Scientific Beta
An SGX Group Company

Does ESG Information Deliver Investment Value? A High-Dimensional Portfolio Perspective

Research Overview

June 2025

About the Authors

Giovanni Bruno is a Senior Quantitative Analyst at Scientific Beta. His research focuses on asset pricing. He earned his PhD in finance at the Norwegian School of Economics, where he also worked as a Teaching Assistant delivering courses on Investments, Derivatives and Risk Management and Quantitative Investment. He holds a Master's Degree from LUISS Guido Carli University (Italy), where he obtained First-Class Honours in Quantitative Finance. Previously to his PhD, he held roles as a Consultant and a Quantitative Analyst with PricewaterhouseCoopers and Altran, respectively, where he specialised in Financial Risk Management.

Felix Goltz, PhD, is Research Director at Scientific Beta. He has been with Scientific Beta since inception. He carries out research in empirical finance and asset allocation, with a focus on alternative investments and indexing strategies. His work has appeared in various international academic and practitioner journals and handbooks, including the Journal of Portfolio Management, the Financial Analysts Journal, the Journal of Index Investing, the Journal of Investment Management and the Handbook of Finance (Wiley). He obtained an MSc and a PhD in finance from the University of Nice Sophia-Antipolis after studying economics and business administration at the University of Bayreuth and EDHEC Business School.

Antoine Naly is a Senior Quantitative Researcher at Scientific Beta. His research focuses on the interactions between equity markets, ESG, and macroeconomic variables. He holds a Master's degree in Financial Economics from HEC Montréal, and previously studied finance and macroeconomics at Lyon 3 University and the University of British Columbia. Prior to joining Scientific Beta, he successively worked as an investment analyst in advisory services at Mercer, as a quantitative macro analyst in global macro research at PGM Global, and as a Macro ESG strategist at Sustainable Market Strategies. He is also a CFA Charterholder.

We tested whether ESG data improves investment performance by analyzing 200+ ESG metrics. Our results show that ESG information allows to increase performance in traditional backtests but adds no value to portfolio construction when assessed in out-of-sample tests. In addition, financially optimal use of ESG data requires taking tilts towards ESG laggards on several dimensions. For investors, this means they need to set realistic expectations about what ESG integration can achieve. And when faced with claims about performance benefits, they need to require out-of-sample testing, not just traditional backtests.

Institutional investors and asset managers use an increasing number of ESG metrics when constructing their portfolios, yet a fundamental question remains unanswered: does all this sustainability data help improve portfolio performance? ESG information is inherently high-dimensional, with hundreds of metrics spanning climate impact, labour practices, board diversity, and more. We test whether this vast information set helps investors build better portfolios beyond what traditional financial data already provides.

Our approach is deliberately focussed on financial objectives. We adopt the perspective of "ESG aware" investors who care only about risk-adjusted returns, not about the ESG score of their portfolio. As one pension fund manager put it: "Having more information when making investment decisions can be informative and better than having less information." But is it? Our study finds that adding 200+ ESG metrics to standard financial data does not improve portfolio performance out of sample.

While ESG information is notoriously high-dimensional - encompassing countless issues measured in countless ways - empirical studies typically treat ESG as a monolith. They rely on either a single aggregate ESG score or focus on individual issues like carbon emissions or employee satisfaction. This not only ignores the rich, multifaceted information available to investors, but also opens the door to cherry picking metrics that showed high performance in the past. In contrast, we harness the full set of high-dimensional ESG data and test its value using out-of-sample analysis.

Why does out-of-sample testing matter? Traditional studies examine whether ESG variables earned a premium by looking backward over the entire history of returns - essentially asking, with perfect hindsight, whether an ESG tilt would have improved performance. But this tells us little about whether ESG data helps investors make better decisions going forward. We take a different approach: we evaluate ESG metrics using only the information that would have been available to investors at the time they constructed their portfolios.

1. A New Way to Assess the Benefits of ESG Information

Our analysis relies on two main ingredients: a high-dimensional dataset and robust portfolio methods.

Our high-dimensional dataset

To capture high-dimensional ESG information, we build a dataset of 222 characteristics from various sources. We include both traditional ESG ratings based on analyst opinions, and alternative ESG characteristics that rely on systematic textual analysis of news or firm documents. The ESG metrics we use span a wide array of sustainability issues in the environmental, social and governance domain, as well as various aggregate measures. We assess the incremental value of ESG information when used together with standard financial characteristics, reflected by 130 characteristics measuring valuation, momentum, quality and similar financial ratios.

Robust portfolio methods

Optimized portfolios are notoriously sensitive to noise in the risk and return parameters they use as inputs - a problem that is especially severe in high-dimensional settings. We use robust portfolio methods that reduce the impact of noise in our high-dimensional dataset by imposing sufficient diversification across the signals we use. We first build a portfolio for each characteristic, weighting stocks in proportion to their rank. These single factor portfolios form the basic building blocks of our analysis. We then construct a multi-factor portfolio that combines single factors in a way that maximises risk-adjusted returns, while ensuring portfolios remain diversified over multiple factors.

We use three different methods to make sure that our results do not hinge on any single technique. For each method, we consider three levels of diversification constraints: weak, moderate or stringent. These constraints allow us to use the high-dimensional information efficiently rather than overfit portfolio weights to noise. For comparison, we also consider a naive optimization where we maximise the Sharpe ratio without imposing diversification constraints.

2. New Evidence on the Value of ESG Information

Our analysis delivers three key results:

Apparent performance benefits vanish when removing hindsight on risk and return parameters

First, we show that ESG information is not redundant with respect to financial information, delivering an increase in Sharpe ratio of about 25%. However, these apparent benefits of using ESG information disappear out of sample, suggesting that the additional estimation risk in the expanded information sets offsets any information advantage.

Our main findings on the incremental value of ESG information are shown in Figure 1. The figure shows the difference in Sharpe ratio between the multifactor portfolio that combines ESG and financial information and the corresponding portfolio that uses only Financials, across the ten different portfolio methods we employ. The figure shows that any incremental benefits that appear in sample, disappear completely in the out-of-sample setting, even when using robust portfolio methods that avoid overfitting. This key result challenges the idea that integrating ESG information adds value.

Figure 1: Increase in Risk-adjusted returns when adding ESG Information

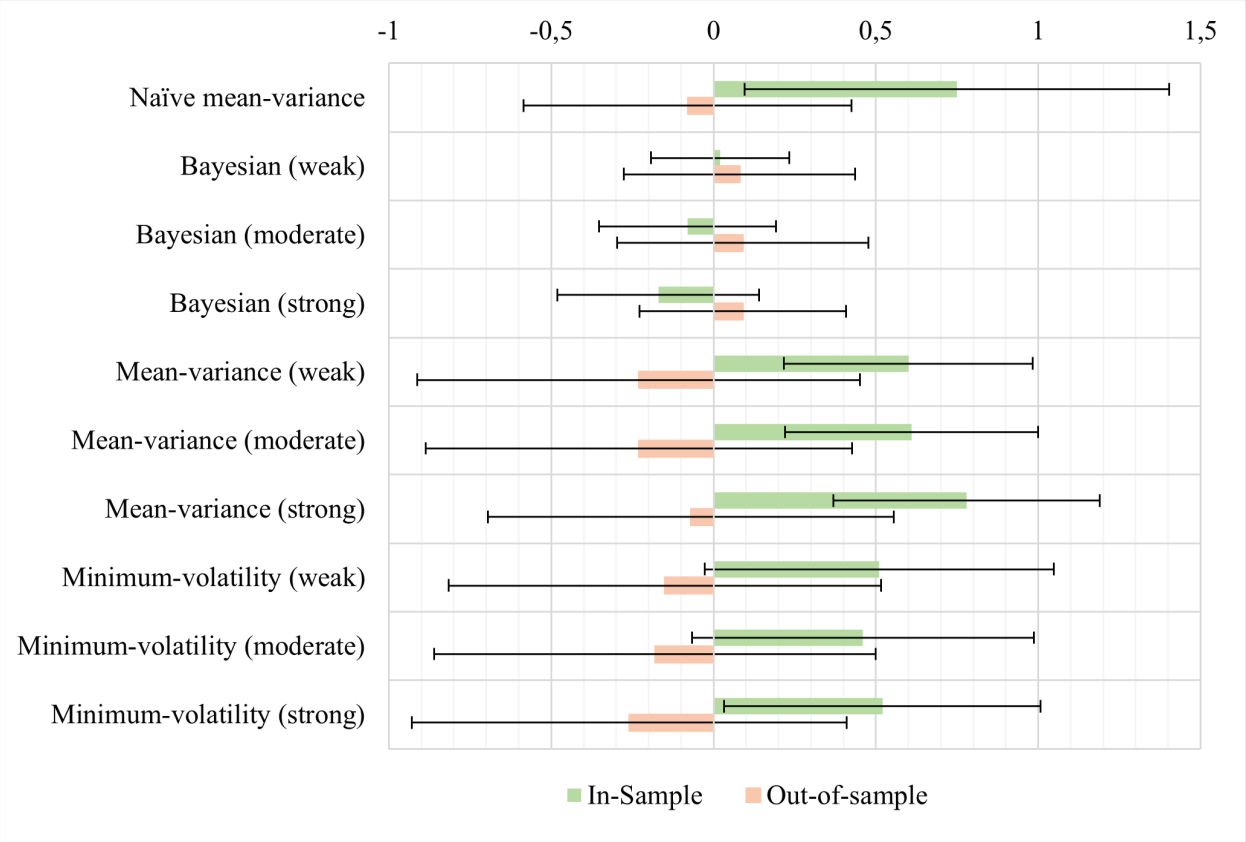

The chart reports differences in annualized Sharpe ratios (ΔSR) between various multifactor portfolios and their respective benchmarks, along with the bootstrapped 90%-confidence intervals of ΔSR. The multifactor portfolios are constructed using the set of factors that contains both ESG and Financial factors and 10 portfolio construction approaches. As for the multifactor portfolios' construction approaches, we report results for the Naïve mean-variance portfolio, three mean-variance approaches that employ Bayesian shrinkage with different priors (weak, moderate, and strong), three constrained mean-variance approaches with different levels of diversification constraint (weak, moderate, and strong), and three constrained minimum-volatility approaches with different levels of diversification constraint (weak, moderate, and strong). For each multifactor portfolio the benchmark is constructed using the same portfolio construction approach but using only the financial factors. Green (orange) bars represent in-sample (out-of-sample) results. We compute 90%-confidence intervals via a stationary block bootstrap.

Financially optimal use of ESG information does not lead to positive sustainability

Second, we show that optimal use of ESG information involves both "green" tilts emphasizing ESG leaders and "brown" tilts emphasizing ESG laggards. Figure 2 shows that ESG factors obtain substantial weight relative to financial factors, with brown tilts roughly as important as green tilts. For example, portfolios consistently tilt to sin stocks across different portfolio construction methods. On the other hand, human rights issues and water resource use are examples of consistent green tilts. Purely considering ESG as an information advantage thus requires anti-ESG investing on some dimensions. This result shows that using ESG information does not necessarily imply positive sustainability. We also find that alternative ESG metrics receive substantial weights relative to traditional analyst ratings.

Figure 2: Factor Tilts in Portfolios that Optimally Combine ESG and Financial Factors

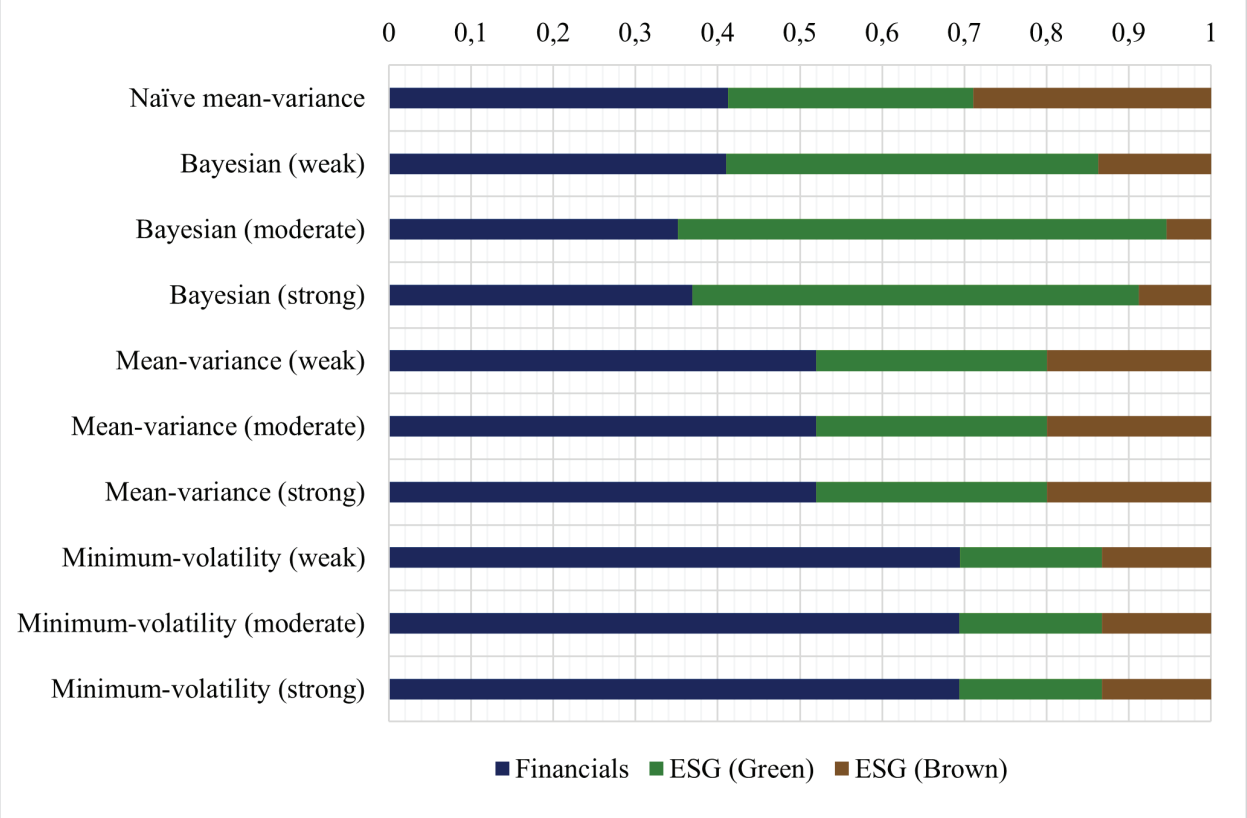

This figure displays the in-sample normalized factor weights across five factor categories for 10 multifactor portfolios. For each factor category, normalized weights are computed as the ratio of the sum of absolute weights within that category to the portfolio's total absolute factor weight (leverage). Factor categories comprise Financials (blue), ESG-Green (green) and ESG-Brown (brown). For ESG factors, green (brown) designation indicates that the factor weight in the multifactor portfolio corresponds to a long (short) position in the most sustainable leg paired with a short (long) position in the least sustainable leg of the factor. As for the multifactor portfolios' construction approaches, we report results for the Naïve mean-variance portfolio, three mean-variance approaches that employ Bayesian shrinkage with different priors (weak, moderate, and strong), three constrained mean-variance approaches with different levels of diversification constraint (weak, moderate, and strong), and three constrained minimum-volatility approaches with different levels of diversification constraint (weak, moderate, and strong). Factor weights are estimated using the full-sample covariance matrix and mean factor returns from August 2008 through June 2023.

ESG information does not provide risk reduction to portfolios

Third, we find that a focus on risk mitigation does not yield stronger benefits of using ESG information. Motivated by investor beliefs that ESG information should be useful for risk mitigation (Larcker et al., 2024), we test strategies that focus solely on volatility reduction. Similar to the Sharpe ratio results, we find that adding ESG characteristics does not lead to significant volatility reduction out of sample.

3. Beyond Traditional Backtests: The Need for Out-of-Sample Validation

Our analysis delivers a clear message: despite the explosion of ESG data, more information does not mean better portfolios. The 200+ ESG metrics we tested expand the information set without adding value. For investors, this means they need to set realistic expectations about what ESG integration can achieve. The lesson extends beyond ESG: Any claim of informational advantage must survive out-of-sample testing, not just traditional backtests.

Reference to full paper:

Bruno, Giovanni and Goltz, Felix and Naly, Antoine. 2025. "Does ESG Information Deliver Investment Value? A High-Dimensional Portfolio Perspective". *Scientific Beta working paper*, available at SSRN: https://ssrn.com/abstract=5280378

Disclaimer

The information contained on the Scientific Beta website (the "information") has been prepared by Scientific Beta Pte solely for informational purposes, is not a recommendation to participate in any particular trading strategy and should not be considered as an investment advice or an offer to sell or buy securities. All information provided by Scientific Beta Pte is impersonal and not tailored to the needs of any person, entity or group of persons. The information shall not be used for any unlawful or unauthorised purposes. The information is provided on an "as is" basis. Although Scientific Beta Pte shall obtain information from sources which Scientific Beta Pte considers to be reliable, neither Scientific Beta Pte nor its information providers involved in, or related to, compiling, computing or creating the information (collectively, the "Scientific Beta Pte Parties") guarantees the accuracy and/or the completeness of any of this information. None of the Scientific Beta Pte Parties makes any representation or warranty, express or implied, as to the results to be obtained by any person or entity from any use of this information, and the user of this information assumes the entire risk of any use made of this information. None of the Scientific Beta Pte Parties makes any express or implied warranties, and the Scientific Beta Pte Parties hereby expressly disclaim all implied warranties (including, without limitation, any implied warranties of accuracy, completeness, timeliness, sequence, currentness, merchantability, quality or fitness for a particular purpose) with respect to any of this information. Without limiting any of the foregoing, in no event shall any of the Scientific Beta Pte Parties have any liability for any direct, indirect, special, punitive, consequential or any other damages (including lost profits), even if notified of the possibility of such damages.

All Scientific Beta Indices and data are the exclusive property of Scientific Beta Pte.

Information containing any historical information, data or analysis should not be taken as an indication or guarantee of any future performance, analysis, forecast or prediction. Past performance does not guarantee future results. In many cases, hypothetical, back-tested results were achieved by means of the retroactive application of a simulation model and, as such, the corresponding results have inherent limitations. The Index returns shown do not represent the results of actual trading of investable assets/securities. Scientific Beta Pte maintains the Index and calculates the Index levels and performance shown or discussed, but does not manage actual assets. Index returns do not reflect payment of any sales charges or fees an investor may pay to purchase the securities underlying the Index or investment funds that are intended to track the performance of the Index. The imposition of these fees and charges would cause actual and back-tested performance of the securities/fund to be lower than the Index performance shown. Back-tested performance may not reflect the impact that any material market or economic factors might have had on the advisor's management of actual client assets.

The information may be used to create works such as charts and reports. Limited extracts of information and/or data derived from the information may be distributed or redistributed provided this is done infrequently in a non-systematic manner. The information may be used within the framework of investment activities provided that it is not done in connection with the marketing or promotion of any financial instrument or investment product that makes any explicit reference to the trademarks licensed to Scientific Beta Pte (SCIENTIFIC BETA, SCIBETA and any other trademarks licensed to Scientific Beta Pte) and that is based on, or seeks to match, the performance of the whole, or any part, of a Scientific Beta index. Such use requires that the Subscriber first enters into a separate license agreement with Scientific Beta Pte. The Information may not be used to verify or correct other data or information from other sources.

The terms contained in this Disclaimer are in addition to the Terms of Service for users without a subscription applicable to the Scientific Beta website, which are incorporated herein by reference.

For more information, please contact:
Séverine Cibelly on: +33 4 93 18 78 51 or by e-mail to: severine.cibelly@scientificbeta.com

Scientific Beta HQ & Asia-Pacific
2 Shenton Way
#02-02
SGX Centre I
Singapore 068804
Tel: +65 6713 7578

Scientific Beta R&D
470 Promenade des Anglais
06200 Nice
France
Tel: +33 493 187 851

Scientific Beta—Europe
77 Leadenhall Street
London EC3A 3DE London
United Kingdom

Scientific Beta—North America
1450 Broadway
12th floor
New York, NY 10018
United States

Scientific Beta—Sydney
Level 35, 100 Barangaroo Avenue
Sydney NSW 2000
Australia
Tel: +61 2 8114 4588

www.scientificbeta.com

Scientific Beta
An SGX Group Company