Balancing Returns and Responsibility: Evidence from Shrinkage-based Portfolios

Christos A. Makridis* Majeed Simaan†

January 16, 2024

Abstract

We introduce environmental, social, and governance (ESG) scores into the portfolio selection framework using shrinkage estimators. We study nine linear shrinkage techniques for estimating stock returns' covariance matrix and eight ESG-based rules, documenting large variations in performance with two rules - stock factor-related - emerging as top performers in terms of risk-adjusted returns, outperforming ESG-based as well as the naive and market portfolios. While integrating ESG scores mitigates downside risk, it leads to high turnover and mixed results in out-of-sample (OOS) ESG performance. Additionally, the market portfolio shows the highest OOS overall ESG score but is not the most efficient in risk-adjusted terms. A field survey further supports our findings, reflecting preferences for profit maximization over an exclusive focus on ESG. Our results offer guidance to asset managers in balancing returns, risk, and social considerations, suggesting that ESG information may have content, but targeting it over other factors may distort organizational roles and yield suboptimal outcomes.

Keywords: ESG-Efficiency, Bias-Variance Trade-Off, Covariance Shrinkage, Portfolio Selection

JEL Codes: C13, C61, G11

*Arizona State University, USA, and Stanford University, USA; cmakridi@stanford.edu. Christos thanks the Institute of Humane Studies for their support.

†School of Business, Stevens Institute of Technology, USA.

1 Introduction

There is an enormous portfolio selection literature that studies how investors should allocate their wealth across different assets to balance their expected returns and associated risks. Dating back to Markowitz (1952) and Merton (1973), there has been a general recognition that investors aim for the highest possible return for a given level of risk smoothed intertemporally. However, what information is considered, and how are different pieces of information weighted in the portfolio selection process? This creates a friction that is known as *estimation risk*. Commonly, estimation risk in portfolio selection refers to the uncertainty and potential errors associated with the estimation of key parameters, such as expected returns, volatilities, and correlations. It is well-recognized that such estimation error leads to poor ex-post investment decisions (Michaud, 1989; Best and Grauer, 1991; DeMiguel et al., 2009b). Nonetheless, a common remedy involves the application of shrinkage techniques, leading to improved efficiency and stability of portfolio weights.[1]

In mean-variance (MV) portfolios, model's inputs are given by the mean vector and the covariance matrix of asset returns. However, estimating mean returns has been recognized as a challenging task (Merton, 1980). Consequently, the literature has placed a greater emphasis on the covariance matrix and has developed various shrinkage-based techniques to mitigate estimation risk.[2] To motivate the idea behind shrinkage-based techniques, consider the sample covariance matrix. Under certain assumptions, it can be shown that the estimator is unbiased; however, its performance is subject to variability. A shrinkage-based estimator, on the other hand, imposes a bias on the estimator in favor of reducing its variance. Such a trade-off can be formulated using a mean-squared error (MSE) metric. Traditionally, common approaches designed to mitigate estimation risk rely on stock returns only, upon which an optimal decision rule is attained. In this research, we move beyond asset returns to construct shrinkage-based rules by integrating information covering firm characteristics based on environmental, social, and governance (henceforth ESG) scores.

Why focus on ESG scores? There has long been a recognition that corporate governance (Gompers et al., 2003) and culture (Edmans, 2011; Graham et al., 2022) matter for understanding asset prices. Nonetheless, the emergence of ESG indicators has raised new questions in portfolio

[1]These techniques are equivalent to imposing additional structure/constraints into portfolio optimization (Jagannathan and Ma, 2003; DeMiguel et al., 2009a).

[2]For an extensive overview of these approaches, readers can refer to the recent survey by Ledoit and Wolf (2022).

theory, particularly around how to balance new data on the performance of organizations across alternative metrics. On the one hand, some have argued that investors are willing to accept lower returns in exchange for "impact investing" (Barber et al., 2021). On the other hand, others have argued that ESG scores contain fundamental information about an organization and can play a role in identifying the efficient frontier (Pedersen et al., 2021). However, recent research has pointed out that ESG scores vary markedly across rating agencies due to differences in measurement, scope, and weight across indicators (Berg et al., 2022). Consequently, the question of whether ESG scores can meaningfully guide portfolio selection remains unsettled, particularly in light of recent findings indicating the noise and disagreement inherent in such scores. The answer will play a major role in determining where and how much capital flows, which is especially relevant for new business formation given the role of capital constraints on entrepreneurs.

In this regard, our paper aims to answer the above question by utilizing shrinkage-based techniques that integrate both stock returns and ESG scores in constructing out-of-sample (henceforth OOS) robust portfolios. To the best of our knowledge, the shrinkage literature has not considered the impact of ESG on covariance matrix shrinkage and, consequently, the resulting risk-based rules. Specifically, our research seeks to examine the marginal value of such scores relative to market (priced-in) information. This is highly relevant as ESG scores potentially convey relevant information about an organization's broader impact in ways that are not reflected in returns - potentially for behavioral reasons or other frictions in the market. Therefore, understanding how ESG scores play a role in portfolio selection with shrinkage estimators is integral, given what asset managers are responsible for managing and seeking guidance on incorporating non-market factors into their portfolio choices. Furthermore, disentangling the role of preferences versus returns in the debate about ESG has become a pivotal issue in the finance literature (Giglio et al., 2023).

The first part of our paper establishes a basic framework for our shrinkage-based rules. Our approach is inspired by the literature on the shrinkage of covariance matrix (Ledoit and Wolf, 2022) and its link to the generalized-norm constraints in portfolio selection by DeMiguel et al. (2009a). The ESG scores can be incorporated as a constraint on the objective function, which is penalized by lower ratings. This formulation can be viewed from a generalized norm point of view, which is consistent with the covariance shrinkage-based approach. In this regard, we leverage state-of-the-art shrinkage-based tools to conduct our main investigation. In particular, we consider nine different linear shrinkage techniques commonly used in the literature to estimate the covariance matrix of stock returns and eight different ESG-based rules (Ardia and Boudt, 2015; Ardia et al., 2017a). The ESG-based rules shrink the covariance matrix towards a diagonal ESG penalty (reciprocal rating) matrix, penalizing loading too low on ESG scores. We use these estimators to construct different risk-managed portfolios and evaluate their OOS performance net of transaction costs while considering the equally and value-weighted portfolios as additional robustness check.

The second part of our paper presents these results. First, we find a large variation in performance across the nine shrinkage rules that rely on stock returns only. The top two performers in terms of risk-adjusted returns are factor-based, denoted by **LW** and **Factor** (Ledoit and Wolf, 2003).[3] Additionally, both rules beat the naive (DeMiguel et al., 2009b) and market portfolios. We note that even though the **Factor** does not rely on ESG data, it results in the highest OOS ESG performance among all other rules. This result sheds an interesting light on filtering information from market prices about ESG ratings. Second, the rules that rely on both stock returns and ESG ratings outperform the naive and market portfolios in terms of the Sortino ratio but not the Sharpe ratio. This result indicates that shrinkage towards ESG mitigates downside risk. However, their turnover is relatively high compared to the other rules. Additionally, these rules do not fare well in terms of OOS ESG scores compared to others. On the other hand, the ESG rules that rely on ESG scores alone fare better in terms of ESG OOS performance. Third, the market portfolio results in the highest OOS overall ESG score. This result should not be surprising, given the increased focus on these scores by asset managers. However, the market portfolio is not the most efficient when it comes to risk-adjusted returns. Investors with a mean-variance preference would gain lower utility by investing in the market portfolio, even though the market attains a high OOS ESG score (roughly the same as the **Factor** rule). Nonetheless, given the long-term implications of ESG, one should consider intertemporal utility in drawing conclusions. Overall, our portfolio analysis demonstrates that rules that do not rely on ESG directly, such as the **LW** and **Factor**, could enhance OOS ESG scores and risk-adjusted returns simultaneously.

The third part of our paper provides reduced-form evidence consistent with our main results on why market-based portfolios outperform the ESG rules. In particular, we delve deeper into in-

[3]These two rules filter out idiosyncratic risk from the target matrix using a factor model, where the factor from the **Factor** (respectively **LW**) rule is constructed using the first-principal (respectively equally weighted) portfolio.

vestors' preferences to better understand how they form their preferences and, hence, their portfolio choices. To do so, we field a survey (1,500 observations) to gauge preferences for ESG investing and assess the role of information in forming preferences for ESG. Our field study is also inspired by recent work by Giglio et al. (2023), who document "four facts about ESG beliefs and investor portfolios." Overall, we establish three main results from our survey. First, in terms of focus, individuals say that paying a living wage ranks as not only their highest priority but also the highest priority that they believe an organization should have, followed by a close tie between issues relating to climate change and human trafficking. Second, after controlling for individual preferences for different ESG factors, there is no correlation between preferences for non-market factors and organizational ESG priorities. In other words, individual preferences for ESG investing do not necessarily translate into organizational priorities. Third, providing people with information about the actual costs of renewable energy production reduces the degree of support for such policies. In sum, our results provide perspective on the distinction between personal and organizational support for ESG policies and the important role of information in shaping attitudes.

Our paper also contributes to a growing empirical literature on ESG investing.[4] While there is already voluminous evidence that high ESG firms have lower expected returns (Hong and Kacperczyk, 2009; Christensen et al., 2022; Pastor et al., 2022; Bolton and Kacperczyk, 2022), and others find statistically insignificant differences (Hartzmark and Sussman, 2019; Pedersen et al., 2021), there have been few studies incorporating ESG into portfolio choice. One reason for the conflicting results in the literature stems from the substantial differences in ESG scores among rating agencies, as Berg et al. (2022) point out and decompose into three categories: measurement (56%), scope (38%), and weight (6%). Recent work by Lindsey et al. (2023) builds upon the instrumented principal components analysis (IPCA) methodology from Kelly et al. (2019), showing that optimal stock portfolios can be adjusted to achieve ESG goals without sacrificing returns. That is possible, however, precisely because of the vast disagreement on "how to measure, weight, and combine ESG information," so investors could hold varying portfolios and still satisfy their ESG preferences.

Our results also speak to an important tension in the literature about the distinction between individual preferences and organizational priorities. For example, Giglio et al. (2023) present novel evidence on ESG investing and preferences, showing that ESG beliefs are important drivers of actual

[4]See Giglio et al. (2021) for a recent survey of the literature.

portfolio allocation decisions and that non-pecuniary factors play a major role after accounting for financial factors. Our results lend additional perspective to these patterns by showing that the bulk of the variation in preferences for non-pecuniary organizational objectives stems from individual preferences and that additional information about ESG costs can tilt support against ESG policies. These empirical results from our survey also explain the intuition behind our quantitative result that the market achieves a better allocation than explicit ESG-based portfolios.

The rest of the paper proceeds as follows. In Section 2, we demonstrate the portfolio framework studied in our research and discuss the shrinkage-based methodology that allows us to incorporate ESG ratings into the portfolio choice. Section 3 covers the empirical design in terms of data and the OOS performance evaluation, whereas we devote Section 4 to the discussion of our portfolio results and findings. In Section 5, conduct a survey and provide some insights that help deepen our understanding of the mechanisms of our baseline results. Finally, Section 6 concludes.

2 Framework

In this section, we briefly describe the idea of covariance matrix shrinkage and its link to imposing ESG constraints on portfolio optimization. Regarding notations, we use bold lowercase letters for vectors and bold uppercase letters for matrices. Greek letters are used to denote parameters, whereas Latin letters are used to denote random variables. A non-bold typeface denotes a univariate variable/parameter.

2.1 Global Minimum Variance Portfolio

We consider a special portfolio rule known as the global minimum variance portfolio (GMVP). The GMVP relies on a single input, which is the covariance matrix denoted by $\mathbf{\Sigma}$. Note that $\mathbf{\Sigma}$ is a non-singular $N \times N$ matrix, with N as the number of assets in the portfolio. Let $\boldsymbol{\xi} \in \mathbb{R}^{N \times 1}$ denote the portfolio weights. Under no position constraints, the GMVP is the solution to the following optimization problem:

$$\begin{aligned} \min_{\boldsymbol{\xi}} \quad & \boldsymbol{\xi}^\top \mathbf{\Sigma} \boldsymbol{\xi} \\ \text{s.t.} \quad & \mathbf{1}^\top \boldsymbol{\xi} = 1, \end{aligned} \tag{2.1}$$

with $\mathbf{1}$ denotes a $N \times 1$ vector of ones. Using the Lagrangian multiplier, it can be shown that the GMVP is given:

$$\boldsymbol{\xi}_0 = \frac{\boldsymbol{\Sigma}^{-1}\mathbf{1}}{\mathbf{1}^\top \boldsymbol{\Sigma}^{-1}\mathbf{1}}. \tag{2.2}$$

Intuitively, if assets were uncorrelated, the decision rule allocates wealth to each asset in the portfolio proportionally to the reciprocal of variance. For instance, when $\boldsymbol{\Sigma}_{ij} = 0$ for $i \neq j$, then the weight allocated to asset $i = 1, ..., N$ is given by

$$\boldsymbol{\xi}_{0,i} = \frac{1}{\sigma_i^2} \left[\sum_{k=1}^{N} \frac{1}{\sigma_k^2} \right]^{-1}.$$

Typically, the covariance matrix measures the riskiness of the assets since it represents the assets' individual risk as well as its co-risk with other assets, where the latter is captured by the covariances. In general terms, the covariance matrix can be viewed as a loss function, and the decision-maker may have different views about what constitutes such risk/loss. For instance, in the case of ESG, one potential loss is that the stock exhibits a low rating. In this case, one can map such scores into a risk metric and solve for GMVP that minimizes the risk of attaining a low ESG portfolio score. We cover this issue further in Section 2.3.

2.2 Generalized-norm-constrained GMVP

Before we introduce how to integrate ESG rating into the portfolio problem, it is worthwhile discussing the idea of the generalized-norm-constrained GMVP briefly. DeMiguel et al. (2009a) propose a generalized-norm-constrained GMVP portfolio. This constraint can be viewed as a covariance matrix shrinkage approach. Specifically, according to Proposition 2 from DeMiguel et al. (2009a), if one imposes an additional constraint of $\boldsymbol{\xi}^\top \boldsymbol{\Omega} \boldsymbol{\xi} \leq \delta$ on the optimization problem from Equation (2.1) for some feasible point $\delta \geq [\mathbf{1}^\top \boldsymbol{\Omega}^{-1}\mathbf{1}]^{-1}$ and positive definite matrix $\boldsymbol{\Omega} \in \mathbb{R}^{N \times N}$, then it can be shown the resulting GMVP is similar to Equation (2.2), whereas the main difference is that the covariance matrix is replaced by a convex combination between $\boldsymbol{\Sigma}$ and $\boldsymbol{\Omega}$.

To relate to the above result, consider the $\boldsymbol{\xi}^\top \boldsymbol{\Omega} \boldsymbol{\xi} \leq \delta$ constraint on the optimization problem from Equation (2.1), such that

$$\begin{aligned}
\min_{\xi} \quad & \boldsymbol{\xi}^\top \boldsymbol{\Sigma} \boldsymbol{\xi} \\
\text{s.t.} \quad & \xi^\top \mathbf{1} = 1 \\
& \boldsymbol{\xi}^\top \boldsymbol{\Omega} \boldsymbol{\xi} \leq \delta.
\end{aligned} \tag{2.3}$$

Using the Lagrangian, we can reformulate the optimization problem as

$$\begin{aligned}
\min_{\boldsymbol{\xi}} \quad & \boldsymbol{\xi}^\top \left[\boldsymbol{\Sigma} + \lambda \boldsymbol{\Omega} \right] \boldsymbol{\xi} \\
\text{s.t.} \quad & \mathbf{1}^\top \boldsymbol{\xi} = 1,
\end{aligned} \tag{2.4}$$

where λ depends on δ. With a small adjustment, the objective function can be reformulated as

$$(1+\lambda)\, \boldsymbol{\xi}^\top \left[\frac{1}{1+\lambda} \boldsymbol{\Sigma} + \frac{\lambda}{1+\lambda} \boldsymbol{\Omega} \right] \boldsymbol{\xi}.$$

Hence, solving the optimization problem under the $\boldsymbol{\Omega}$-norm-constraint is consistent with solving the following optimization problem

$$\begin{aligned}
\min_{\boldsymbol{\xi}} \quad & \boldsymbol{\xi}^\top \boldsymbol{\Sigma}_S`\boldsymbol{\xi} \\
\text{s.t.} \quad & \mathbf{1}^\top \boldsymbol{\xi} = 1,
\end{aligned} \tag{2.5}$$

where

$$\boldsymbol{\Sigma}_S = \frac{1}{1+\lambda} \boldsymbol{\Sigma} + \frac{\lambda}{1+\lambda} \boldsymbol{\Omega} \tag{2.6}$$

denotes the shrank covariance matrix. In this regard, Proposition 2 from DeMiguel et al. (2009a) states that there exists a δ (hence λ) value, for which the coefficient $\alpha = \lambda/(1+\lambda)$ corresponds to the shrinkage intensity from Ledoit and Wolf (2004). As a result, the generalized-norm-constrained GMVP is given by

$$\boldsymbol{\xi}_S = \frac{\boldsymbol{\Sigma}_S^{-1} \mathbf{1}}{\mathbf{1}^\top \boldsymbol{\Sigma}_S^{-1} \mathbf{1}} \tag{2.7}$$

where the S subscript denotes that the portfolio is based on a shrank covariance matrix (alternatively, constrained).

2.3 Application

The portfolio rule from Equation (2.7) relies on two inputs: the covariance matrix, $\mathbf{\Sigma}$, and the constraint matrix, $\mathbf{\Omega}$. In reality, neither matrix is known, except in the cases when the elements of $\mathbf{\Omega}$ are set to fixed values, such as in the case of the identity matrix. Therefore, one needs to rely on a model to map stock-related data to form estimates of either matrix. In the shrinkage literature, typically, the covariance matrix $\mathbf{\Sigma}$ is estimated using the sample covariance denoted by $\mathbf{S}$, whereas the $\mathbf{\Omega}$ matrix is either pre-determined by fixed values such as the case of the identity matrix (Ledoit and Wolf, 2004) or follows a more parsimonious representation of the covariance matrix, e.g., a diagonal matrix.

In our analysis, we fix the estimate of the covariance matrix to be $\mathbf{S}$ and consider different models to construct the target matrix, $\mathbf{\Omega}$, and the shrinkage intensity, α. To be more specific, let $m \in \mathcal{M}$ denote a specific model, where $\mathcal{M}$ is the set of possible models. For a given model m, the estimated covariance matrix is given by

$$\hat{\mathbf{\Sigma}}_S(m) = (1 - \alpha_m)\mathbf{S} + \alpha_m \hat{\mathbf{\Omega}}_m \tag{2.8}$$

where $\hat{\mathbf{\Omega}}_m$ (α_m) is the target matrix (shrinkage intensity) based on model m. Depending on the model used, the matrix $\hat{\mathbf{\Omega}}_m$ could be data-driven or determined according to a well-conditioned matrix, such as the identity matrix that requires no data for construction. At the same time, the shrinkage intensity is generally determined by minimizing a loss function of $\hat{\mathbf{\Sigma}}_S(m)$. In the linear shrinkage literature (Ledoit and Wolf, 2004), the loss function denotes the mean-squared error of the $\hat{\mathbf{\Sigma}}_S(m)$, which captures the mean deviation of the shrank estimate from the population (unknown) matrix. In this regard, the shrinkage approach aims to balance between bias and variance. On the one hand, introducing a well-conditioned matrix, such as the identity matrix, as a target matrix induces greater bias into the covariance matrix. On the other hand, such a target has zero variance, resulting in a lower variance for the shrank estimate. Therefore, there is an optimal convex combination that attains the best bias-variance trade-off, which the shrinkage intensity tries to attain.

Motivated by the covariance shrinkage literature, we consider several portfolio rules that utilize the closed-form solution of the GMVP using different shrinkage models. Specifically, for a given data set $\mathcal{D}$ and model $m \in \mathcal{M}$, the constrained GMVP is given by

$$\mathbf{w}_S(m) = \frac{\hat{\mathbf{\Sigma}}_S(m)^{-1}\mathbf{1}}{\mathbf{1}^\top \hat{\mathbf{\Sigma}}_S(m)^{-1}\mathbf{1}} \tag{2.9}$$

Our framework utilizes the closed-form solution from Equation (2.9) using different estimation models. In the following discussion, we describe the various models used in our main analysis and demonstrate the incorporation of the ESG rating into the portfolio decision.

2.4 Portfolio Rules

We evaluate 17 shrinkage models to determine the constrained GMVP. We refer to these rules as shrinkage-based portfolios. The first nine models rely on stock returns data only, where both the shrinkage intensity and target matrix are determined by the stock returns alone. The next four rules rely on both stock returns and ESG ratings, where the target matrix penalizes portfolios for having low ESG scores, and the shrinkage intensity relies on both data sources. The following four are ad-hoc that fully shrink the covariance matrix based on the ESG data. Finally, we consider two shrinkage-free rules: one is the naive portfolio that equally allocates wealth among all assets (DeMiguel et al., 2009b), and the other is based on the market capitalization of each asset, i.e., value-weighted portfolio. Below, we discuss each rule accordingly:

1. **Sample** denotes the sample covariance matrix that is estimated using T historical daily returns. In our case, we set $T \approx 5 \times 252$, representing five years of daily data, on average. When relevant, we set the same sample size for all other models.

2. **EWMA** denotes the exponentially weighted moving average dynamic estimate that attributes θ to the previous estimate and $1-\theta$ to recent innovations in squared returns. We set $\theta = 0.94$ in line with the literature (see, e.g., Longerstaey and Spencer (1996); Ardia and Boudt (2015)).

3. **LW** corresponds to the shrinkage estimator, where the covariance matrix is a convex combination between the sample covariance matrix and the shrinkage target (also referred to as prior). In this model, the prior is given by a one-factor model, where the factor is the equally weighted portfolio of all asset returns, and the shrinkage intensity (i.e., optimal convex com-

bination) is determined following Ledoit and Wolf (2003).

4. **Factor** is similar to the previous model, whereas the main difference is the calculation of the factor. In this case, the factor is extracted using factor analysis (Harman and Harman, 1976), and by default, we set the analysis to a single factor.[5]

5. **Constant** is similar to the last two models, whereas the main difference is the choice of the target matrix. In this model, all off-diagonal elements of the target matrix are equal to a single value denoting the average correlation among the stocks (Ledoit and Wolf, 2003). This representation reduces the number of parameters in the target matrix to $N + 1$, where N values for the diagonal elements represent each stock variance and one parameter for the average correlation.

6. **Cor** is similar to the **Constant** model, with the main difference being that the diagonal elements of the target matrix are set to one, such that the target matrix denotes a correlation matrix with a single constant correlation value in the off-diagonal elements.

7. **OnePar** follows from Ledoit and Wolf (2004), whereas the model shrinks the covariance matrix towards a diagonal matrix, and the diagonal elements are uniform and equal to the average stock variance.

8. **Diag** is similar to the **OnePar**, whereas the main difference is that the diagonal matrix preserves the diagonal elements of the sample covariance.

9. **Large** follows suit with the **LW** rule, with the main difference being the estimation of the shrinkage intensity estimator is derived under the case of a large number of assets/sample size, i.e., when $N \to \infty$, $T \to \infty$, and $N/T \to \rho \in (0, 1)$

10. **ESG** follows suit with the **Diag** rule, whereas the target matrix is defined using the reciprocal of the ESG ratings. We determine the shrinkage intensity using the closed-form solution from Ledoit and Wolf (2004).

11. **Econ** denotes the economic performance based on the ESG rating. Similar to the previous **ESG** rule, the target matrix is a diagonal matrix where element $i = 1, ..., N$ is the reciprocal of the score.[6] We determine the shrinkage intensity using the closed-form solution from Ledoit

[5]As a robustness check, we evaluate the sensitivity of this model using a larger number of factors.

[6]The reason we utilize the economic aspect from the ESG data is that most of the ratings of the environmental

and Wolf (2004).

12. **Corp** is similar to the above **ESG** rule, where the target matrix is a diagonal matrix where element $i = 1, ..., N$ is the reciprocal of the governance score based on the ESG rating. We determine the shrinkage intensity using the closed-form solution from Ledoit and Wolf (2004).

13. **Social** is similar to the above **ESG** rule, where the target matrix is a diagonal matrix with element $i = 1, ..., N$ stands for the reciprocal of the social score based on the ESG rating. Same as with the other ESG rules, we determine the shrinkage intensity using the closed-form solution from Ledoit and Wolf (2004).

14. **ESG_F** is similar to the **ESG** rule, whereas the covariance matrix is fully shrunk towards the target, i.e., $\alpha = 1$.

15. **Econ_F** is similar to the **Econ** rule, whereas the covariance matrix is fully shrunk towards the target, i.e., $\alpha = 1$

16. **Corp_F** is similar to the **Corp** rule, whereas the covariance matrix is fully shrunk towards the target, i.e., $\alpha = 1$

17. **Social_F** is similar to the **Social** rule, whereas the covariance matrix is fully shrunk towards the target, i.e., $\alpha = 1$

18. **Naive** denotes the equally weighted portfolio (DeMiguel et al., 2009b) that requires no data or estimation.

19. **Market** corresponds to the rule that allocates weight to the underlying assets based on total market capitalization, which is computed as the number of shares outstanding times the stock price.

For the first nine rules, we employ the `R` package developed by Ardia et al. (2017b). This package provides many features that are essential for optimal risk management. Additionally, researchers have already demonstrated its effectiveness in practice. For instance, Ardia and Boudt (2015) conducted research that employed this tool to investigate the implied expected returns of stocks using different efficient portfolio rules. Similarly, Ardia et al. (2017a) utilized the same package to

are missing over the full sample period of data. Nonetheless, by utilizing the overall score of the ESG, our analysis does not completely omit the environmental dimension from our investigation.

examine the impact of covariance matrix specification in risk-based rules. For ESG shrinkage rules 10 through 13, we utilize the linear shrinkage intensity manually using the closed-form solution proposed by Ledoit and Wolf (2004).

We note that for all ESG rules, the target matrix is determined by the diagonal matrix of the reciprocal of ESG scores. Such reciprocal values denote a penalty/risk in terms of ESG standing. Nonetheless, the values of these penalty metrics have different units than the ones of the stock volatility. Additionally, they exhibit larger values than stock volatility and greater heterogeneity, as we demonstrate later in Section 3.1. Such a difference could potentially result in a very small shrinkage intensity towards the target and, hence, ignore any value from the ESG penalties. To mitigate this issue, we rescale the ESG penalty scores such that they exhibit the same cross-sectional standard deviation as stock volatility.

3 Empirical Design

In this section, we describe our main empirical design, which covers the data sources, the backtesting procedure, and the performance evaluation metrics.

3.1 Data

Our data comes from two main sources. The first one is the Center for Research in Security Prices (CRSP), from which we collect stocks traded on the main three exchanges with common shares. The second one is ESG data from Refintiv, a global financial markets data and infrastructure provider. The data corresponds to a collection of environmental, economic, social, and governance scores and covers a broad range of metrics that reflect a company's performance and practices related to sustainability and responsible business conduct. The scores cover over 450 metrics across 10 themes, such as carbon emissions, labor standards, board diversity, and business ethics.

The Refinitiv data dates between 2002 and 2022, whereas the stock returns are daily and cover a larger time period. We allow the CRSP data to date back to 1998 to have an initial sample of five years, which is used to estimate the sample covariance matrix and, hence, perform different shrinkage techniques. Since the ESG data covers fiscal years, we shift the data by one year to make sure the data is available by the time the portfolio rules are constructed. We focus on all stocks

Figure 1: **Correlation**
This figure demonstrates the pairwise Pearson correlation coefficient among the variables from the final merged panel data summarized in Table 1. The × denotes the coefficients that are deemed statistically insignificant at the 5% level according to the Pearson statistical test.

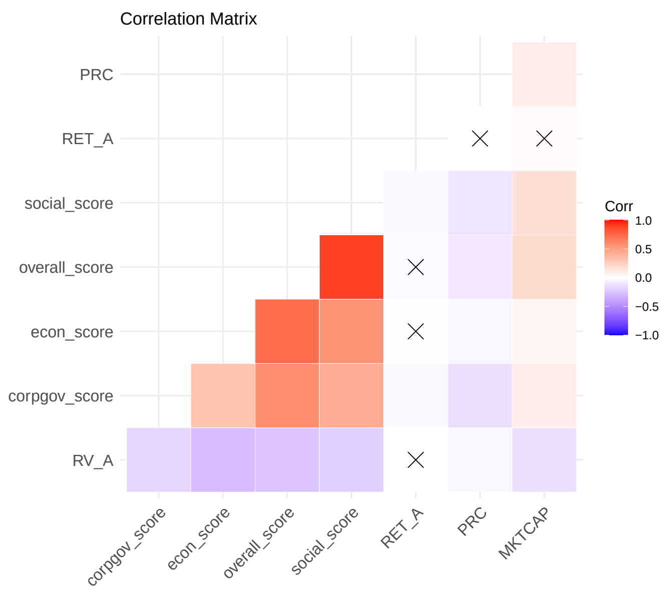

that were available during the sample period between 1998 and 2022, resulting in 235 unique stocks covering more than 6K trading days. In the same manner, all 235 stocks have data available for four ESG scores covering (i) overall score, (ii) economic score, (iii) corporate governance, and (iv) social score. We note that the Refinitiv data has a lot of missing data points for the environmental score. For this reason, we utilize the overall score of the ESG such that our analysis does not completely omit the environmental dimension, whereas the economic rating serves as a complement.

In Table 1, we report basic summary statistics using the merged annual panel data between the stock data and the ESG scores. Additionally, in Figure 1, we visualize the correlation among these variables using a heatmap that distinguishes between the significant and the insignificant coefficients using a basic Pearson test. Overall, we observe that the ESG ratings are highly correlated. For stock returns, we do not find statistical evidence on the panel level between annual returns and different ESG scores. The only exception is the corporate governance score; however, the correlation coefficient is relatively small (-3%). On the other hand, we find that higher ESG ratings are associated with lower realized volatility. The correlation coefficient for these scores is about -20%.

To motivate the appeal of the ESG scores as a covariance matrix target, we compare stock volatility with an ESG penalty metric. Specifically, let X denote a given ESG score, such that the ESG penalty is the reciprocal of X. As an illustration, Figure 2 plots the stock volatility versus the ESG penalty metrics. The former is computed using annual stock returns, which are annualized using $\sqrt{252}$. The latter is computed as a time series average of $\log(1/X)$ for a given ESG score X. In total, there are 235 observations for volatility and the ESG-related penalty. The observations are sorted into deciles based on the corresponding ESG penalty and aggregated using equal weighting. Across all panels, we observe a positive relationship between the two, i.e., higher stock volatility is associated with a higher ESG penalty metric from a cross-sectional perspective. This illustration is consistent with the one from Figure 1, which demonstrates a negative correlation between ESG scores and realized volatility.

In unreported results, we take a look at the variance of the ESG penalty scores and observe a greater heterogeneity among these scores compared to stock volatility, even though both are positively correlated. For instance, the cross-sectional median (volatility) of the ESG penalty metric using the overall score is 0.61 (1.70), whereas the corresponding statistic for stock volatility is 0.35 (0.12). This raises the question of whether integrating these penalty scores into the shrinkage of the covariance results in more or less robust OOS portfolios (Berg et al., 2022). We come back to this question later in Section 4 when that summarizes the main findings of our empirical design.

3.2 Performance Evaluation

Next, we discuss the performance evaluation of the 19 portfolio rules studied in our research. Specifically, standing at the end of the year 2002, we have five years of stock returns as well as the recent ESG scores. We shift ESG data by one a year ahead to avoid any look-ahead bias. For model m, we use the recent five years of stock returns to estimate the sample covariance and shrink it towards the relevant target, which may or may not depend on the stock returns data depending on the model. For instance, in the case of shrinking the covariance towards the overall ESG score metric, the target is constructed using the recent ESG score data.

Given the data at the end of year t and model m, we determine the shrank covariance matrix $\hat{\mathbf{\Sigma}}_{t,m}$ and solve for the GMVP portfolio, $\mathbf{w}_{t,m}$, according to Equation (2.9). We then compute the cumulative return on a single dollar invested in this portfolio over the next year. In the following

Figure 2: **Stock Volatility versus ESG Penalty**
This figure demonstrates the stock volatility versus the ESG penalty metrics. The former is computed using annual stock returns, which are annualized using $\sqrt{252}$. The latter is computed as a time series average of $\log(1/X)$, where X denotes the corresponding ESG score. In total, there are 235 observations for volatility and the penalty metric. For a cross-sectional illustration, the observations are sorted into deciles based on the corresponding ESG penalty score and aggregated by taking the average. The y-axis (x-axis) denotes the average decile volatility (ESG penalty score). Panels (a), (b), (c), and (d) refer to the cases where the ESG penalty is computed using the overall score, economic score, social score, and corporate governance, respectively.

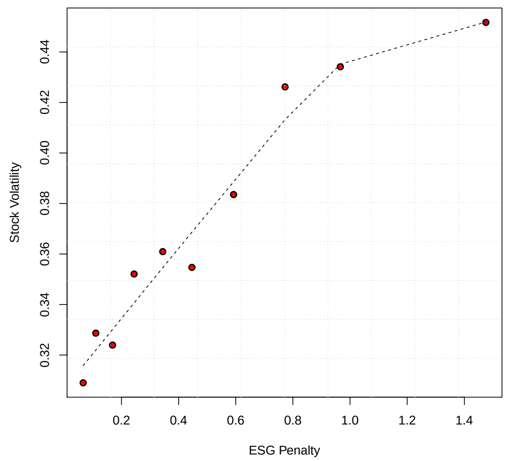

(a) Overall Score

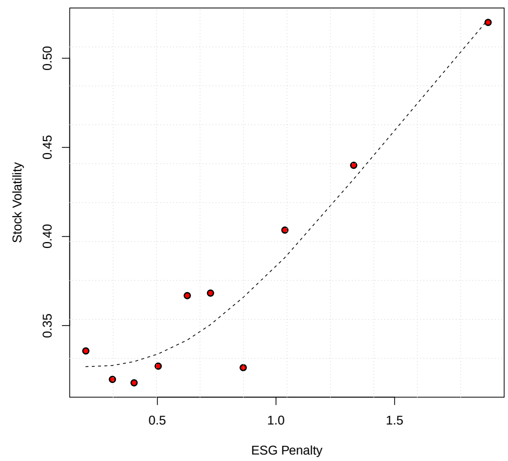

(b) Econ Score

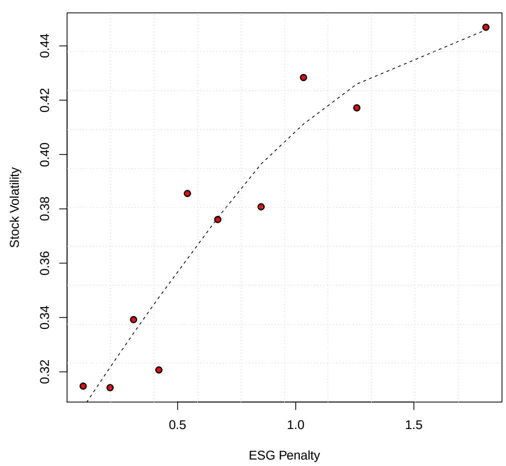

(c) Social Score

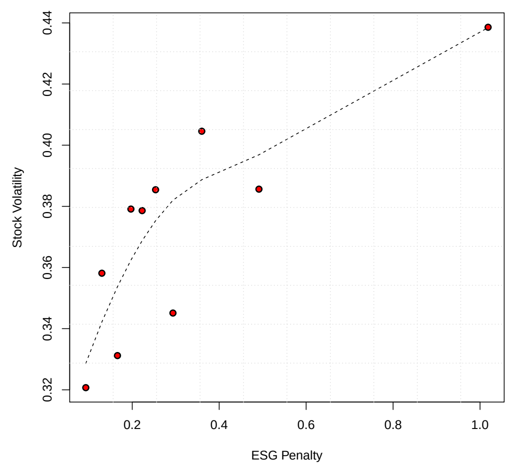

(d) Governance Score

year, we normalize the wealth of the portfolio to 1 and repeat the experiment. The experiment is repeated until we construct the last portfolio at the end of 2021. For a single model, the experiment results in 20 portfolio weights and 20 OOS portfolio annual returns. Based on these 20 observations, we compute a number of performance metrics:

Economic Performance Metrics

1. **TO** denotes the portfolio turnover, which takes into consideration the change in the target weight as well as the change in the previous portfolio due to stock returns. Consistent with the literature (see, e.g., Kan et al. (2022); Lassance et al. (2023)), we compute portfolio turnover as

$$TO_{t+1} =\mid \mathbf{w}_{t+1,m} - \mathbf{w}_{t,m} \circ (\mathbf{1} + \mathbf{r}_{t+1}) \mid_1 \tag{3.1}$$

   where $\circ$ denotes the Hadamard product, i.e., the element-by-element product between two vectors, $\mid \mathbf{v} \mid_1$ is the first norm of vector $\mathbf{v}$, and $\mathbf{r}_{t+1}$ is the annual vector of returns. For consistent comparison, we normalize the vector $\mathbf{w}_{t,m} \circ (\mathbf{1} + \mathbf{r}_{t+1})$ such that it weights sum to one. We report the average TO_{t+1} over the 20 years of data we have.

2. **TC** denotes transaction costs, which we set to 20 basis points in line with the literature (Kan et al., 2022; Lassance et al., 2023). Given the portfolio's TO and a fixed value of TC, we penalize the portfolio return as

$$\tilde{r}_{p,t+1} = r_{p,t+1} - TO_{t+1} \times TC \tag{3.2}$$

   where $r_{p,t+1}$ denotes the gross portfolio return over year $t + 1$.

3. **Mean** denotes the average mean return of the portfolio net of the transaction costs.

4. **Stdev** computes the annual standard deviation of the portfolio return over the whole sample period

5. **Sharpe** stands for the Sharpe ratio computed as the ratio between **Mean** and **Stdev**.

6. **Sortino** is the ratio between the **Mean** and the semi-standard deviation, which is computed as the conditional volatility of the portfolio return when its return is negative (Sortino and Van Der Meer, 1991).

7. **VaR** measures the 95% value-at-risk, computed as **Mean** minus the bottom 5% percentile of the OOS portfolio return. Since we have 20 annual observations, the bottom percentile stands for the worst portfolio return during the full OOS testing period.

ESG-Related Metrics

In addition to the above performance metrics, we evaluate how the above rules fare OOS in terms of ESG efficiency. Let $\mathbf{e}_{t+1}$ denote the vector of ESG ratings for all stocks in the portfolio at the end of year $t+1$. Given the portfolio $\mathbf{w}_{t,m}$ constructed at the end of year t based on model m, we compute its ex-post ESG performance using the dot product between the two vectors, i.e., $\mathbf{e}_{t+1}^{\top}\mathbf{w}_{t,m}$. We compute this metric for each year of the testing period and for each one of the four ESG scores we have, covering overall, economic, corporate, and social. As a summary, we report the average value of this metric.

4 Results and Discussion

In this section, we run a battery of tests given our empirical design to evaluate the OOS performance of the 19 rules covered in Section 2.4, taking into consideration different robustness checks.

4.1 Baseline

As baseline analysis, we consider the full universe of assets (covering 235 stocks) and perform a backtesting procedure to compute the performance metrics discussed in Section 3. The baseline results are summarized in Table 2, which covers three main panels. Panel (a) reports the results for the non-ESG rules that utilize different shrinkage techniques based on stock returns only. Panel (b) covers the eight ESG-related rules, where the first four are shrinkage-based that utilize both stock returns and ESG ratings and the other four rely on the ESG scores only. Finally, Panel (c) covers the equally (naive) and value-weighted (market) portfolios.

A number of important observations follow from Table 2. First, there is a large variation in performance across the nine shrinkage rules that rely on stock returns only. The two top performers in terms of risk-adjusted returns are **LW** and **Factor**. Both rules beat the naive and market portfolios. Interestingly, even though the **Factor** does not rely on ESG data, it results in the highest ESG performance among all other rules. This result sheds an interesting light in terms of filtering information from market prices about ESG ratings. Second, the rules that rely on both stock returns and ESG ratings outperform the naive and market portfolios in terms of Sortino ratio but not Sharpe ratio. This result indicates that shrinkage towards ESG mitigates downside risk.

However, their turnover is relatively high compared to the other rules. Additionally, these rules do not fare well in terms of OOS ESG scores compared to others. On the other hand, the ESG rules that rely on ESG scores alone fare better in terms of ESG OOS performance.

Third, the market portfolio results in the highest OOS overall ESG score. This result should not be surprising, given the increased focus on these scores by asset managers. For instance, Pedersen et al. (2021) state that "over \$30 trillion invested with explicit ESG goals as of the beginning of 2018," according to the Global Sustainable Investment Review reports. At the same time, we find that the market portfolio is not the most efficient when it comes to risk-adjusted returns. For investors with mean-variance preference, we note that this could result in lower utility. However, given the long-term implications of ESG, one should take into consideration intertemporal utility in drawing conclusions. Regardless, our findings overall state that rules such as the **LW** and **Factor**, which do not rely on ESG per se, could enhance ESG scores and risk-adjusted returns. This sheds interesting implications in terms of how asset managers should process stock information in forming portfolios.

4.2 Bootstrap

Next, we perform the following bootstrap experiment where we make 1000 draws from the full universe of stocks. At each draw, we randomly select $N = 50$ stocks and repeat the backtesting analysis to compute the performance metrics. For each performance metric, the bootstrap experiment results in 1000×19 values. For each rule, we compute the average performance metric and summarize the results in Table 3 in line with the baseline table. Similar to the findings of the baseline analysis, the **LW** and **Factor** rank among the top performers in terms of OOS risk-adjusted returns and ESG rating. Additionally, the bootstrap results reaffirm the evidence from the market portfolio that suggests that investors' valuations are tilted toward higher ESG scores but not higher risk-adjusted returns, measured using Sharpe and Sortino ratios. This result highlights a drawback in terms of market efficiency. Investors could utilize higher risk-adjusted returns by employing shrinkage techniques using market data alone while maintaining a relatively high ESG overall score.

4.3 Panel Regression Result

Given the bootstrap analysis, we utilize the 1000 experiments to study the relationship between economic metrics and ESG ratings. From policy implications, we are interested in investigating the contribution between these ratings and OOS performance. For instance, what is the impact of a one standard deviation increase in overall ESG rating on the portfolio mean-variance efficiency OOS? To answer this question, we run a panel regression with fixed effects. The panel corresponds to the bootstrap experiment with 19 models and 1000 repetitions summarized in Table 3. The panel covers in total $19 \times 1000 = 19,000$ model-experiment observation. For each observation, we focus on the OOS portfolio's economic performance as well as its ex-post ESG rating. Given the bootstrapped panel, we regress the six economic performance metrics on each ESG rating, capturing four aspects: overall, economic, corporate governance, and social. In all cases, we control for both model and experiment fixed effects. Table 4 documents these results.

We observe that all ESG ratings are associated with negative returns OOS. At the same time, these ratings are associated with lower portfolio volatility, where the coefficient of the ESG rating is negative across all four attributes. For the Sharpe ratio, the economic score is the only one that is positive and statistically significant compared to the other ESG ratings. For instance, a one standard deviation increase in the economic rating is associated with a 15% increase in the annual Sharpe ratio. Additionally, the econ score has the highest impact on portfolio downside risk, as evident by the Sortino ratio and VaR metrics. The corporate governance score shows the most significant improvement in portfolio turnover. Increased turnover could be attributed to different aspects, such as increased estimation risk or changes in hedging needs. Regardless, this result indicates that aligning portfolio preference with corporate governance score results in greater stability, better re-balancing, and enhanced lower transaction costs.

4.4 Additional Results

As robustness checks, we consider additional tests to address the sensitivity of our findings to different specifications. Due to space considerations, the tables from this section are located in the appendix.

4.4.1 Bootstrap using $N = 100$ Assets

The bootstrap considers $N = 50$ assets. The appeal of shrinkage-based techniques is more significant when considering high-dimensional settings. In this regard, we repeat the bootstrap analysis by drawing $N = 100$ asset rather than $N = 50$. Similar to Table 3, Table IA.1 reports the average for each performance metric and portfolio rule. The results affirm earlier conclusions drawn from the baseline and the previous bootstrap analysis. Additionally, the **Factor** rule remains among the top performers in terms of OOS risk-adjusted returns and ESG ratings. In fact, this rule outperforms the market by a large extent in terms of the Sortino ratio while attaining a relatively higher overall ESG score. Our analysis further stresses the appeal of shrinkage techniques using market data alone in terms of mean-variance-ESG efficiency.

4.4.2 Sensitivity to Different Factors

Given the consistent outperformance of the factor-based rule across our experiments (**Factor**), we proceed to investigate the sensitivity of this rule with respect to the number of components needed for implementation. We recall that this shrinkage model utilizes K components to identify common risk factors from which the target matrix is constructed. Therefore, we expect different shrunk covariance matrices and, hence, different portfolio rules for different values of K. At the same time, increasing the number of components could potentially result in greater noise since the resulting covariance matrix would conform to the sample matrix, which is less robust. As a result, it is unclear whether higher K results in better OOS performance.

We summarize the results for different components $K = 1, ..., 10$ in Table IA.2. Interestingly, we find that the rule based on the first component results in the highest overall ESG score. On the other hand, we observe that adding additional components results in a monotonic decrease in the overall score. Additionally, there is a jump in the portfolio risk-adjusted returns as we increase the number of components. However, this relationship is non-monotonic. While we consider arbitrary numbers of K, decision-makers need to choose the optimal value of K ex-ante. In this regard, portfolio managers could determine such value using cross-validation, which we leave for future research.

5 Understanding the Mechanisms

Our previous analysis provides several important observations in terms of the value of ESG ratings in portfolio selection. In the following section, we dig deeper into the mechanism behind our findings. Inspired by recent work by Giglio et al. (2023), we field a nationally representative survey of 1,500 respondents on ESG-related questions to identify consumer preferences for ESG versus value creation. On the one hand, an ESG portfolio would outperform a market allocation if consumers value non-pecuniary amenities and lack the means to express their preferences fully. On the other hand, it would underperform a market allocation if consumers appreciate value creation over non-pecuniary amenities. To answer this question, we launch a survey that allows us to distinguish between individual preferences and individual beliefs about organizational priorities - that is, what an individual may value versus their beliefs about what an organization should value.

We field our survey through Prolific.co, a leading survey platform that recruits participants and redirects them to Qualtrics-designed surveys by researchers. We obtain a nationally representative sample of 1,500 respondents based on sex, age, and ethnicity from the Census Bureau. We gather standard demographics, including age, education, race, sex, marital status, income, political affiliation, and other characteristics.

Crucially, we also ask several questions relating to ESG and investing. First, we ask respondents about their investment attitudes, allowing them to respond in one of five possible ways: "I do not invest," "I focus on financial performance," "I focus on environmental and social indicators," "I focus on industry and macroeconomic conditions," and "I am a passive investor." We find that 20% are passive investors (i.e., buy-and-hold portfolio strategy with a long-term investment horizon and minimal trading in the market), 25.5% do not invest, 5.5% focus on ESG indicators, 41% focus on financial performance, and 7.7% focus on industry and macroeconomic factors. Second, we ask the following ranked on a 1-5 scale: "How strongly do you agree/disagree with the following statement: a corporation that abides by the law and already maximizes shareholder value should consider additional objectives."[7]

[7] We believe that these scores are likely inflated because there is no forced trade-off; the tendency is to respond to such survey questions with a higher value, in general. Nonetheless, there is useful information in the response that we leverage below. We find that 27% report strongly agree, 45% report agree, 23% report neutral, 3.4% report disagree, and 0.93% report strongly disagree.

Third, we ask the following on a ranked scale of priorities: (i) “Below is a list of issues that are faced by businesses. For each one, please indicate the relative degree to which *you personally worry* about the following issues. Rank the most worrisome issue with a value of 1 and the least worrisome issue with a value of 10” and “Below is a list of issues that are faced by businesses. For each one, please indicate the relative degree to which *you believe that corporations* should be doing more versus less. Rank the most worrisome issue with a value of 1 and the least worrisome issue with a value of 10.” These issues include:

- Mitigating climate change and reducing carbon emissions.
- Promoting biodiversity and reducing land exploitation.
- Improving air quality and reducing waste production.
- Adopting energy-efficient technologies, such as renewable energy and sustainable designs.
- Ensuring supply chains are free of human exploitation and child labor.
- Ensuring supply chains are insulated from strategic adversaries like China and Russia.
- Promoting employee well-being and engagement.
- Advocating for workplace diversity for minorities and LBTQ+.
- Creating spaces for free thought and expression.
- Reducing income inequality and pay gaps between executives and non-managers.
- Paying employees a living wage.
- Investing in the local community, particularly by providing jobs at home.
- Investing in the well-being of citizens of other countries.

We begin by reporting the means and standard deviations of the ranking in Table 5. We find that “paying employees a living wage,” “mitigating climate change and reducing carbon emissions,” and “ensuring supply chains are free of human exploitation and child labor” receive the highest personal priorities with scores of 3.17, 4.58, and 4.82, respectively, and “creating spaces for free thought and expression” and “advocating for workplace diversity for minorities and LGBTQ” receive the lowest priorities with scores of 7.75 and 7.22, respectively. We also see large partisan differences, with liberals ranking climate change much higher - although, surprisingly, biodiversity is more important for conservatives - whereas conservatives rank paying and living wage and employee well-being much higher. We see broad similarities with what they believe companies should be doing.

Next, motivated by the result that only 5.5% focus on ESG indicators when investing, coupled with our ranking across ten different dimensions of ESG factors, we now examine how preferences for considering other objectives as a company are related to preferences for various dimensions of ESG. We juxtapose the focus on environmental versus governance factors by reporting the relationship between two environmental indicators - “mitigating climate change and reducing carbon emissions” and “promoting biodiversity and reducing land exploitation” - and two governance factors - “promoting employee well-being and engagement” and “creating spaces for free thought and expression.” Importantly, we ask how preferences for other organizational objectives are related to perceived company priorities over these ESG dimensions with and without controlling for individual preferences over these priorities. In particular, we examine whether there exists a residual correlation for the consideration of additional priorities after controlling for individual preferences:

$$y_i^{CORP} = \gamma OTHEROBJ_i + \phi y_i^{OWN} + \beta X_i + \varepsilon_i, \tag{5.1}$$

where y^{CORP} denotes an individual's belief about how a company should prioritize issue y, $OTHEROBJ$ denotes an indicator for whether an individual responds “strongly agree” or “agree” to the statement that a company should hold other objectives besides maximizing shareholder value, y^{OWN} denotes an individual's belief about how they prioritize issue y, and X denotes a vector of controls, including age, sex, employment status, race (White, Black, Asian), marital status, education (high school or less, some college, more than college, normalized to college), and political affiliation (normalized to moderates).

Table 6 documents the results of the regression model from (5.1). Column 1 shows that individuals who think that companies should consider other objectives place a 0.75 unit lower ranking (i.e., higher priority) on mitigating climate change, conditional on controls. Including political affiliation lowers the coefficient to -0.38, but it remains statistically significant. However, once their preferences are included in column 3, the statistical and economic significance of considerations for other priorities fades. There is an identical phenomenon in columns 4-6 when we focus on rankings for biodiversity. Next, we turn to more governance-related measures in the remaining columns and find a very different phenomenon. In particular, column 7 shows that individuals who believe companies should consider other objectives raise their rating on employee wellbeing-related matters by 0.81 units. That declines to 0.43 in column 8 with political affiliation as a control and down to 0.186 in column 9. However, now it becomes statistically insignificant at the 10% level (p-value of 0.15). We see the same phenomenon in columns 10-12, where we focus on ranking over free thought and expression.

These results highlight two corollaries. First, while the raw data suggests individuals want companies to prioritize environmental aims, these attitudes are driven by their own individual preferences that get super-imposed onto a company; after taking into account standard controls and their own preferences, there is no statistically or economically significant association with the prioritization of other objectives. Second, whereas those who want companies to focus on additional objectives rank environmental factors as more pressing, they rank governance-related factors as less pressing. That is informative because these governance-related factors are arguably pre-conditions for a robust corporate culture that would permit innovation and execution of environmental objectives, among others (i.e., lack of employee engagement undermines the delivery of core company objectives). In this sense, tilting towards ESG-related aims can disrupt core company priorities.

As a final diagnostic, we examine how the provision of information about the production of renewable energies affects support for them. In particular, we randomize two treatments and one control group across our sample. The first information treatment reads: "According to a 2021 investigation conducted by researchers at Sheffield University in the United Kingdom and reported on by the BBC, the global production of solar panels is conducted using forced labor from China's Uyghur Muslims in the Xinjiang province." The second information treatment reads: "According to a 2021 study published in Harvard Business Review, the waste produced by solar panels will make electricity from solar panels four times more expensive than the world's leading energy analysts thought. By 2035, discarded panels would outweigh new units sold by 2.56 times, making the levelized cost of energy - a measure of the overall cost of an energy-producing asset over its lifetime - four times the current projection." Next, we ask respondents to rate the following two statements on a 1-5 scale, ranging from strongly agree (value = 5) to strongly disagree (value = 1): "The installation and use of solar panels represent a key solution to addressing our current problems with environmental pollution and climate change" and "Electrical vehicles represent a key solution to addressing our current problems with environmental pollution and climate change."

We find that the provision of information, specifically around the second treatment pertaining to the Harvard Business Review (HBS) study exposure, is linked with declines in environmental attitudes: a 0.56 unit decline in attitudes about solar panels as a key part of the solution (p-value = 0.00), and a 0.092 decline in attitudes about electric vehicles (p-value = 0.19). However, we do not find statistically significant results for those who receive the other treatment with reference to the study by Sheffield University. One reason could be an insufficient sample size. Nevertheless, we view this as evidence that steady increases in the degree of information will mitigate support for ESG-tilted portfolios. Overall, the results of the information experiment suggest that exposure to additional and well-understood information about the typically unobserved costs of broader climate-mitigating policies can reduce support for ESG investing.

6 Conclusion

Our research examines the use of different shrinkage techniques and ESG-based rules in estimating the covariance matrix of stock returns. Based on these rules, we are able to reconcile risk-adjusted performance and ESG efficiency. Our analysis shows that factor analysis-based shrinkage techniques result in consistently high risk-adjusted returns and ESG scores out-of-sample, suggesting that filtering information from the covariance matrix has forward-looking implications for stock's related ESG performance. Moreover, we find that the calibration of the shrinkage intensity based on ESG metrics shows some merit of improvement rather than building rules based on ESG targets only. Our results also show that the market portfolio attains high ESG scores but at the cost of lower efficiency. Finally, the panel regression analysis demonstrates a positive relationship between all ESG scores regarding enhanced risk-adjusted returns and reduced downside risk. Overall, our study highlights the importance of incorporating ESG considerations into portfolio construction and provides insights for investors to design ESG-friendly portfolios that balance risk-adjusted performance and sustainability criteria.

References

Ardia, D., Bolliger, G., Boudt, K., Gagnon-Fleury, J.-P., 2017a. The impact of covariance misspecification in risk-based portfolios. Annals of Operations Research 254, 1–16.

Ardia, D., Boudt, K., 2015. Implied expected returns and the choice of a mean-variance efficient portfolio proxy. The Journal of Portfolio Management 41, 68–81.

Ardia, D., Boudt, K., Gagnon-Fleury, J.-P., 2017b. Riskportfolios: Computation of risk-based portfolios in r. J. Open Source Softw. 2, 171.

Barber, B. M., Morse, A., Yasuda, A., 2021. Impact investing. Journal of Financial Economics 139, 162–185.

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

Best, M. J., Grauer, R. R., 1991. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results. The review of financial studies 4, 315–342.

Bolton, P., Kacperczyk, M., 2022. Do investors care about carbon risk? NBER working paper 142, 517–549.

Christensen, D. M., Serafeim, G., Sikochi, A., 2022. Why is corporate virtue in the eye of the beholder? the case of esg ratings. The Accounting Review 97, 147–175.

DeMiguel, V., Garlappi, L., Nogales, F. J., Uppal, R., 2009a. A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. Management science 55, 798–812.

DeMiguel, V., Garlappi, L., Uppal, R., 2009b. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? The review of Financial studies 22, 1915–1953.

Edmans, A., 2011. Does the stock market fully value intangibles? employee satisfaction and equity prices. Journal of Financial Economics 101, 621–640.

Giglio, S., Kelly, B., Stroebel, J., 2021. Climate finance. Annual Review of Financial Economics 13, 15–36.

Giglio, S., Maggiori, M., Stroebel, J., Tan, Z., Utkus, S., Xu, X., 2023. Four facts about esg beliefs and investor portfolios. Tech. rep., National Bureau of Economic Research.

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

Graham, J. R., Grennan, J., Campbell, H. R., Rajgopal, S., 2022. Corporate culture: Evidence from the field. Journal of Financial Economics 146, 552–593.

Harman, H. H., Harman, H. H., 1976. Modern factor analysis. University of Chicago press.

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

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

Jagannathan, R., Ma, T., 2003. Risk reduction in large portfolios: Why imposing the wrong constraints helps. The journal of finance 58, 1651–1683.

Kan, R., Wang, X., Zhou, G., 2022. Optimal portfolio choice with estimation risk: No risk-free asset case. Management Science 68, 2047–2068.

Kelly, B., Pruitt, S., Su, Y., 2019. Characteristics are covariances: A unified model of risk and return. Journal of Financial Economics 134, 501–524.

Lassance, N., Martín-Utrera, A., Simaan, M., 2023. The risk of out-of-sample portfolio performance. Management Science (forthcoming) .

Ledoit, O., Wolf, M., 2003. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. Journal of empirical finance 10, 603–621.

Ledoit, O., Wolf, M., 2004. A well-conditioned estimator for large-dimensional covariance matrices. Journal of multivariate analysis 88, 365–411.

Ledoit, O., Wolf, M., 2022. The power of (non-) linear shrinking: A review and guide to covariance matrix estimation. Journal of Financial Econometrics 20, 187–218.

Lindsey, L., Pruitt, S., Schiller, C., 2023. The cost of esg investing. working paper .

Longerstaey, J., Spencer, M., 1996. Riskmetricstm-technical document. Morgan Guaranty Trust Company of New York: New York 51, 54.

Markowitz, H., 1952. Portfolio selection. The Journal of Finance 7, 77–91.

Merton, R. C., 1973. An intertemporal capital asset pricing model. Econometrica: Journal of the Econometric Society pp. 867–887.

Merton, R. C., 1980. On estimating the expected return on the market: An exploratory investigation. Journal of financial economics 8, 323–361.

Michaud, R. O., 1989. The markowitz optimization enigma: Is 'optimized' optimal? Financial analysts journal 45, 31–42.

Pastor, L., Stambaugh, R. F., Taylor, L. A., 2022. Dissecting green returns. Journal of Financial Economics 146, 403–424.

Pedersen, L. H., Fitzgibbons, S., Pomorski, L., 2021. Responsible investing: The esg-efficient frontier. Journal of Financial Economics 142, 572–597.

Sortino, F. A., Van Der Meer, R., 1991. Downside risk. Journal of portfolio Management 17, 27.

Tables

Table 1: **Summary Statistics**
The following table covers basis summary statistics for the final panel data that combines the CRSP and the Refinitiv ESG data sets. From top to bottom, the variables correspond, respectively to the year observation, stock price, market capitalization (in millions), annual return (cumulative over one year), realized annual volatility (computed as $\sqrt{252}\sigma_d$ with σ_d is the daily volatility), overall ESG score, economic score, corporate governance score, and social score.

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
year	4,935	2,012.000	6.056	2,002	2,007	2,017	2,022
PRC	4,935	931.991	15,618.200	0.067	28.455	74.930	528,921.000
MKTCAP	4,935	44,615.400	97,506.510	7.611	6,486.191	41,117.730	2,525,084.000
RET_A	4,935	0.140	0.456	−0.973	−0.060	0.290	17.743
RV_A	4,935	0.321	0.191	0.090	0.201	0.379	1.851
overall_score	4,935	0.732	0.252	0.016	0.581	0.932	0.987
econ_score	4,935	0.658	0.269	0.010	0.463	0.890	0.989
corpgov_score	4,935	0.795	0.162	0.012	0.734	0.910	0.983
social_score	4,935	0.651	0.269	0.016	0.462	0.880	0.994

Table 2: **Baseline Results**
The following table reports the baseline results using the full universe of assets, consisting of 235 stocks listed on the U.S. main three exchanges between 1998 and 2022. Panel (a) reports the results for the non-ESG rules that utilize different shrinkage techniques based on stock returns only. Panel (b) covers the eight ESG-related rules, where the first four the shrinkage-based that utilize both stock returns and ESG ratings, while the other four depend on the ESG scores only. Finally, Panel (c) covers the equally (naive) and value-weighted (market) portfolios. For more information about the rules see Section 2.4. For the performance metrics, see Section 3.2.

	Economic Performance Metrics						ESG Metrics			
Model	Mean	Stdev	Sharpe	Sortino	VaR	TO	Overall	Econ	Corporate	Social
Panel (a) Return-Based Rules										
Sample	0.090	0.141	0.638	1.501	0.213	3.107	0.721	0.733	0.771	0.660
EWMA	0.076	0.305	0.248	0.418	0.434	31.035	0.766	0.648	0.784	0.774
LW	0.095	0.120	0.789	2.285	0.150	2.321	0.731	0.734	0.775	0.669
Factor	0.097	0.114	0.857	2.407	0.139	0.764	0.854	0.752	0.837	0.803
Constant	0.043	0.176	0.244	0.281	0.225	1.063	0.890	0.802	0.834	0.824
Cor	0.084	0.139	0.609	1.264	0.202	2.607	0.729	0.744	0.772	0.662
OnePar	0.093	0.131	0.708	1.853	0.208	2.706	0.725	0.735	0.773	0.661
Diagonal	0.092	0.133	0.692	2.092	0.200	2.855	0.723	0.732	0.772	0.660
Large	0.095	0.120	0.789	2.285	0.150	2.321	0.731	0.734	0.775	0.669
Panel (b) ESG-Based Rules										
ESG	0.093	0.131	0.714	2.216	0.193	2.781	0.738	0.739	0.780	0.675
Econ	0.091	0.130	0.701	1.965	0.205	2.827	0.729	0.739	0.774	0.667
Corp	0.093	0.131	0.708	2.137	0.192	2.761	0.731	0.736	0.778	0.668
Social	0.095	0.128	0.738	2.382	0.163	2.795	0.749	0.743	0.784	0.689
ESG_F	0.145	0.180	0.808	0.850	0.234	0.221	0.811	0.715	0.825	0.729
Econ_F	0.141	0.175	0.810	0.917	0.231	0.303	0.797	0.733	0.815	0.711
Corp_F	0.149	0.191	0.780	0.920	0.238	0.215	0.767	0.679	0.822	0.682
Social_F	0.148	0.182	0.812	0.949	0.236	0.254	0.817	0.715	0.823	0.749
Panel (c) Other Rules										
Naive	0.153	0.199	0.770	0.949	0.242	0.188	0.744	0.665	0.800	0.663
Market	0.111	0.142	0.783	0.665	0.156	0.038	0.842	0.724	0.832	0.764

Table 3: **Bootstrap Results**

The following table conducts a similar experiment from Table 2. The main difference is the experiment is repeated 1000 times, where 50 stocks out of the 235 are sampled randomly at each experiment. As a summary, the table reports the average performance metrics based on the 1000 samples. Panel (a) reports the results for the non-ESG rules that utilize different shrinkage techniques based on stock returns only. Panel (b) covers the eight ESG-related rules, where the first four the shrinkage-based that utilize both stock returns and ESG ratings, while the other four depend on the ESG scores only. Finally, Panel (c) covers the equally (naive) and value-weighted (market) portfolios. For more information about the rules see Section 2.4. For the performance metrics, see Section 3.2.

	Economic Performance Metrics						ESG Metrics			
Model	Mean	Stdev	Sharpe	Sortino	VaR	TO	Overall	Econ	Corporate	Social
Panel (a) Return-Based Rules										
Sample	0.107	0.117	0.922	2.171	0.168	0.927	0.794	0.737	0.810	0.729
EWMA	0.085	0.217	0.408	1.034	0.307	7.761	0.777	0.708	0.797	0.722
LW	0.106	0.114	0.935	2.980	0.164	0.833	0.799	0.739	0.812	0.734
Factor	0.101	0.122	0.845	3.458	0.177	0.623	0.842	0.743	0.834	0.785
Constant	0.065	0.164	0.444	0.912	0.214	0.990	0.866	0.780	0.833	0.798
Cor	0.100	0.117	0.873	2.204	0.171	0.846	0.801	0.743	0.811	0.734
OnePar	0.108	0.116	0.940	2.576	0.165	0.852	0.795	0.736	0.811	0.729
Diagonal	0.107	0.116	0.932	2.206	0.165	0.887	0.795	0.737	0.810	0.730
Large	0.106	0.114	0.935	2.980	0.164	0.833	0.799	0.739	0.812	0.734
Panel (b) ESG-Based Rules										
ESG	0.107	0.116	0.936	2.211	0.166	0.878	0.801	0.740	0.814	0.735
Econ	0.107	0.115	0.937	2.218	0.165	0.889	0.798	0.741	0.811	0.732
Corp	0.107	0.116	0.933	2.336	0.166	0.876	0.798	0.738	0.814	0.732
Social	0.107	0.116	0.938	2.183	0.165	0.881	0.804	0.741	0.815	0.740
ESG_F	0.145	0.183	0.797	0.937	0.234	0.220	0.811	0.715	0.824	0.729
Econ_F	0.141	0.178	0.798	0.927	0.232	0.302	0.797	0.732	0.815	0.710
Corp_F	0.149	0.195	0.768	0.942	0.238	0.213	0.768	0.680	0.822	0.683
Social_F	0.147	0.185	0.800	0.954	0.236	0.253	0.817	0.715	0.822	0.749
Panel (c) Other Rules										
Naive	0.153	0.203	0.757	0.975	0.243	0.186	0.745	0.666	0.799	0.663
Market	0.112	0.150	0.750	0.755	0.180	0.039	0.840	0.724	0.830	0.762

Table 4: **Panel Regression with Fixed Effects**
The following table reports the results of a panel regression with fixed effects. The panel corresponds to the bootstrap experiment with 19 models and 1000 repetitions summarized in Table 3. The panel covers in total $19 \times 1000 = 19,000$ model-experiment observation. For each observation, we observe the out-of-sample portfolio economic performance as well as its ex-post ESG rating. We regress six performance metrics on each ESG rating, capturing four aspects: overall, economic, corporate governance, and social. In all cases, we control for the model-experiment's fixed effects.

	Dependent variable:					
	Mean	Stdev	Sharpe	Sortino	VaR	TO
	(1)	(2)	(3)	(4)	(5)	(6)
overall_score	-0.270^{***}	-0.216^{***}	-0.714^{***}	2.945^{**}	-0.243^{***}	-6.945^{***}
	(0.005)	(0.009)	(0.040)	(1.407)	(0.011)	(0.402)
Observations	19,000	19,000	19,000	18,541	19,000	19,000
R^2	0.124	0.033	0.018	0.0002	0.028	0.016
econ_score	-0.408^{***}	-0.595^{***}	0.575^{***}	10.852^{***}	-0.575^{***}	-2.419^{***}
	(0.005)	(0.008)	(0.043)	(1.518)	(0.011)	(0.438)
Observations	19,000	19,000	19,000	18,541	19,000	19,000
R^2	0.242	0.217	0.010	0.003	0.133	0.002
social_score	-0.308^{***}	-0.271^{***}	-0.530^{***}	4.958^{***}	-0.267^{***}	-0.623^{*}
	(0.004)	(0.007)	(0.035)	(1.234)	(0.009)	(0.356)
Observations	19,000	19,000	19,000	18,541	19,000	19,000
R^2	0.209	0.068	0.013	0.001	0.044	0.0002
corpgov_score	-0.055^{***}	-0.068^{***}	-0.562^{***}	4.106	-0.199^{***}	-23.091^{***}
	(0.011)	(0.016)	(0.075)	(2.650)	(0.020)	(0.741)
Observations	19,000	19,000	19,000	18,541	19,000	19,000
R^2	0.001	0.001	0.003	0.0001	0.005	0.051
Note:						*p<0.1; **p<0.05; ***p<0.01

Table 5: **Descriptive Results on Ranking of ESG Preferences**
The following table reports the means and standard deviations across different partitions of a nationally representative sample of 1,500 respondents fielded through Prolific in March 2023, reporting their rankings in response to the following two questions: "Below is a list of issues that are faced by businesses. For each one, please indicate the relative degree to which *you personally worry* about the following issues. Rank the most worrisome issue with a value of 1, and the least worrisome issue with a value of 10" and "Below is a list of issues that are faced by businesses. For each one, please indicate the relative degree to which *you believe that corporations* should be doing more versus less. Rank the most worrisome issue with a value of 1, and the least worrisome issue with a value of 10." The possible rankings include: Mitigating climate change and reducing carbon emissions; Promoting biodiversity and reducing land exploitation; Improving air quality and reducing waste production; Adopting energy efficient technologies, such as renewable energy and sustainable designs; Ensuring supply chains are free of human exploitation and child labor; Ensuring supply chains are insulated from strategic adversaries like China and Russia; Promoting employee well-being and engagement; Advocating for workplace diversity for minorities and LBTQ+; Creating spaces for free thought and expression; Reducing income inequality and pay gaps between executives and non-managers; Paying employees a living wage; Investing in the local community, particularly by providing jobs at home; Investing in the well-being of citizens of other countries.

	Pooled		Conservative		Liberal		Moderate	
	(1)		(2)		(3)		(4)	
	mean	sd	mean	sd	mean	sd	mean	sd
				Panel (a) Individual				
Your ranking of free throught/expression	7.75	2.58	6.85	2.76	8.30	2.31	7.51	2.62
Your ranking of LGBT/minority diversity	7.22	2.59	8.42	2.11	6.36	2.61	7.76	2.32
Your ranking of biodiversity	6.07	2.46	5.75	2.31	6.28	2.55	5.93	2.39
Your ranking of pay gaps among executives	5.59	2.92	6.15	2.95	5.19	2.82	5.85	2.98
Your ranking of energy efficiency/renewables	5.52	2.41	5.58	2.20	5.49	2.49	5.50	2.44
Your ranking of employee well-being	5.20	2.58	4.16	2.34	5.81	2.52	4.98	2.57
Your ranking of air quality/waste	5.09	2.33	4.72	2.18	5.44	2.34	4.75	2.34
Your ranking of anti-trafficking	4.82	2.47	4.07	2.36	5.13	2.46	4.93	2.47
Your ranking of climate change	4.58	2.85	6.04	2.77	3.81	2.64	4.71	2.76
Your ranking of paying a living wage	3.17	2.68	3.24	2.75	3.18	2.68	3.06	2.61
				Panel (a) Company				
Company ranking of free throught/expression	7.86	2.51	6.82	2.69	8.47	2.21	7.65	2.52
Company ranking of LGBT/minority diversity	7.11	2.61	8.26	2.21	6.40	2.63	7.43	2.48
Company ranking of biodiversity	5.83	2.42	5.42	2.37	6.20	2.39	5.47	2.42
Company ranking of energy efficiency/renewables	5.47	2.37	5.49	2.15	5.52	2.46	5.33	2.39
Company ranking of air quality/waste	5.30	2.30	4.80	2.21	5.63	2.27	5.14	2.35
Company ranking of pay gaps among executives	5.19	3.02	5.91	3.06	4.69	2.90	5.49	3.04
Company ranking of employee well-being	4.94	2.62	3.86	2.39	5.64	2.60	4.59	2.45
Company ranking of anti-trafficking	4.70	2.55	4.10	2.55	4.87	2.47	4.94	2.64
Company ranking of climate change	4.37	2.75	5.53	2.91	3.77	2.48	4.41	2.73
Company ranking of paying a living wage	4.23	3.22	4.82	3.20	3.80	3.06	4.55	3.42
Observations	1498		375		760		363	

Table 6: **Do Preferences for Non-market Factors Predict ESG Priorities?**

The following table reports the coefficients associated with regressions of ratings associated with a given ESG issue on an indicator for whether the individual believes that organizations should have other objectives besides value maximization for shareholders, conditional on controls and, in some specifications, their own preferences over the associated ESG issue. Controls include: political affiliation (conservative and liberal normalized to moderate), age, male, race (White, Black, and Asian), employment status, marital status, and education (high school or less, some college, more than college, normalized to college). The outcome variables are answers in response to the following questions: "Below is a list of issues that are faced by businesses. For each one, please indicate the relative degree to which *you personally worry* about the following issues. Rank the most worrisome issue with a value of 1, and the least worrisome issue with a value of 10" and "Below is a list of issues that are faced by businesses. For each one, please indicate the relative degree to which *you believe that corporations* should be doing more versus less. Rank the most worrisome issue with a value of 1, and the least worrisome issue with a value of 10." The specific rankings in this table include: Mitigating climate change and reducing carbon emissions; Promoting biodiversity and reducing land exploitation; Promoting employee well-being and engagement; Creating spaces for free thought and expression.

Dep. var.	Climate change			Biodiversity			Employee well-being			Free thought		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Other objectives	-.753***	-.383**	-.036	.192	.014	-.023	.812***	.428***	.186	.802***	.394***	.143
	[.164]	[.165]	[.131]	[.139]	[.146]	[.125]	[.150]	[.147]	[.132]	[.148]	[.149]	[.131]
Conservative		1.082***	.266		-.012	.101		-.684***	-.299*		-.971***	-.611***
		[.212]	[.166]		[.185]	[.159]		[.181]	[.160]		[.192]	[.167]
Liberal		-.597***	-.057		.709***	.532***		.984***	.601***		.843***	.473***
		[.172]	[.134]		[.156]	[.133]		[.161]	[.143]		[.156]	[.138]
Your ranking			.633***			.535***			.508***			.484***
Age	.015***	.010**	.000	-.005	-.002	-.004	-.013***	-.008*	-.004	.010**	.016***	.008**
	[.005]	[.005]	[.004]	[.004]	[.004]	[.004]	[.004]	[.004]	[.004]	[.004]	[.004]	[.004]
Male	-.195	-.276**	-.081	-.467***	-.426***	-.179*	-.318**	-.233*	-.205*	-.169	-.079	.021
	[.143]	[.139]	[.108]	[.126]	[.125]	[.105]	[.136]	[.132]	[.114]	[.128]	[.123]	[.105]
White	.395	.177	.173	.080	.165	.068	-.922**	-.712*	-.572	-.208	.024	.093
	[.369]	[.355]	[.274]	[.335]	[.338]	[.240]	[.403]	[.389]	[.351]	[.361]	[.341]	[.276]
Black	.339	.335	.186	.036	.060	.048	-1.127**	-1.104***	-.649*	-1.085***	-1.070***	-.721**
	[.417]	[.404]	[.313]	[.376]	[.378]	[.278]	[.440]	[.427]	[.386]	[.406]	[.387]	[.327]
Asian	.248	.064	-.002	.029	.130	.088	-1.326***	-1.125**	-.830**	-.454	-.246	-.155
	[.463]	[.452]	[.354]	[.425]	[.424]	[.326]	[.469]	[.456]	[.403]	[.438]	[.424]	[.350]
Fully employed	.111	-.019	-.024	-.023	.010	.031	-.227	-.117	-.007	-.116	.015	.014
	[.146]	[.144]	[.111]	[.130]	[.130]	[.109]	[.140]	[.136]	[.118]	[.134]	[.129]	[.110]
Married	.064	-.102	.016	-.263*	-.188	-.136	.032	.200	.253**	.006	.187	.073
	[.159]	[.154]	[.119]	[.135]	[.135]	[.114]	[.148]	[.143]	[.121]	[.140]	[.135]	[.116]
High school or less	.652	-.039	.573	.647	.963*	-.032	-1.216*	-.511	-.547	-.147	.608	.060
	[.840]	[.907]	[.541]	[.581]	[.584]	[.715]	[.621]	[.675]	[.570]	[.754]	[.669]	[.737]
Some college	.174	.169	.021	.117	.113	.135	.241	.241	.328**	.092	.095	.092
	[.177]	[.172]	[.136]	[.152]	[.150]	[.125]	[.162]	[.156]	[.136]	[.155]	[.149]	[.127]
More than college	-.160	-.047	-.043	.325*	.261	.200	.144	.019	.055	-.542***	-.670***	-.558***
	[.189]	[.183]	[.139]	[.175]	[.176]	[.151]	[.196]	[.190]	[.160]	[.187]	[.178]	[.158]
R-squared	.03	.08	.46	.02	.04	.33	.04	.10	.33	.05	.12	.35
Sample Size	1485	1485	1483	1485	1485	1483	1485	1485	1483	1485	1485	1483

Appendix

Table IA.1: **Bootstrap Results using $N = 100$ Assets**
The following table conducts a similar experiment from Table 3. The main difference is the analysis samples $N = 100$ stocks out of 235 randomly at each experiment. Overall, the analysis conducts 1000 experiments, and the table below reports the average performance metrics. Panel (a) reports the results for the non-ESG rules that utilize different shrinkage techniques based on stock returns only. Panel (b) covers the eight ESG-related rules, where the first four the shrinkage-based that utilize both stock returns and ESG ratings, while the other four depend on the ESG scores only. Finally, Panel (c) covers the equally (naive) and value-weighted (market) portfolios. For more information about the rules see Section 2.4. For the performance metrics, see Section 3.2.

	Economic Performance Metrics						ESG Metrics			
Model	Mean	Stdev	Sharpe	Sortino	VaR	TO	Overall	Econ	Corporate	Social
Panel (a) Return-Based Rules										
Sample	0.100	0.116	0.870	2.238	0.168	1.546	0.766	0.732	0.794	0.703
EWMA	0.073	0.278	0.286	0.783	0.379	16.213	0.760	0.703	0.771	0.713
LW	0.101	0.110	0.921	2.287	0.158	1.299	0.774	0.736	0.798	0.711
Factor	0.099	0.117	0.856	3.465	0.161	0.706	0.847	0.747	0.835	0.794
Constant	0.053	0.168	0.336	0.576	0.197	0.959	0.878	0.791	0.833	0.812
Cor	0.093	0.115	0.813	2.294	0.171	1.374	0.775	0.741	0.796	0.709
OnePar	0.101	0.113	0.903	2.111	0.163	1.401	0.769	0.734	0.797	0.705
Diagonal	0.100	0.114	0.888	2.941	0.164	1.464	0.768	0.733	0.795	0.705
Large	0.101	0.110	0.921	2.287	0.158	1.299	0.774	0.736	0.798	0.711
Panel (b) ESG-Based Rules										
ESG	0.101	0.114	0.897	2.197	0.164	1.441	0.777	0.738	0.801	0.714
Econ	0.100	0.113	0.900	2.252	0.163	1.460	0.772	0.738	0.797	0.709
Corp	0.101	0.114	0.891	2.214	0.164	1.435	0.773	0.735	0.801	0.710
Social	0.102	0.113	0.905	2.990	0.163	1.447	0.783	0.740	0.803	0.722
ESG_F	0.145	0.181	0.804	0.913	0.234	0.221	0.811	0.715	0.824	0.729
Econ_F	0.142	0.176	0.805	0.904	0.231	0.303	0.797	0.732	0.815	0.711
Corp_F	0.149	0.193	0.775	0.923	0.237	0.214	0.767	0.679	0.822	0.683
Social_F	0.148	0.183	0.807	0.925	0.235	0.254	0.817	0.715	0.823	0.749
Panel (c) Other Rules										
Naive	0.153	0.201	0.765	0.961	0.241	0.187	0.744	0.665	0.800	0.663
Market	0.112	0.146	0.770	0.710	0.168	0.038	0.841	0.724	0.831	0.763

Table IA.2: **Baseline Results using Different Factors**

The following table conducts a similar experiment from Table 2. The main difference is the analysis focuses on the factor model. In the baseline results, the factor analysis utilizes the first component ($K = 1$) to determine the target covariance matrix. In the following, we repeat the experiment for different values of $K = 1, 2, ..., 10$. The first row of the table is consistent with the **Factor** rule from Table 2.

	Economic Performance Metrics						ESG Metrics			
Model	Mean	Stdev	Sharpe	Sortino	VaR	TO	Overall	Econ	Corporate	Social
Factor_1	0.097	0.114	0.857	2.407	0.139	0.764	0.854	0.752	0.837	0.803
Factor_2	0.096	0.118	0.813	3.015	0.169	1.425	0.790	0.727	0.810	0.736
Factor_3	0.112	0.114	0.984	5.462	0.166	1.505	0.785	0.738	0.802	0.738
Factor_4	0.065	0.101	0.641	1.038	0.130	1.658	0.793	0.758	0.800	0.751
Factor_5	0.085	0.102	0.827	1.367	0.118	1.871	0.776	0.749	0.804	0.724
Factor_6	0.093	0.105	0.892	8.015	0.154	1.967	0.746	0.724	0.794	0.690
Factor_7	0.092	0.112	0.824	1.799	0.168	1.995	0.741	0.728	0.793	0.681
Factor_8	0.095	0.111	0.855	2.003	0.164	1.976	0.732	0.722	0.790	0.673
Factor_9	0.100	0.128	0.779	1.542	0.168	2.012	0.716	0.713	0.784	0.653
Factor_10	0.159	0.147	1.078	2.380	0.238	0.859	0.732	0.643	0.795	0.711