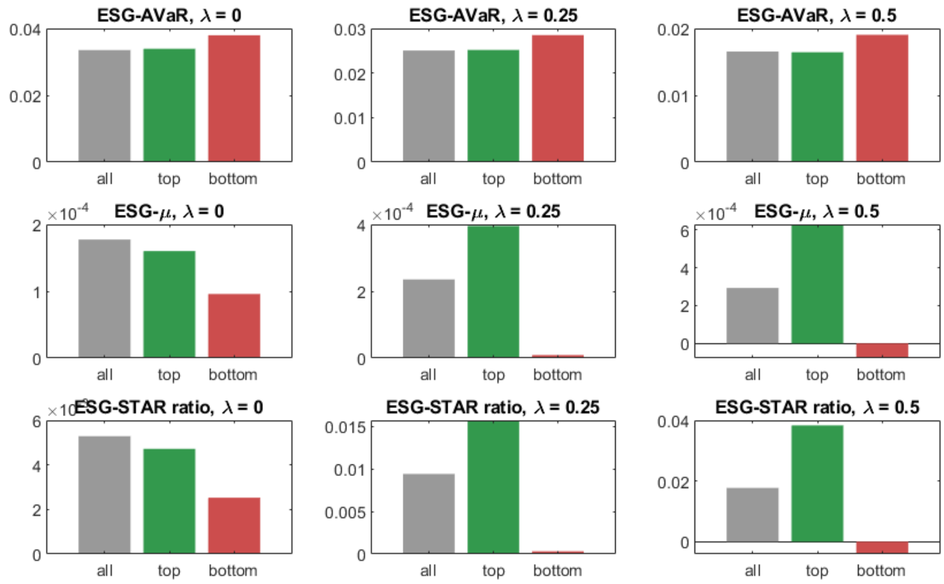

Figure 2 - ESG-risk and reward measures for different values of λ.

---

6. Policy & Practical Implications

Below are key takeaways and recommendations for financial regulators, institutional investors, ESG standard setters, and portfolio managers.

A. Toward standardized ESG risk metrics

- The framework offers a **principled basis** to standardize ESG-inclusive risk measures — a sorely needed objective given the fragmented ESG metrics landscape.
- ESG regulators or taxonomy bodies (e.g. EU SFDR, EU taxonomy) might consider endorsing or referencing axiomatically founded ESG risk measures as benchmarks or guidelines.
- Encouraging data providers to publish **high frequency ESG scores** (not just static scores) becomes crucial for implementing such methods.

B. Investor disclosure & transparency

- Institutional investors that market themselves as "ESG" or "sustainable" may disclose which λ (preference weight) underlies their risk-return calculations. This avoids **"black-box" ESG branding**.
- Mandates or fiduciary guidelines could require that ESG funds provide **ESG-reward-risk ratio metrics** alongside classical performance indicators.

C. Portfolio construction and regulation

- Asset managers could embed **ESG-coherent risk constraints** or ESG RRR **thresholds** into portfolio optimization algorithms, thus internalizing ESG tradeoffs rather than applying ESG as an external filter.

- Regulators might mandate **stress testing** of ESG portfolios under "adverse ESG shocks" (e.g. regulatory shocks, reputational risks), using the bivariate risk framework.

D. Data & infrastructure support

- ESG data providers (e.g. MSCI, Sustainalytics, Truvalue) should move toward **higher-frequency**, time-series ESG disclosures, to enable the modeling of stochastic ESG dynamics.
- Financial infrastructure (e.g. risk management systems) must evolve to accept **multivariate risk inputs** and compute dual representations, scenario analyses, etc.

E. Encouraging research and adoption

- Encourage academic and industry research to refine estimation techniques, e.g. how to estimate joint distributions of (X, Y), especially in presence of sparse ESG data.
- Pilot programs in sovereign wealth funds, pension funds, or development finance institutions (DFIs) could test the approach on real portfolios.

---

7. Limitations, Challenges & Future Directions

While promising, the framework also faces several caveats and areas needing further work:

1. **Quality and comparability of ESG data**. ESG ratings lack standardization across providers. Differences in methodology, sector adjustment, and reporting frequency may bias risk estimates.
2. **Modeling correlation structure**. The dependency between financial returns and ESG flows is complex and may be nonlinear; simple linear correlation assumptions may misestimate tail risks.
3. **Choice of λ is subjective**. The investor must choose λ, but its calibration is nontrivial. Too much sensitivity to λ may lead to unstable ranking or perverse incentives.
4. **Dynamic, multi-period extension**. The present work is mostly static (single horizon). Extending to time-dynamic portfolio allocation with stochastic ESG over multiple periods is a natural next step.
5. **Behavioral or regulatory constraints**. Some ESG preferences reflect normative or regulatory mandates rather than pure risk-return tradeoffs. Integrating constraints or lexicographic preferences (i.e. ESG thresholds that override return) remains to be studied.
6. **Adoption hurdles**. Market participants may resist switching from familiar metrics (e.g. Sharpe ratio) to more complex ESG-aware ones, particularly if interpretability is harder.

---

8. Summary & Recommendations for Policymakers

- **Framework contribution**: Torri et al. provide a rigorous, axiomatic extension of risk and reward metrics into a bivariate ESG + financial space, parameterized by investor ESG preference λ.
- **Practical utility**: The method is implementable (as shown in DJIA empirical test) and yields different stock rankings than purely financial metrics.
- **Policy relevance**: Adopting or endorsing ESG-coherent risk measures can improve transparency, comparability, and credibility of ESG investing.
- **Key actions** for policymakers / regulators / ESG standard bodies:
    1. Encourage or require the publication of stochastic ESG time-series data.

2. Promote (or mandate) disclosure of ESG-RRR alongside classical performance metrics.
3. Support pilot adoption of ESG-coherent risk metrics in institutional portfolios.
4. Fund further research into dynamic, multi-period ESG-risk modeling, robust λ calibration, and scenario stress testing of ESG shocks.