Benchmarking with Sustainability Mandates *

Jason Cen Linquan Chen Yao Chen

Abstract

We develop a novel non-parametric approach to constructing bespoke fund benchmarks that incorporate sustainability or ESG mandates into the actual mutual fund portfolios. By counterfactually evaluating the difference in return between the actual fund and its ESG-mandated benchmarks, we examine whether sustainable investing achieves its dual objectives of superior financial and ESG performances. Our main empirical findings are: i) U.S. active equity funds are concentrated in the *average* or BBB band in ESG rating; ii) on average, the pursuit of higher ESG rating is monotonically associated with higher fund returns; and iii) the implementation of ESG mandates in either direction entails the cost of increasing idiosyncratic risk.

Keywords: Mutual Fund Performance, Sustainable Investing, Benchmark, ESG, Greenwashing

JEL Codes: G11, G23

*This version: March 2024. Please send correspondence to Jason Cen, University of Exeter Business School, Rennes Drive, Exeter EX4 4PU, United Kingdom. Jason Cen (Email: j.cen2@exeter.ac.uk), Linquan Chen (Email: l.chen2@exeter.ac.uk), and Yao Chen (Email: y.chen5@exeter.ac.uk) are all with University of Exeter Business School. The authors thank Daniele Bianchi, Francesco Celentano, Alok Kumar, Ian Marsh, Desislava Vladimirova, and participants at Exeter Finance Brownbag Seminar, the 2023 CEQURA Conference, and the 2023 SESTEF Conference, Quoniam Asset Management Research Seminar for helpful comments. The authors gratefully acknowledge the financial support from the British Academy through the BA/Leverhulme Small Research Grants (SG2122\210971).

1 Introduction

The recent decade has witnessed continuing growth in sustainable investing that incorporates environmental, social, and governance (ESG) factors into portfolio choice decisions (Matos, 2020). According to Global Sustainable Investment Alliance (2020), global sustainable investment reached $35.3 trillion in 2020, a 55% increase from 2016. In promoting responsible business practices, institutional investors assume a pivotal role by employing sustainable investment strategies, including screening (e.g., Hong and Kacperczyk, 2009; Renneboog, Horst, and Zhang, 2011), engagement (e.g., Dimson, Karakaş, and Li, 2015; He, Kahraman, and Lowry, 2023), and ESG integration (e.g., Chen, Chen, Kumar, and Leung, 2023). Despite the growing prevalence of sustainable investing, the extant literature falls short of robustly validating its dual objectives of achieving both superior financial and sustainability performances due to the lack of proper benchmarks for fund-specific performance evaluation. Consequently, a fundamental question remains unanswered: Does sustainable investing benefit investors financially?[1]

This paper bridges this gap by examining the economic trade-off in the concurrent pursuit of both financial and sustainability goals. To this end, we develop a novel methodology for non-parametrically constructing bespoke fund benchmarks with ESG mandates, which enables us to counterfactually assess the impact of ESG mandates on fund performance. The rationale behind our methodology is to perturb a fund's original portfolio to the least extent necessary in terms of loss in diversification relative to the original portfolio, precisely reaching a specified target ESG rating. Constructed from the actual stock holdings of a fund, our bespoke ESG-mandated benchmarks incorporate the existing investment objectives, preferences, and constraints faced by the fund manager, which are otherwise

[1] Understanding financial gains from the ESG pursuit is of vital importance because intuitively the aim of sustainable investing is to include profit and enhance value rather than sacrifice profit and destroy value (Edmans, 2021).

unobservable. The resulting benchmarks feature intuitive ESG-tilts: A minimum-divergence ESG-enhanced (moderated) benchmark increases (reduces) holdings in high-ESG stocks while reducing (increasing) holdings in low-ESG stocks relative to a fund's original portfolio. Our ESG-tilting advocates a more constructive sustainable investing strategy by promoting high-ESG stocks while refraining from the outright exclusion of low-ESG stocks. Empirically, an outperformance (underperformance) of the ESG-enhanced (moderated) benchmark relative to a fund's original portfolio serves as supporting evidence for the compatibility of the dual-objectives in sustainable investing.

Employing our methodology on a sample of actively managed equity funds in the United States spanning from 2007 to 2022, our empirical analysis reveals three main findings:

First, virtually all funds in our sample appear *average* in their ESG ratings. Specifically, we compute fund-level ESG ratings based on firm-level ESG ratings using the relevant portfolio weights. We find that the overwhelming majority of U.S. active equity funds have a BBB rating in contrast to the wider dispersion of the unconditional distribution of firm-level ESG ratings.[2] The concentration of fund-level ESG rating distribution is nearly time-invariant and it is rare to see a fund experiencing any change in its ESG rating. These initial results are crucial as they establish the foundation for our subsequent analyses and results. On the one hand, given the prevalence of *average* ESG ratings among funds, there exists ample opportunity for them to bolster their sustainability profile thanks to the more dispersed nature of firm-level ESG ratings. On the other hand, it appears puzzling why almost every fund settles for an *average* ESG profile, leaning towards neither higher nor lower ratings. Our third main finding below helps rationalize this puzzle.

Second, higher ESG ratings are monotonically associated with higher fund returns

[2]See Section 3.1 for details on the ESG rating scheme.

on average. Specifically, we focus on funds with a BBB ESG rating, constituting the majority within our sample (67%), and aggregate these BBB-rated funds into a fund of funds (FoF). We find that the BBB-FoF could achieve a 3% per annum increase in expected return with an information ratio of 0.5 per annum by tilting its original portfolio toward an AAA rating. In contrast, the BBB-FoF could face a potential loss of over 2% per annum if it opts to lower its ESG target to a CCC rating. These results are not subsumed by the impact of ESG mandates on fund exposures to systematic risk factors.

Third, the superior performance in both ESG rating and fund return is not obtained without entailing associated costs. Specifically, we show that the BBB-FoF would have to take into account the key trade-off between achieving higher ESG targets and heightening idiosyncratic risk. Interestingly, the BBB-FoF would also increase its idiosyncratic risk by downgrading the ESG criteria of its portfolio. This finding provides justification for the concentration of fund-level ESG ratings within the *average* or BBB range since it minimises the idiosyncratic risk.

Taken together, our primary findings answer the fundamental question with important policy implications: sustainable investing does benefit investors financially, though they would have to bear more idiosyncratic risk in the attainment of the dual-objectives. What's more, maintaining at least their current ESG standards is prudent for investors to steer clear of an all-lose scenario where a lower ESG rating, diminished returns, and the burden of additional idiosyncratic risk converge.

Finally, we conduct a battery of additional analyses to scrutinize the robustness and heterogeneity of our results. Specifically, our results cannot be solely explained by the recent asset price bubble during the ESG investment boom. Moreover, there exists substantial heterogeneity across funds in the impact of ESG mandates on fund performance. Last but not least, pillar-specific mandates should be considered separately from the ESG mandate and among the three pillars, only the enhancement of governance (G) results in a significantly positive effect on fund returns.

Related literature The prior literature on sustainable investing centers on evaluating the performance of ESG funds in reality using non-ESG passive indices (e.g., Cremers, Riley, and Zambrana, 2023) and peer groups (e.g., Hartzmark and Sussman, 2019; Pástor and Vorsatz, 2020) as benchmarks. The limitation of these approaches lies in their narrow focus on ESG funds, henceforth unable to address how a fund, especially non-ESG, could benefit from adopting an ESG mandate. This paper presents a methodological contribution through a novel approach for non-parametrically constructing ESG-mandated benchmarks applicable to any fund, irrespective of its current ESG label. Our approach enables the integration of continuous, multi-dimensional ESG targets, providing flexibility to assess the impact of both ESG enhancement and moderation on fund performance.

Our methodology also contributes broadly to the literature on benchmarking fund performance (e.g., Berk and van Binsbergen, 2015; Daniel, Grinblatt, Titman, and Wermers, 1997; Pástor, Stambaugh, and Taylor, 2015). Our bespoke benchmarks enable more precise performance evaluations tailored to each fund, preventing inappropriate comparisons between fundamentally different funds. Unlike Beber, Brandt, Cen, and Kavajecz (2021), who use a parametric portfolio approach to constructing bespoke benchmarks, our non-parametric bespoke benchmarking method derives from each fund's actual stock holdings, henceforth capturing skill, preferences, and constraints implicitly.

In addition to our methodological contributions, we also empirically contribute to the growing literature examining the relation between ESG and financial performance.[3] Empirical evidence documented so far in the existing literature has

[3] Pástor, Stambaugh, and Taylor (2021) show that ESG performance is expected to be negatively associated with average future stock returns based on an equilibrium model that incorporates ESG criteria in portfolio decisions. On the other hand, based on an ESG-efficient frontier, Pedersen, Fitzgibbons, and Pomorski (2021) find that ESG ratings could either enhance or deteriorate financial performance, depending on whether the ratings contain material information.

been mixed. At the stock level, for instance, Khan, Serafeim, and Yoon (2016) find a positive relation between stock returns and firm-level performance on material ESG issues, while Hong and Kacperczyk (2009) show that sin stocks (i.e., stocks with low ESG ratings) earn higher returns. In contrast, Avramov, Cheng, Lioui, and Tarelli (2022) claim no significant difference between high- and low-ESG stock returns[4]. At the fund level, the empirical evidence documented in the existing literature has not reached an agreement either. Among others, Cremers, Riley, and Zambrana (2023) develop an *Active ESG Share* metric that captures fund managers' activeness in using ESG information. They document a positive relation between Active ESG Share and future fund performance among ESG funds. In contrast, Orlov, Ramelli, and Wagner (2023) find a negative impact of mutual fund managerial ownership on portfolio sustainability, indicating that fund managers do not consider ESG enhancement as a way to improve risk-adjusted returns. Further, Hartzmark and Sussman (2019) and Geczy, Stambaugh, and Levin (2021) find no significant return difference between high and low sustainability funds. Through a comprehensive performance evaluation analysis using our novel fund-specific benchmarking method, our empirical findings shed light on the feasibility of achieving both improved financial and sustainability outcomes, along with the associated costs linked to attaining these dual objectives.

The paper proceeds as follows. Section 2 introduces our methodology including our non-parametric portfolio approach. Section 3 describes the data. Section 4 presents results of our empirical analysis. Section 5 concludes.

[4]For more examples, see Masulis and Reza (2015) and Cheng, Hong, and Shue (2023) who show that corporate involvement in ESG activities is driven by agency problem and reduces firm value in contrast to Ferrell, Liang, and Renneboog (2016) and Albuquerque, Koskinen, and Zhang (2019) who find a positive relation between ESG and firm value. Also see Friede, Busch, and Bassen (2015) for a review.

2 Methodology

In this section, we first introduce a non-parametric approach to constructing bespoke fund benchmarks that incorporate sustainability mandates. We then present our methodology of using these bespoke ESG-mandated benchmarks to estimate the impact of potential sustainability mandates on fund performance.

2.1 Minimum-divergence benchmarks with ESG mandates

The key idea of our approach is to create a fund's ESG-mandated benchmark as close to its original portfolio as possible only to achieve a prescribe ESG target. Since the investment objectives, constraints, and manager skill and preference of each fund are not directly observable[5], we build the benchmark portfolio from each fund's actual stock holdings which should have encoded such information as a result of the fund manager's optimal portfolio choice.

We start by establishing the notations. A universe of active equity mutual funds are indexed by $j = 1, 2, \ldots, J_t$, where J_t is the total number of funds at time t. The original long-only portfolio held by fund j at time t is represented by an $N_t \times 1$ vector of weights $w_t^{jo} = (w_{1,t}^{jo}, w_{2,t}^{jo}, \ldots, w_{N_t,t}^{jo})'$. Stocks in the time-$t$ investment universe are indexed by $i = 1, 2, \ldots, N_t$, where N_t is the total number of investible stocks at time t. Let $X_{i,t}$ denote the firm-level ESG rating associated with stock i at time t. Hence, the original fund-level ESG rating for fund j is $x_t^{jo} = \sum_{i=1}^{N_t} w_{i,t}^{jo} X_{i,t}$.

Next, given a fund's original portfolio, we formally define its benchmarks with ESG mandates as follows:

Definition 2.1 (ESG-enhanced and ESG-moderated benchmarks)**.** At time t and for a fund j, the ESG-enhanced benchmark is a set of portfolios that target at fund-level ESG rating of x higher than the fund's original rating of x_t^{jo} and the set of long-only

[5]The prospectus of a fund does not necessarily reveal the complete and precise description of the true risk preference and investment objectives of the fund manager.

portfolio weights $w_t^j = (w_{1,t}^j, w_{2,t}^j, \ldots, w_{N_t,t}^j)'$ is defined as

$$\mathcal{B}^{+}(x, x_t^{jo}) = \left\{ w_t^j \,\middle|\, w_{i,t}^j > 0 \text{ for } \forall i,\ \sum_{i=1}^{N_t} w_{i,t}^j = 1,\ \sum_{i=1}^{N_t} w_{i,t}^j X_{i,t} \geq x, \text{ for } x \in (x_t^{jo},\, \overline{x}) \right\}, \quad (1)$$

On the other hand, the fund's ESG-moderated benchmark is a set of portfolios that target at a fund-level ESG rating of x lower than the fund's original rating x_t^{jo} and the set of portfolio weights is defined as

$$\mathcal{B}^{-}(x, x_t^{jo}) = \left\{ w_t^j \,\middle|\, w_{i,t}^j > 0 \text{ for } \forall i,\ \sum_{i=1}^{N_t} w_{i,t}^j = 1,\ \sum_{i=1}^{N_t} w_{i,t}^j X_{i,t} \leq x, \text{ for } x \in (\underline{x},\, x_t^{jo}) \right\}. \quad (2)$$

Evidently, the above-defined fund benchmarks for any given ESG target of x are generally not unique, which renders them limited usefulness for evaluating the impact of potential sustainability mandates on fund performance. It is therefore vital to refine the characterization of the sets by, for instance, restricting our attention to the ESG-enhanced and ESG-moderated benchmarks, respectively, in the nearest vicinity of each fund's original portfolio. It is worth noting that the weights of any long-only portfolio are consistent with the structure of a probability distribution as they are non-negative and sum up to one. Enlightened by this fact, we measure the distance between each fund's original portfolio and its benchmark for a given ESG target x by the Kullback-Leibler Information Criterion divergence (I-divergence) (Kullback and Leibler, 1951):

$$\mathcal{D}_{KL}\left(w_t^j \| w_t^{jo}\right) \equiv \sum_{i=1}^{N_t} w_{i,t}^j \log\left(\frac{w_{i,t}^j}{w_{i,t}^{jo}}\right), \quad (3)$$

which has been broadly applied in statistics and information theory to quantify how one probability distribution diverges from the other or the information loss of using one probability distribution to approximate the other[6]. The I-divergence is always non-

[6]The asset pricing literature has adopted the relative entropy minimization approach to recover the risk-neutral probability distribution and the stochastic discount factor (SDF) (e.g., Stutzer, 1995; Ghosh, Julliard, and Taylor, 2019). On the other hand, Bera and Park (2008) has introduced a portfolio choice problem based on the maximum entropy principle.

negative and it attains the zero lower bound if and only if the fund's original portfolio and its ESG-mandated benchmark are identical. To see this, one can simply rearrange Eq. (3) and proceed using Jensen's inequality as follows:

$$\sum_{i=1}^{N_t} w_{i,t}^j \log\left(\frac{w_{i,t}^j}{w_{i,t}^{jo}}\right) = -\sum_{i=1}^{N_t} w_{i,t}^j \log\left(\frac{w_{i,t}^{jo}}{w_{i,t}^j}\right) \geq -\log\left(\sum_{i=1}^{N_t} w_{i,t}^j \frac{w_{i,t}^{jo}}{w_{i,t}^j}\right) = 0\,,$$

where the equality holds if and only if $w_{i,t}^j = w_{i,t}^{jo}$ for $\forall\, i$. The I-divergence defined in Eq. (3) has an appealing economic interpretation: the loss in diversification of the fund-specific benchmark w_t^j relative to the original portfolio w_t^{jo} for fund j at time t. This interpretation is motivated by the fact that the entropy measure of the weights of a long-only portfolio indicates its level of diversification. To see this, consider a general long-only portfolio $w = (w_1, \ldots, w_N)'$. Its diversification can be measured by the entropy of its weights, $\mathsf{Entropy}(w) = -\sum_{i=1}^N w_i \ln(w_i)$. Intuitively, for an extremely diversified portfolio such as an equally-weighted portfolio, $\mathsf{Entropy} = \ln(N)$ while for an extremely concentrated portfolio such as one that invests in only one stock, $\mathsf{Entropy} \to 0$. Furthermore, note that since

$$\begin{aligned}\mathsf{Entropy}(w) = -\sum_{i=1}^N w_i \ln(w_i) &= -\sum_{i=1}^N w_i \left[\ln(w_i) - \ln\left(\frac{1}{N}\right)\right] + \ln(N) \\ &= -\mathcal{D}_{KL}\left(w \middle\| \frac{1}{N}\right) + \ln(N)\,,\end{aligned}$$

the maximization of entropy or diversification for a general long-only portfolio is equivalent to the minimization of the I-divergence or loss in diversification of the portfolio relative to an equally-weighted portfolio.

Motivated by the economic interpretation of I-divergence above, we now formally define the minimum-divergence benchmarks with ESG mandates as follows:

Definition 2.2 (Minimum-divergence benchmarks with ESG mandates)**.** At time t, and with a prescribed ESG target of $x \in (\underline{x}, \overline{x})$, the minimum-divergence benchmark for fund j is a portfolio whose weights $w_t^{j\star} = (w_{1,t}^{j\star}, w_{2,t}^{j\star}, \ldots, w_{N_t,t}^{j\star})'$, are defined as

$$w_t^{j\star}(x, x_t^{jo}) = \begin{cases} \underset{w_t^j \in \mathcal{B}^+(x, x_t^{jo})}{\arg\min} \quad \mathcal{D}_{KL}\left(w_t^j \| w_t^{jo}\right), & \text{for} \quad x > x_t^{jo} \\ \\ \underset{w_t^j \in \mathcal{B}^-(x, x_t^{jo})}{\arg\min} \quad \mathcal{D}_{KL}\left(w_t^j \| w_t^{jo}\right), & \text{for} \quad x < x_t^{jo}\,. \end{cases} \tag{4}$$

Specifically, when $x > (<)\, x_t^{jo}$, $w_t^{j\star}(x, x_t^{jo})$ represents the minimum-divergence ESG-enhanced (ESG-moderated) benchmark for fund j at time t.

What makes the I-divergence more appealing than other (pseudo) distance metrics? First, the logarithm functional form preserves the long-only nature of the original portfolio in the benchmarks. This property ensure the investment rationale that sustainable investing be focused on the allocation of capital on assets rather than resorting to short-sale or other speculative techniques to achieve a sustainability objective. In contrast, other distance metrics such as the sum of squared portfolio weight differentials do not necessarily guarantee the non-negativity of portfolio weights. Second, given that the entropy measure of the weights of a long-only portfolio indicates diversification, the minimization of the I-divergence offers the most parsimonious solution that merely requires the minimum amount of loss in diversification relative to the original portfolio to fulfill any prescribed ESG target. Last but not least, the I-divergence encodes information about all the moments of cross-sectional distribution of the relative portfolio weight differentials $\xi_{i,t}^j = \frac{w_{i,t}^j - w_{i,t}^{jo}}{w_{i,t}^{jo}}$ for $\forall\, i = 1, 2, \ldots, N_t$. To see this, one can apply the Taylor expansion to the I-divergence with regard to the relative weight differentials

$$\mathcal{D}_{KL}\left(w_t^j \| w_t^{jo}\right) \approx \frac{1}{2}\sum_{i=1}^{N_t} w_{i,t}^{jo}\,(\xi_{i,t}^j)^2 - \frac{1}{6}\sum_{i=1}^{N_t} w_{i,t}^{jo}\,(\xi_{i,t}^j)^3 + \frac{1}{12}\sum_{i=1}^{N_t} w_{i,t}^{jo}\,(\xi_{i,t}^j)^4 - \cdots\,. \tag{5}$$

Essentially, the I-divergence penalizes the variance (represented by the first term on the right-hand side of Eq. (5)), negative skewness (the second term), and kurtosis (the third term) of portfolio weight differentials. Therefore, the I-divergence minimization aims not only to reduce the dispersion of portfolio perturbation across stocks but also to mitigate the need of extreme rebalancing. In contrast, the sum of squared portfolio weight differentials captures only the variance and therefore are more susceptible to drastic portfolio reshuffling.

At first sight, the optimal control in Eq. (4) is a high-dimensional ($N_t \geq 3000$) constrained minimization problem that would be computationally intensive to solve. However, we formally show that the problem can be transformed so that the solution is quasi-analytical subject to a one-dimensional unconstrained minimization problem as follows:

Proposition 1. *At time t, and with a prescribed ESG target of $x \in (\underline{x}, \overline{x})$, the minimum-divergence benchmark for fund j defined in Eq. (4), is solved as*

$$w_{i,t}^{j\star}(x, x_t^{jo}) = \frac{w_{i,t}^{jo} \exp\left(\lambda^{j\star}(x, x_t^{jo})\,(X_{i,t} - x)\right)}{\sum_{i=1}^{N_t} w_{i,t}^{jo} \exp\left(\lambda^{j\star}(x, x_t^{jo})\,(X_{i,t} - x)\right)} \tag{6}$$

where the Lagrangian multiplier $\lambda^{j\star}$ solves the following unconstrained minimization problem

$$\lambda^{j\star}(x, x_t^{jo}) = \arg\min_{\lambda^j} \sum_{i=1}^{N_t} w_{i,t}^{jo} \exp\left(\lambda^j\,(X_{i,t} - x)\right) . \tag{7}$$

Specifically, when $x \in (x_t^{jo}, \overline{x})$, we have $\lambda^{j\star} > 0$ and the solution (Eq. (6)) represents a minimum-divergence ESG-enhanced benchmark; when $x \in (\underline{x}, x_t^{jo})$, we have $\lambda^{j\star} < 0$ and the solution (Eq. (6)) represents a minimum-divergence ESG-moderated benchmark.

Corollary 1. *In the special case where the ESG target coincides with the original ESG rating ($x = x_t^{jo}$), all ESG-mandated benchmarks of fund j converges to its original portfolio (w_t^{jo}). Formally, $w_{i,t}^{j\star}(x_t^{jo}, x_t^{jo}) \equiv w_{i,t}^{jo}(x_t^{jo}) \equiv w_{i,t}^{jo}$, for $\forall i$, and $\lambda^{j\star}(x_t^{jo}, x_t^{jo}) = 0$.*

Intuitively, the minimum-divergence ESG-enhanced benchmark has a tilt toward stocks with higher ESG ratings whereas the minimum-divergence ESG-moderated benchmark has a tilt toward stocks with lower ESG ratings. The main properties of the ESG-tilt are illustrated in Figure 1. The minimum-divergence benchmark tends to perturb the portfolio by increasing holdings of high-ESG stocks more aggressively than selling low-ESG stocks, in stark contrast to more radical methods such as negative screening that excludes low-ESG stocks entirely. Furthermore, the solution also suggests that our minimum-divergence benchmarks are more robust when the cross-sectional distribution of firm-level ESG ratings deviates from normality because the I-divergence summarizes information regarding all the moments of the cross-sectional distribution of firm-level ESG ratings.

2.2 Modification of the observed fund portfolio

A technical issue arises when in practice a fund's observed actual portfolio weights, denoted by $w_{i,t}^{j,Raw}$, is used to represent the fund's original portfolio $w_{i,t}^{jo}$. The reason is that the portfolio perturbation in our methodology relies on the assumption that the fund manager makes non-zero investment in all stocks. However, in reality, the observed actual portfolio of the fund may not cover the entire cross section of U.S. stocks, i.e., it contains zero weights so that $w_{i,t}^{j,Raw} = 0$ for at least some stocks. The fund's benchmark with any ESG mandates would never be able to invest in stocks not currently held in the actual observed portfolio of the fund because for these zero-weight stocks, we have $\log\left(w_{i,t}^{j,Raw}\right) = -\infty$. The fund would have to incur an infinite amount of cost of deviation from its orignial optimal portfolio choice when altering the weights from $w_{i,t}^{j,Raw} = 0$ to $w_{i,t}^{j\star} > 0$. Therefore, the optimal solution to estimating the ESG-mandated benchmark in this case would have to be keeping these zero-weight stocks as they are, effectively restricting the investment universe to include only stocks currently held. Such restrictions might render some ESG targets unattainable when, for example, the target is AAA whereas the top-rated stock currently held by the fund has an ESG rating of A.

To overcome this technical issue, we transform the observed actual portfolio of each fund $w_{i,t}^{j,Raw}$ to an original portfolio $w_{i,t}^{jo}$ by a small modification $0 < \kappa \ll 1$ such that

$$w_{i,t}^{jo} = (1-\kappa)w_{i,t}^{j,Raw} + \kappa\, w_{i,t}^{j,Not}\,, \tag{8}$$

where $w_{i,t}^{j,Not}$ represents the weights of a value-weighted portfolio of stocks not currently held in the fund's observed actual portfolio. In our subsequent empirical analysis, we set $\kappa = 0.01$. Through the untabulated sanity check, we verify that the returns of the actual observed portfolio and the original portfolio modified as above for a fund are virtually perfectly correlated; moreover, the correlation is not sensitive to the choice of the parameter κ as long as $0 < \kappa \ll 1$. Hence, we use the original portfolio $w_{i,t}^{jo}$ as the input of our estimation of the minimum-divergence benchmarks with ESG mandates as it not only effectively represents the actual observed fund portfolio $w_{i,t}^{j,Raw}$ but also addresses the technical issue abovementioned.

2.3 Performance evaluation of sustainable investing

Combining portfolio weights for each fund and its benchmarks with stock returns, we obtain the time series of excess returns

$$R_{t+1}^{j\star}(x, x_t^{jo}) - R_{t+1}^{jo}(x_t^{jo}) = \sum_{i=1}^{N_t} w_{i,t}^{j\star}(x, x_t^{jo})R_{i,t+1} - \sum_{i=1}^{N_t} w_{i,t}^{jo}(x_t^{jo})R_{i,t+1}\,, \tag{9}$$

where $R^{jo}(x_t^{jo})$ is the gross return of fund j's original portfolio with an original fund-level ESG rating of x_t^{jo}, and $R^{j\star}(x, x_t^{jo})$ is the gross return of fund j's benchmark if the fund's ESG rating target is set at $x \neq x_t^{jo}$. Another technical issue emerges from the fact that the original fund-level score, x_t^{jo}, is likely to vary through time for a fund: it would therefore be inappropriate to estimate the impact of ESG mandates on fund performance by directly taking a sample average of (9) because we aim to evaluate the impact of a potential ESG mandate (indicated by a given ESG target of x) for a fixed original fund-level ESG rating of x^o. To address this issue, we proceed to aggregate funds to an asset under management (AUM)-weighted portfolio, or a fund of funds (FoF), for any given time-invariant original fund-level ESG rating of x^o. The gross return of this x^o-FoF is denoted by

$$\mathcal{R}^{o}_{t+1}(x^o) = \sum_{j=1}^{J_t} \omega_t^j R^{jo}_{t+1}(x_t^{jo}) \, \mathbb{I}\left\{x_t^{jo} = x^o\right\} \,, \tag{10}$$

where $\mathbb{I}\left\{x_t^{jo} = x^o\right\}$ is a dummy variable equal to one if $x_t^{jo} = x^o$ and zero otherwise. Similarly, by aggregating the benchmarks targeting at an ESG rating of x for all funds with an original ESG rating of x^o, we obtain the x-benchmark for the x^o-FoF whose gross return is denoted by

$$\mathcal{R}^{\star}_{t+1}(x, x^o) = \sum_{j=1}^{J_t} \omega_t^j R^{j\star}_{t+1}(x, x_t^{jo}) \, \mathbb{I}\left\{x_t^{jo} = x^o\right\} \,. \tag{11}$$

Ultimately, this aggregation enables us to evaluate the unconditional impact of an ESG target x given a fixed original ESG rating of x^o by the difference in expected return between the x^o-FoF and its x-benchmark:

$$\underbrace{\mathsf{E}\left[\mathcal{R}^{\star}_{t+1}(x, x^o) - \mathcal{R}^{o}_{t+1}(x^o)\right]}_{\text{Unconditional Performance}} = \mathsf{E}\left[\sum_{j=1}^{J_t} \omega_t^j \underbrace{\mathsf{E}_t\left[R^{j\star}_{t+1}(x, x_t^{jo}) - R^{jo}_{t+1}(x_t^{jo})\right]}_{\text{Conditional Performance}} \mathbb{I}\left\{x_t^{jo} = x^o\right\}\right] \,. \tag{12}$$

.

3 Data

In this section, we describe the datasets used in the subsequent empirical analysis.

3.1 Firm-level ESG ratings

We collect firm-level ESG ratings from MSCI. MSCI reports industry-adjusted ESG scores using both numeric scores and letter-based ratings. The numeric ESG score ranges from 0 to 10 while the letter-based rating of a firm takes one of the following values: AAA, AA, A, BBB, BB, B, and CCC (see Figure 2, source: MSCI ESG Research (2023, p. 7)). Firms rated as A or above (i.e., with an ESG score no lower than 7.143) are labeled as ESG leaders while firms with ratings lower than BB (i.e., scores below 2.857) are labeled as ESG laggards. A firm whose ratings lies between A and BB is labeled to deliver an average ESG performance.

Our key variable is the industry-adjusted ESG score. For firms not covered by MSCI, we proxy for their ESG ratings using the following steps. First, following the latest MSCI methodology, we manually assign each stock in our sample to an ESG sub-industry, an ESG rating industry, and an ESG sector using the GICS classification obtained from Compustat. For firms with retired GICS codes, we use the GICS historical files obtained from S&P to assign corresponding ESG industries. In addition, for firms with missing GICS classification, we use firm names or main businesses to identify their ESG industries. Next, in each month, we calculate the average ESG scores at the sub-industry, industry, and sector levels. For firms not covered by MSCI, we proxy for their ESG ratings using the average ESG score of the corresponding sub-industry. For firms without sub-industry classification, we use the average ESG score at a broader (i.e., industry or sector) level to proxy their ESG performance.

3.2 Mutual fund data

We obtain mutual fund returns and other fund characteristics from the CRSP survivorship-bias-free U.S. mutual fund dataset. We focus on actively managed equity funds and exclude international, index, balanced, sector, target-date, and hedge funds.

We obtain quarterly mutual fund stock holdings from Thomson/Refinitiv and merge the two data sets using the MFLINKS provided by WRDS. We calculate the fund-level ESG rating as an weight average of ESG ratings of individual stocks held in the portfolio using the corresponding portfolio weights.

3.3 Other data

We obtain monthly stock returns from CRSP and accounting information from Compustat. We focus on common stocks (share code equals 10 or 11) that are listed on NYSE, NYSE American, or NASDAQ. We exclude penny stocks whose share price is below $1. Finally, we obtain risk factors from Ken French's data library. Our final sample covers 3,182 unique equity funds holding more than 7,000 unique stocks from January 2007 to December 2022.

4 Empirical analysis

4.1 Preliminaries

Before evaluating the performance of ESG enhancement and moderation strategies, we first present the ESG characteristics of all common stocks and equity funds in our sample.

Figure 3a exhibits the unconditional distribution of firm- and fund-level ESG scores. Specifically, the region between the two dotted lines indicates the letter-based rating of BBB, while the region between the two solid lines indicates firms or funds with average ESG performance (i.e., with ESG ratings of A, BBB, or BB). We find that the majority of U.S. active equity funds have a BBB letter rating, and almost all funds appear *average* in light of ESG performance. In contrast, the unconditional distribution of firm-level ESG scores shows remarkably wider dispersion. While the majority of firms fall in the average group, the mass of firms as ESG leaders or laggards are non-trivial, which makes feasible both ESG enhancement and moderation strategies in constructing benchmarks with ESG mandates.

Next, to have a closer examination of fund-level ESG scores in contrast to firm-level ESG scores, we present the evolution of cross-sectional distributions of ESG scores in Figure 3b. We find an upward trend in both fund- and firm-level ESG performance. Since larger stocks tend to have better ESG ratings and higher portfolio weights, the average fund-level ESG rating is always higher than the average stock-level rating. In addition, while all funds fall within the average category (i.e., within the two solid lines), a large proportion of stocks belong to the leader or laggard group (i.e., above or below the two solid lines).

Finally, we show how the ESG scores vary through time for individual funds in contrast to individual firms. Specifically, we compute the time-series standard deviation of ESG scores for each individual funds and firms and plot the distributions of these time-series standard deviations for individual funds (in red shades) and individual firms (in blue shades), respectively, in Figure 3c. The vertical dotted line indicates that a fund or firm has to increase its numeric ESG score by at least 1.429 to raise its letter-based rating by one notch (i.e., from BBB to A). Similarly, the vertical solid line indicates that a fund or a firm has to increase its numeric ESG score by 2.857 to promote its ESG profile by a category (i.e., from Average to Leader). We show that the time-series standard deviation of number ratings is smaller than 1 for almost all funds, which suggests that fund-level letter-based ratings are unlikely to change over time. In contrast, the time-series standard deviation of ESG scores is larger for stocks, and a small fraction of such changes are big enough to shift the letter rating or even the ESG leader/laggard status.

Taken together, Figure 3 shows that U.S. active equity funds tend to look average in terms of their ESG performance and their ESG ratings remains nearly constant over time. In contrast, stock-level ESG ratings feature a more dispersed distribution in terms of variations both in the time-series and in the cross-section. Therefore, it is feasible for these funds to achieve different ESG objectives by tilting portfolios toward stocks in the leader or laggard group.

4.2 Impact of ESG mandates on fund returns

In this section, we examine how shifts in ESG targets affect mutual fund returns. Since the majority of funds have BBB ratings, we form a value weighted mutual fund portfolio of all BBB funds using asset-under-management (AUM) as portfolio weights. This BBB-rated fund of funds (FoF) ensures that all funds included in this portfolio meet the BBB rating requirement. We then examine the difference in performance between the BBB-FoF and its benchmarks with enhanced or moderated ESG mandates. Figure 4a presents the results. We find that expected return increases monotonically with ESG targets. For instance, when the BBB FoF introduced an ESG-enhanced mandate that increases its rating from BBB to AAA, it would earn an additional 3% per annum in expected returns which is statistically significant based on Newey-West standard errors (Newey and West, 1987) with a lag of 3 months. This result suggests that on average a BBB-rated fund could achieve the dual objective of better financial and sustainability performances. On the contrary, the BBB FoF would lose more than 2% per annum in expected return for an ESG-moderation that downgrades its ESG rating from BBB to CCC.

To adjust for the effect of volatility on performance measurement, we also report the information ratios associated with the difference in return between the BBB FoF and its ESG-mandated benchmarks in Figure 4b. We find that the implementation of ESG-enhanced mandates would result in consistently higher information ratio. For instance, an ESG-enhanced mandate that increases the FoF's rating from BBB to AAA would significantly increase the annualized information ratio by 0.5, which is statistically significant based on stationary bootstrapping standard errors (Politis and Romano, 1994) with an expected lag length of 3 months. It is worth noting that the marginal effect is strongest when the BBB FoF improves its ESG performance by only one notch, i.e., from BBB to A. This finding suggests that mutual funds do not have to pursue aggressive ESG enhancement strategies to achieve the dual objective of sustainable investing. Even a marginal ESG enhancement could be economically meaningful to improve a fund's financial performance on average. In contrast, targeting a slightly lower ESG criteria (i.e., from BBB to BB) will significantly lower the information ratio. Overall, these findings demonstrate that on average funds can achieve higher returns when implementing ESG-enhanced mandates.

One potential explanation on the positive relation between ESG mandates and fund performance could be that implementing different ESG targets are simply repacking exposures to risk factors. To address this potential concern, we examine loadings of the BBB FoF returns to standard risk factors by estimating the following Fama-French-Carhart 4-factor model (Carhart, 1997):

$$\mathcal{R}_t^\star(x, BBB) - \mathcal{R}_t^o(BBB) = \alpha + \boldsymbol{\beta}'\mathbf{f}_t + \epsilon_t\,, \quad \text{for any given ESG target } x \tag{13}$$

where the vector f_t comprises the market factor (MKT), the size factor (SMB), the value factor (HML), and the momentum factor (WML). Figure 5 presents the results. We find that the impact of ESG mandates on exposures to the market and momentum factors are statistically insignificant (Panels (a) and (d)). On the other hand, we find that ESG-enhanced mandates tend to reduce the BBB FoF's exposures to the size and value factors while ESG-moderated mandates tend to increase exposures to these two factors (Panels (b) and (c)). These effects are statistically significant, though their magnitudes are economically small. To isolate the effect of ESG mandates on factor exposures, we report the difference in abnormal return (α) between the BBB FoF and its ESG-mandated benchmarks in Figure 6. Consistent with the evidence regarding expected return differences documented in Figure 4a, alpha increases monotonically with ESG targets in virtually the same magnitude as expected return. For instance, the BBB FoF would earn an extra 2% per annum in abnormal return, after controlling for size and value factor exposures. Therefore, the positive relation between ESG score and fund performance remains robust despite the effect of ESG mandates on factor exposures.

4.3 Cost of ESG mandates

Our findings so far suggest that a fund's financial performance seems aligned with its ESG rating in that fund managers on average could achieve the dual objective of higher return and better ESG profile by altering their existing portfolios. The implication of these findings appears puzzling. Why do funds still appear *average* in ESG rating in reality if they could improve their financial performance by promoting their ESG criteria? In this section, we investigate the potential costs associated with implementing ESG mandates.

Specifically, we examine the impact of ESG mandates on the idiosyncratic risks of the BBB FoF. The results are reported in Figure 7a. We find that both the ESG-enhanced and ESG-moderated mandates would increase the idiosyncratic volatility of returns significantly. This finding is intuitive since most stocks belong to the average ESG rating group (i.e., A, BBB, or BB). The ESG-enhanced (moderated) mandate would restrict the investment universe toward ESG leaders (laggards) while the original BBB FoF could potentially invest in any firms. Figure 7b presents the changes in R-squared when deviating from the BBB FoF original portfolio. Consistent with the effect on idiosyncratic volatility, we show that the R-squared would decrease when a fund deviates from BBB rating, indicating an increase in idiosyncratic risk. Overall, we show that although tilting portfolios toward high ESG firms could improve risk-adjusted fund performance, such enhancement strategies would also increase portfolio idiosyncratic risk.

Consistent with our findings that implementing ESG mandates that target at different ESG scores would increase the idiosyncratic risks in the BBB FoF's returns, we also find that more turnover of the BBB FoF's original portfolio is required when the ESG mandate aims for a score farther away from BBB, regardless of the direction of the mandate (see Figure 8).

4.4 Additional analysis

In this section, we provide additional analysis that examines the robustness and heterogeneity of our results on the effect of ESG mandates on fund performance.

4.4.1 ESG bubble

One potential critique of our results on the positive relation between ESG targets and fund performance is that they are soley driven by the ESG investment boom in the recent decade. Such a critique implies that the positive relation between ESG target and fund performance would only exist in the recent decade during the ESG investment boom. On the contrary, we find that the outperformance of ESG-enhanced benchmarks have existed before the recent decade's ESG investment boom. In fact, the AAA-benchmark outperforms the BBB FoF while the CCC-benchmark underperforms in our entire sample period except for a brief period during the peak of the 2008 Global Financial Crisis (Figure 9a).

To further address the critique that our results are solely driven by asset price bubbles during the recent ESG investment boom, we compare the cumulative returns of the ESG-enhanced benchmarks for the BBB FoF and those of the ESG fund in reality, the aggregate ESG portfolio documented by van der Beck (2023) in Figure 9b. Consistent with van der Beck (2023) who documents a significant increase in ESG fund flows (accompanied by a growing number of ESG funds, and more frequent fund name changes) in the 2016-2022 period, we find that ESG funds in reality, represented by the aggregate ESG portfolio, closely tracked our AAA-benchmark before 2017 and started to diverge more and more remotely from our AAA-benchmark afterwards. In fact, the aggregate ESG portfolio have not even remained close to our single A-benchmark since 2017. This finding not only suggests noteworthy greenwashing in the recent ESG investment boom but also rule out the ESG-flow driven asset price bubble as the sole explanation for our findings on the positive relation between ESG targets and fund performance.

4.4.2 Heterogeneity

Our empirical analysis has focused on the aggregate or average fund level by examining returns of the BBB FoF and its benchmarks mandated with different levels of ESG rating. Another potential critique emerges from this analysis that our results are entirely driven by an unconditionally positive correlation between firm-level ESG scores and stock expected returns. In this sense, it would be irrelevant to examine fund-specific benchmarks with ESG mandates. We argue that this critique is invalid for our findings for three reasons:

To begin with, the literature has yet to reach a consensus whether stocks issued by firms with higher ESG rating indeed deliver higher expected return. This fact motivates us to shift the focus from stock-level analysis to fund-level analysis in this paper.

To further address concerns arising from this critique, we show the cross-sectional distribution of the difference in expected return between each BBB-rated fund and its ESG enhanced (AAA) and ESG-moderated (CCC) benchmarks in Figure 10a. Based on this potential critique, the effect of ESG mandates on fund performance would be homogeneous across funds. Consistent with our previous results, a BBB-rated funds on average tends to earn an increase in expected return by 2% per annum when pursuing an ESG-enhanced (AAA) mandate while it tends to lose 2% per annum when pursuing an ESG-moderated (CCC) mandate. However, there exists substantial heterogeneity of this effect across funds, which suggests that the dual objective of both higher ESG rating and higher returns is feasible for some funds and not for others. We also estimate the abnormal return (alpha) for each BBB-rated fund based on the Fama-French-Carhart four-factor model and present the cross-sectional distribution of the difference in alpha between each BBB-rated fund and its ESG-enhanced (AAA) and ESG-moderated (CCC) benchmarks in Figure 10b. The result remains the same.

4.4.3 Pillar-specific mandates

The concept of sustainable or ESG investing encompasses three pillars: Environmental (E), Social (S), and Governance (G). Given our findings on the impact of ESG mandates on fund performance, it is legitimate to extend our analysis to the impact of pillar-specific mandates on fund performance. Figure 11 summarizes the results: While a BBB-rated FoF in E, S, or G tends to deliver higher alpha if it pursues a higher target for the specific pillar, such effect is only statistically significant for the G-pillar. Moreover, the effect of G-enhanced mandates is also economically significant. For example, a FoF with G-pillar rating of BBB would earn an extra alpha of more than 5% per annum if it implements a G-mandate targeting at the AAA rating.

Note that it would be a misinterpretation to claim that our main finding on the positive relation between ESG targets and average fund performance is driven by the governance rating. The reason is that neither is the ESG rating a linear function of the three underlying pillars, nor targeting at ESG can be linearly attributed to targeting at each individual pillars. Hence, ESG mandates and pillar-specific mandates should be considered and analyzed separately.

5 Conclusion

In this paper, we develop a novel non-parametric approach to constructing fund benchmarks that incorporate sustainability or ESG mandates into the actual mutual fund portfolios. Using this new method, we perform counterfactual evaluation that compares the actual fund returns with ESG-mandated benchmark portfolio returns to examine whether sustainable investing achieves its dual objective of superior financial and ESG performances.

Our empirical analysis leads to the following main findings. First, all funds in our sample appear *average* in terms of their fund-level ESG ratings. We find that the vast majority of U.S. active equity funds have a BBB rating in contrast to the wider dispersion of the unconditional distribution of firm-level ESG ratings. Second, the dual objectives of improving ESG rating and fund returns are aligned with each other on average. A BBB FoF would increase expected return by 3% per annum at an information ratio of 0.5 per annum by improving its ESG rating to AAA. Third, the attainment of superior fund performance in both ESG rating and return entails the cost of higher idiosyncratic risks.

Through a series of additional analyses, we show that our results cannot be solely explained by the recent asset price bubble from the ESG investment boom. Moreover, there exists substantial heterogeneity across funds in the impact of ESG mandates on fund performance. Last but not least, pillar-specific mandates should be considered separately from the ESG mandate and only the Governance (G)-enhanced mandates result in higher fund returns.

References

Albuquerque, Rui, Yrjo Koskinen, and Chendi Zhang, 2019, Corporate social responsibility and firm risk: Theory and empirical evidence, *Management Science* 65, 4451–4469.

Avramov, Doron, Si Cheng, Abraham Lioui, and Andrea Tarelli, 2022, Sustainable investing with esg rating uncertainty, *Journal of Financial Economics* 145, 642–664.

Beber, Alessandro, Michael W. Brandt, Jason Cen, and Kenneth A. Kavajecz, 2021, Mutual fund performance: Using bespoke benchmarks to disentangle mandates, constraints and skill, *Journal of Empirical Finance* 60, 74–93.

Bera, Anil, and Sung Y. Park, 2008, Optimal portfolio diversification using the maximum entropy principle, *Econometric Reviews* 27, 484–512.

Berk, Jonathan B., and Jules H. van Binsbergen, 2015, Measuring skill in the mutual fund industry, *Journal of Financial Economics* 118, 1–20.

Carhart, Mark M., 1997, On the persistence of mutual fund performance, *Journal of Finance* 52, 57–82.

Chen, Linquan, Yao Chen, Alok Kumar, and Woo Sau Leung, 2023, Firm-level esg information and active fund management, Working Paper. University of Miami.

Cheng, Ing-Haw, Harrison Hong, and Kelly Shue, 2023, Do managers do good with other people's money?, *Review of Corporate Finance Studies* 12, 443–487.

Cremers, Martijn, Timothy B. Riley, and Rafael Zambrana, 2023, The complex materiality of ESG ratings: Evidence from actively managed ESG funds, Harvard Law School Forum on Corporate Governance.

Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristic-based benchmarks, *Journal of Finance* 52, 1035–1058.

Dimson, Elroy, Oğuzhan Karakaş, and Xi Li, 2015, Active ownership, *Review of Financial Studies* 28, 3225–3268.

Edmans, Alex, 2021, The social responsibility of business includes profits, available at https://promarket-org.cdn.ampproject.org/c/s/promarket.org/2021/10/19/social-responsibility-business-profits-pieconomics/?amp.

Ferrell, Allen, Hao Liang, and Luc Renneboog, 2016, Socially responsible firms, *Journal of Financial Economics* 122, 585–606.

Friede, Gunnar, Timo Busch, and Alexander Bassen, 2015, ESG and financial performance: Aggregated evidence from more than 2000 empirical studies, *Journal of Sustainable Finance and Investment* 5, 210–233.

Geczy, Christopher C., Robert F. Stambaugh, and David Levin, 2021, Investing in socially responsible mutual funds, *Review of Asset Pricing Studies* 11, 309–351.

Ghosh, Anisha, Christian Julliard, and Alex P. Taylor, 2019, An information-theoretic asset pricing model, Working Paper. London School of Economics.

Global Sustainable Investment Alliance, 2020, Global sustainable investment review.

Hartzmark, Samuel M., and Abigail B. Sussman, 2019, Do investors value sustainability? A natural experiment examining ranking and fund flows, *Journal of Finance* 74, 2789–2837.

He, Yazhou, Bige Kahraman, and Michelle Lowry, 2023, ES risks and shareholder voice, *Review of Financial Studies* 36, 4824–4863.

Hong, Harrison, and Marcin Kacperczyk, 2009, The price of sin: The effects of social norms on markets, *Journal of Financial Economics* 93, 15–36.

Khan, Mozaffar, George Serafeim, and Aaron Yoon, 2016, Corporate sustainability: First evidence of materiality, *The Accounting Reivew* 91, 1697–1724.

Kullback, Solomon, and Richard A. Leibler, 1951, On information and sufficiency, *Annuals of Mathematical Statistics* 22, 79–86.

Masulis, Ronald W., and Syed Walid Reza, 2015, Agency problems of corporate philanthropy, *Review of Financial Studies* 28, 592–636.

Matos, Pedro, 2020, ESG and responsible institutional investing around the world: A critical review, CFA Institute Research Foundation.

MSCI ESG Research, 2023, *ESG Ratings Methdology*, available at https://www.msci.com/esg-and-climate-methodologies.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703.

Orlov, Vitaly, Stefano Ramelli, and Alexander F. Wagner, 2023, Revealed beliefs about responsible investing: Evidence from mutual fund managers, Working Paper. Swiss Finance Institute and ECGI.

Pástor, L'uboš, Robert F. Stambaugh, and Lucian A. Taylor, 2015, Scale and skill in active management, *Journal of Financial Economics* 28, 23–45.

Pástor, L'uboš, Robert F. Stambaugh, and Lucian A. Taylor, 2021, Sustainable investing in equilibrium, *Journal of Financial Economics* 142, 550–571.

Pástor, L'uboš, and M. Blair Vorsatz, 2020, Mutual fund performance and flows during the COVID-19 crisis, *Review of Asset Pricing Studies* 10, 791–833.

Pedersen, Lasse Heje, Shaun Fitzgibbons, and Lukasz Pomorski, 2021, Responsible investing: The ESG-efficient frontier, *Journal of Financial Economics* 142, 572–597.

Politis, Dimitris N., and Joseph P. Romano, 1994, The stationary bootstrap, *Journal of American Statistical Association* 89, 1303–1313.

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

Stutzer, Michael, 1995, A Bayesian approach to diagnosis of asset pricing models, *Journal of Econometrics* 68, 367–397.

van der Beck, Philippe, 2023, Flow-driven ESG returns, Working Paper. Harvard Business School.

Figures

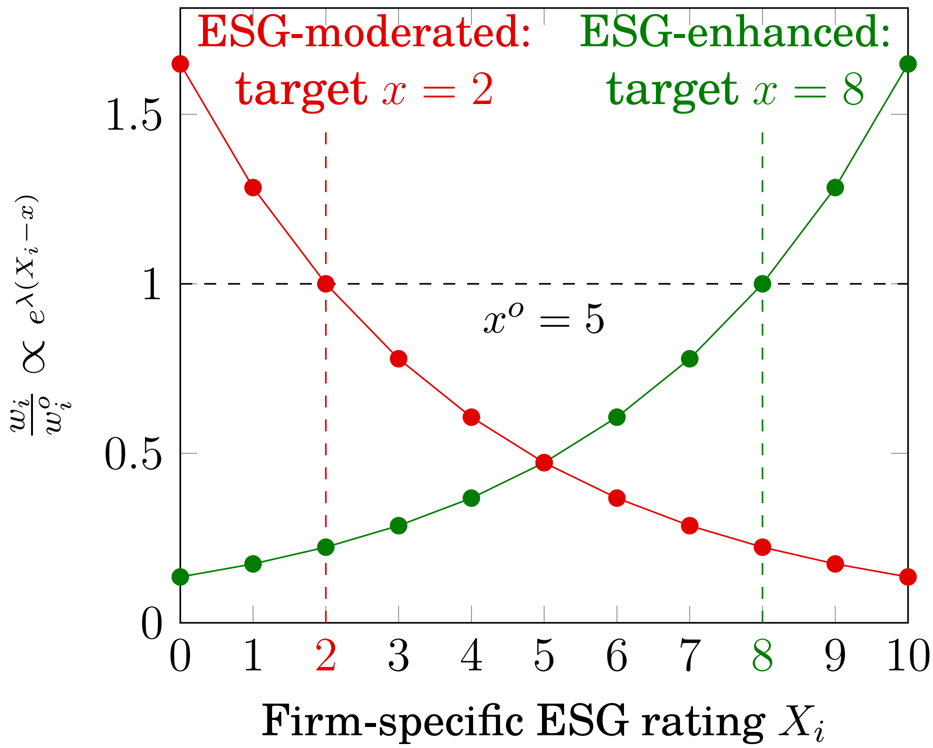

Figure 1: This figure illustrates the allocation of an exponentially-tilted portfolio with regard to stock-level ESG ratings. Note that though in our actual estimation, λ is not a free parameter and is instead determined endogenously by the unconstrained minimization problem in Eq. (7), it is set at $\lambda = 0.25$ merely to illustrate how the exponential tilting makes the portfolio allocation.

Letter Rating	Leader/Laggard	Final Industry-Adjusted Company Score
AAA	Leader	8.571* - 10.0
AA	Leader	7.143 – 8.571
A	Average	5.714 – 7.143
BBB	Average	4.286 – 5.714
BB	Average	2.857 – 4.286
B	Laggard	1.429 – 2.857
CCC	Laggard	0.0 – 1.429

Figure 2: MSCI firm-level ESG scores and letter-based ratings (MSCI ESG Research, 2023, p. 7)

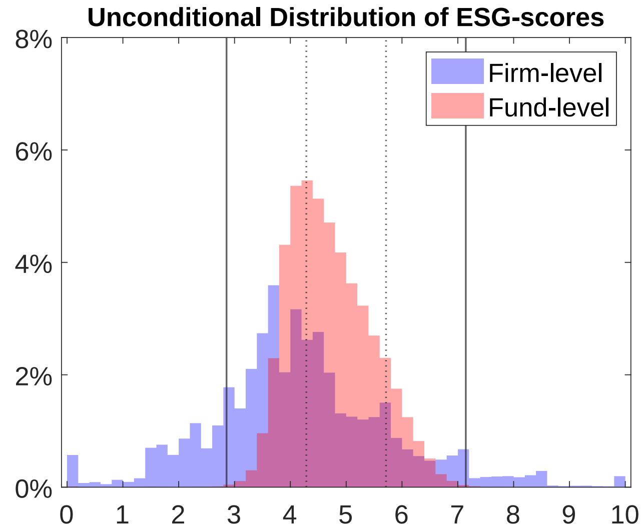

(a) Unconditional distributions

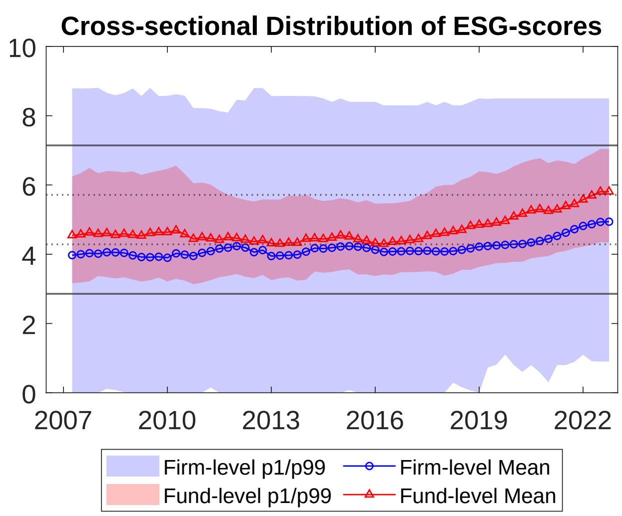

(b) Cross-sectional distributions

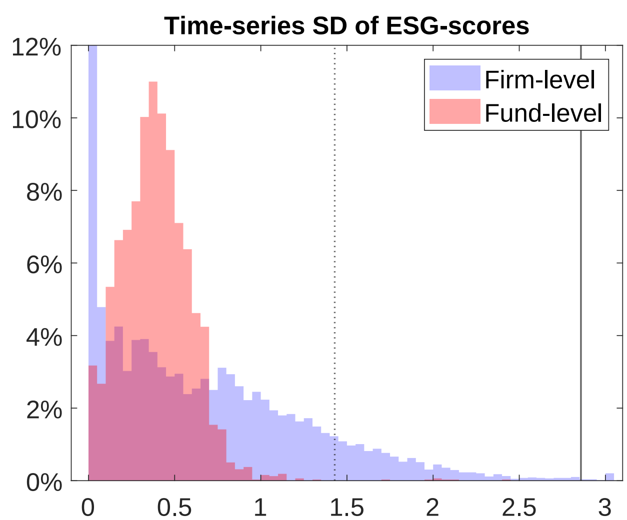

(c) Time-series standard deviations

Figure 3: Distribution of firm-level and fund-level ESG ratings.

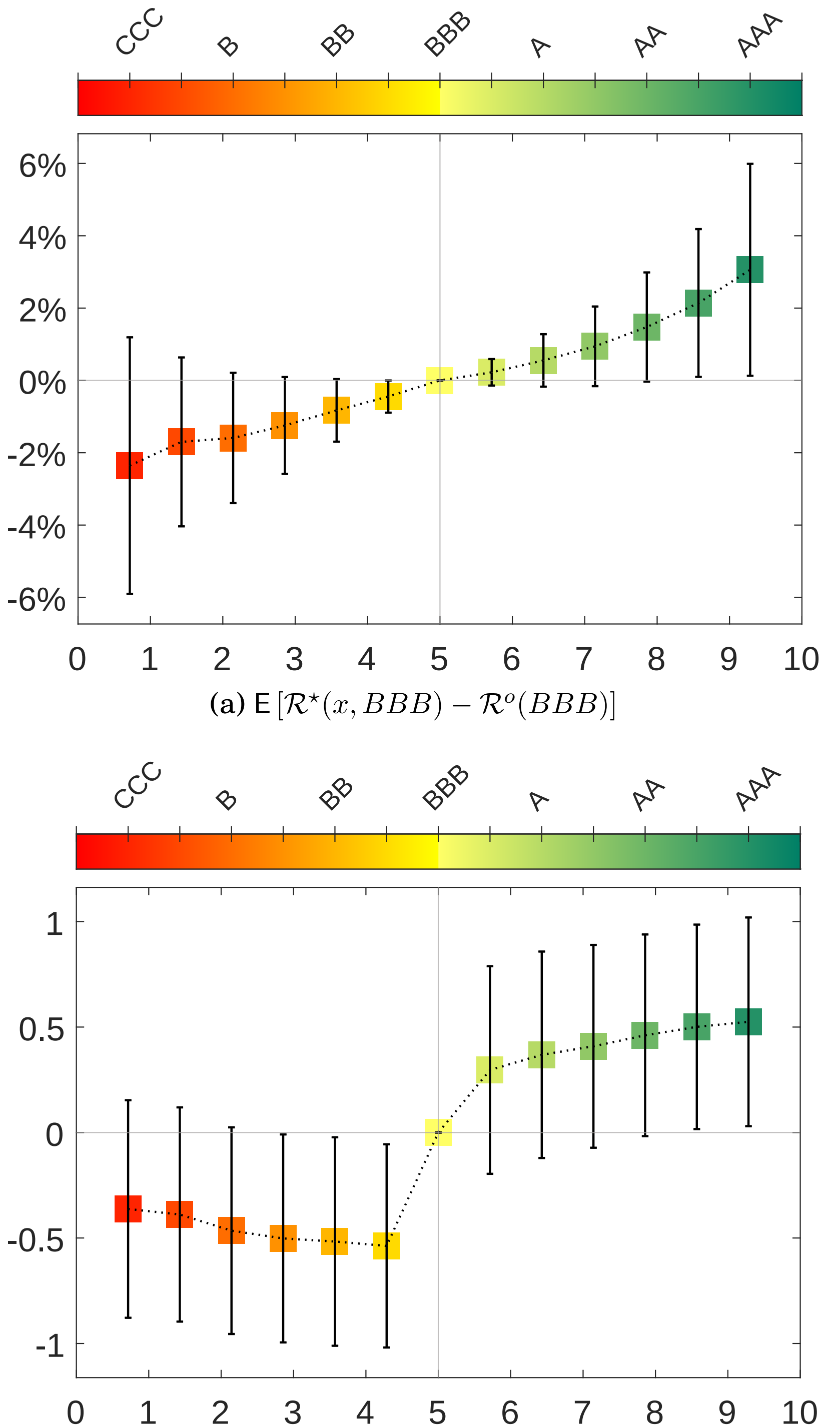

(a) $\mathsf{E}\left[\mathcal{R}^{\star}(x, BBB) - \mathcal{R}^{o}(BBB)\right]$

(b) Information Ratio for $\mathcal{R}^{\star}(x, BBB) - \mathcal{R}^{o}(BBB)$

Figure 4: Impact of ESG mandates on the BBB-FoF's performance.

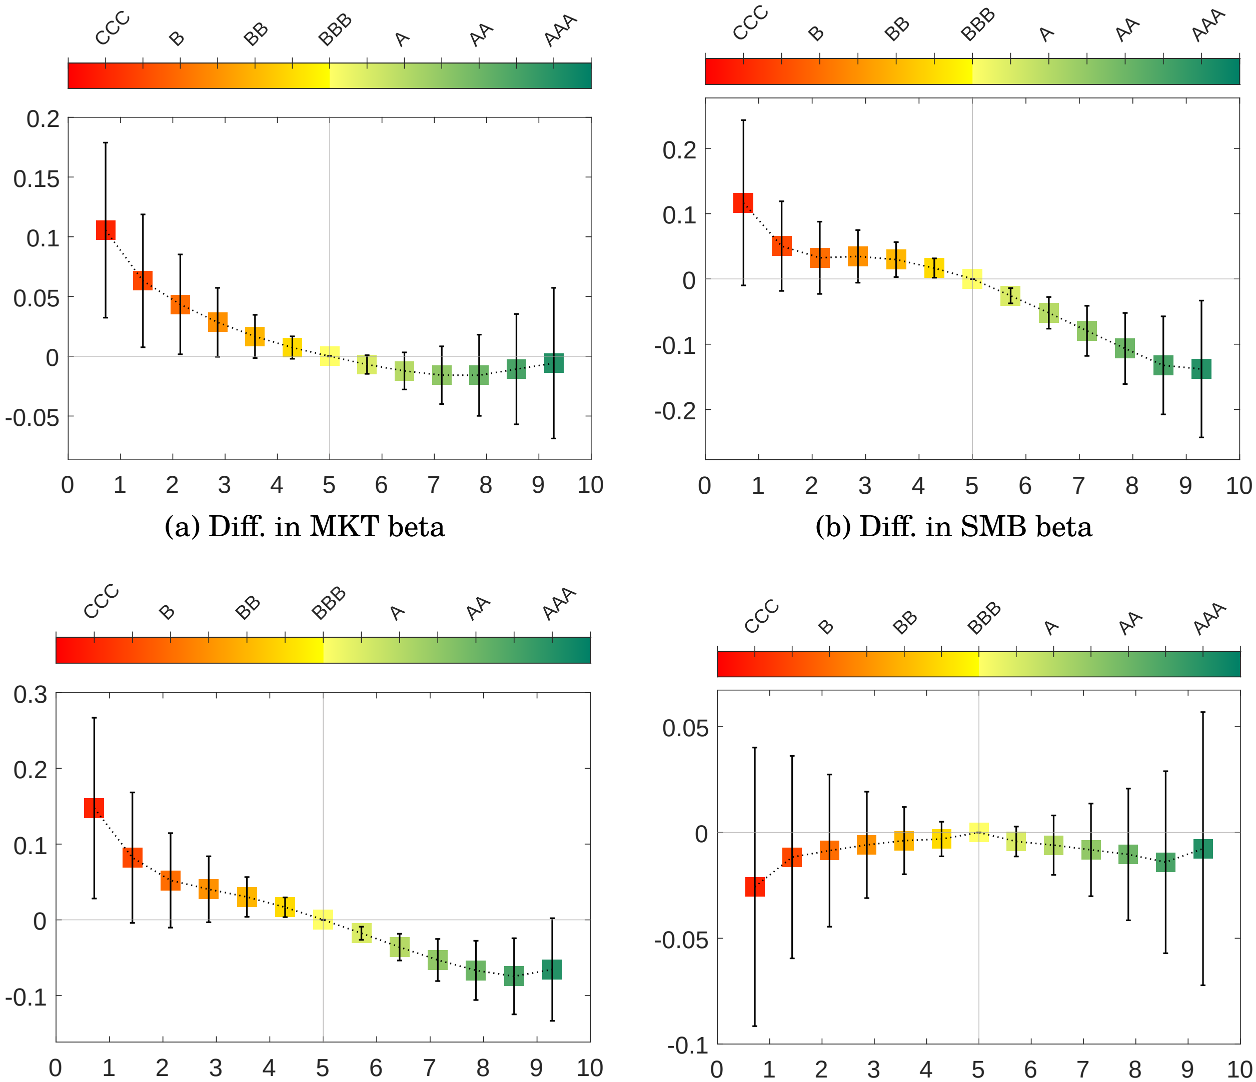

(a) Diff. in MKT beta

(b) Diff. in SMB beta

(c) Diff. in HML beta

(d) Diff. in WML beta

Figure 5: Impact of ESG mandates on the BBB-FoF's risk exposures

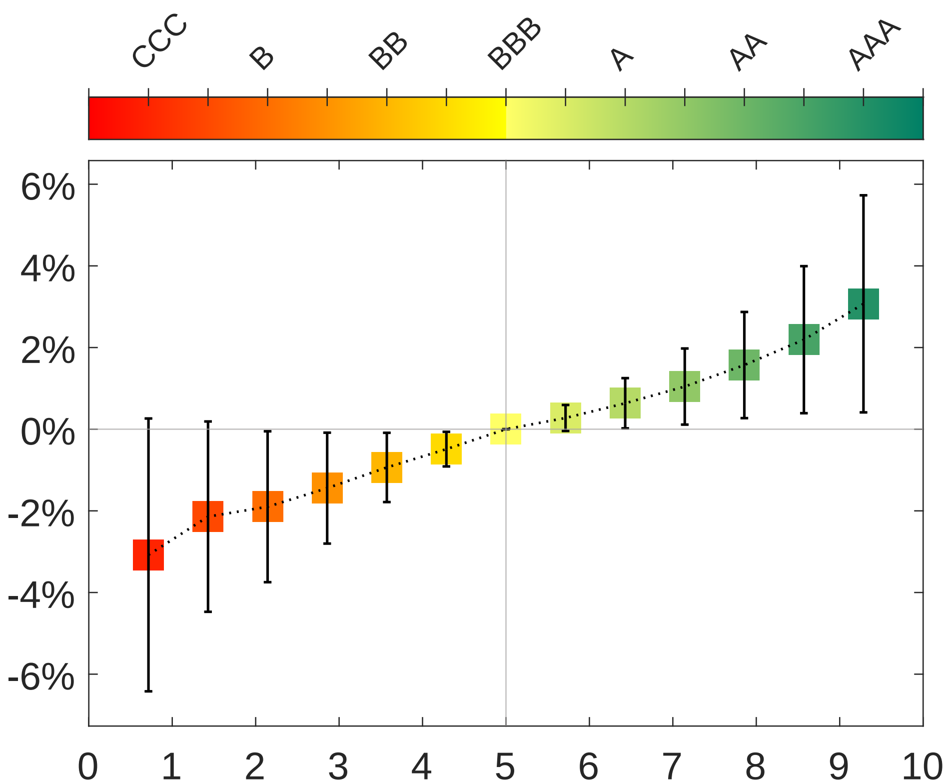

Figure 6: Impact of ESG mandates on the BBB-FoF's alpha.

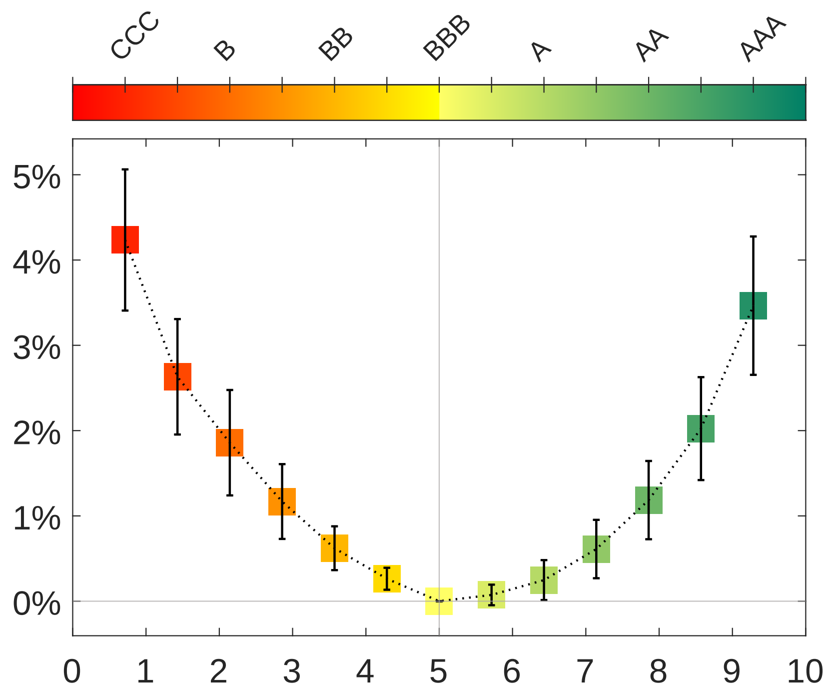

(a) Idiosyncratic volatility differentials

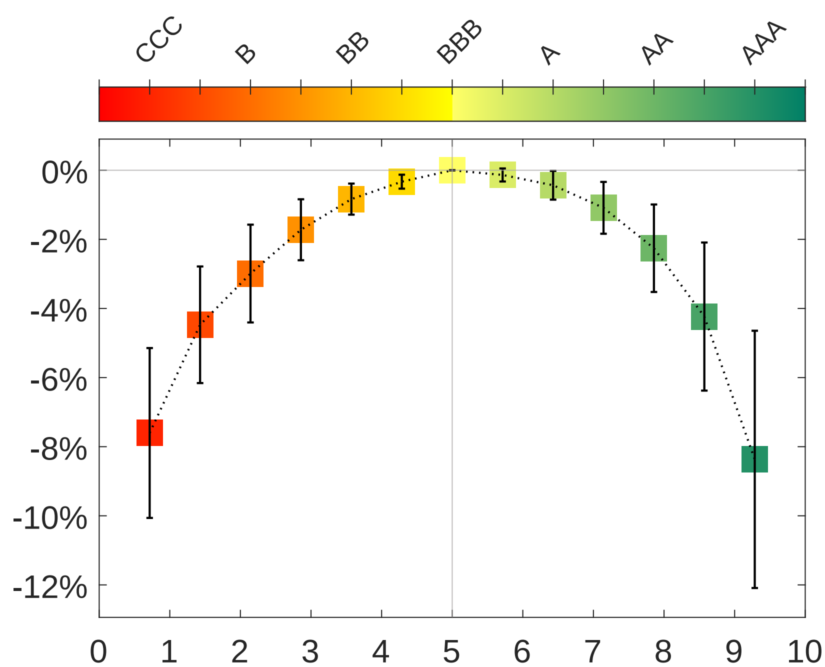

(b) R-squared differentials

Figure 7: Impact of ESG mandates on the BBB-FoF's risks.

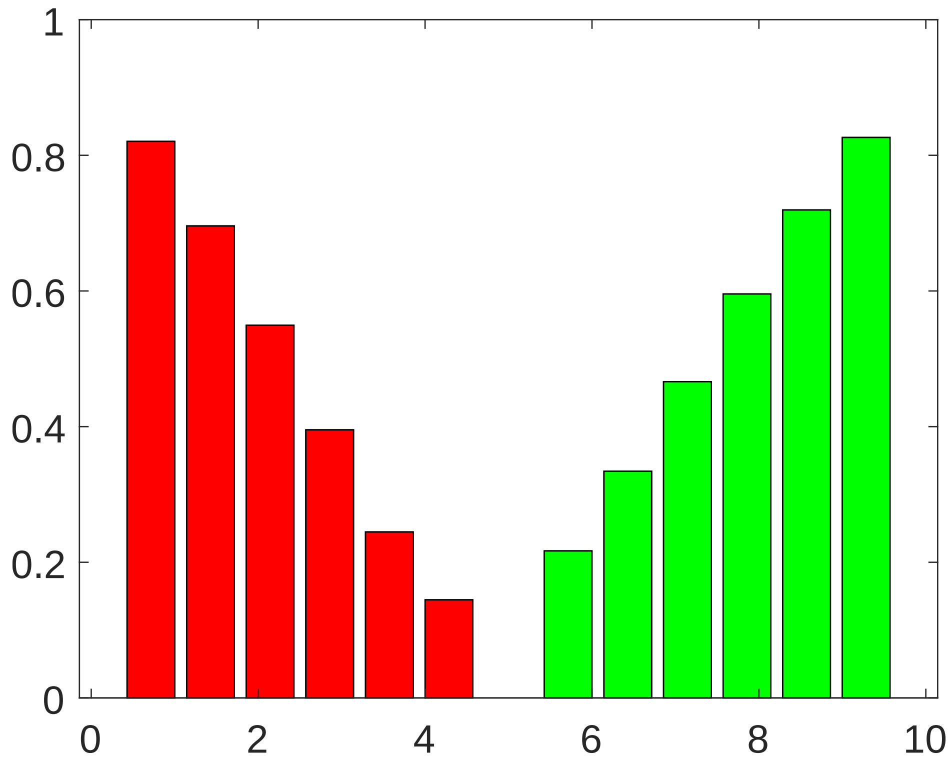

Figure 8: Portfolio turnover for a BBB-fund to fulfill the potential ESG mandates

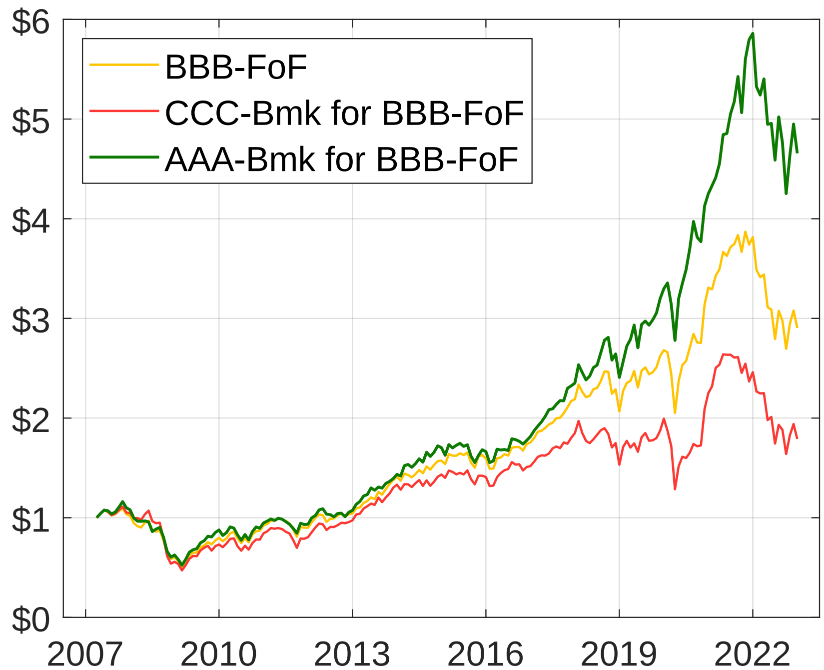

(a) BBB-FoF and its benchmarks

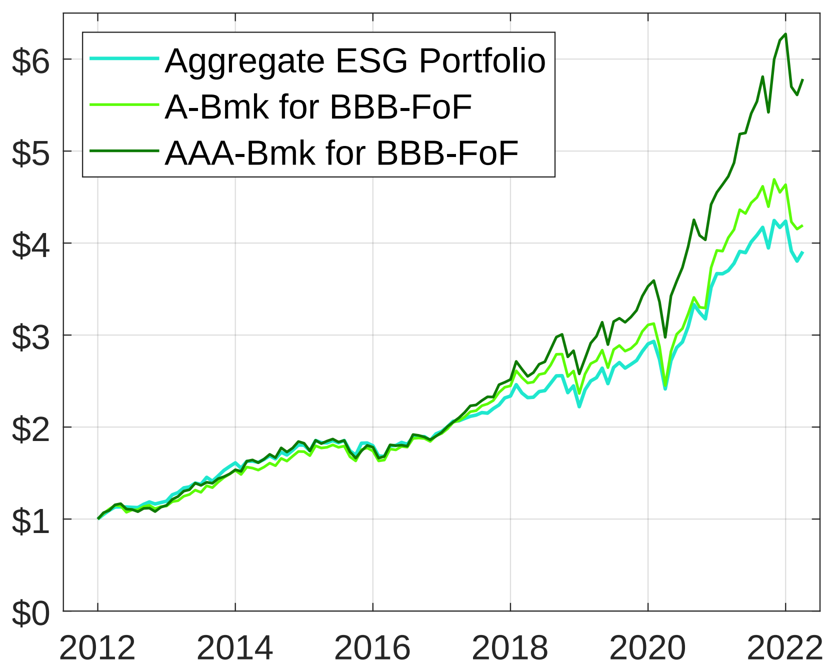

(b) Benchmarks for the BBB-FoF and the aggregate ESG portfolio

Figure 9: Comparison of cumulative returns of the BBB-FoF, its ESG-mandated benchmarks, and the aggregate ESG portfolio.

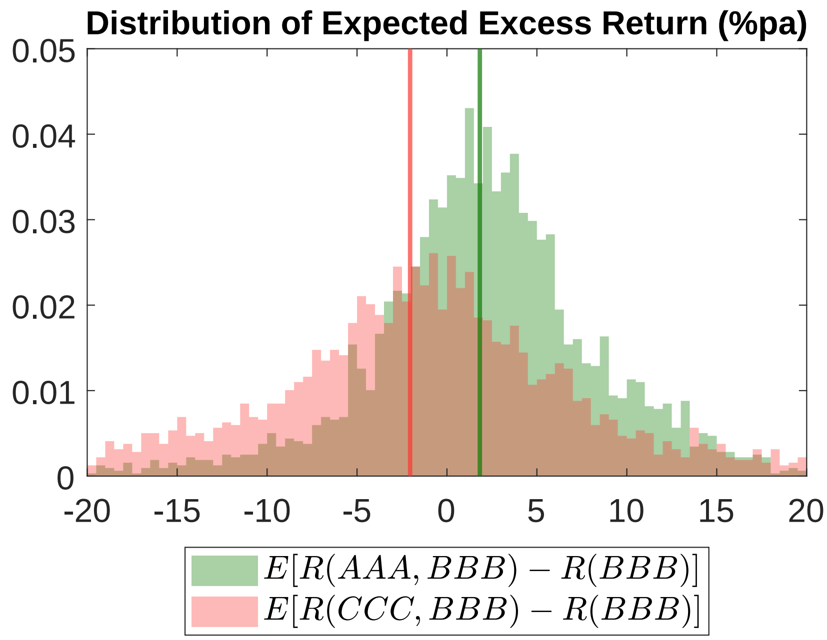

(a) Expected excess returns

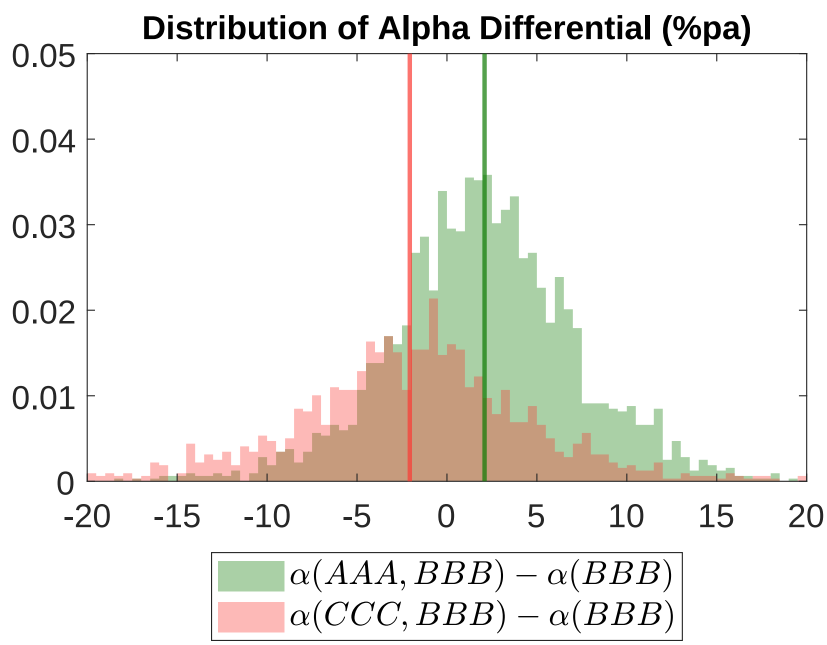

(b) Alpha spreads

Figure 10: Distribution of difference in performance between a BBB-rated fund and its ESG-mandated benchmarks.

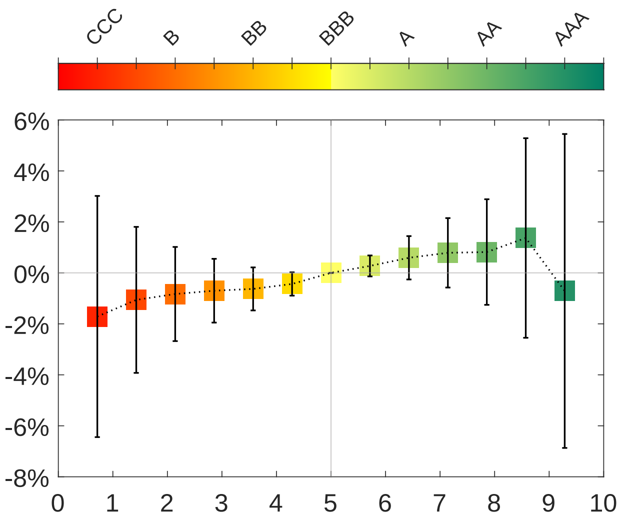

(a) Environmental (E) mandates

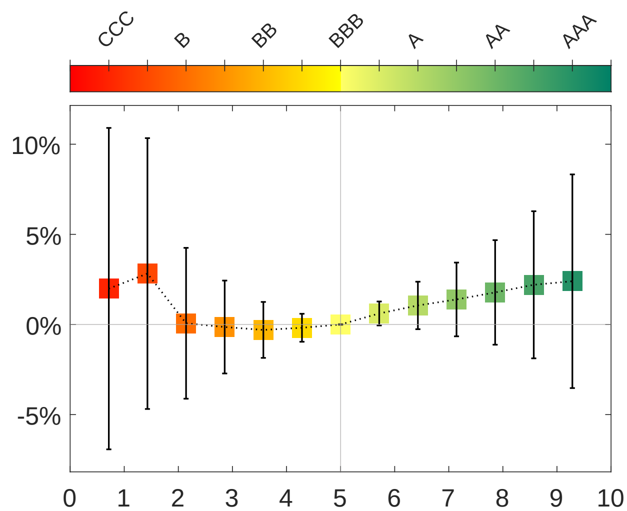

(b) Social (S) mandates

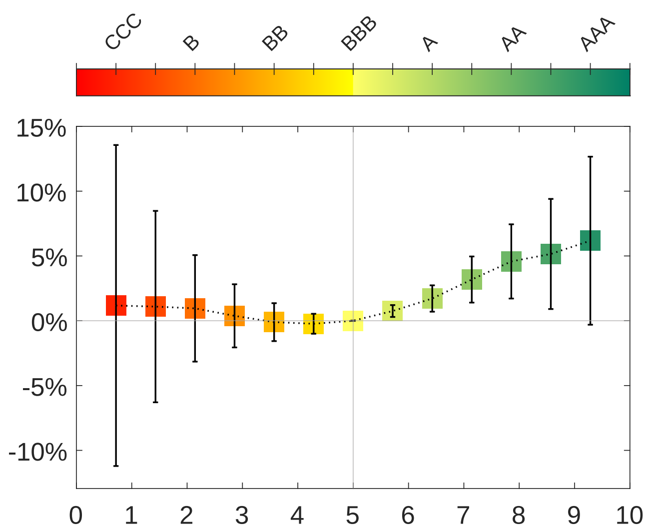

(c) Governance (G) mandates

Figure 11: Impact of pillar-specific mandates on fund performance.

Appendix A Proof of Proposition 1

For brevity, the fund index j and time index t are both omitted herein. Now we restate the I-divergence minimization problem for an ESG-enhanced benchmark given a target $x > x^o$ as follows:

$$\min_{w} \sum_{i=1}^{N} w_i^o \, \theta_i \log(\theta_i) \quad \text{s.t.} \; \sum_{i=1}^{N} w_i^o \, \theta_i \geq x, \quad \sum_{i=1}^{N} w_i^o \, \theta_i = 1, \quad w_i^o \, \theta_i > 0, \quad \text{for } \forall \, i,$$

where $\theta_i = \frac{w_i}{w_i^o}$. We then write the Lagrangian as

$$\mathcal{L} = \sum_{i=1}^{N} w_i^o \, \theta_i \log(\theta_i) - \lambda \sum_{i=1}^{N} w_i^o \, \theta_i \, (X_i - x) - \mu \left(1 - \sum_{i=1}^{N} w_i^o \, \theta_i\right) - \gamma_i \, w_i^o \, \theta_i \, .$$

The first-order condition (FOC) with respect to θ_i, $\frac{\partial \mathcal{L}}{\partial \theta_i}(\theta_i^\star, \lambda^\star, \mu^\star, \gamma_i^\star) = 0$, yields

$$\log(\theta_i^\star) + 1 - \lambda^\star (X_i - x) + \mu^\star - \gamma_i^\star = 0 \, .$$

Since the ESG-enhancement constraint and the budget constraint are both binding whereas the logarithm objective function guarantees the strict positiveness of portfolio weights, it follows that, at the optimum,

$$\lambda^\star > 0, \quad \mu^\star > 0, \quad \text{and} \quad \gamma_i^\star \equiv 0 \quad \text{for } \forall \, i.$$

Multiplying both sides of the FOC by w_i and summing up across i, we obtain

$$\mu^\star = -1 - \sum_{i=1}^{N} w_i \log(\theta_i^\star) \, .$$

On the other hand, the FOC leads to

$$\theta_i^\star = \exp\left(\lambda^\star (X_i - x)\right) \, \exp\left(-1 - \mu^\star\right) \, .$$

Combining the above with the budget constraint $\sum_{i=1}^{N} w_i = \sum_{i=1}^{N} w_i^o \, \theta_i^\star = 1$, we obtain

$$1 = \exp(-1 - \mu^\star) \sum_{i=1}^{N} w_i^o \exp\left(\lambda^\star (X_i - x)\right) \implies \exp(-1 - \mu^\star) = \left[\sum_{i=1}^{N} w_i^o \exp\left(\lambda^\star (X_i - x)\right)\right]^{-1} .$$

Hence, the solution for θ_i is

$$\theta_i^\star = \frac{\exp\left(\lambda^\star\left(X_i - x\right)\right)}{\sum_{i=1}^N w_i^o \exp\left(\lambda^\star\left(X_i - x\right)\right)} = \frac{e^{\lambda^\star X_i}}{\sum_{i=1}^N w_i^o e^{\lambda^\star X_i}}\,,$$

and the solution for μ is

$$\mu^\star = -1 - \sum_{i=1}^N w_i \log\left(\theta_i^\star\right) = -1 + \log\left(\sum_{i=1}^N w_i^o \exp\left(\lambda^\star\left(X_i - x\right)\right)\right)\,.$$

Finally, the Lagrangian multiplier associated with the ESG-enhancement constraint (λ) is the solution to

$$\sum_{i=1}^N w_i^o \theta_i^\star\left(X_i - x\right) = 0 \implies \sum_{i=1}^N w_i^o\, e^{\lambda^\star X_i}\left(X_i - x\right) = 0\,.$$

This equation coincides with the first-order condition of the following unconstrained problem

$$\min_{\lambda > 0} \sum_{i=1}^N w_i^o\, e^{\lambda(X_i - x)}\,,$$

which completes the proof for the ESG-enhanced benchmark with a given target of $x > x^o$. For the ESG-moderated benchmark with a given target of $x < x^o$, everything remains the same except for $\lambda^\star < 0$.