Financial Analysts Journal

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/ufaj20

3D Investing: Jointly Optimizing Return, Risk, and Sustainability

David Blitz, Mike Chen, Clint Howard & Harald Lohre

Financial Analysts Journal | A Publication of **CFA Institute**
https://doi.org/10.1080/0015198X.2024.2335142

Research

3D Investing: Jointly Optimizing Return, Risk, and Sustainability

David Blitz, Mike Chen, Clint Howard, and Harald Lohre

David Blitz is Chief Researcher at Robeco, Rotterdam, the Netherlands. Mike Chen is Head of Next Gen Research at Robeco, Rotterdam, the Netherlands. Clint Howard is a Quantitative Researcher at Robeco, Rotterdam, the Netherlands. Harald Lohre is Head of Quant Equity Research at Robeco, Rotterdam, the Netherlands, and an Honorary Researcher at Lancaster University Management School, Lancaster, United Kingdom. Send correspondence to Clint Howard at c.howard@robeco.com.

Traditional mean-variance portfolio optimization is based on the premise that investors only care about risk and return. However, some investors also have non-financial objectives such as sustainability goals. We show how the traditional approach can readily be extended to mean-variance-sustainability optimization and explain why this 3D investing approach is ex-ante Pareto-optimal. We illustrate its efficacy empirically in several studies, including carbon footprint and sustainable development goal objectives. Importantly, we highlight conditions under which a 3D optimization approach is superior to a naïve 2D approach augmented with sustainability constraints.

Keywords: carbon footprint; ESG; factor investing; portfolio optimization; sustainable development goals (SDG); sustainable investing

Disclosure: The authors disclose that they are employed by Robeco, a firm that offers various investment products. The construction of these products may, at times, draw on insights related to this research. The views and results presented in this article were not driven by the views or interests of Robeco and are not a reflection of its points of view.

PL Credits: 2.0

Introduction

The standard risk and return portfolio framework has been challenged by numerous approaches that all focus on making investment decisions based on objectives that are not strictly risk- or return-based, such as impact investing, socially responsible investing (SRI), or environmental, social, and corporate governance (ESG) investing. The proliferation of these sustainable investing (SI) philosophies highlights how the standard mean-variance framework of Markowitz (1952) is no longer sufficient. Accordingly, investment practice has evolved to incorporate sustainability objectives into the investment problem, with salient examples being metrics related to carbon footprint, ESG, and sustainability development goals (SDG). In this article, we bring together the multi-objective portfolio optimization framework with the real-world implementation of alternative investment objectives. Specifically, we contrast the common practice of incorporating sustainability objectives into a portfolio using constraints with the use of objective function targets and discuss when one may be preferable to the other.

Sandberg et al. (2009) and Horan et al. (2022) highlight the heterogeneity of SRI, where there is no one-size-fits-all approach to developing sustainability-oriented investment portfolios. Many approaches have been proposed in the literature that strive to incorporate sustainability objectives into a portfolio. These include excluding undesirable stocks from the investment universe (Diltz 1995; Kinder and Domini 1997; Naber 2001), constraining the portfolio's exposure to such objectives (Boudt, Cornelissen, and Croux 2013), and incorporating sustainable targets into the return/alpha component of the objective function (Bilbao-Terol, Arenas-Parra, and Cañal-Fernández 2012; Hirschberger et al. 2013; Utz et al. 2014; Chen and Mussalli 2020). However, a

The authors thank two anonymous reviewers as well as participants at the Robeco research seminar and EQDerivatives Europe SRI workshop. The views expressed in this paper are not necessarily shared by Robeco or its subsidiaries.

core question around the optimal way to incorporate sustainability objectives into investment portfolios is underserved.

Investors often target a portfolio carbon footprint lower than some benchmark or achieving a higher sustainability score than the benchmark.[1] These objectives are naturally suited toward a constraint-based framework, as benchmark-relative constraints can readily cater to these investment desires. At low constraint levels, such as simply being better than the benchmark, this approach may work well. However, under more ambitious targets, constraint-based approaches face several challenges. Blitz and Hoogteijling (2022) highlight how a constraint on carbon footprint can be interpreted as an implicit carbon tax (on expected return). When the constraint is non-binding (i.e., redundant), then the tax is effectively zero. However, when the constraint is binding this implies a non-zero tax on lowering carbon footprint. As the constraint becomes more binding, it is implied that high–carbon footprint stocks are highly attractive from an expected return perspective and thus the constraint must apply a large implicit carbon tax on these stocks to prevent the optimization algorithm from purchasing them. The empirical question we aim to answer is whether better ex-post portfolio return and sustainability characteristics are achieved when adding the sustainability metric in the objective function (using a multi-objective optimization framework) or when applying an implicit penalty using portfolio-level constraints. In other words, is 3D investing targeting a sustainability objective alongside risk and return objectives superior to traditional 2D investing augmented by sustainability constraints?

We answer this question based on two relevant practical examples of integrating sustainability objectives into a portfolio. Figure 1 shows how the break-even transaction cost[2] of realistic enhanced indexing portfolios varies when aiming to reduce carbon footprint or improve SDG scores relative to the MSCI World benchmark index. We find that for both carbon footprint reduction and SDG score improvement, the 3D investing approach is the superior solution, especially at more ambitious sustainability targets. Taken together, our results show that for portfolios that seek to track the benchmark closely while outperforming the benchmark, ambitious sustainability goals are better implemented using a direct objective function term, rather than a portfolio-level constraint. The objective function term allows for a rewarded time-varying tradeoff of a stock's expected return and the stock's contribution toward the sustainability objective. It is this flexibility to decide at the portfolio construction's run time when it might be better to go for expected return vis-à-vis sustainability that gives the superior result of the objective function approach.

Such an outcome is not surprising, as the 3D investing optimization framework that we use is ex-ante Pareto-optimal (Zadeh 1963). That is, for a given level of expected return, sustainability performance, and risk level, our approach achieves maximum expected sustainability performance or expected return, respectively. Improving one characteristic

Figure 1. The Cost of Implementing an Alternative Objective

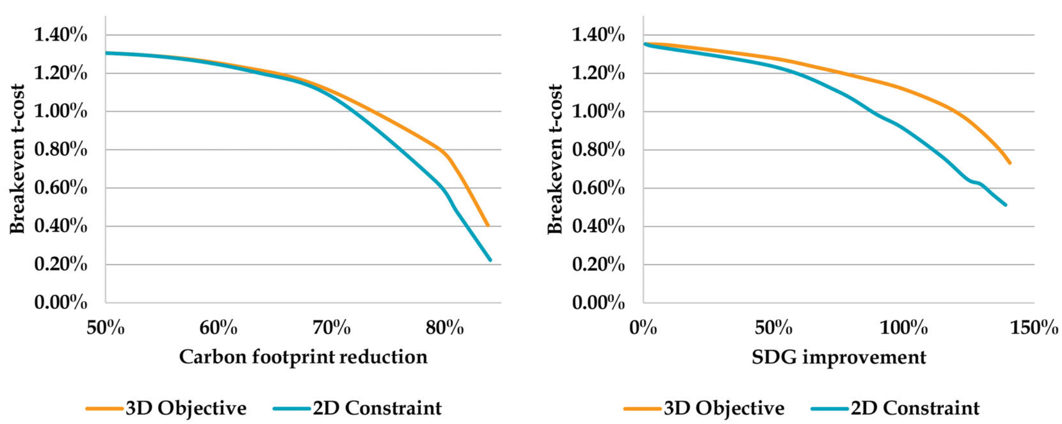

Notes: This figure plots the time-series average break-even transaction cost associated with different levels of benchmark-relative carbon footprint reduction and sustainable development goals (SDG) score improvement under an objective function term (3D Objective) or constraint-based (2D Constraint) portfolio-construction approach. The sample runs from December 1989 to December 2022.

(expected return, sustainability, or risk) requires accepting a worse result in one or more of the other characteristics. Pareto optimality adds to the theoretical appeal of our approach and helps rationalize its empirical value-add.

Notwithstanding the abovementioned, one may argue that sustainability is simply a risk that can be measured and incorporated into the risk dimension of the standard mean-variance optimization framework. This is certainly a possibility wherein some element of sustainability could be incorporated into a factor-based risk model. However, as sustainability definitions are extremely varied and diverse (Berg, Kölbel, and Rigobon 2022), certain sustainability views do not neatly fit into traditional risk or return considerations. Having a flexible framework where such views can easily be incorporated and attributed, apart from risk and return, is valuable.

This article relates to an extensive literature that extends the standard mean-variance optimization framework to incorporate alternative investment objectives. Hallerbach et al. (2004) introduce a multi-decision investment framework that incorporates SRI preferences into the risk-return portfolio-construction process. Other early efforts focused on the optimization approaches required for integrating non-financial objectives into the portfolio optimization process (Bilbao-Terol, Arenas-Parra, and Cañal-Fernández 2012; Ballestero et al. 2012; Dorfleitner and Utz 2012; Utz et al. 2014; Calvo, Ivorra, and Liern 2015; Calvo, Ivorra, and Liern 2016). In recent years, the focus shifted to the construction of ESG-efficient portfolios and how these portfolios relate to the standard mean-variance efficient frontiers (Chen and Mussalli 2020; Geczy, Stambaugh, and Levin 2021; Pedersen, Fitzgibbons, and Pomorski 2021; Schmidt 2020; Shushi 2022; Steuer and Utz 2023; Wu et al. 2022; Xidonas and Essner 2022; Alessandrini and Jondeau 2021; Coqueret et al. 2021). We contribute to this literature by demonstrating the effectiveness of the practical implementation of these methodologies. Specifically, we show how the desired portfolio characteristics interact with the different methods of integrating sustainability characteristics into the portfolio.

Ultimately, portfolio constraints are still (and will continue to be) relevant in the portfolio-construction paradigm. There are scenarios where minimum portfolio exposures or sustainability profiles must always be maintained, and this can only be guaranteed by a constraint. However, if an investor is targeting long-run average sustainability objectives and deviations around this average are acceptable, our results show that 3D investing can deliver portfolios that satisfy this requirement at lower levels of turnover and hence higher after-cost performance. Furthermore, the flexibility of our approach enables taking advantage when a given security's expected return is particularly high with respect to sustainability or vice versa. Thus, when both expected return and sustainability characteristics are important to an investor, our multi-objective 3D investing framework ensures the joint optimality of expected return and sustainability.

Multi-Objective Optimization Framework

Standard Mean-Variance Optimization.

The classic mean-variance optimization problem can be written as:

$$\begin{aligned} \max_{w} \quad & \lambda w'\mu - \frac{\gamma}{2} w'\Sigma w \\ s.t.\ & w'e = 1, \end{aligned} \tag{1}$$

where w is an $N \times 1$ vector of asset weights, μ is an $N \times 1$ vector of expected returns, Σ is the $N \times N$ variance-covariance matrix, e is an $N \times 1$ vector of ones, and λ and γ are scalar coefficients. Portfolios generated under Eq. (1) are mean-variance–optimal in that they achieve the maximum expected return for a given level of risk. This framework can be extended to include additional dimensions, such as constraining the portfolio relative to some benchmark (Jorion 2003), incorporating transaction cost penalties (Taksar, Klass, and Assaf 1988; Ledoit and Wolf 2022), penalizing turnover (Hautsch and Voigt 2019), or enforcing positive asset weights (Jagannathan and Ma 2003).

A Multi-Objective Optimization Framework.

It is straightforward to extend the mean-variance optimizer from Eq. (1) to construct portfolios on an efficient frontier surface in three (or more) dimensions. In the case of additional sustainability considerations, Eq. (1) can be extended to three dimensions as follows:

$$\begin{aligned} \max_{w} \quad & \lambda w'\mu + (1-\lambda) w'\mu_{SI} - \frac{\gamma}{2} w'\Sigma w \\ s.t.\ & w'e = 1, w \in \Omega, \end{aligned} \tag{2}$$

where μ_{SI} is an $N \times 1$ vector of any (discrete or continuous) sustainability metric, λ becomes the relative preference between the return and sustainability objectives, and Ω is the set of feasible solutions, which includes any portfolio constraints. This formulation is general and can accommodate the incorporation of common sustainability characteristics. These include commercial ESG metrics from vendors such as MSCI and Sustainalytics, carbon footprint, SDG scores, and climate transition scores. The only requirement here is that the sustainability metric is ordinal.[3]

This multi-objective optimization technique is called the weighted-sum method (Marler and Arora 2010; Stanimirović, Zlatanović, and Petković 2011), and the resulting solutions can be shown to be Pareto-optimal (Zadeh 1963). This technique allows the construction of portfolios on a multidimensional efficient frontier surface. Previously, this type of portfolio construction has been applied in investment examples such as Gintschel and Scherer (2004), O'Cinneide, Scherer, and Xu (2006), Ballestero et al. (2012), Dorfleitner and Utz (2012), Calvo, Ivorra, and Liern (2015), Chen and Mussalli (2020), and Steuer and Utz (2023). We adopt this 3D investing framework in our subsequent empirical analysis, where we focus on the incorporation of carbon footprint reduction and SDG score improvements in a benchmark-relative portfolio-optimization setting. Although we zoom in on two specific applications, the proposed framework generalizes to any ordinal measure that can be expressed as a series of discrete or continuous values.

Data and Methodology

Data. Our sample consists of MSCI World constituents at the end of every month from December 1989 to December 2022.[4] We source stock returns and fundamental data from Refinitiv. Following Blitz and Hoogteijling (2022), we calculate the carbon footprint of a stock by dividing scope 1 and scope 2 carbon emissions (sourced from Trucost)[5] by enterprise value including cash (EVIC).[6] As highlighted in Busch, Johnson, and Pioch (2020), corporate carbon data can vary in quality across data vendors, especially concerning estimated data points. As we are constructing portfolios from the MSCI World universe, we need sufficient data coverage. Thus, we use the company-reported data where available and use the Trucost-estimated data where data are missing. EVIC is calculated as the market value of a firm's shares plus the book value of its debt. We source stock-level ESG scores from MSCI, where our data begins in January 2009. We source SDG scores from Robeco (Van Zanten and Huij 2022); these are a set of seven discrete variables between −3 and +3 that measure a companies' contribution to the SDGs.[7] For carbon footprint, ESG scores, and SDG scores, missing data are filled with the cross-sectional Global Industry Classification Standard (GICS) subsector (level 2) median.

In our subsequent empirical analyses, we present results contrasting portfolio-construction approaches via multi-objective optimization versus constraints, where we target carbon footprint reductions and SDG improvements. Note that we deliberately choose to not report results for MSCI ESG improvements. This decision is motivated by the skewed distribution and concentration of MSCI ESG scores in recent years. Such a changing distribution, whether empirically warranted or not, makes it increasingly challenging to incorporate the respective sustainability objective into a benchmark-relative portfolio-optimization framework. A changing distribution requires changing parameters for both constraints and objective function terms, thus adding further complexity to the optimization problem. We will revisit the investment implications of such evolving sustainability data following our presentation of the 3D investing outcomes for carbon footprint and SDG scores.

Portfolio Optimization. We use a portfolio-optimization setting that mimics the construction of a real-life investment portfolio applying realistic portfolio constraints and settings. We seek to construct portfolios with tracking errors in the range of 0.5% to 1.0%, as this represents the challenging multi-objective scenario of delivering high expected returns and sustainability goals with a limited risk budget. Thus, the design parameters used are reflective of these lower–tracking error portfolio targets. The portfolio exposure to regions (defined as North America, Europe, and Asia Pacific) and GICS level-one sectors are restricted to ±0.5% of the benchmark market-capitalization weighted value. Portfolio weights must be non-negative (i.e., long-only). The maximum trade size is limited to 25% of a stock's average daily volume over the past 65 trading days. The maximum stock weight relative to the benchmark (i.e., active weight) is ±0.5%. The maximum active share of the portfolio is 40%. The gross exposure of the portfolio must be 100% (i.e., fully invested). We assume that the funds under management grow with the realized market return, and we design the simulations such that the final fund size at the end of 2022 is EUR 4 billion. We incorporate a turnover penalty into the objective function, which is the sum of the squared absolute trade sizes.

As we target specific tracking errors, we transform the weight vector of Eq. (2) from absolute asset weights to benchmark relative weights[8]:

$$w_{new} = w_p - w_{bm}.$$

Our portfolio-optimization problem for a single time-step is then given by:

$$\max_w \quad \lambda_1 w'_{new}\mu + \lambda_2 w'_{new}\mu_{SI} - \frac{\gamma}{2} w'_{new}\Sigma w_{new} \\ -\kappa \| w_{new} - w_{old} \|, \tag{3}$$

where w_{old} is the portfolio weights immediately before the rebalance and κ is a scaling parameter for the turnover penalty (we set $\kappa = 1$), and we incorporate the previously described constraints. We use a base set of portfolio-construction constraints and settings across our simulations, and then we permute the expected return coefficient λ_1, risk aversion coefficient γ, and sustainability coefficient λ_2 in each different optimization. Last, we introduce an additional optional constraint on either carbon footprint or SDG scores (e.g., portfolio carbon footprint must be less than or equal to the benchmark carbon footprint).

Expected Returns and Risk. As inputs of expected returns μ, we use a simple equal-weighted multifactor score (denoted QMV) consisting of quality, momentum, and value signals. For value, we use an equal-weighted combination of book to price and 12-month forward earnings to price, ranked within GICS sectors. For quality, we use an equal-weighted combination of return on equity and debt-to-assets. For momentum, we use the previous twelve-minus-one month return. Each of the four underlying signals is first rank–standardized between -1 and $+1$. The signals are then combined into a single multifactor score. We do not aim to construct the best multifactor score, but rather a simplified score that is representative of common choices and implementations of multifactor investment strategies.

Table 1 presents the stand-alone results from conducting portfolio sorts on the sustainability scores; quality, momentum, and value signals; and the multifactor score. At the end of each month, we sort our stock universe into quintile portfolios and present the return spread between the top and bottom quintile portfolios earned from holding the portfolios for one month. Panel A of Table 1 presents standard portfolio statistics, and panel B presents the correlation between these top-minus-bottom portfolio returns. We can observe that the sustainability scores

Table 1. Stand-Alone Performance of Single-Factor and Multifactor Quintile Portfolios

A. Performance statistics

	Carbon footprint	ESG	SDG	Quality	Value	Momentum	QMV
Mean return (%)	–0.03	1.00	0.58	6.29	6.71	7.67	11.35
Volatility (%)	9.07	4.61	6.27	9.73	12.78	17.87	14.00
Sharpe ratio	0.00	0.22	0.09	0.65	0.52	0.43	0.81
Mean CAPM alpha (%)	0.86	1.80	0.15	7.63	5.77	11.22	13.28
Alpha volatility (%)	8.80	4.50	6.17	9.26	12.58	16.37	13.41
Alpha ratio	0.10	0.40	0.02	0.82	0.46	0.69	0.99
Beta	–0.13	–0.08	0.07	–0.19	–0.19	–0.46	–0.26
Turnover (%)	217.5	145.3	99.0	217.2	340.7	643.8	511.9

B. T-B return correlations

SDG	–43%					
ESG	–11%	20%				
Quality	–18%	–12%	7%			
Value	3%	9%	–8%	2%		
Momentum	12%	–13%	3%	55%	–27%	
QMV	1%	–16%	8%	78%	27%	78%

Notes. This table presents the univariate top-minus-bottom (T-B) portfolio statistics of the sustainability measures, investment factors, and multifactor portfolio. The sample for ESG is from December 2009 to December 2022. The sample for all other results is from December 1989 to December 2022. Stocks are sorted based on each characteristic into quintile portfolios that are rebalanced monthly and held for one month. Panel A reports the annualized performance statistics of the T-B portfolio. Alpha is calculated by regressing the T-B portfolio return on the market return in excess of the risk-free rate. Turnover is the annualized one-way portfolio turnover (e.g., a value of 2,400% per year corresponds to fully replacing the top and bottom portfolio each month). Panel B reports the correlation of the T-B portfolio return series.
ESG = environmental, social, and corporate governance; QMV = quality, momentum, and value multifactor strategy; SDG = sustainable development goals; CAPM = capital asset pricing model.

tend to have low top-minus-bottom returns and alphas over the market return, while we see that the common quality, momentum, and value factors have significant capital asset pricing model (CAPM) alphas with high Sharpe ratios. From a correlation perspective, we observe consistent and positive correlation among quality, momentum, and value, while the correlation between sustainability scores and QMV is close to zero. This highlights the differentiating nature that these selected sustainability measures can have in contrast to typical measures of expected returns.

As for expected risk, we use a standard variance-covariance matrix (Σ) that follows a latent factor model approach where we apply PCA with 20 components to the sample variance-covariance matrix estimated using 60 months of daily returns data. We use five-day overlapping returns to account for market asynchronicity (Burns, Engle, and Mezrich 1998; Martens and Poon 2001).

Empirical Results

Mean-Variance-Sustainability Frontier.

Before constructing the fully optimized portfolios across an extended period, we first compute an ex-ante view on expected return, risk, and sustainability. Consequently, the traditional *2D efficient frontier* between risk and return transforms into a *3D efficient surface*. Figure 2 presents the ex-ante 3D efficient surface among tracking error, expected return, and benchmark-relative carbon footprint as of December 2022. The surface is colored from green (lower carbon footprint than benchmark) to magenta (higher carbon footprint than benchmark). We additionally plot the simple 2D mean-variance efficient frontier (solid black line). We observe the first important result, which is the divergence between the performance of the "green" region (high values of carbon footprint reduction) and the "magenta" region. In the green region, for a given tracking error, higher expected returns require sacrificing carbon footprint reduction (i.e., the surface is coming toward the reader). In the magenta region, for a given tracking error, we typically observe that the carbon footprint reduction remains constant as the expected return increases (i.e., the surface is shaped like a canonical mean-variance efficient frontier). This result highlights how the expected return-tracking error efficient frontiers change for a given level of carbon footprint reduction.

While Figure 2 depicts a 3D surface, Figure 3 presents the ex-ante risk-return-sustainability efficient surfaces as a collection of topographical lines for both carbon footprint and SDG scores. Panel A shows the standard "risk-return efficient frontiers," where each additional frontier away from the traditional "maximum risk-return efficient frontier" corresponds to a higher sustainability target.[9] Panel B presents the "risk-sustainability efficient frontiers"[10] where each additional frontier corresponds to a higher expected return target. In general, for both carbon footprint reduction and SDG improvement, as the desired sustainability or expected return goals increase, the achievable efficient frontiers move further away from the maximal risk-return or risk-sustainability efficient frontiers, respectively. There is thus room to reduce the portfolio carbon footprint without incurring significant tracking error increases or expected return decreases. For example, at a tracking error of 2%, there is effectively no expected return difference between 0% reduction and 60% reduction. However, for the SDG score improvement, any increase above the benchmark level typically requires sacrificing expected return or increasing ex-ante tracking error.

The results in Figures 2 and 3 demonstrate how, at lower tracking errors, one typically needs to sacrifice expected return if one wants to meaningfully improve portfolio sustainability. As the tracking error increases, the upper bound of the risk-return-sustainability efficient surface tends to be closer to the maximum risk-return or risk-sustainability efficient frontiers, reflecting the larger available opportunity set. An alternative way of viewing this is through the lens of the minimum ex-ante tracking error required to implement a given carbon footprint reduction. For example, in December 2022, a 50% carbon footprint reduction required a minimum ex-ante tracking error of around 0.75%, while a 70% carbon footprint reduction required a minimum ex-ante tracking error of 1.75%. Ignoring any correlation between carbon footprint and expected returns, if you have a tracking error budget of 1.00% and you use 0.75% of your risk budget to achieve the desired carbon footprint reduction, there is less risk budget available to take expected return exposure. However, if your tracking error budget is 5.00%, you have relatively more risk budget available to increase your exposure toward expected returns and thus are more likely to be able to jointly satisfy your sustainability and expected return objectives.

Thus, while the multi-objective optimization framework we propose and implement is generalizable

Figure 2. Ex-Ante Mean-Variance-Carbon Footprint Reduction Efficient Surface

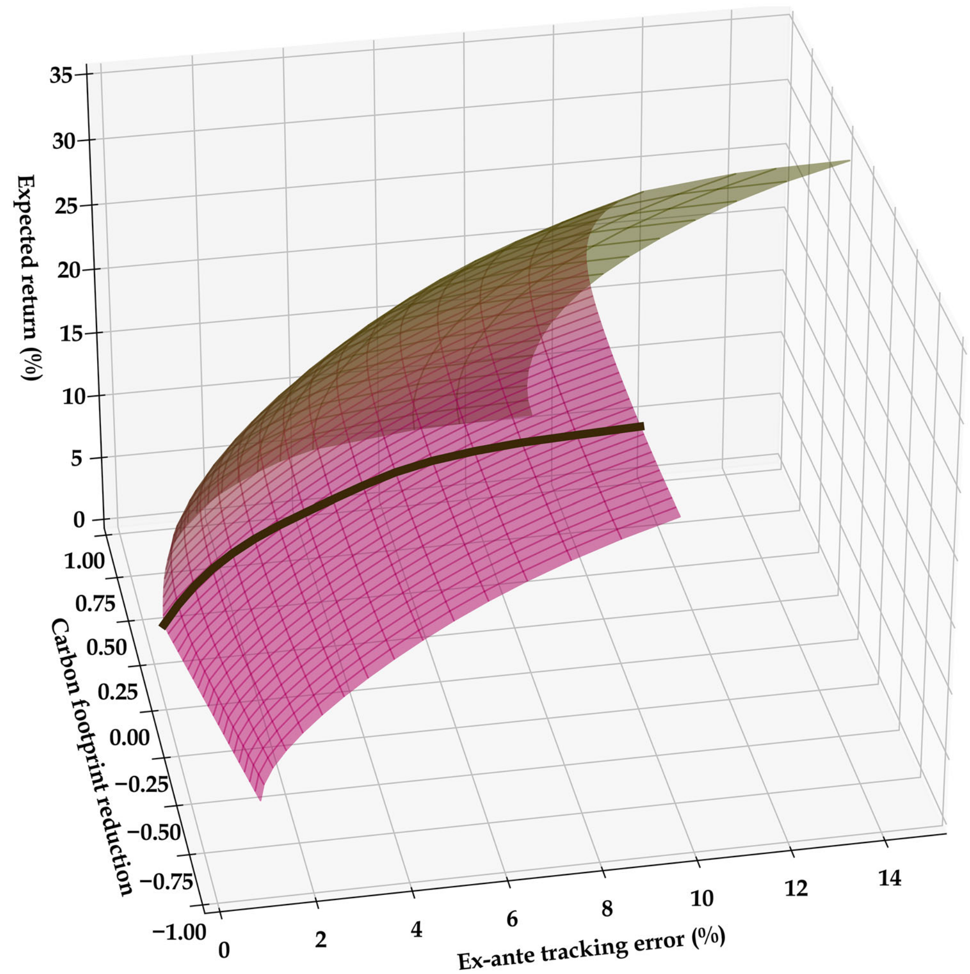

Notes: This figure plots the ex-ante expected return-tracking error-sustainability surface for carbon footprint reduction. The solid black line corresponds to the ex-ante expected return-tracking error efficient frontier. The surface is shaded based on the y-axis variable (carbon footprint reduction relative to the benchmark), where green corresponds with a higher reduction and magenta with a lower reduction. This surface was calculated using data as of December 2022.

across all portfolios, it is of particular relevance for lower–tracking error portfolios that wish to achieve ambitious sustainability targets. Hence, our empirical use cases will focus on low–tracking error portfolios.

Reducing Carbon Footprint. A common sustainability objective of an investment portfolio is to reduce the carbon footprint relative to some benchmark, with the aim of steering the portfolio away from carbon-emitting companies. A basic way to achieve this objective is to enforce a portfolio constraint, such that the portfolio's carbon footprint must always be at least y% better than the benchmark. Although such an approach will guarantee adherence to this requirement, it can lead to suboptimal performance.

A constraint-based approach implies a time-varying carbon tax on expected return, as when a constraint on carbon footprint becomes more binding the optimization algorithm will impose a larger tax on stocks with higher carbon footprints; see Blitz and Hoogteijling (2022). Therefore, we propose the usage of security-level carbon footprint in the objective function as an alternative mechanism for reducing the portfolio's overall carbon footprint, while jointly considering the risk versus expected return versus sustainability tradeoff. Such an approach applies a more stable tax on expected return and provides greater

Figure 3. Ex-Ante Mean-Variance-Sustainability Efficient Frontiers

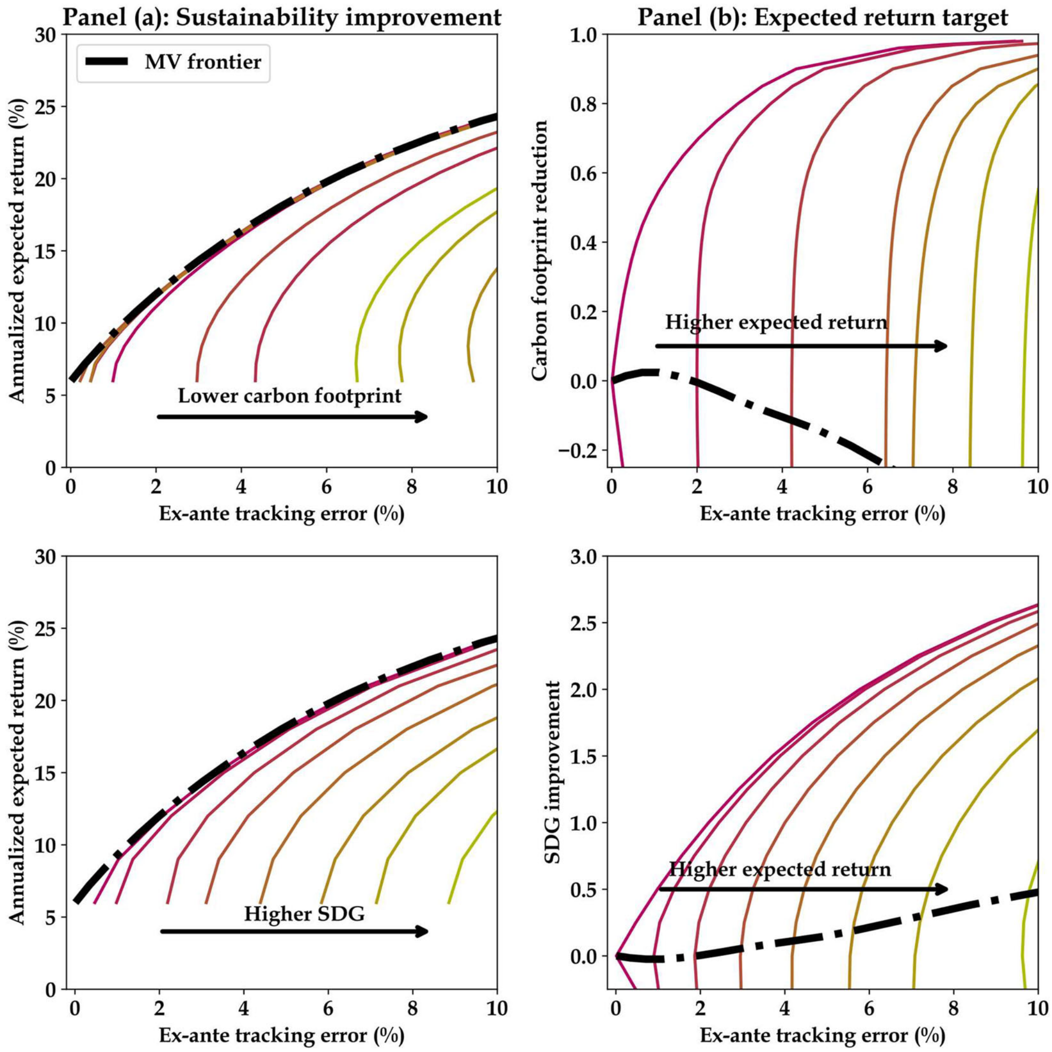

Notes: This figure plots ex-ante expected return-tracking error-sustainability frontiers for benchmark-relative carbon footprint reduction (top row) and sustainable development goals (SDG) improvement (bottom row). Panel A shows the efficient frontiers when changing the sustainability target of the portfolio. Panel B plots the sustainability measure on the y-axis for different levels of expected return. These plots were calculated using data as of December 2022.

scope for the optimization algorithm to trade off expected return and carbon footprint.

To evaluate the practical implications of both approaches, we run a series of simulations where we alter selected portfolio-construction parameters to explore the impact on a portfolio's carbon footprint via constraint and optimization approaches. We run the following simulations:

1. Unconstrained simulation (denoted UC)
2. Constrained simulation with a carbon constraint at y% (denoted Cy)
3. Unconstrained simulation with a carbon metric in the objective function (denoted P)[11]
4. Constrained simulation with a carbon metric in the objective function and a carbon constraint at y% (denoted PCy)

Figure 4 depicts the portfolio's carbon footprint relative to the benchmark under four different scenarios. The unconstrained portfolio has a carbon exposure

Figure 4. Carbon Footprint Reduction Relative to MSCI World under Different Optimization Scenarios

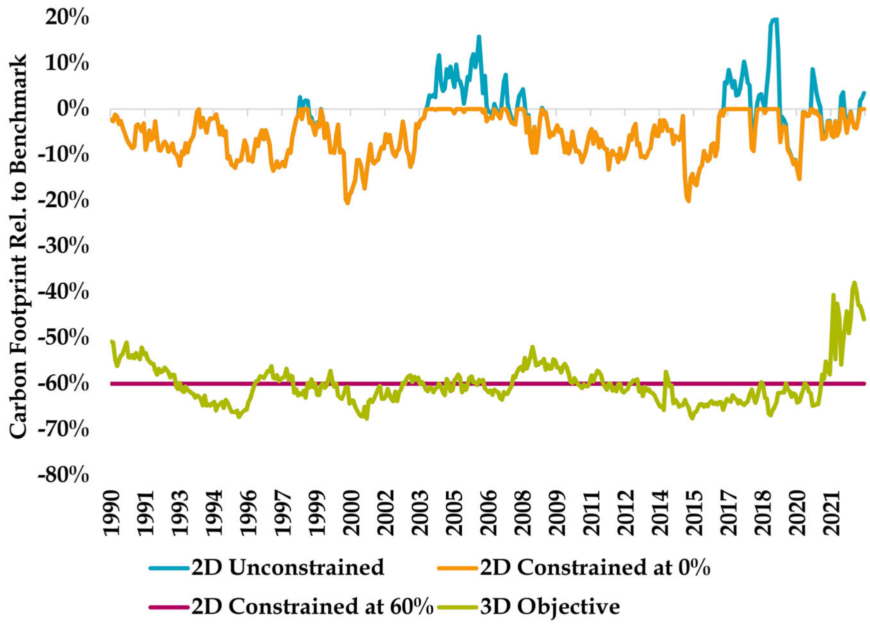

Notes: This figure plots the percentage improvement of the portfolio's carbon footprint over the MSCI World carbon footprint using different 2D and 3D portfolio-construction approaches. We report results for a portfolio with a tracking error target of 0.5%.

that deviates around 0% versus the benchmark, which is expected given the lower tracking error of the portfolio. The portfolio with a carbon constraint at 0% (i.e., better than benchmark) produces the same result as the unconstrained case, except the carbon footprint can only be lower than the benchmark's footprint. This result demonstrates the time-varying nature of the carbon footprint of a portfolio. Such a constraint is not always binding and thus is relatively "cheap" to implement from an expected return perspective. This is contrasted directly with the 60% carbon constraint, where the portfolio's carbon exposure has a static exposure of −60% versus the benchmark and thus is a case where this constraint is always binding. Finally, the carbon metric in the objective function scenario highlights how the portfolio can achieve an average carbon footprint reduction of 60%, but in a time-varying nature. Thus, the portfolio optimizer has more flexibility to trade off carbon footprint reduction with risk and return objectives.

The results in Figure 4 can also be linked to underlying economic phenomena. In recent years, oil and gas stocks have had strong price momentum while still being cheap from a valuation perspective and thus have become attractive from a factor perspective. Hence, under the 3D investing approach, the optimizer elects to purchase these stocks, which thus results in an increase in the carbon footprint of the portfolio. This can then be contrasted with earlier periods, such as 2012 to 2020, when such stocks were relatively unattractive from a factor perspective and thus the objective term approach can produce a larger-than-average reduction of the portfolio's carbon footprint.

The differences between the constrained optimization and 3D investing approach are generally driven by the time-series variation in expected returns. A stock's carbon footprint does not significantly change month-to-month, but expected return forecasts can. Thus, at each monthly rebalance the 3D investing approach is trading off expected returns and carbon footprint.

Table 2 presents the detailed performance statistics over the December 1989 to December 2022 period. First, we note how in both panel A and panel B, applying constraints for carbon footprint reduction reduces gross outperformance and hence IR (as tracking error is relatively constant). Second, in both

Table 2. Portfolio Simulation Results with Different Carbon Footprint Construction Approaches

A. Ex-post tracking error target 0.5%

	2D			3D	
	UC	C0	C60	P	PC60
Gross outp. (%)	0.70	0.70	0.57	0.61	0.56
Tracking error (%)	0.53	0.53	0.52	0.53	0.52
Information ratio	1.32	1.32	1.10	1.15	1.08
Turnover one-way ann. (%)	30.18	30.19	34.38	32.66	33.98
Carbon footprint reduction (%)	−4.4	−5.7	−59.1	−59.2	−60.9
Ann. alpha (%)	0.44	0.45	0.42	0.44	0.40
Alpha t-stat	(5.25)	(5.28)	(5.19)	(5.38)	(5.03)
Mkt-RF t-stat	(−2.14)	(−2.14)	(−3.26)	(−3.36)	(−3.28)
Quality t-stat	(4.65)	(4.65)	(3.55)	(4.29)	(3.51)
Value t-stat	(1.49)	(1.47)	(0.80)	(0.95)	(0.82)
Momentum t-stat	(5.68)	(5.68)	(4.77)	(4.69)	(4.71)
Size t-stat	(0.04)	(0.04)	(-0.05)	(-0.04)	(0.04)
Carbon footprint coeff.	−0.20	−0.20	−0.73	−0.66	−0.73
Carbon footprint t-stat	(−1.01)	(−0.99)	(−2.37)	(−2.32)	(−2.42)
R-squared (%)	41.6	41.6	28.1	31.0	27.1

B. Ex-post tracking error target 1.0%

	2D			3D	
	UC	C0	C75	P	PC75
Gross outp. (%)	1.30	1.30	1.00	1.06	0.98
Tracking error (%)	0.97	0.97	0.96	0.98	0.96
Information ratio	1.34	1.34	1.04	1.08	1.02
Turnover one-way ann. (%)	47.81	47.85	53.75	51.44	53.67
Carbon footprint reduction (%)	−1.4	−6.6	−74.1	−73.2	−75.0
Ann. alpha (%)	0.79	0.80	0.61	0.66	0.60
Alpha t-stat	(5.08)	(5.16)	(4.11)	(4.32)	(4.05)
Mkt-RF t-stat	(−1.88)	(−1.92)	(−2.45)	(−2.69)	(−2.44)
Quality t-stat	(5.55)	(5.56)	(4.89)	(5.45)	(4.86)
Value t-stat	(1.08)	(1.01)	(1.02)	(0.78)	(0.91)
Momentum t-stat	(5.55)	(5.54)	(4.57)	(4.47)	(4.45)
Size t-stat	(2.33)	(2.30)	(2.55)	(2.50)	(2.56)
Carbon footprint coeff.	−0.25	−0.25	−1.32	−1.24	−1.34
Carbon footprint t-stat	(−0.71)	(−0.72)	(−2.34)	(−2.36)	(−2.36)
R-squared (%)	40.7	40.7	28.8	31.3	28.1

Notes. This table presents the performance and spanning regression results for fully invested long-only portfolios optimized using a multifactor expected return target, variance-covariance matrix, and either a constraint on benchmark relative carbon footprint (2D) or directly in the objective function (3D). Our sample runs from December 1989 to December 2022 using constituents of the MSCI World universe. Portfolios are rebalanced monthly. Panel (A) targets a 0.5% tracking error portfolio. Panel (B) targets a 1.0% tracking error portfolio. The spanning regression regresses the gross outperformance of the optimized portfolio on the outperformance of the top-minus-bottom portfolios of the different factors (Quality, Value, Momentum, Size, carbon footprint). UC denotes unconstrained. Cx denotes a constraint at x% lower than the benchmark carbon footprint. P denotes a term in the objective function. PCy denotes a term in the objective function and a constraint at y% lower than the benchmark carbon footprint. R-squared is calculated in a regression excluding the SI-regressor.

scenarios, the 3D-objective function approach (denoted P) outperforms the 2D-constraint approach (C60%/C75%). The 3D-objective function approach delivers an increased gross outperformance and a lower one-way turnover, while maintaining a similar tracking error and carbon footprint reduction. These results are reflected in the spanning alpha regressions we run over the underlying factors we use to construct the portfolio (QMV) as well as log market capitalization (size) and a factor constructed from carbon footprint scores. Across all cases, there are similar exposures to the targeted factors, highlighting how the 3D investing incorporates the carbon footprint objective while maintaining similar factor exposures in a more efficient manner.

We also explore the combination of the objective function approach and the constraint approach. We find that such an approach underperforms the isolated approaches, since it achieves larger carbon footprint reductions in both cases but often overshoots the targeted reduction level. As we are operating with ambitious carbon reduction targets, any increase can significantly impact the gross outperformance of the portfolio. Nevertheless, such a combination highlights how one can ensure a base level of reduction, while also doing better than this target if the risk-return-sustainability tradeoff is appropriate.

We have denoted the "2D constraint" as that where the objective is to maximize expected returns and minimize risk subject to some carbon footprint constraint. It is also possible to frame this as maximizing expected returns and minimizing carbon footprint subject to a maximum tracking error constraint. In unreported robustness tests, we find that the 3D investing approach is superior to this "2D TE constraint" approach, but this "2D TE constraint" approach is superior to the "2D constraint" approach. This suggests that our approach benefits relatively more from introducing the carbon footprint term into the objective function and allowing for this tradeoff to occur.

Improving SDG Score. An alternative application of our proposed framework is improving a portfolio's exposure to a positive measure of sustainability. We use the Robeco SDG scores, which assign a discrete score between −3 (poor) and +3 (good) on how a company is contributing to the UN SDG goals. Instead of adding a negative coefficient on a stock's carbon footprint, we can add a reward/tax "refund" to the objective function, which encourages the optimization algorithm to hold stocks with positive SDG scores.

Figure 5 presents the portfolio's SDG score relative to the benchmark under five different scenarios, as in Figure 4. We find qualitatively similar results here as for the carbon footprint reduction exercise. Importantly for the objective function scenario, we

Figure 5. SDG Score Improvement over MSCI World under Different Optimization Scenarios

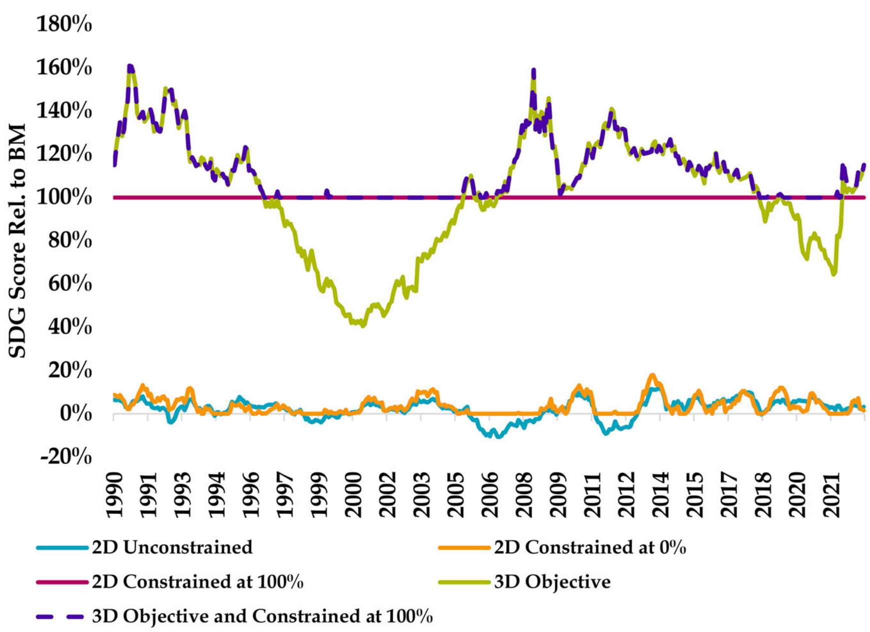

Notes: This figure plots the percentage improvement of the portfolio's sustainable development goals (SDG) score over the MSCI World SDG score using different 2D and 3D portfolio-construction approaches. We report results for a portfolio with a tracking error target of 1.0%.

observe a time-variation in the portfolio's SDG score around the constrained average of 100% improvement. This result again highlights the dynamic nature of the 3D-objective function approach to targeting sustainability objectives, as opposed to the fixed nature of a constraint.

Table 3 presents the detailed performance statistics over the December 1989 to December 2022 period. Following the same notation as in Table 2, we document quantitatively similar results as in the carbon footprint scenario. In panel A, the objective function scenario (P) outperforms the constraint scenario (C70%) across most metrics: higher gross outperformance (0.49% versus 0.46%), higher gross IR (0.98 versus 0.88), and lower one-way turnover (26.7% versus 31.4%). In panel B, for the 1.0% tracking error scenario, we find comparable results, with the objective function scenario outperforming the constraint scenario.

Implications and Outlook

The Impact of Evolving Sustainability Data.

Through these two examples, we have shown how using an SI metric in the objective function is generally superior to simply imposing a fixed constraint at the portfolio level. The 3D investing framework is generalizable to any sustainability metric that can be expressed as a discrete or continuous series (if it is ordinal). The empirical examples correspond to commonly explored sustainability measures in investment management, and we observe how there can be structurally different outcomes dependent on the measure itself (in conjunction with expected return and risk models used). To this end, the broader question remains: What is the best way to construct portfolios that satisfy sustainability desires going forward? Figure 6 presents the MSCI World value-weighted carbon footprint, ESG, and SDG scores over time. We observe time-variation in the benchmark sustainability scores, but it is not always trending in one direction. Such changes have significant implications for how portfolios that target these measures should be constructed.

In particular, targeting relative improvements of MSCI ESG scores over the benchmark becomes increasingly challenging, as the average benchmark score has increased from 2015 to 2022.[12] It is not strictly the increasing benchmark score that is problematic, but rather the skewness and concentration of scores which have an upper bound of 10.0. For example, suppose that the benchmark score is 8.0 and a relative improvement of 20% is desired. This corresponds to a portfolio score of 9.6, which would require holding a large number of stocks with an ESG score of 10.0. Thus, the portfolio construction is going to be heavily driven by ESG scores, while expected return and risk concerns become secondary. By using the objective function approach, the impacts of such benchmark changes can be less impactful, and the optimization algorithm will be able to better trade off expected return, risk, and sustainability objectives. However, in both scenarios, the metric of choice has considerable influence on the optimization algorithm, and thus it is important to select measures that have desirable properties when targeting them in a portfolio-optimization algorithm.

When to Use One Approach Versus the Other.

Portfolio constraints are the most common way to ensure compliance with sustainability objectives. Another popular way to ensure portfolio sustainability compliance is via universe exclusion, for example, excluding names that are considered "sin stocks." In this article, 3D investing emerges as an effective way to improve portfolio sustainability that offers Pareto optimality and more flexibility. This result, however, does not mean that the traditional constraint- and exclusion-based approaches are without merit and should be discarded.

A constraint-based approach to portfolio sustainability is suitable when one wants portfolio-level sustainability goals to be achieved *at all times*. Similarly, an exclusion-based approach ensures that individual stock–level sustainability goals are achieved *at all times*.[13] This is because constraints and exclusions are hard criteria and, thus, the portfolio optimizer must satisfy these objectives for all proposed portfolios. On the other hand, an optimization-based approach as discussed in this paper represents a soft criterion. It is more flexible, as it enables the optimizer to trade off among sustainability, risk, and expected return. This tradeoff ensures a superior sustainability profile versus those portfolios without sustainability in the objective *on average*, but it does not guarantee a specific sustainability profile at any given point in time.

These hard and soft approaches both have their use cases in portfolio construction. If one wants to always ensure a certain level of guaranteed sustainability profile or ensure that certain names will not be held in the portfolio, constraints and exclusions should be used, respectively. On the other hand, if the portfolio manager wants to achieve a better

Table 3. Portfolio Simulation Results with Different SDG Construction Approaches

A. Ex-post tracking error target 0.5%

	2D			3D	
	UC	C0	C70	P	PC70
Gross outp. (%)	0.70	0.70	0.46	0.49	0.42
Tracking error (%)	0.53	0.53	0.52	0.50	0.51
Information ratio	1.32	1.32	0.88	0.98	0.82
Turnover one-way ann. (%)	30.2	30.3	31.4	26.7	30.1
SDG improvement (%)	0.2	3.1	69.6	70.1	78.6
Ann. alpha (%)	0.44	0.44	0.30	0.31	0.28
Alpha t-stat	(5.19)	(5.18)	(3.38)	(4.14)	(3.15)
Mkt-RF t-stat	(−2.23)	(−2.21)	(−2.60)	(−2.91)	(−2.64)
Quality t-stat	(4.67)	(4.65)	(5.72)	(5.35)	(5.46)
Value t-stat	(1.48)	(1.49)	(1.10)	(0.84)	(1.04)
Momentum t-stat	(5.56)	(5.59)	(1.22)	(3.24)	(0.99)
Size t-stat	(−0.01)	(0.03)	(0.88)	(−0.24)	(0.82)
SDG coeff.	0.41	0.43	1.39	1.58	1.56
SDG t-stat	(1.32)	(1.40)	(4.42)	(5.61)	(5.10)
R-squared (%)	41.6	41.5	17.0	25.1	13.9

B. Ex-post tracking error target 1.0%

	2D			3D	
	UC	C0	C100	P	PC100
Gross outp. (%)	1.30	1.29	0.94	0.98	0.87
Tracking error (%)	0.97	0.97	1.02	0.96	1.01
Information ratio	1.34	1.33	0.92	1.02	0.86
Turnover one-way ann. (%)	47.8	48.0	51.3	44.3	49.5
SDG improvement (%)	−0.2	3.4	99.5	101.4	111.9
Ann. alpha (%)	0.78	0.77	0.59	0.59	0.56
Alpha t-stat	(5.03)	(4.95)	(3.62)	(4.12)	(3.45)
Mkt-RF t-stat	(−2.02)	(−1.98)	(−2.66)	(−3.22)	(−2.78)
Quality t-stat	(5.62)	(5.60)	(5.83)	(5.85)	(5.51)
Value t-stat	(1.11)	(1.10)	(1.01)	(0.54)	(1.03)
Momentum t-stat	(5.47)	(5.50)	(1.55)	(3.61)	(1.36)
Size t-stat	(2.38)	(2.41)	(2.87)	(1.89)	(2.72)
SDG coeff.	0.69	0.73	2.61	2.87	2.89
SDG t-stat	(1.28)	(1.35)	(3.74)	(5.09)	(4.25)
R-squared (%)	40.7	40.9	18.1	27.5	15.8

Notes. This table presents the performance and spanning regression results for fully invested long-only portfolios optimized using a multifactor expected return target, variance-covariance matrix, and either a constraint on benchmark relative SDG (2D) or directly in the objective function (3D). Our sample runs from December 1989 to December 2022 using constituents of the MSCI World universe. Portfolios are rebalanced monthly. Panel A targets a 0.5% tracking error portfolio. Panel B targets a 1.0% tracking error portfolio. The spanning regression regresses the gross outperformance of the optimized portfolio on the outperformance of the top-minus-bottom portfolios of the different factors (quality, value, momentum, size, SDG). Cx denotes a constraint at x% higher than the benchmark SDG score. P denotes a term in the objective function. PCy denotes a term in the objective function and a constraint at y% higher than the benchmark SDG score. R-squared is calculated in a regression excluding the SI regressor.
SDG = sustainable development goals; SI = sustainable investing; UC = unconstrained.

sustainability profile on average but, when conditions are right, may go for higher expected returns instead of a better sustainability profile, then the multi-objective optimization approach is appropriate. For example, when oil and gas companies are so cheap that their expected future returns are very high, the optimization-based approach allows for temporary sacrifice of sustainability for higher expected return.

Figure 6. MSCI World Value-Weighted Carbon Footprint, ESG, and SDG Scores over Time

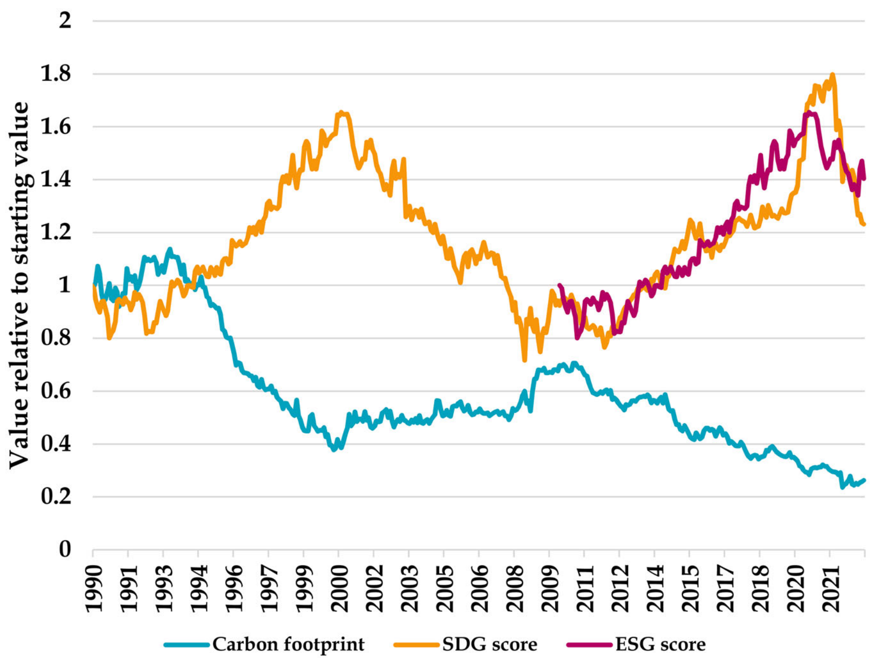

Notes: This figure plots the value-weighted carbon footprint, sustainable development goals (SDG) score, and environmental, social, and corporate governance (ESG) score, for the MSCI World benchmark. The sample for carbon footprint and SDG runs from December 1989 to December 2022. The sample for ESG runs from December 2009 to December 2022. Each series is reported as the ratio of the current value to the first value in the time series.

Finally, with new regulatory developments, such as the European Green Deal, the necessity to integrate alternative objectives into the investment paradigm will continue to grow. Our methodology provides a framework that is compatible with these regulatory environments and allows investors to make more transparent decisions around the integration of sustainability objectives alongside explicit risk and return considerations. Ultimately, at the heart of our framework is the stock-selection model wherein we select stocks with the highest expected returns. This continues to remain front and center in the investment paradigm, and our framework allows for a more flexible approach for incorporating alternative objectives, such as sustainability.

Conclusions

Investing has historically been a multidimensional endeavor, but portfolio-construction approaches have most often been considered 2D. Sustainable investing is the latest example of multi-objective investing in an extensive line of examples. We demonstrate a 3D investing framework that results in the "best possible" solution when jointly considering more than two portfolio objectives. Historical simulations highlight the superiority of this approach versus the traditional constraint-based approach for sustainable investing (in the context of carbon footprint reduction and attaining higher SDG scores). 3D investing achieves, on average, higher sustainability characteristics and expected returns when compared with a pure constraint-based approach.

Notably, constraints are not without their use in sustainable investing. In practice, a mixed approach, with a nonbinding sustainability constraint in conjunction with incorporating the sustainability criteria into the objective function, may be preferred. Such an approach guarantees a basic level of sustainability targeting while allowing the optimization algorithm to make opportunistic tradeoffs among return, risk, and sustainability. If aggressive sustainability objectives are desired, the 3D investing approach where sustainability is explicitly targeted alongside alpha and risk is optimal.

Editor's Note
Submitted 20 December 2023
Accepted 22 March 2024 by William N. Goetzmann

Notes

1. The reduction of the carbon footprint of a portfolio (as given by CO_2 emissions scaled by some measure of a company's size) is one of the most common sustainability objectives. For examples, see Andersson, Bolton, and Samama (2016), Hao, Soe, and Tang (2018), Görgen, Jacob, and Nerlinger (2020), Roncalli et al. (2020), Bender et al. (2020), Benz et al. (2020), Bolton and Kacperczyk (2021), Bolton, Kacperczyk, and Samama (2022), and Kolle et al. (2022).

2. This is the transaction cost at which the outperformance of the portfolio would be zero.

3. For practical considerations on the sustainability metric μ_{SI}, see Chen and Mussalli (2020).

4. Prior to 2001, we use constituents of the FTSE Development Markets index as a proxy for MSCI World constituents.

5. In unreported results, we find qualitatively similar results when using raw scope 1 and scope 2 emissions and carbon intensity (scaled by revenue instead of EVIC) and when incorporating scope 3 emissions into carbon footprint.

6. We additionally use the data-simulation approach of Blitz and Hoogteijling (2022) to produce a longer history of carbon footprint data and SDG data. Note that any potential forward information leakage is of little concern as we are comparing two portfolio-construction approaches on the same data. We aim to illustrate the broad application of our methodology on a representative set of sustainability data.

7. Examples can be found at the Robeco SI open-access page: https://www.robeco.com/en-int/sustainable-investing/how-do-companies-and-countries-score-on-sustainability.

8. We use the same benchmark, MSCI World, when constructing portfolios and evaluating financial and sustainability objectives.

9. We define the "risk-return efficient frontier" to be the traditional efficient frontier for a constant level of portfolio sustainability and the "maximum risk-return efficient frontier" to be the risk-return efficient frontier when sustainability considerations are dropped. That is, the maximum risk-return efficient frontier is the efficient frontier in the traditional sense.

10. Similarly, we define the "risk-sustainability efficient frontier" to be the risk versus sustainability efficient frontier for a constant level of expected return and the "maximum risk-sustainability efficient frontier" to be the risk-sustainability efficient frontier when expected return considerations are dropped.

11. For the carbon footprint scenario, we use coefficients of (–0.016 and –0.020) and for the SDG scenario we use coefficients of (2.0 and 1.5) for the 0.5%/1.0% tracking error targets, respectively.

12. For more discussions on the skewed distribution of MSCI ESG scores, see Chen, von Behren, and Mussalli (2021).

13. This approach is not extensively discussed in this paper, as it is common and straightforward.

References

Alessandrini, F., and E. Jondeau. 2021. "Optimal Strategies for ESG Portfolios." *The Journal of Portfolio Management* 47 (6): 114–138. doi:10.3905/jpm.2021.1.241.

Andersson, M., P. Bolton, and F. Samama. 2016. "Hedging Climate Risk." *Financial Analysts Journal* 72 (3): 13–32. doi:10.2469/faj.v72.n3.4.

Ballestero, E., M. Bravo, B. Pérez-Gladish, M. Arenas-Parra, and D. Plà-Santamaria. 2012. "Socially Responsible Investment: A Multicriteria Approach to Portfolio Selection Combining Ethical and Financial Objectives." *European Journal of Operational Research* 216 (2): 487–494. doi:10.1016/j.ejor.2011.07.011.

Bender, J., C. He, C. Ooi, and X. Sun. 2020. "Reducing the Carbon Intensity of Low Volatility Portfolios." *The Journal of Portfolio Management* 46 (3): 108–122. doi:10.3905/jpm.2020.46.3.108.

Benz, L., A. Jacob, S. Paulus, and M. Wilkens. 2020. "Herds on Green Meadows: The Decarbonization of Institutional Portfolios." *Journal of Asset Management* 21 (1): 13–31. doi:10.1057/s41260-019-00147-z.

Berg, F., J. F. Kölbel, and R. Rigobon. 2022. "Aggregate Confusion: The Divergence of ESG Ratings." *Review of Finance* 26 (6): 1315–1344. doi:10.1093/rof/rfac033.

Bilbao-Terol, A., M. Arenas-Parra, and V. Cañal-Fernández. 2012. "A Fuzzy Multi-Objective Approach for Sustainable Investments." *Expert Systems with Applications* 39 (12): 10904–10915. doi:10.1016/j.eswa.2012.03.034.

Blitz, D., and T. Hoogteijling. 2022. "Carbon-Tax-Adjusted Value." *The Journal of Portfolio Management* 48 (5): 121–137. doi:10.3905/jpm.2022.1.343.

Bolton, P., and M. Kacperczyk. 2021. "Do Investors Care about Carbon Risk?" *Journal of Financial Economics* 142 (2): 517–549. doi:10.1016/j.jfineco.2021.05.008.

Bolton, P., M. Kacperczyk, and F. Samama. 2022. "Net-Zero Carbon Portfolio Alignment." *Financial Analysts Journal* 78 (2): 19–33. doi:10.1080/0015198X.2022.2033105.

Boudt, K., J. Cornelissen, and C. Croux. 2013. "The Impact of a Sustainability Constraint on the Mean-Tracking Error Efficient Frontier." *Economics Letters* 119 (3): 255–260. doi:10.1016/j.econlet.2013.03.020.

Burns, P., R. F. Engle, and J. J. Mezrich. 1998. "Correlations and Volatilities of Asynchronous Data." *The Journal of Derivatives* 5 (4): 7–18. doi:10.3905/jod.1998.408000.

Busch, T., M. Johnson, and T. Pioch. 2020. "Corporate Carbon Performance Data: Quo Vadis?" *Journal of Industrial Ecology* 26 (1): 350–363. doi:10.1111/jiec.13008.

Calvo, C., C. Ivorra, and V. Liern. 2015. "Finding Socially Responsible Portfolios Close to Conventional Ones." *International Review of Financial Analysis* 40: 52–63. doi:10.1016/j.irfa.2015.03.014.

Calvo, C., C. Ivorra, and V. Liern. 2016. "Fuzzy Portfolio Selection with Non-Financial Goals: Exploring the Efficient Frontier." *Annals of Operations Research* 245 (1-2): 31–46. doi:10.1007/s10479-014-1561-2.

Chen, M., and G. Mussalli. 2020. "An Integrated Approach to Quantitative ESG Investing." *The Journal of Portfolio Management* 46 (3): 65–74. doi:10.3905/jpm.2020.46.3.065.

Chen, M., R. Von Behren, and G. Mussalli. 2021. "The Unreasonable Attractiveness of More ESG Data." *The Journal of Portfolio Management* 48 (1): 147–162. doi:10.3905/jpm.2021.1.281.

Coqueret, G., S. Stiernegrip, C. Morgenstern, J. Kelly, J. Frey-Skött, and B. Österberg. 2021. "Boosting ESG-Based Optimization with Asset Pricing Characteristics." SSRN Working Paper, no. 3877242.

Diltz, J. D. 1995. "Does Social Screening Affect Portfolio Performance?" *The Journal of Investing* 4 (1): 64–69. doi:10.3905/joi.4.1.64.

Dorfleitner, G., and S. Utz. 2012. "Safety First Portfolio Choice Based on Financial and Sustainability Returns." *European Journal of Operational Research* 221 (1): 155–164. doi:10.1016/j.ejor.2012.02.034.

Geczy, C. C., R. F. Stambaugh, and D. Levin. 2021. "Investing in Socially Responsible Mutual Funds." *Review of Asset Pricing Studies* 11 (2): 309–351. doi:10.1093/rapstu/raab004.

Gintschel, A., and B. Scherer. 2004. "Currency Reserve Management by Dual Benchmark Optimisation." In *Risk Management for Central Bank Foreign Reserves*, edited by C. Bernadell, P. Cardon, J. Coche, F. X. Diebold, and S. Manganelli, 137–151. Frankfurt am Main, Germany: European Central Bank.

Görgen, M., A. Jacob, and M. Nerlinger. 2020. "Get Green or Die Trying? Carbon Risk Integration into Portfolio Management." *The Journal of Portfolio Management* 47 (3): 77–93. doi:10.3905/jpm.2020.1.200.

Hallerbach, W., H. Ning, A. Soppe, and J. Spronk. 2004. "A Framework for Managing a Portfolio of Socially Responsible Investments." *European Journal of Operational Research* 153 (2): 517–529. doi:10.1016/S0377-2217(03)00172-3.

Hao, B., A. M. Soe, and K. Tang. 2018. "Carbon Risk Integration with Factors." *The Journal of Index Investing* 9 (2): 53–65. doi:10.3905/jii.2018.1.061.

Hautsch, N., and S. Voigt. 2019. "Large-Scale Portfolio Allocation under Transaction Costs and Model Uncertainty." *Journal of Econometrics* 212 (1): 221–240. doi:10.1016/j.jeconom.2019.04.028.

Hirschberger, M., R. E. Steuer, S. Utz, M. Wimmer, and Y. Qi. 2013. "Computing the Nondominated Surface in Tri-Criterion Portfolio Selection." *Operations Research* 61 (1): 169–183. doi:10.1287/opre.1120.1140.

Horan, S. M., E. Dimson, C. Emery, K. Blay, G. Yelton, and A. Agarwal. 2022. "ESG Investment Outcomes, Performance Evaluation, and Attribution." CFA Institute Research Foundation. doi: 10.56227/22.1.14.

Jagannathan, R., and T. Ma. 2003. "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps." *The Journal of Finance* 58 (4): 1651–83. doi:10.1111/1540-6261.00580.

Jorion, P. 2003. "Portfolio Optimization with Tracking-Error Constraints." *Financial Analysts Journal* 59 (5): 70–82. doi:10.2469/faj.v59.n5.2565.

Kinder, P. D., and A. L. Domini. 1997. "Social Screening: Paradigms Old and New." *The Journal of Investing* 6 (4): 12–19. doi:10.3905/joi.1997.408443.

Kolle, J., H. Lohre, E. Radatz, and C. Rother. 2022. "Factor Investing in Paris: Managing Climate Change Risk in Portfolio Construction." *Journal of Investment Management* 20 (4): 35–51.

Ledoit, O., and M. Wolf. 2022. "Markowitz Portfolios under Transaction Costs." University of Zurich, Department of Economics, Working Paper Series, no. 420.

Markowitz, H. 1952. "Portfolio Selection." *The Journal of Finance* 7 (1): 77–91. doi:10.2307/2975974.

Marler, R. T., and J. S. Arora. 2010. "The Weighted Sum Method for Multi-Objective Optimization: New Insights." *Structural and Multidisciplinary Optimization* 41 (6): 853–862. doi:10.1007/s00158-009-0460-7.

Martens, M., and S.-H. Poon. 2001. "Returns Synchronization and Daily Correlation Dynamics between International Stock Markets." *Journal of Banking & Finance* 25 (10): 1805–1827. doi:10.1016/S0378-4266(00)00159-X.

Naber M. 2001. "Catholic Investing: The Effects of Screens on Financial Returns." *The Journal of Investing* 10 (4): 58–65. doi:10.3905/joi.2001.319487.

O'Cinneide, C., B. Scherer, and X. Xu. 2006. "Pooling Trades in a Quantitative Investment Process." *The Journal of Portfolio Management* 32 (4): 33–43. doi:10.3905/jpm.2006.644191.

Pedersen, L. H., S. Fitzgibbons, and L. Pomorski. 2021. "Responsible Investing: The ESG-Efficient Frontier." *Journal of*

Financial Economics 142 (2): 572–597. doi:10.1016/j.jfineco.2020.11.001.

Roncalli, T., T. Le Guenedal, F. Lepetit, T. Roncalli, and T. Sekine. 2020. "Measuring and Managing Carbon Risk in Investment Portfolios." SSRN Working Paper, no. 3681266.

Sandberg, J., C. Juravle, T. M. Hedesström, and I. Hamilton. 2009. "The Heterogeneity of Socially Responsible Investment." *Journal of Business Ethics* 87 (4): 519–533. doi:10.1007/s10551-008-9956-0.

Schmidt, A. B. 2020. "Optimal ESG Portfolios: An Example for the Dow Jones Index." *Journal of Sustainable Finance & Investment* 12 (2): 529–535. doi:10.1080/20430795.2020.1783180.

Shushi, T. 2022. "The Optimal Solution of ESG Portfolio Selection Models that are Based on the Average ESG Score." *Operations Research Letters* 50 (5): 513–516. doi:10.1016/j.orl.2022.07.008.

Stanimirović, I. P., M. L. J. Zlatanović, and M. D. Petković. 2011. "On the Linear Weighted Sum Method for Multi-Objective Optimization." *Facta Universitatis* 26: 49–63.

Steuer, R. E., and S. Utz. 2023. "Non-Contour Efficient Fronts for Identifying Most Preferred Portfolios in Sustainability Investing." *European Journal of Operational Research* 306 (2): 742–753. doi:10.1016/j.ejor.2022.08.007.

Taksar, M., M. J. Klass, and D. Assaf. 1988. "A Diffusion Model for Optimal Portfolio Selection in the Presence of Brokerage Fees." *Mathematics of Operations Research* 13 (2): 277–294. doi:10.1287/moor.13.2.277.

Utz, S., M. Wimmer, M. Hirschberger, and R. E. Steuer. 2014. "Tri-Criterion Inverse Portfolio Optimization with Application to Socially Responsible Mutual Funds." *European Journal of Operational Research* 234 (2): 491–498. doi:10.1016/j.ejor.2013.07.024.

Van Zanten, J. A., and J. Huij. 2022. "ESG to SDG: Do Sustainable Investing Ratings Align with the Sustainability Preferences of Investors, Regulators, and Scientists?." SSRN Working Paper, no. 4186680.

Wu, Q., X. Liu, J. Qin, L. Zhou, A. Mardani, and M. Deveci. 2022. "An Integrated Multi-Criteria Decision-Making and Multi-Objective Optimization Model for Socially Responsible Portfolio Selection." *Technological Forecasting and Social Change* 184: 121977. doi:10.1016/j.techfore.2022.121977.

Xidonas, P., and E. Essner. 2022. "On ESG Portfolio Construction: A Multi-Objective Optimization Approach." *Computational Economics* 63 (1): 21–45. doi:10.1007/s10614-022-10327-6.

Zadeh L. 1963. "Optimality and Non-Scalar-Valued Performance Criteria." *IEEE Transactions on Automatic Control* 8 (1): 59–60. doi:10.1109/TAC.1963.1105511.