JOIM
www.joim.com

Journal Of Investment Management, Vol. 21, No. 3, (2023), pp. 42–60

ASSET ALLOCATION WITH NON-PECUNIARY ESG PREFERENCES: EFFICIENTLY BLENDING VALUE WITH VALUES

Douglas M. Grim[a,b,*], *Giulio Renzi-Ricci*[a,c] *and Anna Madamba*[a,d]

The explosion of interest in ESG investing has yielded several quantitative frameworks that seek to incorporate non-pecuniary ESG preferences into conventional multi-asset portfolio optimization models. In this article, the authors specify an accessible approach that allows investors to simultaneously optimize for both pecuniary preferences (such as systematic, factor, and active risk aversion) and non-pecuniary ESG tastes in a way that avoids "one size fits all" solutions and arbitrary portfolio decisions. Using case studies, they demonstrate that the strength of non-pecuniary desires along with both pecuniary expectations and risk preferences are important determinants of the optimal portfolio choice.

Interest in environmental, social, and governance (ESG) investments has grown dramatically over the last decade. According to the latest report from Global Sustainable Investment Alliance (GSIA), global assets in ESG investments at the end of 2019 surpassed $35 trillion with cumulative growth since 2014 in all major regions included in the report (GSIA, 2020).[1] This growth has coincided with a proliferation of investment products and other strategies available to investors for use when implementing different ESG investment approaches, including a substantial increase in commingled fund options and direct indexing/separately managed account (SMA) offers. At the same time, published research on ESG-related investment topics has increased dramatically (Whelan *et al.*, 2021).

This increase in both attention and implementation options has sparked more industry debate about how to determine whether, or to what extent, investors should include ESG investment strategies in their portfolio. For typical risk/return focused investors who are unable or unwilling to sacrifice financial utility, different ESG strategies are dealt with in a conventional

[a]The Vanguard Group.
[b]Email: douglas_m_grim@vanguard.com.
[c]Email: giulio.renzi.ricci@vanguard.co.uk.
[d]anna_b_madamba@vanguard.com
*Corresponding author.

portfolio optimization framework. An advisor can set investment performance expectations for assets being considered for inclusion in a client's portfolio and input the investor's personal risk preferences for those assets. If the financial return, volatility, and correlation properties of the ESG strategies are attractive enough in comparison to the conventional asset options being considered, then a portfolio optimization model will allocate weight to them without any special treatment.[2]

Some investors derive non-pecuniary social and/or private benefits from owning assets with ESG features they find personally appealing. This taste-based utility they earn can rationally influence the process for choosing an optimal portfolio (e.g., Fama and French, 2007; Dorfleitner and Nguyen, 2017; Gomes, 2020). For example, holding financial assumptions constant, they may earn extra satisfaction from what they believe is altruistic behavior (Lewis and Cullis, 1990) if they take on outsized positions in securities of businesses that are delivering positive societal change by contributing toward a more sustainable and inclusive world through their standard operations.[3] Alternatively, they could earn warm glow benefits (Andreoni, 1989) by avoiding investment in securities of firms that have business operations that do not align with their personal values (e.g., tobacco firms, oil and gas companies). Consequently, determining the prudent portfolio investment mix for these investors should formally consider more than just the risk/return profile of assets and their attitude toward such risks.[4]

Academics and practitioners have struggled to solve this unconventional portfolio choice problem. Existing studies propose approaches that attempt to account for an investor's non-pecuniary preferences, but to our knowledge there is no widely-accepted technique that allows investors to fully integrate pecuniary and personal non-pecuniary motives into a quantitative, multi-asset, portfolio construction process without requiring multiple steps or ad-hoc choices which can lead to suboptimal decisions. Also, there is no commonly agreed upon methodology that investors can use to define their non-pecuniary preferences. We fill this gap in the literature by proposing a method that can be used to build an optimal portfolio for an investor with a combination of pecuniary and non-pecuniary preferences in a single, efficient step using an accessible way of measuring the non-pecuniary benefit. The framework builds upon Aliaga-Diaz *et al.* (2020) by enabling explicit, customized, and intuitive investor choices for ESG taste alongside systematic, factor, and active risk aversion. We produce a few case studies to test it and use different assumptions to assess how varying parameter values impacts the optimal portfolio mix.

Our results are in line with our expectations and suggest that there can be a trade-off between pecuniary and non-pecuniary preferences, along with a positive relationship between the weight of an ESG investment in a portfolio and an investor's level of ESG-driven non-pecuniary utility (hereafter, NPU) for that particular asset. In addition, exploratory interviews with actual investors suggest a range of NPU values and a willingness by a subset to accept ESG strategy underperformance if necessary, indicating that required investor inputs are accessible and can be identified. These findings are consistent with theory and provide the evidence necessary to believe that the framework can be utilized in practice for institutional and retail investors, regardless of their personal motives or ESG preferences.

In the remaining sections, we provide a brief summary of the literature on NPU and asset allocation frameworks, explain our methodology, and present case study results that examine our model's sensitivity to different levels of investor-specific ESG taste, asset-level ESG "quality", and risk aversion.

1 Literature Review

There is a growing body of research that supports the hypothesis that a subset of investors explicitly and rationally consider more than just the pecuniary features of assets when making allocation decisions.[5] This is consistent with evidence that some consumers factor in NPU in other product choice decisions—for example, choosing a car, apparel, food, or a home electricity provider. Researchers have found that the root cause of this NPU is typically induced by altruism, warm-glow, and/or social signaling (e.g., Hartzmark and Sussman, 2019). If someone prefers to invest only in assets that align with their personal values, this can be driven by warm-glow feelings or social signaling (Becker, 1974; Andreoni, 1989; Hong and Kacperczyk, 2009; Riedl and Smeets, 2017). Alternatively, someone preferring to invest in assets they believe will positively impact broader society in some way, often referred to as generating a positive externality (e.g., a business focused on producing low carbon technology to help other businesses reduce carbon emissions), could be generating NPU through a form of altruism (Brodback *et al.*, 2019).

Conventional frameworks for multi-asset portfolio construction, like traditional mean–variance optimization (Markowitz, 1952, 1959), Black–Litterman (Black and Litterman, 1991), and expected utility ("full-scale") optimization (Adler and Kritzman, 2007; Sharpe, 2007) assume that investors seek to maximize their future expected wealth while accounting only for investment-related risks and preferences. As a result, if an investor derives NPU that could impact portfolio choice, their advisor must either design a portfolio construction methodology from scratch or augment one of these conventional methods to account for their client's unique ESG tastes. This has led to greater attention toward developing portfolio construction frameworks that allow for non-pecuniary preferences. We cover this strand of literature in the section that describes our methodology.

2 Incorporating ESG NPU into a Full-Scale Portfolio Construction Framework

2.1 Defining and quantifying an investor's non-pecuniary ESG preferences

Investors can have different attitudes and preferences toward certain ESG-related features of assets, and assets can have materially different ESG attributes (e.g., Capelle-Blancard and Monjon, 2014; Nofsinger and Varma, 2014) which can lead investors to differentiate between such assets (hereafter, ESG funds).[6] Given these facts, we first define a generalizable ESG *fund-specific* parameter derived from one or more third-party sources. This parameter represents the asset's level of ESG "quality" (hereafter, ESG Quality Score). We also define a propensity parameter (hereafter, ESG Propensity Score) that is *investor-specific* and scales to the unique level of non-pecuniary "value" that an investor obtains from an ESG fund with that specific ESG Quality Score.

Our objective in this section is to quantify the level of NPU, which is the added utility score value that makes an investor indifferent between investing in conventional funds versus a specific set of ESG funds. We do this by finding the optimal asset allocation for a certain investor, such that the utility derived from investing in ESG funds is equal to the utility derived from investing in traditional, non-ESG focused assets (hereafter, conventional funds), adjusted for the amount of NPU. For this calculation, we assume that the ESG funds being considered do not outperform or underperform conventional funds to more easily isolate the non-pecuniary effect. Mathematically:

$$U(r_i, x_i)^{ESG}_{i=1\ldots n} = U(r_i, x_i)^{no\_ESG}_{i=1\ldots n} + f(S)$$
$$S = \sum_{i=1}^{n} \lambda_i x_i s_i \tag{1}$$

where U^{no_ESG} represents the utility of an optimized portfolio with just n conventional funds included, whereas U^{ESG} represents the utility of a portfolio using the same weights as U^{no_ESG} but with at least some investments considered being ESG funds and assume that they have the same pecuniary expectations as the conventional funds. x_i denotes the portfolio weight for asset i, r_i denotes the return for asset i, s_i denotes the ESG Quality Score for asset i, and λ_i represents the investor's personal ESG Propensity Score for asset i. Finally, f is the ESG preference function that quantifies the level of NPU and translates the investor-adjusted ESG Quality Scores, S, into a utility score equivalent. f is positive and increasing in S, and captures the additional utility that the ESG investor derives from holding the ESG funds.

So far, the framework we propose is flexible and permits the ESG Quality Score to be represented by *any* type or number of ESG attributes and/or from *any* third-party source. For instance, it could be represented by a specific single attribute such as a carbon intensity value.[7] In the following sections, we explain how we can define f and the ESG Quality Score in a way that would allow us to integrate the NPU into a unique utility function accounting for both pecuniary and non-pecuniary preferences.

2.2 Options for integrating non-pecuniary preferences into the optimization process

In this section, we describe how the non-pecuniary preferences and ESG Quality Scores can be integrated into a multi-asset portfolio optimization framework, highlighting the key characteristics of each approach. To our knowledge, the proposed approaches in the literature rely on standard utility maximization theory and most of them assume mean–variance optimization.

We can identify three different approaches that have been suggested, which under certain assumptions, are equivalent: (i) adding a constraint based on a specific ESG variable, such as a score based on a third-party rating agency and/or proprietary data; (ii) adding a non-pecuniary preference/penalty on top of the pecuniary utility; and (iii) integrating a non-pecuniary preference within a unique utility function and therefore defining an "ESG-adjusted" level of wealth.

2.2.1 Utility function with explicit ESG variable constraint

This approach adds a constraint to the portfolio optimization problem where an investor maximizes his or her expected utility subject to meeting a certain portfolio ESG Quality Score. Similar to Branch *et al.* (2019) and Kilmurray *et al.* (2021), the optimization problem can be defined as:

$$\max_{\{x_i\}} \ \mathbb{E}[U(r_{i,t}, x_i)]$$
$$s.t. \ \{x_i \in \mathbb{R} | 0 \le x_i \le 1\} \wedge \sum_i x_i = 1 \tag{2}$$
$$\sum_i s_i x_i = K$$

where $r_{i,t}$ represents the return for asset i at time t, x_i denotes the portfolio weight for asset i, s_i refers to the ESG Quality Score for each asset i, and K is the asset-weighted ESG value for the portfolio.

With this approach, the ESG Quality Scores can be defined in many ways and the method is simple to implement because some ESG variables (e.g., Scope 1 carbon intensity) are easy to measure. However, there are a few shortcomings that limit its attractiveness. First, it forces the investor to set an arbitrary constraint level (e.g., asset-weighted ESG Quality Score). Second, imposing an exogenous constraint on the portfolio can lead to unfeasible or unreasonable allocations (i.e., corner solutions), because it subordinates pecuniary preferences (e.g., Dorfleitner and Utz, 2012) and there is a disconnect between the dual objective of maximizing a pecuniary utility while meeting an ESG constraint that is defined in a different search space.[8]

2.2.2 *Adding an ESG preference/penalty to the pecuniary utility*

This approach assumes that the impact of the pecuniary component (i.e., utility function of wealth) on the total utility score is kept separate from the component that expresses non-pecuniary preferences. Depending on how the ESG Quality Score is defined and measured, this approach can be seen as adding an element of non-pecuniary preference or correcting the utility function for the disutility (i.e., penalty) from assets that do not have attractive non-pecuniary ESG features.[9] Similar to Pedersen *et al.* (2021), we can express this approach as:

$$\begin{aligned}&\max_{\{x_i\}}\ \mathbb{E}[U(r_{i,t}, x_i)] + f(x_i, s_i, \lambda_i)\\ &s.t.\ \{x_i \in \mathbb{R} \,|\, 0 \le x_i \le 1\} \wedge \sum_i x_i = 1\end{aligned} \tag{3}$$

where, consistent with Equation (1), s_i denotes the ESG Quality Score for any asset i and λ_i is the degree to which an investor desires to own certain ESG funds (i.e., ESG Propensity Score) for taste reasons. Function f is the ESG preference function that translates the NPU into a utility score equivalent. Mercereau and Melin (2020) follow a similar approach assuming mean–variance preferences and adjusting utility for a climate temperature constraint.

This approach has the advantage of being very generalizable and keeping pecuniary utility and non-pecuniary utility separate. However, the pecuniary utility part ($\mathbb{E}[U(r_{i,t}, x_i)]$) and the non-pecuniary utility part ($f(x_i, s_i, \lambda_i)$) ultimately need to have the same unit. This is the objective of the preference function (f): converting any type of ESG scoring system into a utility-consistent measure. Defining the preference function can therefore be challenging. Some approaches simply assign a relative importance weight to the pecuniary and non-pecuniary components (e.g., Bollen, 2007; Jessen, 2012). Unfortunately, determining preferred weights does not address the unit of measure issue and is not an intuitive exercise for many investors, which could lead to an inaccurate estimation of their personal level of NPU.

2.2.3 *Integrating non-pecuniary preferences by defining an "ESG-adjusted" level of wealth*

In this approach, an investor maximizes the expected utility of the level of wealth at maturity adjusted for non-pecuniary ESG preferences.[10] This is the approach followed by Zerbib (2022), and Pastor *et al.* (2021) who use an exponential utility function and Dorfleitner and Nguyen (2017), Klement (2018), and Idzorek and Kaplan (2022) who use mean–variance preferences. Mathematically, we can express this optimization as:

$$\begin{aligned}&\max_{\{x_i\}}\ \mathbb{E}[U(r_{i,t} + \lambda_i s_i, x_i)]\\ &s.t.\ \{x_i \in \mathbb{R} \,|\, 0 \le x_i \le 1\} \wedge \sum_i x_i = 1\end{aligned} \tag{4}$$

where λ_i is still the ESG Propensity Score. This is also defined by Pastor *et al.* (2021) as a "scalar measuring the degree of [an agent's] 'ESG taste."' Also, since the asset returns and ESG Quality Scores are directly added together, they need to be expressed in the same unit (Zerbib, 2022).

This method has the advantage of being mathematically pure and very direct. However, ESG Quality Scores can only be expressed following a definition which makes them directly comparable with asset returns. Also, this approach is equivalent to the second approach if two conditions hold: (i) the utility function can be expressed as just a function of expected returns and volatility (i.e., mean–variance utility); and (ii) the preference function is defined such that $f(x_i, s_i, \lambda_i) = \sum_{i=1}^{n} \lambda_i x_i s_i$. In that case, we can show that, similar to the approach followed by Hayes *et al.* (2015) and Idzorek and Kaplan (2022), the objective function becomes:

$$\max_{\{x_i\}} \sum_i x_i \mathbb{E}[r_i] - \frac{\gamma}{2} \sum_i \sum_j x_i x_j \sigma_i \sigma_j \rho_{i,j} + \sum_i \lambda_i x_i s_i \tag{5}$$

where σ and ρ refer to the volatility and correlation levels, respectively.

We consider this third approach the most appropriate to integrate into an expected utility optimization model and the most effective for application with investors. Therefore, we propose a framework that expands on this methodology.

2.3 *Our approach: Solve for "ESG-adjusted" level of wealth*

Our approach is to integrate pecuniary and non-pecuniary preferences in multi-asset portfolios that augments the full-scale multi-asset optimization model proposed by Aliaga-Diaz *et al.* (2020) which systematically considers active, systematic, and factor exposures.

We start from a simpler derivation of the asset allocation problem where we have only one risk bucket and risk aversion coefficient:

$$\begin{aligned} &\max_{\{x_i\}} \mathbb{E}[U(W_T)] = \max_{\{x_i\}} \mathbb{E}\left[\frac{W_{Tot}^{1-\gamma}}{(1-\gamma)}\right] \\ &s.t.\ \{x_i \in \mathbb{R} \mid 0 \le x_i \le 1\} \wedge \sum_i x_i = 1 \\ &\qquad \sum_i C \cdot x_i \le b \end{aligned} \tag{6}$$

where γ is the investor's relative risk aversion and C and b refer to the set of linear constraints. Also, for ease of explanation, we further simplify the optimization problem by assuming that the investor's ESG Propensity Scores are identical across all ESG funds considered ($\lambda_i = \lambda, \forall i$) and the ESG Quality Scores for each fund do not change over time.[11] Finally, W_{Tot} is the level of terminal wealth adjusted for non-pecuniary preferences and ESG Quality Scores and defined as:

$$W_{Tot} = \prod_{t=1}^{T} \left(1 + \sum_{i=1}^{N} x_i (r_{i,t} + \lambda s_i)\right) \tag{7}$$

The framework defined by Aliaga-Diaz *et al.* (2020) is characterized by splitting the total level of wealth coming from different sources of risk (active, factors, and systematic) into separate buckets. Our objective is then to redefine the wealth coming from the systematic (W_p), factor (W_f) and active (W_a) risk sources to also account for non-pecuniary preferences. Once we appropriately update the wealth components for considering ESG fund exposures and preferences, we can rely on the same expected utility maximization framework. To accomplish that and show how the model specification changes once we add ESG Quality and ESG Propensity Scores for each asset, we recast the portfolio return at any given point in time:

$$\sum_{i=1}^{N} x_i(r_{i,t} + \lambda s_i) = \sum_{i=1}^{N} x_i^p\left(r_{i,t}^p + \lambda s_i^p\right) + \sum_{i=1}^{N}\sum_{f=1}^{F} x_i^f\left(r_{i,t}^f + \lambda s_i^f\right) + \sum_{i=1}^{N} x_i^a(r_{i,t}^a + \lambda s_i^a) \quad (8)$$

Here we do not assume that the ESG Quality Scores are identical across the active, factor, and systematic risk exposures for each asset class. This allows a higher degree of generalization, and we acknowledge the fact that, for example, an active U.S. ESG equity fund does not necessarily have the same ESG Quality Score of its passive ESG benchmark. However, to keep the algebra simple, we assume that λ is the same across the different risk buckets ($\lambda^k = \lambda, \forall\, k \in \{p, f, a\}$). In the Appendix, we report the additional steps and details for the derivation of the objective function.

2.4 *Willingness-to-pay as an NPU measure*

Our framework requires that NPU is measured in a unit that can be combined with a financial measure of utility. But unlike a financial outcome, a non-pecuniary benefit is subjective and not directly observable, and therefore difficult to value (e.g., Cooper *et al.*, 2016; Dorfleitner and Nguyen, 2017). There is no widely accepted methodology to capture investors' non-pecuniary preferences. As a result, making an accurate measurement and mapping it into something universally understandable and practical for inclusion in a portfolio optimization framework is inherently challenging. More precisely, we need to provide a definition and measure for the ESG Quality Scores (s_i) and the investor's ESG Propensity Score (λ).

Contingent valuation methods are a widely accepted way to attach an economic value to a nonmarket resource.[12] It involves a survey-based methodology that requires survey respondents to make an economic decision based on a real or hypothetical scenario. Economic value can be assessed through a willingness-to-pay (WTP) or willingness-to-accept (WTA) measure.[13] Measuring an investor's WTP can be revealed through actual (or simulated) investor purchases or by asking their purchase intention through surveys. Valuation can be captured through different techniques, from an open-ended question format to more complex methodologies such as the design of choice experiments (Breidert *et al.*, 2006; Miller *et al.*, 2011).

However, the reliability and validity of contingent valuation methods are subject to debate, due to numerous issues, including: (i) hypothetical bias, reflecting a difference between what respondents say and what they do, (ii) unstable preferences, evidenced by the documented gap between WTP and WTA measures, and, (iii) a scope sensitivity problem, where values vary depending on whether goods are valued separately or as part of a package (Oerlemans *et al.*, 2016). For the most part, these known issues can be mitigated by a well-designed survey-based choice elicitation technique (see, for example, McFadden, 2001; Plott and Zeiler, 2005; Mas and Pallais, 2017; Schmidt and Bijmolt, 2020).

For our purpose, the WTP corresponds to the monetary value of the NPU that the investor derives from a particular ESG fund. Specifically, it is the maximum performance shortfall, relative to a conventional fund in the same asset category, that an investor would be willing to accept. Therefore, all else equal, the stronger the investor's non-pecuniary preference for a specific ESG fund, the larger their WTP will be.

We are agnostic as to the source of NPU—whether from altruism, warm glow, or social signaling—and treat it as the aggregate of all three sources. The source of NPU may be reflected in the strength of investors' desire for ESG. For example, an investor motivated by altruism may have a stronger desire for ESG and require a larger portfolio share than one motivated by warm glow or social signaling. Someone who would rather exercise their altruistic motives through different channels (such as volunteering or donations to charity) would be expected to express low propensity for ESG funds. Either way, the investor's propensity for ESG is captured in the model.

For the purposes of the model specification defined in the previous section, we define the ESG Quality Scores (s_i) as the average WTP value assigned to an ESG fund based on a large survey of investors. This survey would be conducted to capture the distribution of WTPs across investors, by fund profile. One approach that can be used for the estimation of the entire distribution of WTPs is through a discrete choice model (DCM) as specified by McFadden (1984).

In the survey, investors would be presented with a hypothetical scenario where they are shown a series of choice tasks and are asked to choose between an ESG fund and a conventional fund. These funds will vary based on pre-established attributes, such as management strategy, fund return, and a vector of ESG characteristics. Each of the attributes have different levels – for example, active and index for management strategy. The objective is that at the end of the exercise, enough investors would have made a choice between various types of ESG funds and traditional funds to have a sufficiently large sample of preferences. At the same time, the return differences between the conventional and ESG funds that are embedded in the exercise would allow the calculation of WTP post-survey. More precisely, through a binary regression (e.g., logit model), one would estimate the probability for a given investor to pick the ESG fund, over the conventional fund equivalent, given their respective attributes, for any given level of ESG fund underperformance. A cumulative distribution function for any given combination of ESG funds attributes can be created from this to then calculate the probability density function (i.e., distribution of WTPs) for a wide range of different ESG funds. This process would ultimately lead to specifying a unique ESG Quality Score metric for each ESG fund.

If a survey-based WTP value is not used as the proxy, the advisor could develop a conversion function to translate the chosen proxy or proxy combination into WTP equivalents in a way that keeps the ordinal level of ESG "quality" across assets. For example, an investor could decide that asset-weighted carbon intensity is how they personally measure the non-pecuniary quality of an ESG fund. They would express their personal WTP level for each ESG fund being considered based on each fund's ESG Quality Score, evaluated in terms of carbon intensity. In effect, this specification approach collapses the ESG Quality Score and ESG Propensity Score into a single measure. If the investor derives NPU from more than just a single attribute like carbon intensity in our example (e.g., workforce gender diversity), then the ordinal level for ESG quality depends on the relative importance they assign to each ESG attribute of interest (Cooper *et al.*, 2016). The investor would take this into account when expressing their unique WTP value for each ESG fund to the advisor.

The ESG Propensity Score (λ), which is based on an investor's specific input, would then act as a scaling factor, adjusting an investor's personal

WTP to its position in the survey-based distribution. The actual estimation of the hyper-parameter λ is outside the scope of this article. However, one approach that can be used for the estimation of λ is through the DCM mentioned previously, where the ESG Propensity Score would be derived based on a specific percentile of the distribution of WTPs.

For example, imagine a hypothetical investor who, after reviewing the attributes provided by her advisor, considers her WTP to be 1.04% for a specific ESG fund. Based on the distribution of WTPs for that fund from the survey, her personal WTP is higher than the *average* WTP of 0.8% (i.e., ESG Quality Score). For instance, her value of 1.04% could fall in the top 30% of the distribution of WTPs. Therefore, we observe that the associated ESG Propensity Score (λ) scaling factor for her profile would be consistent with a level of 1.3 ($1.04/0.8\% = 1.3$). Alternatively, if she indicated that her WTP was zero which would translate into $\lambda = 0$, that would be equivalent to a portfolio for a typical investor who considers only pecuniary goals. Figure 1 shows an illustrative example for two ESG funds where ESG Fund B has a higher average WTP compared to ESG Fund A, presumably from having more attractive ESG attributes (e.g., lower carbon intensity value). This setup allows the investor to build portfolios where different ESG funds, with more/less appealing ESG features, would have higher/lower average WTPs. The ESG Propensity Score, which is investor specific, scales the average WTP to the personalized WTP.[14] Also, this approach provides for an intuitive and straightforward application of our framework to efficiently build model portfolios and single fund solutions, where different profiles of ESG investors or personas can be assumed.

An important consideration for our portfolio construction framework is whether the required investor input can be easily and accurately captured. Through exploratory interviews, we confirmed that ESG investors and those seriously considering an ESG fund can understand and relate to the concepts of NPU and WTP. We found that some investors are willing to give up conventional fund returns to invest in an ESG fund

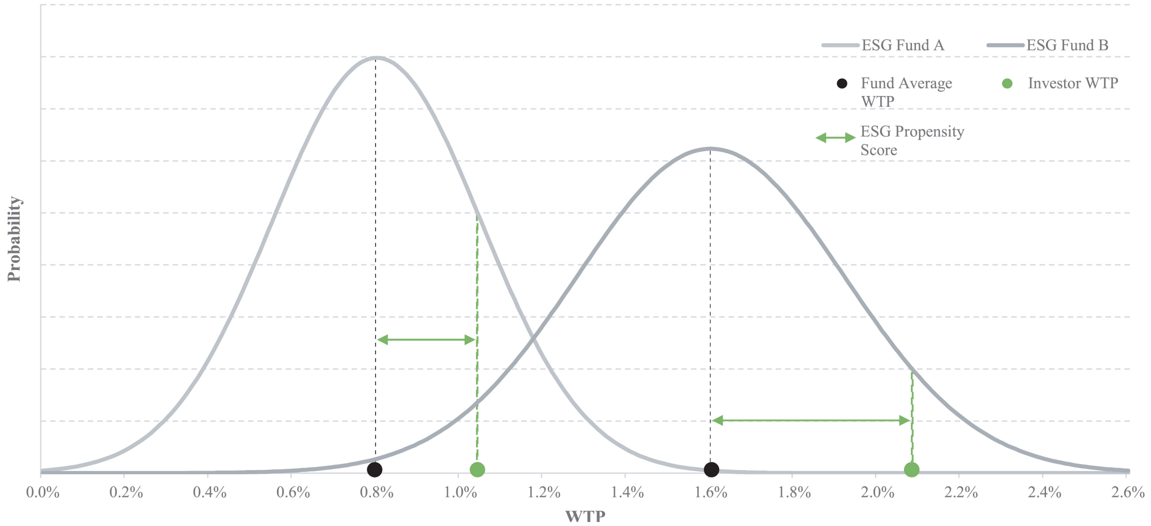

Figure 1 ESG fund WTP distribution and ESG propensity score.

Notes: Hypothetical and for illustrative purposes only. In this example, the ESG Quality Score for each fund is proxied by the median WTP from a distribution based on a survey of investors. The ESG Propensity Score is set at a value of 1.3 and is assumed to be the same for both ESG Fund A and ESG Fund B.

and can provide an estimate of their WTP for ESG.[15]

Besides displaying the attributes in a two-fund (ESG versus conventional) side-by-side comparison format, we have found that there are a few other best practices for an advisor who is either constructing survey questions for a large group of investors or creating a simple questionnaire template with the objective of gathering an investor's WTP for a given ESG fund. First, make sure that the horizon used for the hypothetical fund returns is long term, because the goal is to measure an investor's WTP over a period reasonably similar to how long they could be in the fund for standard financial reasons. Second, ensure that the ESG fund performance they can choose is always lower than the conventional fund and that the performance difference is certain. Recall for this activity, we are attempting to measure their point of utility indifference between the two funds solely for NPU measurement purposes. Therefore, the fund return expectations are hypothetical and do not have to match actual financial performance expectations that will be embedded in the portfolio optimization process.[16] Third, include the end wealth WTP difference, so all investors understand how return differences translate into dollar wealth differences. Lastly, given we have found that an investor's WTP can vary based on the expected absolute return of the conventional fund (with some willing to pay less when absolute returns are lower), test their WTP using different absolute return assumptions for the conventional fund. Collectively, these practices can help improve the accuracy of the parameter estimates.

2.5 *Methodological features lead to differentiated portfolio construction benefits*

This utility-based portfolio construction framework that explicitly accounts for NPU exhibits features that make it desirable for practitioner usage. In particular, the approach:

- **Builds off a strong utility-based foundation.** It augments a traditional, utility-based portfolio modeling approach ensuring that the desirable pecuniary features are retained, while allowing for the inclusion of NPU for the subset of investors who prefer to choose a portfolio that maximizes a more comprehensive measure of utility (Jessen, 2012; Cooper *et al.*, 2016; Dorfleitner and Nguyen, 2017).
- **Enables gathering of more accurate inputs from investors.** Our approach uses an intuitive and accessible method, reducing the risk of "garbage in, garbage out" which can result when an investor does not understand how to estimate an important model input.
- **Utilizes a well-supported proxy for an investor's personal NPU.** WTP, our proxy for NPU, has robust theoretical support in the literature across different product choice domains. This proxy can capture NPU derived from the most common sources (i.e., altruism, warm-glow, and social signaling).
- **Matches unit for measurement of asset returns and investor-specific NPU.** Our approach allows for a direct unit correspondence between assets' returns and investor-specific non-pecuniary preferences, as both are defined in the "return space" when scoring for utility purposes. In this regard, our proxy for NPU can be seen as a form of "social alpha". This helps with NPU measurement accuracy and permits full interaction effects with pecuniary features since it integrates non-pecuniary preferences into a one-step framework that is built on pecuniary preferences.[17]
- **Accommodates broad flexibility and applicability for many investor types**. The framework allows for customized financial expectations for ESG funds (concessionary or non-concessionary), personalized systematic,

factor, and active risk aversion, investor-specific degree of non-pecuniary propensity for each ESG fund, and ESG funds with different levels of ESG "quality."[18] This permits a highly bespoke process and provides for more transparent trade-off discussions between the investor and the advisor.

- **Avoids requiring arbitrary weight constraints and two-step portfolio construction.** Weight constraints are suboptimal because the "trading-off" of pecuniary and non-pecuniary attributes of the assets being considered in the optimization is not occurring simultaneously, leading to the subordination of one goal over the other (e.g., Dorfleitner and Utz, 2012; Kilmurray *et al.*, 2021).

3 Case Studies

In this section, we present three hypothetical case studies and the associated results from incorporating the aforementioned methodology into a multi-asset portfolio construction process with active and indexed strategies. The intent of these examples is to assess the sensitivity of the optimal asset allocation to the ESG Propensity Score, the ESG Quality Score, and the active risk aversion, using some hypothetical values. Calibrating the ESG Quality Scores (s_i) and ESG Propensity Scores (λ) or the utility risk aversion coefficients is out of scope.

For simplicity and interpretability, we consider just six investment categories: four conventional-only index funds (U.S. equity, U.S. bonds, global ex-U.S. equity, and global ex-U.S. bonds) and two ESG index funds (U.S. equity and U.S. bonds). For both the ESG U.S. equity index and ESG U.S. bond index funds, we assume no factor exposures and an alpha of 0% and a tracking error of 3% and 2% versus a conventional total market index, respectively. Also, we construct a baseline portfolio with no ESG funds exposure. The baseline portfolio is used for comparison once non-pecuniary preferences are included. The systematic and active risk aversions are calibrated such that the baseline portfolio has a total 60% equity and 40% bond allocation. Lastly, we will assume that the ESG Quality Score is measured using only a fund-specific carbon intensity metric.[19]

We perform a series of three case studies: (i) for increasing values of ESG Propensity Scores, we examine the impact on the optimal allocation and its carbon intensity level; (ii) for increasing values of the U.S. equity fund ESG Quality Score, we look at the active-passive allocation split; and (iii) for increasing values of ESG Propensity Scores and a range of active risk aversions, we compare the allocation with ESG funds versus the baseline portfolio.[20]

In Table 1 we report the first case study, where different ESG Quality Scores are assumed for the ESG U.S. equity and the U.S. bond funds. The carbon intensity levels are then translated into a return or wealth-based measure, the hypothetical investors' average WTP in this case, in order to be seamlessly integrated into the utility function. In practice, this step would rely on survey-derived investor WTPs. The first column represents the baseline portfolio as the ESG Propensity Score is set equal to zero. In the other columns, we gradually increase the ESG Propensity Score up to a value of 2.0.

The allocation to the ESG U.S. equity and U.S. bond funds increases as the ESG Propensity Score increases. In fact, for the largest ESG Propensity Score that we chose to use in this example, the weight in ESG funds represent a significant portion of the portfolio (54%). Since in this case study the ESG funds being considered are expected to generate zero alpha, we notice how the portfolio's expected return does not change when the ESG Propensity is set at higher levels, but the portfolio's expected tracking error increases up

Table 1 Optimal asset allocation responds to different levels of taste for ESG funds.

ESG Input Assumptions						
ESG Propensity Score	0.0	0.4	0.8	1.2	1.6	2.0
U.S. Equities ESG Quality Score	1.14%	1.14%	1.14%	1.14%	1.14%	1.14%
U.S. Bonds ESG Quality Score	1.33%	1.33%	1.33%	1.33%	1.33%	1.33%
Asset Class Weights						
U.S. Equities	36%	36%	36%	36%	36%	36%
Conventional	36%	36%	36%	31%	19%	6%
ESG	0%	0%	0%	5%	17%	30%
Global ex-U.S. Equities	24%	25%	25%	24%	24%	24%
U.S. Bonds	36%	35%	35%	36%	36%	36%
Conventional	36%	10%	3%	3%	7%	12%
ESG	0%	25%	32%	33%	29%	24%
Global ex-U.S. Bonds	4%	4%	4%	4%	4%	4%
	100%	100%	100%	100%	100%	100%
Summary Analytics						
Total Equity Allocation	60%	61%	61%	60%	60%	60%
Additional Allocation to ESG Funds	0%	25%	32%	38%	46%	54%
Portfolio Carbon Intensity	121.9	110.1	106.6	102.6	97.2	91.7
Expected Return	4.2%	4.2%	4.2%	4.2%	4.2%	4.2%
Expected Volatility	9.6%	9.8%	9.8%	9.7%	9.7%	9.7%
Tracking Error	n.a.	0.4%	0.6%	0.7%	1.0%	1.2%

Notes: Portfolios have been optimized over a 10-year investment horizon. Expected returns, volatility (standard deviation), and tracking error are median values from a distribution of 10,000 simulations. We use each hypothetical fund's carbon intensity level as the ESG Quality Score. The carbon intensity levels for the ESG U.S. equity and U.S. bond funds are 81.0 and 42.4, respectively. The conventional index fund carbon intensity level is 142.0 for equity and 91.7 for bonds. The ESG Quality Scores have been translated into an average WTP of 1.14% and 1.33%, respectively, and the scores are kept constant in each scenario. The investor's ESG Propensity Scores are assumed to be identical for both U.S. equity and U.S. bonds funds. The portfolio assumes six asset categories: four conventional index strategies covering U.S. equity, global ex-U.S. equity, U.S. bonds, and global ex-U.S. bonds and two indexed ESG strategies covering U.S. equity and U.S. bonds. The baseline portfolio, for an investor who does not generate NPU (i.e., ESG Propensity Score equals zero), is used for comparison once non-financial ESG preferences are included. The systematic and active risk aversions are calibrated such that the baseline portfolio has a total of 60% equity and 40% bond allocation. The optimal asset allocation options are also constrained such that global ex-U.S. equity is lower or equal to 40% of total equity, and global ex-U.S. bonds is lower or equal to 50% of total bonds.
Sources: Vanguard calculations, using asset returns' projections as of December 2020, with forecasts in USD.

to 1.2% with larger ESG fund exposures. Finally, having a larger fraction of the portfolio allocated to ESG funds leads to a reduction in the carbon intensity level (from 121.9 to 91.7). This means that, all else equal, investors with higher ESG Propensity Scores toward carbon intensity will tend to have optimal allocations with lower asset-weighted carbon intensity levels.

The second case study is reported in Table 2. Here we test the sensitivity of the optimal allocation by changing the ESG Quality Score, represented here by investors' average WTP, for the U.S. equity ESG fund. All other inputs, including the ESG Quality Score for the ESG U.S. bond fund, are kept constant, including assuming an investor has an ESG Propensity Score equal to 0.4, our

Table 2 Optimal asset allocation responds to different levels of U.S. equity fund ESG quality.

ESG Input Assumptions						
ESG Propensity Score	0.4	0.4	0.4	0.4	0.4	0.4
U.S. Equities ESG Quality Score	1.14%	1.59%	1.71%	1.82%	2.28%	2.50%
U.S. Bonds ESG Quality Score	1.33%	1.33%	1.33%	1.33%	1.33%	1.33%
Asset Class Weights						
U.S. Equities	36%	36%	36%	37%	36%	36%
Conventional	36%	29%	20%	11%	6%	4%
ESG	0%	7%	16%	26%	30%	32%
Global ex-U.S. Equities	25%	24%	24%	24%	24%	24%
U.S. Bonds	35%	36%	36%	35%	36%	36%
Conventional	10%	17%	26%	35%	36%	36%
ESG	25%	19%	10%	0%	0%	0%
Global ex-U.S. Bonds	4%	4%	4%	4%	4%	4%
	100%	100%	100%	100%	100%	100%
Summary Analytics						
Total Equity Allocation	61%	60%	60%	61%	60%	60%
Additional Allocation to ESG Funds	0%	1%	1%	1%	5%	7%
Portfolio Carbon Intensity	110.1	108.2	107.2	106.5	103.6	102.4
Expected Return	4.2%	4.2%	4.2%	4.2%	4.2%	4.2%
Expected Volatility	9.8%	9.6%	9.7%	9.8%	9.7%	9.7%
Tracking Error	0.4%	0.5%	0.6%	0.7%	0.8%	0.9%

Notes: Portfolios have been optimized over a 10-year investment horizon. Expected returns, volatility (standard deviation), and tracking error are median values from a distribution of 10,000 simulations. We use each hypothetical fund's carbon intensity level as the ESG Quality Score. The carbon intensity levels for the ESG U.S. equity and U.S. bond funds are 81.0 and 42.4, respectively. The conventional index fund carbon intensity level is 142.0 for equity and 91.7 for bonds. The ESG Quality Scores have been translated into an average WTP of 1.14% and 1.33%, respectively. The ESG Quality Score for U.S. equities is then further and gradually increased up to a level of 2.50%. The investor's ESG Propensity Scores are assumed to be identical for both U.S. equity and U.S. bonds funds. The portfolio assumes six asset categories: four conventional index strategies covering U.S. equity, global ex-U.S. equity, U.S. bonds, and global ex-U.S. bonds and two indexed ESG strategies covering U.S. equity and U.S. bonds. The baseline portfolio, for an investor who does not generate NPU (i.e., ESG Propensity Score equals zero), is used for comparison once non-financial ESG preferences are included. The systematic and active risk aversions are calibrated such that the baseline portfolio has a total of 60% equity and 40% bond allocation. The optimal asset allocation options are also constrained such that global ex-U.S. equity is lower or equal to 40% of total equity, and global ex-U.S. bonds is lower or equal to 50% of total bonds. The results are for an investor with an ESG Propensity Score of 0.4.
Sources: Vanguard calculations, using asset returns' projections as of December 2020, with forecasts in USD.

new base case portfolio. Our results show how by increasing the ESG Quality Score associated with the ESG U.S. equity fund by a certain factor (from 1.14 to 2.5), the optimal asset allocation moves from conventional U.S. equity and ESG U.S. bonds into ESG U.S. equity and conventional U.S. bonds. Also, we notice that for high U.S. equity fund ESG Quality Scores, the total allocation to ESG funds is larger than the base case portfolio (column 1), proving that not only does it become optimal to switch from the ESG bond fund to the ESG equity fund, but an investor is also better off having a larger overall exposure to ESG funds.

Table 3 Allocation to ESG funds responds to different levels of active risk aversion.

ESG Propensity Score	Low					Medium					High
0.0	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
0.4	37%	34%	31%	29%	27%	25%	24%	23%	21%	20%	19%
0.8	46%	42%	40%	37%	33%	32%	30%	28%	26%	24%	23%
1.2	59%	53%	47%	44%	40%	38%	35%	34%	32%	31%	29%
1.6	68%	63%	58%	52%	50%	46%	42%	40%	39%	38%	35%
2.0	78%	74%	65%	64%	59%	54%	54%	50%	45%	44%	43%

(Column heading: **Active Risk Aversion**)

Notes: Each percentage represents the total portfolio allocation to ESG funds based on specific ESG Propensity Score and active risk aversion parameter settings. Portfolios have been optimized over a 10-year investment horizon. We use each hypothetical fund's carbon intensity level as the ESG Quality Score. The carbon intensity levels for the ESG U.S. equity and U.S. bond funds are 81.0 and 42.4, respectively. The conventional index carbon intensity level is 142.0 for equity and 91.7 for bonds. The ESG Quality Scores have been translated into an average WTP of 1.14% and 1.33%, respectively, and the scores are kept constant in each scenario. The investor's ESG Propensity Scores are assumed to be identical for both U.S. equity and U.S. bonds funds. The portfolio assumes six asset categories: four conventional index strategies covering U.S. equity, global ex-U.S. equity, U.S. bonds, and global ex-U.S. bonds and two indexed ESG strategies covering U.S. equity and U.S. bonds. The baseline portfolio, for an investor who does not generate NPU (i.e., ESG Propensity Score equals zero), is used for comparison once non-financial ESG preferences are included. The systematic and active risk aversions are calibrated such that the baseline portfolio has a total of 60% equity and 40% bond allocation. The optimal asset allocation options are also constrained such that global ex-U.S. equity is lower or equal to 40% of total equity, and global ex-U.S. bonds is lower or equal to 50% of total bonds.
Sources: Vanguard calculations, using asset returns' projections as of December 2020, with forecasts in USD.

In Table 3, the third case study, we showcase how the allocation to ESG funds changes when we vary active risk aversion, unless the investor does not earn NPU (i.e., ESG Propensity Score is set to zero). The set of portfolios reported in Table 1 correspond to a medium level of active risk aversion (contoured in black in Table 3). Our results show that the total allocation to the two ESG funds increases as the investor's ESG Propensity Score for carbon intensity increases and their active risk aversion decreases. This indicates how a "social alpha," based on the NPU an investor derives, can lead to a higher optimal allocation to ESG funds for investors who are more tolerant with active risk.

In summary, our case study results are consistent with our expectations and suggest that, all else equal, it is prudent for investors with larger ESG Propensity Scores to have a higher allocation in ESG funds. Also, ESG funds that exhibit more attractive ESG characteristics for the subset of investors who have non-pecuniary preferences—such as reduction in carbon emissions—and therefore have higher ESG Quality Scores, see a larger allocation in multi-asset portfolios for such investors. Finally, all else equal, investors with more appetite for active risk (i.e., lower risk aversion) will tend to receive a larger optimal allocation in ESG funds.

4 Conclusion

While the average investor builds a portfolio only considering financial expectations and preferences, it is well-documented that some investors derive non-financial, positive social, and/or private benefits from the ESG-related features of assets which can rationally influence their portfolio construction process. The most common method used to determine this type of investor's policy portfolio only assesses one dimension of risk and arbitrarily applies weight constraints for assets with a certain ESG profile, which subordinates pecuniary goals and can be difficult for an advisor to justify.

In this article, we propose a comprehensive framework that explicitly and simultaneously considers an investor's pecuniary and non-pecuniary preferences in a multi-asset portfolio optimization model. We also consider the "active" risk associated with ESG fund investments, which by construction deviate from a traditional, broad market benchmark. The approach specifies non-pecuniary preferences by defining an "ESG-adjusted" level of wealth that corrects assets' weights based on the level of ESG exposure (i.e., ESG "quality") of any given ESG fund and an investor's unique taste for assets with a certain ESG profile. The quantitative framework measures NPU by obtaining an investor's personal WTP break-even threshold for each of those assets. Willingness-to-pay has substantial theoretical support as an accurate and accessible NPU proxy in a number of research domains and allows for a direct unit correspondence between assets' returns and investor-specific non-pecuniary preferences. The model permits full interaction effects by considering the investor's NPU together with their financial risk preferences when calculating total utility scores for different portfolio mixes.

The strength of non-pecuniary desires along with financial expectations and risk aversion preferences are important determinants of the optimal portfolio choice. This theoretical intuition is supported by the trade-off evidence that we observe in case studies between pecuniary and non-pecuniary preferences. The optimal asset allocation to ESG funds responds to changes in the level of active risk aversion, the non-pecuniary ESG "quality" of the funds, and the degree of investor's desire for the ESG features of the funds. This avoids arbitrary constraint decisions, showcases the importance of being able to fully customize a variety of financial and non-financial inputs, and provides for more transparent trade-off discussions between the investor and their advisor.

Appendix A. Derivation of the Non-Pecuniary Adjusted Objective Function

Following the set of equations from Aliaga-Diaz *et al.* (2020), we know that the market return ($r_{i,t}^M$) is a common variable to all types of investment exposures (i.e., active, factors and passive). Therefore, we can add the non-pecuniary component to $r_{i,t}^M$ and create $r_{i,t}^{M*} = r_{i,t}^M + \lambda s_i^p$. Similarly, we would need to correct $\delta_{i,t}^f$ and α_i with the difference between λs_i^f and λs_i^p and λs_i^a, respectively. We therefore recast Equation (8) and allow for ESG propensities and Quality Scores to adjust the final total return of the portfolio as follows:

$$\sum_{i=1}^{N} x_i(r_{i,t} + \lambda s_i) = \sum_{i=1}^{N} x_i^p r_{i,t}^{p*} + \sum_{i=1}^{N}\sum_{f=1}^{F} x_i^f r_{i,t}^{f*} + \sum_{i=1}^{N} x_i^a r_{i,t}^{a*} \tag{A.1}$$

where:

$$\begin{cases} r_{i,t}^{p*} = r_{i,t}^{M*} \\ r_{i,t}^{f*} = r_{i,t}^{M*} + \delta_{i,t}^{f*} \\ r_{i,t}^{a*} = \alpha_i^* + \beta_i r_{i,t}^{M*} \\ \qquad + \sum_{f=1}^{F} L_i^f \delta_{i,t}^{f*} + \varepsilon_{i,t} \\ \qquad \varepsilon_{i,t} \sim t(v)\sqrt{\sigma_{\alpha_i}^2} \end{cases} \tag{A.2}$$

and where:

$$\begin{cases} r_{i,t}^{M*} = r_{i,t}^{M} + \lambda s_i^p \\ \delta_{i,t}^{f*} = \delta_{i,t}^{f} + \lambda\left(s_i^f - s_i^p\right) \\ \alpha_i^* = \alpha_i + \lambda \left(s_i^a - \beta_i s_i^p \right. \\ \qquad \left. - \sum_{f=1}^{F} L_i^f \left(s_i^f - s_i^p\right) \right) \end{cases} \tag{A.3}$$

At this point, we recast the returns as:

$$\begin{cases} R_{p,t}^* = \sum_{i=1}^{N} \left[\left(x_i^p + \sum_{f=1}^{F} x_i^f + x_i^a \right) r_{i,t}^{M*} \right] \\ R_{f,t}^* = \sum_{i=1}^{N} \left[\left(\sum_{f=1}^{F} L_i^f \delta_{i,t}^{f*} + (\beta_i - 1) r_{i,t}^{M*} \right) x_i^a \right. \\ \qquad \left. + \sum_{f=1}^{F} x_i^f \delta_{i,t}^{f*} \right] \\ R_{a,t}^* = \sum_{i=1}^{N} x_i^a (\alpha_i^* + \varepsilon_{i,t}) \end{cases} \tag{A.4}$$

Finally, assuming log returns, we define the different wealth components as $W_p = e^{\sum_{t=1}^{T} R_{p,t}^*}$, $W_f = e^{\sum_{t=1}^{T} R_{f,t}^*}$ and $W_a = e^{\sum_{t=1}^{T} R_{a,t}^*}$ and rely on the same expected utility maximization setting reported in Aliaga-Diaz *et al.* (2020) in Equation (4).

Disclaimer

The authors report no conflicts of interest.

All investing is subject to risk, including the possible loss of the money you invest. Be aware that fluctuations in the financial markets and other factors may cause declines in the value of your account. There is no guarantee that any particular asset allocation or mix of funds will meet your investment objectives or provide you with a given level of income.

Endnotes

1. GSIA prefers an inclusive definition of ESG assets by counting "investment approaches that consider environmental, social, and governance (ESG) factors in portfolio selection and management across seven strategies of sustainable or responsible investment." See their report for more details. Some practitioners and academics use the term sustainable investments instead of ESG investments to refer to the same concept.
2. This could include cases when the investor has a risk aversion with one or more financially material ESG-related risks (e.g., Andersson *et al.*, 2016; Engle *et al.*, 2020; Pastor *et al.*, 2021).
3. Some investors believe that their investments in certain assets directly cause incremental positive environmental and/or societal change while others disagree (Kölbel *et al.*, 2020). This can impact the level of non-pecuniary utility that they derive from an investment. Whether allocating capital to a particular asset causes incremental change is outside the scope of our research. For the framework proposed in this paper, we assume that the answer to this is known when the portfolio decision is being made.
4. Some investors with non-financial ESG tastes may fully satisfy them outside of their portfolio through alternative methods such as supporting certain non-profit organizations or lobbying efforts or they cannot consider non-financial issues in the asset allocation process for fiduciary reasons. Those investors should use conventional portfolio optimization approaches.
5. The theory of investors deriving NPU from ESG-type strategies and weighing the marginal benefit against the marginal cost has been reported in the academic literature for decades (e.g., Cullis *et al.*, 1992). For revealed preference evidence, see, for instance, Barber *et al.* (2021) and Bialkowski and Starks (2016) and the citations therein. For stated preference evidence, see, for instance, Apostolakis *et al.* (2016); Hartzmark and Sussman (2019); Bauer *et al.* (2021); Brodback *et al.* (2021); and Heeb *et al.* (2022).
6. We use the term "ESG funds" in this article for simplicity purposes. The framework can incorporate non-fund assets, such as separately managed/personalized/direct indexing strategies with a bespoke ESG screen or conventional funds with an ESG attribute profile that an

investor finds desirable even though it is not labeled as an ESG fund.

[7] These can be gathered from different independent providers (e.g., Sustainalytics, MSCI). If the ESG attribute data of investor interest is not available for a particular ESG fund, then there is no ESG Quality Score for the investor to potentially derive added utility from, leading to an ESG Propensity Score and consequently NPU level of zero for that fund.

[8] More precisely, this means that unless an investor has a clear sense of which level he/she wants to achieve on the ESG variable, and whether that is achievable given all other inputs and constraints (e.g., risk aversion, relative weight constraints, etc.), there might be no feasible portfolio meeting all requirements.

[9] Penalty functions are quite popular in the literature of portfolio optimization. See, for example, Serbin *et al.* (2011) and Hayes *et al.* (2015).

[10] We can also consider asset returns to be adjusted for non-pecuniary ESG preferences rather than looking at the level of wealth at maturity.

[11] Just like when the investor's pecuniary expectations or risk aversion changes, if the ESG Quality Score of a fund change significantly in the future, that could trigger the need for a reassessment of the optimal portfolio allocation.

[12] See Carson (2000) for an overview of contingent valuation.

[13] In general, WTP represents the maximum price at which a person is willing to obtain a good while WTA is the minimum price at which a person is willing to give up a good that he or she already owns.

[14] The ESG Propensity Score here is assumed to be a constant scalar. This can be very effective for creating different portfolios associated with different investor personas. However, the ESG Propensity Score can also be calibrated based on a specific percentile of the ESG fund WTP distribution (e.g., an investor who is consistently in the top 10% of the WTP distribution).

[15] Based on a qualitative study conducted by Vanguard (2021).

[16] As a reminder, the framework is flexible and silent on actual expected conventional and ESG fund performance for the purpose of portfolio optimization, leaving those decisions to the advisor.

[17] Although a penalty function approach would not require adjusting the set of return and risk expectations for each asset, finding the function that correctly converts the NPU measure into a utility equivalent can be very challenging. In fact, we need the NPU estimate to have the same order of magnitude/unit of pecuniary utility.

[18] Our framework does consider any ESG index fund to be a rules-based active strategy because they are, by definition, constructed differently than a conventional total market index. This results in performance that varies from the market, generating the need for an active risk aversion setting for the investor. Depending on the specific fund, it may also require a factor risk aversion (Plagge *et al.*, 2022). Therefore, an advisor would input an estimate for the expected alpha, tracking error, and factor loadings (if any) for each ESG fund being considered in the optimization process along with the investor's risk aversion levels.

[19] As described in the previous sections, the ESG Quality Score can by proxied using one or more variables from third-party sources (e.g., Sustainalytics), not necessarily just a carbon intensity metric.

[20] While for simplicity purposes we use only one set of pecuniary assumptions for ESG funds in these case studies, our framework does *not* take a position on whether certain ESG index or active funds will generate positive, zero, or negative alpha or exhibit a certain level of tracking error. It can accommodate any set of pecuniary assumptions for such funds. Theories and empirical evidence on ESG fund performance are mixed (Plagge and Grim, 2020). We ran scenarios using different pecuniary assumptions for ESG funds but did not report those results to conserve space. As one might expect, holding all other assumptions constant, increasing (decreasing) alpha and decreasing (increasing) tracking error led to larger (smaller) optimal allocations to ESG funds.

Acknowledgments

We are grateful for the comments and suggestions from James J. Rowley, Roger Aliaga-Diaz, Dhagash Mehta, and an anonymous referee.

References

Adler, T. and Kritzman, M. (2007). "Mean-Variance Versus Full-Scale Optimization: In and Out of Sample," *Journal of Asset Management* **7**(5), 302–311.

Aliaga-Diaz, R., Renzi-Ricci, G., Daga, A., and Ahluwalia, H. (2020). "Portfolio Optimization with Active, Passive, and Factors: Removing the Ad Hoc Step," *The Journal of Portfolio Management* **46**(4), 39–51.

Andersson, M., Bolton, P., and Samama, F. (2016). "Hedging Climate Risk," *Financial Analysts Journal* **72**(3), 13–32.

Andreoni, J. (1989). "Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence," *Journal of Political Economy* **97**(6), 1447–1458.

Apostolakis, G., Kraanen, F., and van Dijk, G. (2016). "Examining Pension Beneficiaries' Willingness to Pay for a Socially Responsible and Impact Investment Portfolio: A Case Study in the Dutch Healthcare Sector," *The Review of Financial Studies* **34**(8), 3976–4043.

Barber, B. M., Morse, A., and Yasuda, A. (2021). "Impact Investing," *Journal of Financial Economics* **139**(1), 162–185.

Bauer, R., Ruof, T., and Smeets, P. (2021). "Get Real! Individuals Prefer More Sustainable Investments," *The Review of Financial Studies* **34**(8), 3976–4043.

Becker, G. (1974). "A Theory of Social Interactions," *Journal of Political Economy* **82**(6), 1063–1093.

Bialkowski, J. and Starks, L. T. (2016). "SRI Funds: Investor Demand, Exogenous Shocks and ESG Profiles," Working Papers in Economics 16/11, University of Canterbury, Department of Economics and Finance.

Black, F. and Litterman, R. B. (1991). "Asset Allocation: Combining Investor Views with Market Equilibrium," *The Journal of Fixed Income* **1**(2), 7–18.

Bollen, N. (2007). "Mutual fund attributes and investor behavior," *Journal of Financial and Quantitative Analysis* **42**(3), 683–708.

Branch, M., Goldberg, L., and Hand, P. (2019). "A Guide to ESG Portfolio Construction," *The Journal of Portfolio Management* **45**(4), 61–66.

Breidert, C., Hahsler, M., and Reutterer, T. (2006). "A Review of Methods for Measuring Willingness-to-Pay," *Innovative Marketing* **2**(4), 8–32.

Brodback, D., Guenster, N. K., and Mezger, D. (2019). "Altruism and Egoism in Investment Decisions," *Review of Financial Economics* **37**(1), 118–148.

Brodback, D., Guenster, N. K., Pouget, S., and Wang, R. (2021). "The Valuation of Corporate Social Responsibility: A Willingness-to-Pay Experiment," Working Paper.

Capelle-Blancard, G. and Monjon, S. (2014). "The Performance of Socially Responsible Funds: Does the Screening Process Matter?" *European Financial Management* **20**(3), 494–520.

Carson, R. T. (2000). "Contingent Valuation: A User's Guide," *Environmental Science and Technology* **34**(8), 1413–1418.

Cooper, L., Evnine, J., Finkelman, J., Huntington, K., and Lynch, D. (2016). "Social Finance and the Postmodern Portfolio: Theory and Practice," *The Journal of Wealth Management* **18**(4), 9–21.

Cullis, J. G., Lewis, A., and Winnett, A. (1992). "Paying to be Good? U.K. Ethical Investments," *Kyklos International Review of Social Sciences* **45**(1), 3–23.

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

Dorfleitner, G. and Nguyen, M. (2017). "A New Approach for Optimizing Responsible Investments Dependently on the Initial Wealth," *Journal of Asset Management* **18**(2), 81–98.

Engle, R. F., Giglio, S., Kelly, B., Lee, H., and Stroebel, J. (2020). "Hedging Climate Change News," *The Review of Financial Studies* **33**(3), 1184–1216.

Fama, E. F. and French, K. R. (2007). "Disagreement, Tastes, and Asset Prices," *Journal of Financial Economics* **83**(3), 667–689.

"Global Sustainable Investment Review 2020", Global Sustainable Investment Alliance, 2021.

Gomes, O. (2020). "Optimal Growth under Socially Responsible Investment: A Dynamic Theoretical Model of the Trade-off between Financial Gains and Emotional Rewards," *International Journal of Corporate Social Responsibility* **5**(1), 1–17.

Hartzmark, S. M. and Sussman, A. B. (2019). "Do Investors Value Sustainability? A Natural Experiment Examining Ranking and Fund Flows," *The Journal of Finance* **74**(6), 2789–2837.

Hayes, M., Primbs, J. A., and Chiquoine, B. (2015). "A Penalty Cost Approach to Strategic Asset Allocation with Illiquid Asset Classes," *The Journal of Portfolio Management* **41**(2), 33–41.

Heeb, F., Kölbel, J., Paetzold, F., and Zeisberger, S. (2022). "Do Investors Care about Impact?" *The Review of Financial Studies* (forthcoming).

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.

Idzorek, T. M. and Kaplan, P. (2022). "Forming ESG-Oriented Portfolios: A Popularity Approach," *The Journal of Investing* **31**(4), 63–75.

Jessen, P. (2012). "Optimal Responsible Investment," *Applied Financial Economics* **22**(21), 1827–1840.

Kilmurray, D., Melin, L., and Mercereau, B. (2021). "Integrating Impact Funds into Mainstream Portfolios," *The Journal of Impact and ESG Investing* **1**(4), 103–119.

Klement, J. (2018). "Does ESG Matter for Asset Allocation", Fidante Partners.

Kölbel, J. F., Heeb, F., Paetzold, F., and Busch, T. (2020). "Can Sustainable Investing Save the World? Reviewing the Mechanisms of Investor Impact," *Organization & Environment* **33**(4), 554–574.

Lewis, A. and Cullis, J. (1990). "Ethical Investments: Preferences and Morality," *Journal of Behavioral Economics* **19**(4), 395–411.

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

Markowitz, H. (1959). *Portfolio Selection: Efficient Diversification of Investments* (John Wiley & Sons, New York, NY).

Mas, A. and Pallais, A. (2017). "Valuing Alternative Work Arrangements," *American Economic Review* **107**(12), 3722–3759.

McFadden, D. (1984). "Econometric Analysis of Qualitative Response Models," In: *Handbook of Econometrics*, Vol. II. Chap. 24 (Elsevier Science Publishers BV).

McFadden, D. (2001). "Economic Choices," *American Economic Review* **91**(3), 351–378.

Mercereau, B. and Melin, L. (2020). "Optimizing Portfolios across Risk, Return and Climate," *The Journal of Impact and ESG Investing* **1**(1), 115–131.

Miller, K. M., Hofstetter, R., Krohmer, H., and Zhang, Z. J. (2011). "How Should Consumers' Willingness to Pay be Measured? An Empirical Comparison of State-of-the-Art Approaches," *Journal of Marketing Research* **47**, 172–184.

Nofsinger, J. and Varma, A. (2014). "Socially Responsible Funds and Market Crises," *Journal of Banking & Finance* **48**, 180–193.

Oerlemans, L. A., Chan, K. Y., and Volschenk, J. (2016). "Willingness to Pay for Green Electricity: A Review of Contingent Valuation Literature and its Sources of Error," *Renewable and Sustainable Energy Reviews* **66**, 875–885.

Pastor, L., Stambaugh, R., and Taylor, L. (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.

Plagge, J-C. and Grim, D. (2020). "Have Investors Paid a Performance Price? Examining the Behavior of ESG Equity Funds," *The Journal of Portfolio Management* **46**(3), 123–140.

Plagge, J-C., Ertl, M., and Grim, D. (2022). "ESG Equity Index Investing: Don't Forget about Factor Exposures," *The Journal of Beta Investment Strategies* **13**(4), 163–180.

Plott, C. R. and Zeiler, K. (2005). "The Willingness to Pay-Willingness to Accept Gap, the "Endowment Effect", Subject Misconceptions, and Experimental Procedures for Eliciting Valuations," *American Economic Review* **95**(3), 530–545.

Riedl, A. and Smeets, P. (2017). "Why Do Investors Hold Socially Responsible Mutual Funds?" *The Journal of Finance* **72**(6), 2505–2550.

Schmidt, J. and Bijmolt, T. H. A. (2020). "Accurately Measuring Willingness to Pay for Consumer Goods: A Meta-Analysis of the Hypothetical Bias," *Journal of the Academy of Marketing Science* **48**, 499–518.

Serbin, V., Borkovec, M., and Chigirinskiy, M. (2011). "Robust Portfolio Rebalancing with Transaction Cost Penalty: An Empirical Analysis," *Journal of Investment Management* **9**(2), 73–87.

Sharpe, W. (2007). "Expected Utility Asset Allocation," *Financial Analysts Journal* **63**(5), 18–30.

Vanguard, (2021). "Vanguard ESG Investor Research: Full Qualitative Report," [Unpublished manuscript].

Whelan, T., Atz, U., Van Holt, T., and Clark, C. (2021). *ESG and Financial Performance: Uncovering the Relationship by Aggregating Evidence from 1,000 Plus Studies Published Between 2015—2020*. NYU Stern Centre for Sustainable Business.

Zerbib, O. D. (2022). "A Sustainable Capital Asset Pricing Model (S-CAPM): Evidence from Environmental Integration and Sin Stock Exclusion," *Review of Finance* **26**(6), 1345–1388.

Keywords: ESG investing; sustainable investing; portfolio construction with ESG; ESG integration; multi-asset allocation; non-pecuniary preferences; performance measurement