

Statistical Theory and Related Fields

ISSN: 2475-4269 (Print) 2475-4277 (Online) Journal homepage: www.tandfonline.com/journals/tstf20

Optimal multi-period dynamic portfolios for sustainable investing

Danping Li, Ruixun Zhang & Chaoyi Zhao

To cite this article: Danping Li, Ruixun Zhang & Chaoyi Zhao (09 Nov 2026): Optimal multi-period dynamic portfolios for sustainable investing, Statistical Theory and Related Fields, DOI: 10.1080/24754269.2026.2755107

To link to this article: https://doi.org/10.1080/24754269.2026.2755107

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

Published online: 09 Nov 2026.

Submit your article to this journal

Article views: 15

View related articles

View Crossmark data

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=tstf20

STATISTICAL THEORY AND RELATED FIELDS
https://doi.org/10.1080/24754269.2026.2755107

Taylor & Francis
by informa

OPEN ACCESS Check for updates

Optimal multi-period dynamic portfolios for sustainable investing

Danping Li [a,b,c], Ruixun Zhang [d,e,f,g] and Chaoyi Zhao [d,g,h]

[a]School of Statistics, East China Normal University, Shanghai, People's Republic of China; [b]KLATASDS-MOE, Shanghai, People's Republic of China; [c]China Inclusive Aging Finance Research Center, Shanghai, People's Republic of China; [d]School of Mathematical Sciences, Peking University, Beijing, People's Republic of China; [e]Center for Statistical Science, Peking University, Beijing, People's Republic of China; [f]National Engineering Laboratory for Big Data Analysis and Applications, Peking University, Beijing, People's Republic of China; [g]Laboratory for Mathematical Economics and Quantitative Finance, Peking University, Beijing, People's Republic of China; [h]Sloan School of Management and Laboratory for Financial Engineering, MIT, Cambridge, MA, USA

ABSTRACT

Sustainable investing involves balancing traditional investment strategies with environmental, social, and governance-related (ESG) objectives to improve long-term outcomes. We develop a stochastic control framework for constructing optimal multi-period dynamic ESG portfolios, in which investors form dynamic portfolios by maximizing utility functions that account for both profitability and the average ESG scores of their portfolios in the near and distant future. We provide an explicit solution to the investors' utility maximization problem, and show that the optimal weight is a linear combination of three components: a financial performance component, an ESG component, and a component reflecting the weight from the previous period. Using the U.S. stock market data, we find that the multi-period dynamic ESG portfolio outperforms single-period portfolios in terms of returns, Sharpe ratios, alphas, and information ratios. Our article extends dynamic portfolio theory in the context of sustainable investing, and provides guidance to practitioners on tailoring ESG portfolios to their specific preferences and to market conditions related to ESG.

ARTICLE HISTORY
Received 14 December 2025
Revised 19 May 2026
Accepted 20 October 2026

KEYWORDS
Sustainable investing; environmental, social, and governance (ESG); dynamic portfolio theory; transaction costs

1. Introduction

In recent years, sustainable investing – which balances traditional investing with environmental, social, and governance (ESG) objectives to improve long-term outcomes – has gained considerable attention from investors. In line with the increasing concerns about global warming, as of October 2025, more than 130 countries have committed to carbon-neutrality targets in various forms, representing approximately 79% of the world's population and 77% of global GDP, according to a tracker co-led by Oxford Net Zero.[1] As BlackRock CEO Larry

CONTACT Chaoyi Zhao zhaochaoyi@pku.edu.cn School of Mathematical Sciences, Peking University, No. 5 Yiheyuan Road, Haidian District Beijing 100871, People's Republic of China

[1] For example, the United States, the United Kingdom, and Japan have committed to net carbon neutrality by 2050, China by 2060, and India by 2070; see https://zerotracker.net. For more details on net-zero commitments at the firm level, see Bolton & Kacperczyk (2026).

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

Fink wrote in his 2022 letter to CEOs, global warming has led to an accelerating 'tectonic shift of capital', and the transition to carbon neutrality 'will not happen overnight'.[2]

The long journey to carbon neutrality requires sustainable investors to construct long-term green portfolios. However, researchers and investors typically build portfolios based on single-period mean-variance optimization with average ESG or emission constraints, and apply this methodology annually over long horizons (Pástor et al., 2021; Pedersen et al., 2021). Such an approach, while analytically tractable, implicitly assumes that portfolio rebalancing decisions are independent across periods. As a result, it may overlook critical features of real-world investing, such as changes in the underlying market environment, the persistence of firms' ESG performance, and the transaction costs associated with frequent rebalancing (Li et al., 2022).

Instead of repeatedly applying single-period models, can we directly construct a multi-period *dynamic* green portfolio that incorporates the future dynamics of ESG scores and transaction costs? Framing this problem in a multi-period setting is more appealing because, in practice, active portfolio managers make a sequence of rebalancing decisions rather than solving independent one-shot problems each year. Multi-period models offer advantages over single-period models by explicitly accounting for two features that are central in ESG investing. First, firms' ESG performance is highly autocorrelated (Bolton & Kacperczyk, 2021; Lo et al., 2022). Modelling this persistence allows current ESG information to be carried forward in a disciplined way, improving forecasts of future portfolio greenness and enabling investors to plan their ESG exposure over the investment horizon rather than treating each period in isolation. Second, portfolio rebalancing is costly in reality. A multi-period framework incorporates transaction costs directly and therefore chooses not only *what* portfolio to hold, but also *how* to move from current holdings to future desired holdings, optimally trading off expected returns and risks, long-horizon ESG impact, and the costs of adjusting the portfolio over time.

In this article, we develop a stochastic control framework for constructing optimal multi-period dynamic ESG portfolios. We consider a universe of firms with regularly updated ESG scores, and investors form dynamic portfolios by maximizing utility functions that account for both profitability and the average ESG scores of their portfolios in the near and distant future. Transaction costs arising from regular rebalancing are also incorporated. Our framework aligns with a number of multi-factor asset pricing models, such as the Capital Asset Pricing Model (CAPM) (Lintner, 1965; Sharpe, 1964), the Arbitrage Pricing Theory (APT) (Ross, 1976), and the Fama-French factor models (Fama & French, 1993, 2015). In addition, it allows for arbitrary investor preferences regarding the ESG scores of their portfolios, and can describe positive, negative, or zero correlations between assets' ESG scores and expected returns. This generality makes our results widely applicable to different types of markets, investors, and ESG metrics.

Under this framework, we provide an explicit solution to the utility maximization problem. In particular, we find that, in each period, the optimal weight is a linear combination of three components: a pecuniary component, an ESG component, and a component reflecting the weight from the previous period. This result sheds light on the key decision faced by ESG investors – a trade-off between future returns/risks, future greenness, and transaction costs. The first two components are forward-looking, depending on investors'

[2] See https://www.blackrock.com/corporate/investor-relations/larry-fink-ceo-letter.

estimates of future asset returns and the dynamic behavior of ESG scores. The last component is backward-looking, reducing transaction costs by partially retaining previous portfolio weights.

We further demonstrate that the relative proportions of these three components in the entire portfolio vary with investors' profit preferences, ESG tastes, and transaction costs. For example, investors with strong pecuniary preferences increase the proportion of the pecuniary component, while those with strong non-pecuniary preferences increase the proportion of the ESG component. If investors face higher transaction costs, they increase the proportion of the last-period-weight component. These findings provide guidance to practitioners to tailor their portfolios based on specific preferences and market conditions.

Finally, we illustrate our results empirically in the U.S. stock market. We employ the LSEG dataset, which includes annually updated ESG scores for approximately 3,000 U.S. companies from 2002 to 2024. We construct annually rebalanced long/short portfolios based on ESG scores and evaluate their financial performance. We then compare the portfolios formed using our dynamic ESG portfolio framework with their single-period counterparts and with those constructed using traditional single-period mean-variance optimization subject to ESG constraints, as proposed in the recent sustainable investing literature (Lo et al., 2022; Pástor et al., 2021; Pedersen et al., 2021). Our findings reveal that the multi-period dynamic ESG portfolio outperforms single-period portfolios in terms of returns, Sharpe ratios, alphas, and information ratios, thanks to its optimal multi-period trade-off between profitability, greenness, and transaction costs.

Related Literature. Our article is related to several strands of the literature. We focus on sustainable investing from the perspective of portfolio construction. There is a vast literature on green portfolios and their asset pricing implications, including, for example, Albuquerque et al. (2019), Idzorek et al. (2021), Pástor et al. (2021), Pedersen et al. (2021), Zerbib (2022), Flora and Tankov (2023), Lo et al. (2024), Lo and Zhang (2024), Lo and Zhang (2025), Goldstein et al. (2025), and Yang et al. (2025). See Tankov & Zhang (2025) for a recent review. However, this body of work predominantly considers single-period optimal portfolios with green constraints, without incorporating multi-period dynamics.

Jiao and Liu (2025) consider a multi-period portfolio optimization problem with ESG considerations. Our study differs from Jiao and Liu (2025) in three key aspects. First, we adopt a discrete-time framework, which is widely used in portfolio rebalancing practice, whereas they employ a continuous-time framework. Second, our model explicitly incorporates transaction costs. Third, we conduct an empirical analysis using data from approximately 3,000 U.S. companies over more than 20 years.

From a theoretical perspective, our article extends dynamic portfolio theory. Samuelson (1969) formulates and solves a multi-period portfolio selection problem under a utility maximization framework. Li and Ng (2000) make a breakthrough in solving the dynamic mean-variance problem for the multi-period case using dynamic programming. Subsequent research further examines the impact of transaction costs and the uncertainty of market evolution on dynamic portfolio selection under the multi-period case. For example, Gârleanu and Pedersen (2013) derive the optimal dynamic portfolio policy when trading is costly and asset returns are predictable by signals with different mean-reversion speeds. This framework is further extended in Glasserman and Xu (2013) by considering the robustness of the portfolio to the uncertainty of market evolution. Other related literature on dynamic portfolio theory includes Sun et al. (2006), Brown and Smith (2011), Dybvig and Pezzo (2020), and

Li et al. (2022). Our article extends dynamic portfolio theory in the context of sustainable investing.

Our empirical study also contributes to the growing literature on measuring the returns of ESG investing. Bolton and Kacperczyk (2021, 2023) and Bolton et al. (2022) find that higher stock returns and lower price-to-earnings and market-to-book ratios are associated with higher levels and growth rates of carbon emissions, both in the U.S. and internationally. Chen et al. (2025) consider the optimizing portfolios under carbon risk constraints. In the ESG investing literature, several studies report neutral or positive performance for green portfolios (Berg et al., 2024; Lindsey et al., 2025; Lo et al., 2024, 2022; Pástor et al., 2021, 2022). Other empirical literature on ESG includes Bansal et al. (2022), Blitz and Fabozzi (2017), Chan et al. (2020), Galema et al. (2008), Madhavan et al. (2021), and Renneboog et al. (2008).

Notation: In this article, we denote by $A^\top$ the transpose of a vector or matrix A, and by $\mathbb{E}[X]$ the expectation of a random variable or vector X. We use $X_1, X_2, \ldots, X_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mu, \Sigma)$ to represent that random variables $X_1, X_2, \ldots, X_n$ are mutually independent and identically distributed (i.i.d.), following a joint normal distribution with expectation μ and covariance matrix Σ. We use I to denote the identity matrix. All vectors are assumed to be column vectors unless otherwise stated.

Outline. The remainder of this paper is structured as follows. Section 2 introduces the multi-period dynamic ESG framework and provides the explicit solution to the optimization problem. Sections 3 and 4 present numerical experiments and empirical studies, providing intuition and illustrating the model's applicability. We conclude in Section 5 and provide proofs in the Appendix.

2. The multi-period ESG framework

In this section, we introduce our multi-period dynamic ESG investing framework. Section 2.1 introduces our assumptions about asset returns and ESG scores. Section 2.2 describes investors' behavior in forming optimal multi-period portfolios. Section 2.3 provides the explicit solution to our framework.

2.1. Asset returns and ESG scores

Assume that an investor is considering investing in a financial market consisting of $n+1$ assets that includes a risk-free asset and n stocks, over the period from time 0 to time T. For $k = 1, 2, \ldots, T$, let $r_k = (r_k^1, r_k^2, \ldots, r_k^n)^\top$ denote the vector of random returns of the n stocks from time $k-1$ to time k, and let $R_k = (r_k^1 - r^f, r_k^2 - r^f, \ldots, r_k^n - r^f)^\top$ denote the vector of excess returns beyond the risk-free rate, r^f. At time k, for $i = 1, 2, \ldots, n$, the ith stock has an observable ESG score, g_k^i, where a positive value of g_k^i corresponds to a green stock with positive social impact, while a negative value corresponds to a brown stock with negative social impact. Let $g_k = (g_k^1, g_k^2, \ldots, g_k^n)^\top$.

From time $k-1$ to k, assume that the excess return vector follows a multi-factor model with p common factors and the ESG factor:

$$R_k = Bf_k + cg_{k-1} + e_k, \quad \{e_k\}_{k=1}^T \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Sigma_e), \tag{1}$$

where $B \in \mathbb{R}^{n\times p}$ is the constant factor loading matrix, $f_k \in \mathbb{R}^p$ is the vector of common factor returns, $g_{k-1} \in \mathbb{R}^n$ is the vector of observable ESG scores at time $k-1$, and $e_k \in \mathbb{R}^n$ is the residual term. The parameter c is a constant. In particular, $c > 0$ implies that green stocks yield excess future returns, whereas $c < 0$ implies that the positive ESG impact may hurt stock returns. In this setting, f_k represents the collection of common factors, while the ESG scores g_{k-1} can be viewed as observable idiosyncratic factors.

The term cg_{k-1} is intended as a parsimonious reduced-form specification of the empirical relation between ESG characteristics and future stock returns. This linear specification is consistent with the empirical asset-pricing literature, where ESG scores are commonly included as firm-level characteristics in linear predictive regressions, Fama-MacBeth regressions, or multi-factor panel models of subsequent stock returns (Alves et al., 2025; Escobar-Saldívar et al., 2025; Liu et al., 2022). The empirical evidence is mixed: some studies find that higher ESG scores predict neutral or higher subsequent returns (Alves et al., 2025; Berg et al., 2024; Lindsey et al., 2025; Lo et al., 2024, 2022; Pástor et al., 2021, 2022), while others find that higher ESG scores are associated with lower expected returns (Bolton et al., 2022; Bolton & Kacperczyk, 2021, 2023; Escobar-Saldívar et al., 2025; Liu et al., 2022). Therefore, we do not impose any sign restriction on c. The case $c > 0$ corresponds to the view that ESG performance contains favorable information about future fundamentals, whereas $c < 0$ captures the case in which investors are willing to accept lower expected returns for holding greener stocks or in which brown stocks command a compensation premium (Pedersen et al., 2021).

The factor returns are assumed to follow the mean-reverting process:

$$f_k = \bar{f} + (I - \Psi)(f_{k-1} - \bar{f}) + \eta_k, \quad \{\eta_k\}_{k=1}^{T} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Sigma_F), \tag{2}$$

where $\bar{f} \in \mathbb{R}^p$ is the long-run mean of the factor returns and $\Psi \in \mathbb{R}^{p\times p}$ is the mean-reversion coefficient matrix. We define $\mu := B\bar{f} \in \mathbb{R}^n$ and $\widetilde{f}_k := f_k - \bar{f}$. Given all information at time k, the covariance matrix of R_{k+1} is $\Sigma := B\Sigma_F B^\top + \Sigma_e$.

Equation (2) characterizes the dependence of the factors over time. In particular, factor returns, including the Fama-French factors, have been shown to exhibit serial dependence and factor momentum. Ehsani and Linnainmaa (2022) document that most factors are positively autocorrelated and that factor momentum helps explain individual stock momentum. Relatedly, Petkova (2006) shows that the Fama-French factors are correlated with innovations in predictive state variables, and the strategic asset allocation literature commonly models factor variables using vector autoregressions to capture return predictability (Campbell et al., 2003; Gârleanu & Pedersen, 2013). When $\Psi = I$, the factor returns are i.i.d. over time with mean $\bar{f}$ and covariance matrix Σ_F.

As documented by Bolton and Kacperczyk (2021) and Lo et al. (2022), the time series of ESG scores exhibit strong autocorrelations. Therefore, we assume that the dynamics of $\{g_k\}_{k=0}^{T-1}$ follow a mean-reverting process:

$$g_k = (I - \Phi)g_{k-1} + \varepsilon_k, \quad \{\varepsilon_k\}_{k=1}^{T-1} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Omega), \tag{3}$$

where $\Phi \in \mathbb{R}^{n\times n}$ is the constant mean-reversion coefficient matrix, which characterizes the stability of the ESG scores. We further assume that $\{\varepsilon_k\}_{k=1}^{T-1}$, $\{e_k\}_{k=1}^{T}$, and $\{\eta_k\}_{k=1}^{T}$ are mutually independent.

Equation (3) provides a parsimonious way to capture the gradual adjustment of firms' ESG profiles, since ESG ratings are typically updated through changes in corporate policies,

disclosure practices, controversies, and industry-relative assessments rather than through purely transitory shocks. It is also consistent with empirical work that models ESG-score dynamics using autoregressive or rating-migration-type specifications. For example, Wedajo et al. (2024) explicitly study the dynamics of firm-level ESG scores using an ARMA-type specification, while studies on ESG rating migration and ESG rating dynamics document that ESG ratings evolve persistently over time rather than independently across periods (Chen & Lin, 2026). Moreover, major ESG rating providers (e.g., LSEG, S&P Global, and MSCI) construct ESG scores as bounded and often industry-relative measures, which further motivates a stationary mean-reverting specification rather than a nonstationary random walk.

2.2. Investor behavior

An investor manages her investment from time 0 to T. For $k = 0, 1, \ldots, T$, let x_k denote the investor's wealth at time k. Denote by u_k^i the fraction of wealth invested in stock i from time k to $k+1$, and let $u_k = (u_k^1, u_k^2, \ldots, u_k^n)^\top$. The fraction of wealth not invested in stocks, $1 - \sum_{i=1}^n u_k^i$, is allocated to the risk-free asset.[3] Therefore, without considering transaction costs, x_k follows the wealth dynamics given by

$$x_{k+1}^{\text{nocost}} = x_k + x_k r^f \left(1 - \sum_{i=1}^n u_k^i\right) + \sum_{i=1}^n x_k r_{k+1}^i u_k^i = x_k(1 + r^f) + x_k R_{k+1}^\top u_k.$$

However, in practice, trading is associated with transaction costs, including brokerage commissions, bid-ask spreads, and other market frictions such as taxes and regulatory fees. Following Gârleanu & Pedersen (2013), we assume that the transaction costs are quadratic in the trade amount. After incorporating these costs, the wealth process becomes

$$\begin{aligned} x_{k+1} &= x_{k+1}^{\text{nocost}} - \frac{1}{2}(1 + r^f)(x_k u_k - x_{k-1} u_{k-1})^\top \Lambda (x_k u_k - x_{k-1} u_{k-1}) \\ &= x_k(1 + r^f) + x_k R_{k+1}^\top u_k - \frac{1}{2}(1 + r^f)(x_k u_k - x_{k-1} u_{k-1})^\top \Lambda (x_k u_k - x_{k-1} u_{k-1}), \quad (4) \end{aligned}$$

where Λ is a symmetric positive definite matrix measuring the level of trading costs.[4]

We assume that the investor follows a multi-period dynamic ESG investment framework and adopts an exponential utility function as the objective (Jiao & Liu, 2025):

$$\max_{u_0, u_1, \ldots, u_{T-1}} \mathbb{E}\left[-\exp\left(-\gamma x_T - \sum_{k=0}^{T-1} d(1-\rho)^k x_k g_k^\top u_k\right)\right]. \quad (5)$$

The utility function (5) balances financial returns and ESG preferences. The first term, γx_T, represents the terminal wealth, where $\gamma > 0$ is the pecuniary preference parameter. The second term, $\sum_{k=0}^{T-1} d(1-\rho)^k x_k g_k^\top u_k$, represents the discounted cumulative ESG impact of the portfolio, where $d \geq 0$ measures the strength of ESG preference and $\rho \in [0, 1)$ controls the discount rate over time. Different values of γ and d reflect different pecuniary and non-pecuniary preferences of the investor.

[3] A negative value of $1 - \sum_{i=1}^n u_k^i$ implies that the investor borrows the risk-free asset.

[4] As discussed in Gârleanu and Pedersen (2013), Λ can be regarded as a multidimensional version of Kyle's lambda (Kyle, 1985). In addition, the empirical literature generally finds transaction costs to be convex (Lillo et al., 2003), with some researchers explicitly estimating quadratic trading costs (Breen et al., 2002).

2.3. Optimal portfolio weights

The solution to the optimization problem (5) is given in the following theorem.

Theorem 2.1: *The solution to the optimization problem* (5) *is*

$$u_k = \frac{1}{x_k} Q_k^{-1} \left[P_k g_k + H_k \mu + J_k \widetilde{f}_k + x_{k-1} \gamma \, (1 + r^f)^{T-k} \Lambda u_{k-1} \right], \quad k = 1, 2, \ldots, T-1, \tag{6}$$

where $Q_k \in \mathbb{R}^{n \times n}$, $P_k \in \mathbb{R}^{n \times n}$, $H_k \in \mathbb{R}^{n \times n}$, *and* $J_k \in \mathbb{R}^{n \times p}$ *are defined recursively backward by*

$$\begin{aligned}
Q_{k-1} &= \gamma \, (1 + r^f)^{T-k+1} \Lambda + \gamma^{\,2} (1 + r^f)^{2(T-k)} \Sigma_e + 2B_k \\
&\quad + 4U_{z,k}^\top N_k^f U_{z,k} + 4W_{z,k}^\top N_k^g W_{z,k}, && (7) \\
P_{k-1} &= \left[d(1-\rho)^{k-1} + \gamma \, c (1 + r^f)^{T-k} \right] I - 2C_k (I - \Phi) \\
&\quad - 4U_{z,k}^\top N_k^f U_{g,k} - 4W_{z,k}^\top N_k^g W_{g,k}, && (8) \\
H_{k-1} &= \gamma \, (1 + r^f)^{T-k} I - 2F_k - 4U_{z,k}^\top N_k^f U_{\mu,k} - 4W_{z,k}^\top N_k^g W_{\mu,k}, && (9) \\
J_{k-1} &= \gamma \, (1 + r^f)^{T-k} B (I - \Psi) - 2K_k (I - \Psi) \\
&\quad - 4U_{z,k}^\top N_k^f U_{f,k} - 4W_{z,k}^\top N_k^g W_{f,k}, && (10) \\
A_{k-1} &= (I - \Phi)^\top A_k (I - \Phi) + 2U_{g,k}^\top N_k^f U_{g,k} + 2W_{g,k}^\top N_k^g W_{g,k} \\
&\quad - \frac{1}{2} P_{k-1}^\top Q_{k-1}^{-1} P_{k-1}, && (11) \\
S_{k-1} &= (I - \Psi)^\top S_k (I - \Psi) + 2U_{f,k}^\top N_k^f U_{f,k} + 2W_{f,k}^\top N_k^g W_{f,k} \\
&\quad - \frac{1}{2} J_{k-1}^\top Q_{k-1}^{-1} J_{k-1}, && (12) \\
E_{k-1} &= (I - \Phi)^\top E_k (I - \Psi) + 2U_{g,k}^\top N_k^f U_{f,k} + 2W_{g,k}^\top N_k^g W_{f,k} \\
&\quad - \frac{1}{2} P_{k-1}^\top Q_{k-1}^{-1} J_{k-1}, && (13) \\
B_{k-1} &= \frac{1}{2} \left[\gamma \, (1 + r^f)^{T-k+1} \Lambda - \gamma^{\,2} (1 + r^f)^{2(T-k+1)} \Lambda Q_{k-1}^{-1} \Lambda \right], && (14) \\
C_{k-1} &= -\frac{1}{2} \gamma \, (1 + r^f)^{T-k+1} \Lambda Q_{k-1}^{-1} P_{k-1}, && (15) \\
K_{k-1} &= -\frac{1}{2} \gamma \, (1 + r^f)^{T-k+1} \Lambda Q_{k-1}^{-1} J_{k-1}, && (16) \\
D_{k-1} &= (I - \Phi)^\top D_k + 2U_{g,k}^\top N_k^f U_{\mu,k} + 2W_{g,k}^\top N_k^g W_{\mu,k} \\
&\quad - \frac{1}{2} P_{k-1}^\top Q_{k-1}^{-1} H_{k-1}, && (17) \\
L_{k-1} &= L_k (I - \Psi) + 2U_{\mu,k}^\top N_k^f U_{f,k} + 2W_{\mu,k}^\top N_k^g W_{f,k} \\
&\quad - \frac{1}{2} H_{k-1}^\top Q_{k-1}^{-1} J_{k-1}, && (18)
\end{aligned}$$

$$F_{k-1} = -\frac{1}{2}\gamma\,(1+r^f)^{T-k+1}\Lambda Q_{k-1}^{-1}H_{k-1}, \tag{19}$$

$$\begin{aligned} h_{k-1} = h_k &- \frac{1}{2}\ln\left|I - 2\Sigma_F^{1/2}S_k\Sigma_F^{1/2}\right| - \frac{1}{2}\ln\left|I - 2\Omega^{1/2}\widetilde{A}_k\Omega^{1/2}\right| \\ &+ 2\mu^\top U_{\mu,k}^\top N_k^f U_{\mu,k}\mu + 2\mu^\top W_{\mu,k}^\top N_k^g W_{\mu,k}\mu \\ &- \frac{1}{2}\mu^\top H_{k-1}^\top Q_{k-1}^{-1}H_{k-1}\mu, \end{aligned} \tag{20}$$

where

$$\begin{aligned} N_k^f &= \Sigma_F^{1/2}\left(I - 2\Sigma_F^{1/2}S_k\Sigma_F^{1/2}\right)^{-1}\Sigma_F^{1/2}, \\ U_{z,k} &= K_k^\top - \frac{1}{2}\gamma\,(1+r^f)^{T-k}B^\top, \\ U_{g,k} &= E_k^\top(I-\Phi), \\ U_{f,k} &= S_k(I-\Psi), \\ U_{\mu,k} &= L_k^\top, \\ \widetilde{A}_k &= A_k + 2E_kN_k^fE_k^\top, \\ N_k^g &= \Omega^{1/2}\left(I - 2\Omega^{1/2}\widetilde{A}_k\Omega^{1/2}\right)^{-1}\Omega^{1/2}, \\ W_{z,k} &= C_k^\top + 2E_kN_k^fU_{z,k}, \\ W_{g,k} &= A_k(I-\Phi) + 2E_kN_k^fU_{g,k}, \\ W_{f,k} &= E_k(I-\Psi) + 2E_kN_k^fU_{f,k}, \\ W_{\mu,k} &= D_k + 2E_kN_k^fU_{\mu,k}. \end{aligned}$$

Here, $A_k, B_k, C_k, D_k, F_k, Q_k, P_k, H_k \in \mathbb{R}^{n\times n}$, $S_k \in \mathbb{R}^{p\times p}$, $E_k, K_k, L_k, J_k \in \mathbb{R}^{n\times p}$, *and* h_k *satisfy the boundary conditions* $A_T = B_T = C_T = D_T = F_T = Q_T = P_T = H_T = S_T = E_T = K_T = L_T = J_T = 0$ *and* $h_T = 0$.

Theorem 2.1 has three important implications. First, according to (6), the optimal portfolio weights can be divided into three components – an ESG component, $P_k g_k$; a component related to financial performance (the FIN component), $H_k\mu + J_k f_k$; and a component related to the weight in the previous period (the PRE component), $x_{k-1}\gamma\,(1+r^f)^{T-k}\Lambda u_{k-1}$. The ESG and FIN components are prospective, reflecting expectations about future excess returns and the evolving trajectory of ESG scores. The PRE component is retrospective, mitigating trading costs through inertia in portfolio adjustments. This combination of forward-looking and backward-looking considerations illustrates the key advantage of our multi-period framework over static models.

Second, the FIN component itself can be further decomposed into two parts: the long-run component, $H_k\mu$, which depends on the long-run mean of the factor returns, μ, and the short-run component, $J_k\widetilde{f}_k$, which captures the current deviation of the factor returns from their long-run mean, f_k. This decomposition shows how return predictability affects the optimal dynamic portfolio choice. In particular, when $\Psi = I$, the factor returns reduce

to the i.i.d. case, so the current deviation $\widetilde{f}_k$ has no predictive power for future factor returns. In this case, $J_k = 0$, and the FIN component reduces to $H_k\mu$.

Third, Theorem 2.1 provides explicit recursive formulas for calculating the parameters in Equation (6). In practice, investors can calculate the portfolio weights using the backward recursion formulas (7)–(20). To illustrate this result further, we now consider several special cases of our framework.

Special Case 1: No Transaction Cost. If $\Psi = I$ and there are no transaction costs, i.e., $\Lambda = 0$, the optimal investment strategy in Equation (6) reduces to

$$u_k = \frac{1}{\gamma^2(1+r^f)^{2(T-k-1)}x_k} \times \left[\left[d(1-\rho)^k + \gamma\, c(1+r^f)^{T-k-1}\right]\Sigma^{-1}g_k + (1+r^f)^{T-k-1}\gamma\,\Sigma^{-1}\mu\right]. \tag{21}$$

In this case, investors can calculate the optimal weights using Equation (21) directly without relying on the recursive formulas in Theorem 2.1.

Equation (21) has the following implications. First, Equation (21) contains only two components: the ESG component, $\Sigma^{-1}g_k$, and the FIN component, $\Sigma^{-1}\mu$. This is consistent with the results in Jiao and Liu (2025, Remark 3.3). In particular, the third component in Equation (6) – the PRE component – vanishes, thanks to the absence of transaction costs. Second, the FIN component, $\Sigma^{-1}\mu$, coincides with the optimal portfolio weights that maximize the Sharpe ratio; see, for example, Lo et al. (2024).

Third, the ratio of the coefficients before the ESG and FIN components in Equation (21) is

$$\begin{aligned}\frac{\text{coefficient of the ESG component}}{\text{coefficient of the FIN component}} &= \frac{d(1-\rho)^k + \gamma\, c(1+r^f)^{T-k-1}}{(1+r^f)^{T-k-1}\gamma} \\ &= \frac{d\left[(1-\rho)(1+r^f)\right]^k}{\gamma\,(1+r^f)^{T-1}} + c.\end{aligned} \tag{22}$$

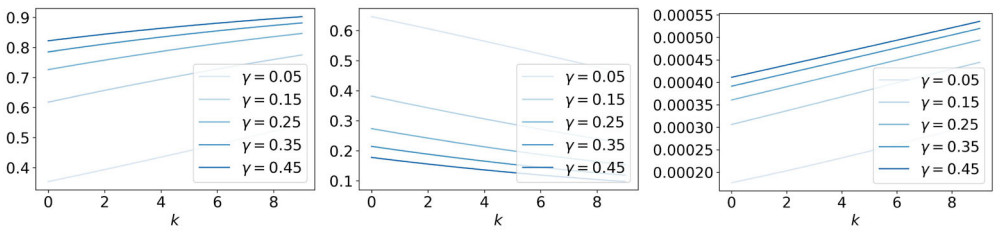
This implies that the relative proportion of the ESG component in the optimal weight increases with respect to the excess return parameter brought by ESG c as well as the ESG preference parameter in the exponential utility function d, while it decreases with respect to the pecuniary preference parameter γ.

Finally, Equation (22) reveals how the proportion of the ESG component changes over time k. In particular, if ρ is sufficiently small such that $(1-\rho)(1+r^f) > 1$, the ratio in Equation (22) increases with k, implying that investors place greater emphasis on ESG over time. By contrast, if ρ is large such that $(1-\rho)(1+r^f) < 1$, the discounted value of future ESG performance in the exponential utility function, Equation (5), diminishes, leading investors to prioritize ESG in the short run, and the ratio in Equation (22) decreases over time.

Special Case 2: Scalar Matrices. Assume that $\Psi = I$ and all of the matrices in our framework are scalar multiples of the identity matrix. That is, there exist positive constants ω, λ, ϕ, and σ such that $\mathbf{\Omega} = \omega I$, $\Lambda = \lambda I$, $\Phi = \phi I$, and $\Sigma = \sigma^2 I$. In this case, the optimal weight given by Theorem 2.1 takes the form

$$u_k \propto w_{\text{ESG},k}g_k + w_{\text{FIN},k}\mu + w_{\text{PRE},k}u_{k-1}, \tag{23}$$

where the scalars $w_{\text{ESG},k}$, $w_{\text{FIN},k}$, and $w_{\text{PRE},k}$ represent the relative proportions of the three components in the optimal weight. Thus, the optimal portfolio weight is a linear combination



(a) $\widetilde{w}_{\mathrm{FIN},k}$. (b) $\widetilde{w}_{\mathrm{ESG},k}$. (c) $\widetilde{w}_{\mathrm{PRE},k}$.

Figure 1. The proportions of the three components in Equation (23) vary with the pecuniary preference (γ). (a) $\widetilde{w}_{\mathrm{FIN},k}$; (b) $\widetilde{w}_{\mathrm{ESG},k}$; and (c) $\widetilde{w}_{\mathrm{PRE},k}$.

of three vectors – g_k, μ, and u_{k-1}, corresponding to the ESG score, the financial performance, and the previous-period weight, respectively.

Although the assumption that all matrices are scalar multiples of the identity is strong, it provides useful intuition about the optimal solution in Theorem 2.1 due to its simplicity. We adopt this assumption in the numerical analysis in the following section for illustrative purposes.[5]

3. Numerical analysis

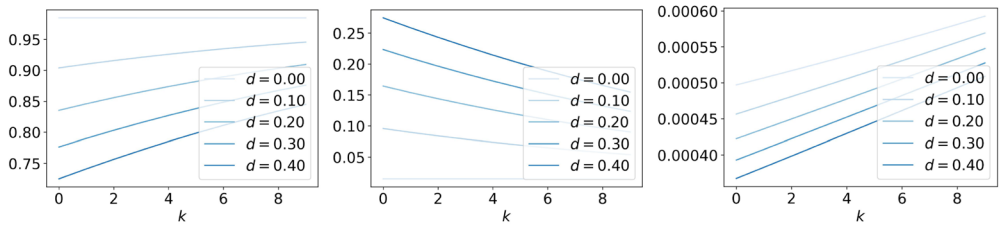
In this section, we conduct a numerical analysis to study the impact of parameters on the solution given by Theorem 2.1. We assume that Special Case 2 described in Section 2.3 holds, and examine how $w_{\mathrm{ESG},k}$, $w_{\mathrm{FIN},k}$, and $w_{\mathrm{PRE},k}$ in Equation (23) change with respect to the model parameters.

In our numerical analysis, we set the parameter values as follows: $\sigma = 0.2$, $\omega = 0.1$, $\phi = 0.3$, $\rho = 0.1$, $c = 0.015$, $r^f = 0.02$, and $T = 10$. We focus on the impact of the pecuniary preference γ, the ESG taste d, and the transaction costs λ. Given these parameters, we use Theorem 2.1 to obtain the optimal portfolio weights for each period. We standardize the three components $w_{\mathrm{ESG},k}$, $w_{\mathrm{FIN},k}$, and $w_{\mathrm{PRE},k}$ so that they sum to one for better illustration. To avoid notational confusion, we denote the standardized component proportions by $\widetilde{w}_{\mathrm{ESG},k}$, $\widetilde{w}_{\mathrm{FIN},k}$, and $\widetilde{w}_{\mathrm{PRE},k}$, respectively.

Figures 1, 2, and 3 illustrate the impact of the pecuniary preference γ, the ESG taste d, and the transaction costs λ, respectively. In Figure 1, we fix $d = 0.1$ and $\lambda = 0.05\%$; in Figure 2, we fix $\gamma = 1$ and $\lambda = 0.05\%$; and in Figure 3, we fix $d = 0.1$ and $\gamma = 1$. The figures show that, as γ, d, and λ change, the allocation of portfolio components adjusts accordingly.

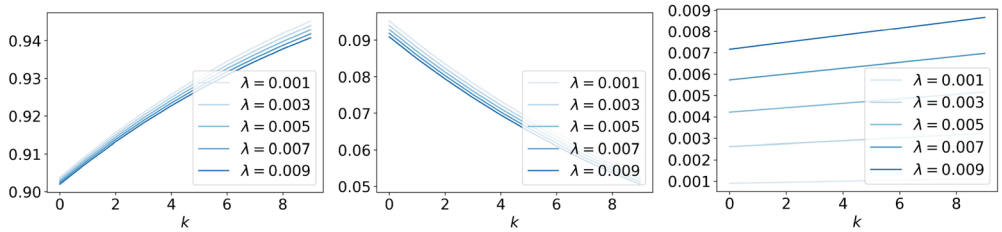
From Figure 1, as the pecuniary preference γ increases, both the FIN component ($\widetilde{w}_{\mathrm{FIN},k}$) and the PRE component ($\widetilde{w}_{\mathrm{PRE},k}$) increase, while the ESG component ($\widetilde{w}_{\mathrm{ESG},k}$) diminishes. Conversely, Figure 2 shows that as the non-pecuniary preference d increases, both the FIN and PRE components decrease, while the ESG component rises. This is intuitive because when the pecuniary preference γ is larger, the investor pays more attention to traditional financial goals, whereas when the non-pecuniary preference d increases, the investor assigns

[5] The assumption $\Lambda = \lambda I$ is used by Gârleanu and Pedersen (2013). They also adopt an assumption similar to $\Phi = \phi I$ for factor returns. The assumption $\Omega = \omega I$ can serve as a good approximation when the dynamics of ESG scores are homogeneous and uncorrelated between different stocks. The last assumption, $\Sigma = \sigma^2 I$, generally does not hold. However, as discussed in Lo et al. (2024, Proposition EC.5) and Lo and Zhang (2024, Proposition A.2), if investors focus only on residual returns beyond the factors, $\Sigma = \sigma^2 I$ approximately holds.



(a) $\widetilde{w}_{\mathrm{FIN},k}$. (b) $\widetilde{w}_{\mathrm{ESG},k}$. (c) $\widetilde{w}_{\mathrm{PRE},k}$.

Figure 2. The proportions of the three components in Equation (23) vary with the ESG preference (d). (a) $\widetilde{w}_{\mathrm{FIN},k}$; (b) $\widetilde{w}_{\mathrm{ESG},k}$; and (c) $\widetilde{w}_{\mathrm{PRE},k}$.



(a) $\widetilde{w}_{\mathrm{FIN},k}$. (b) $\widetilde{w}_{\mathrm{ESG},k}$. (c) $\widetilde{w}_{\mathrm{PRE},k}$.

Figure 3. The proportions of the three components in Equation (23) vary with the transaction costs (λ). (a) $\widetilde{w}_{\mathrm{FIN},k}$; (b) $\widetilde{w}_{\mathrm{ESG},k}$; and (c) $\widetilde{w}_{\mathrm{PRE},k}$.

greater weight to ESG considerations. We also note that the PRE component is relatively small because the transaction costs are modest ($\lambda = 0.05\%$) compared with the return and ESG objectives. Therefore, the incentive to retain last-period weights is weaker, leading to a lower proportion of the PRE component.

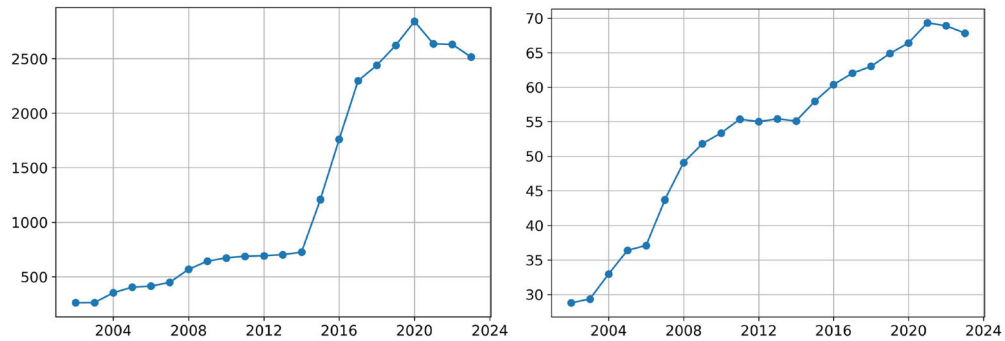
Figure 3 further demonstrates that as the transaction costs λ increase, the optimal portfolio becomes more inclined to retain previous weights, which reduces both the FIN and ESG components. The rationale is that maintaining previous weights lowers portfolio turnover, thereby reducing transaction costs.

Finally, from Figures 1–3, we also observe that the proportions of the three components fluctuate over time, reflecting the dynamic interplay between financial and non-financial considerations in investment decision-making. In particular, under the parameter setup of our numerical analysis, as time k progresses, both the FIN and PRE components increase, while the ESG component decreases. This finding aligns with the results of Special Case 1 in Section 2.3, which indicate that when $(1-\rho)(1+r^f) < 1$, investors focus more on ESG in the short run, but shift their attention to financial returns as the investment horizon approaches its end.[6]

4. Empirical study

In this section, we illustrate our results empirically using data from the U.S. stock market. We obtain ESG scores from the LSEG database, which is one of the most comprehensive ESG

[6] In our numerical analysis, we set $\rho = 0.1$ and $r^f = 0.02$, which implies $(1-\rho)(1+r^f) = 0.918 < 1$.



(a) Number of stocks. (b) Average ESG Combined Scores.

Figure 4. Number of stocks with valid ESG Combined Scores in our dataset and the cross-sectional average values of the ESG Combined Scores for each year. (a) Number of stocks. (b) Average ESG Combined Scores.

datasets in the industry,[7] and construct annually rebalanced portfolios based on ESG scores to evaluate their financial performance. Sections 4.1, 4.2, and 4.3 show the summary statistics of our dataset, the portfolio construction methodology, and the portfolio performance, respectively.

4.1. Data description

The LSEG dataset provides annually updated ESG measures for around 3,000 U.S. companies from 2002 to 2024. We use the ESG Combined Score – an overall company score based on reported information in the environmental, social and corporate governance pillars of the LSEG dataset – as our measure of ESG performance. For each company in each year, the ESG Combined Score ranges from 0 to 100, with higher values implying better ESG performance.

We obtain monthly dividend-adjusted return data for U.S. companies from 2000 to 2024 from the CRSP dataset.[8] We also obtain monthly Fama-French five-factor (Fama & French, 2015) data from Kenneth R. French's website.[9] These Fama-French five factors serve as the factor vector f_k in Equation (1).

For each year k, we match stocks with valid ESG Combined Scores in year k and valid monthly returns in year $k+1$. Figure 4(a,b) show the number of stocks with valid data[10] and the cross-sectional average ESG scores by year k, respectively. Figure 4(a) shows that the number of stocks in our sample generally increases, particularly after 2015 and before 2020, indicating greater coverage over the study period.

Figure 4(b) presents the cross-sectional average ESG Combined Scores for each year. We find that the average ESG Combined Score generally increases over time, likely as a result of improvements in energy efficiency, technological innovation, or a growing reliance on renewable energy sources. This pattern is consistent with observations documented in the literature, for example, Bolton and Kacperczyk (2021) and Lo et al. (2022).

[7] The LSEG database was previously known as Refinitiv. See https://www.lseg.com/en/data-analytics/financial-data/company-data/esg-data.

[8] We obtain the CRSP data through the Wharton Research Data Service.

[9] See https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

[10] The data for 2024 are not included in the figure because, although the ESG Combined Scores for 2024 are available in the LSEG dataset, the monthly returns for 2025 have not yet been updated in the CRSP dataset.

Table 1. Cross-sectional average of ESG Combined Scores for each year in our dataset.

year	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012
mean	28.78	29.35	32.96	36.39	37.11	43.70	49.08	51.80	53.35	55.35	55.00
std	15.05	14.49	15.00	18.03	18.42	18.02	19.09	20.34	19.47	18.80	18.47
first quartile	17.13	18.28	22.41	22.59	22.58	29.93	35.34	34.70	37.91	42.03	41.62
median	26.19	26.09	31.13	32.63	34.27	41.16	48.34	52.00	51.94	55.19	57.15
third quartile	39.47	39.28	41.79	47.31	48.14	56.05	62.39	68.82	69.08	69.33	69.14
year	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
mean	55.40	55.10	57.94	60.39	62.01	63.02	64.87	66.40	69.31	68.90	67.83
std	18.50	18.36	17.73	16.53	16.33	16.19	15.56	15.15	13.32	13.29	13.14
first quartile	42.63	41.16	45.11	47.73	51.20	52.04	55.01	58.46	61.88	62.33	61.34
median	57.18	56.95	60.04	62.03	63.23	64.99	67.48	69.65	70.62	71.32	69.54
third quartile	69.31	69.54	71.70	71.74	74.76	74.47	76.19	77.60	78.90	78.78	77.60

Table 1 shows the annual cross-sectional average (mean), standard deviation (std), first quartile, median, and third quartile over our study period. We observe that both the average and median ESG Combined Scores increase gradually, reflecting steady improvements in ESG performance.

4.2. Portfolio construction

We construct annually rebalanced long/short ESG dynamic portfolios based on ESG Combined Scores, including the multi-period dynamic portfolios proposed in this paper, their single-period counterparts, and constrained mean-variance portfolios commonly used in the sustainable-investing literature (Lo et al., 2022; Pástor et al., 2021; Pedersen et al., 2021). The portfolio construction methods are as follows.

Multi-Period Portfolios. The dynamic long/short portfolios are constructed using the portfolio weights given by Theorem 2.1. We estimate the factor loadings (B), the long-run mean of factor returns ($\bar{f}$), the mean-reversion coefficient matrix of factor returns (Ψ), the covariance matrix of factor innovations (Σ_F), and the diagonal covariance matrix of residuals (Σ_e) based on the monthly return data of the previous three years.[11] Then, we compute $\mu = B\bar{f}$ as described in Section 2. We set $\Phi = 0.3I$ and $\boldsymbol{\Omega} = 0.1I$ based on estimates from ESG scores, and $\Lambda = 0.05\%I$, to account for brokerage commissions, bid-ask spreads, and regulatory fees. In addition, we set $r^f = 2\%$, $\rho = 0.1$, and $c = 0.015$.[12] Because the ESG Combined Score ranges from 0 to 100, for each year, we standardize them by subtracting the cross-sectional average and dividing by the cross-sectional standard deviation. We set g_k equal to these standardized ESG scores.

Single-Period Portfolios. To isolate the contribution of the multi-period dynamic mechanism, we also construct a single-period version of our portfolio. This portfolio is computed using the same model specification, parameter estimates, and ESG scores as the dynamic portfolio, but with the investment horizon set to $T = 1$. Thus, it optimizes only the current-period objective and does not take into account the future evolution of ESG scores, factor returns, or the effect of current trading on future transaction costs. The comparison between

[11] The long-run mean of factor returns ($\bar{f}$), the mean-reversion coefficient matrix of factor returns (Ψ), and the covariance matrix of factor innovations (Σ_F) are estimated from the autoregressive model of factor returns. The factor loadings (B) and the diagonal covariance matrix of residuals (Σ_e) are estimated via ordinary least squares regressions for each stock.

[12] As estimated by Lo et al. (2022), the correlation between ESG scores and stock returns takes values of around 1.5%.

the dynamic and single-period portfolios therefore highlights the incremental value of solving the globally dynamic optimization problem rather than repeatedly applying a one-period rule.

Constrained Portfolios. As an additional benchmark, we follow the constrained long/short optimization framework proposed by Lo et al. (2022). Unlike single-period portfolios, constrained portfolios incorporate ESG preferences through an explicit portfolio-level ESG constraint and impose additional practical constraints, including full investment and position limits. For each year, we consider the following constrained mean-variance optimization problem:

$$\max_{u_k} \quad u_k^\top \mu - \xi u_k^\top \Sigma u_k,$$

$$\text{subject to} \quad \begin{cases} \sum_{i=1}^{n} u_k^i = 1, \\ |u_k^i| \le \tau, \quad i = 1, 2, \ldots, n, \\ \sum_{i=1}^{n} u_k^i g_k^i \ge s, \end{cases}$$

where μ and Σ are estimated in the same way as for the multi-period portfolio; ξ is the risk-aversion tuning parameter, which, without loss of generality, is set to 10; τ is the maximum weight for each stock, set to 10% (Lo et al., 2022; Pysarenko et al., 2019); and s is the threshold that controls the minimum average level of the ESG score of the portfolio, set equal to that of the dynamic portfolio.

By comparing the multi-period portfolio with the single-period portfolio, we can isolate the value of the multi-period optimization mechanism. Since these two portfolios are based on the same objective function, parameter estimates, and ESG information, their performance difference reflects the benefit of taking into account the future evolution of ESG scores and factor returns, as well as the intertemporal impact of current trading decisions on future transaction costs. By further comparing these two portfolios with the constrained portfolio, we can evaluate how our utility-based ESG integration differs from a conventional constrained ESG investment rule. In particular, the constrained portfolio incorporates ESG preferences through an explicit minimum ESG-score constraint in a one-period mean-variance problem, whereas our model incorporates ESG scores directly into the investor's multi-period objective. Therefore, this comparison helps assess whether the proposed framework provides additional economic value relative to both its own single-period counterpart and a widely used constrained mean-variance benchmark.

All model parameters are estimated using a rolling window of the preceding three years of data, and the portfolio investigation period spans from 2003 to 2024. As a robustness check, we repeat the analysis using a five-year estimation window and an investigation period from 2005 to 2024.

Finally, we note that our empirical design explicitly ensures comparability between the multi-period and constrained portfolios in terms of ESG exposure. In particular, the ESG constraint in the static mean-variance optimization is chosen so that the minimal average portfolio ESG score s matches that of the corresponding multi-period portfolio. This design guarantees that the two portfolios are comparable along the ESG dimension, which is observable at the portfolio rebalancing date.

Table 2. Performance measures of the multi-period portfolio, the single-period portfolio, and the constrained portfolio under different parameter setups.

setup	$d = 0.1, \gamma = 1$			$d = 0.1, \gamma = 10$		
portfolio	multi-period	single-period	constrained	multi-period	single-period	constrained
return	9.13%	9.23%	6.01%	10.73%	10.44%	9.17%
std	20.90%	26.39%	19.34%	13.86%	14.04%	13.93%
SR	0.3652	0.2929	0.2333	0.6652	0.6361	0.5500
MDD	48.73%	62.70%	53.90%	26.87%	24.78%	28.58%
α	0.18%	0.40%	−3.12%	0.29%	0.19%	−0.50%
TE	1.24%	1.61%	1.08%	0.61%	0.65%	0.69%
IR	0.0097	0.0160	−0.1880	0.0310	0.0191	−0.0471
ESG	7.6250	8.2525	7.6250	4.9883	5.2038	4.9883
setup	$d = 0.05, \gamma = 1$			$d = 0, \gamma = 1$		
return	9.64%	9.26%	7.56%	11.16%	11.00%	9.49%
std	16.29%	18.15%	16.40%	13.64%	13.61%	13.43%
SR	0.4993	0.4274	0.3695	0.7069	0.6967	0.5945
MDD	25.77%	36.77%	36.26%	31.55%	30.84%	30.73%
α	0.15%	0.02%	−1.73%	0.31%	0.24%	−0.41%
TE	0.89%	1.04%	0.89%	0.54%	0.54%	0.63%
IR	0.0110	0.0013	−0.1259	0.0375	0.0288	−0.0426
ESG	6.4947	7.0274	6.4947	4.4051	4.4311	4.4051

In contrast, it is generally infeasible to impose an ex ante constraint that equalizes realized portfolio risk across models, because future return volatility is not directly observable at the time of portfolio construction and may evolve endogenously with portfolio dynamics. As a result, rather than enforcing equality in realized risk, we evaluate portfolio performance using risk-adjusted measures such as the Sharpe ratio, factor-adjusted alpha, and information ratio. These metrics provide a natural and economically meaningful basis for comparing portfolios with potentially different realized risk profiles.

4.3. Portfolio performance

Table 2 reports the performance of four multi-period portfolios and four single-period portfolios under different parameter setups, as well as the corresponding constrained portfolios with the minimum ESG level s the same as that of the multi-period portfolio. In particular, we report their annualized return (return), standard deviation (std), Sharpe ratio (SR), maximum drawdown (MDD), alpha (α) and tracking error (TE) with the Fama-French five-factor model as the benchmark, information ratio (IR), and the average standardized ESG score (ESG).[13]

For all portfolios, realized portfolio wealth and performance measures are computed after deducting the same quadratic transaction costs specified in Equation (4), using $\Lambda = 0.05\%I$. Thus, although the constrained portfolio and single-period portfolio benchmarks do not internalize transaction costs in their one-period optimization problems, their reported performance is evaluated net of the same transaction costs as the multi-period.

We have the following observations from the results of the four parameter setups in Table 2. First, the multi-period portfolios generally outperform the corresponding single-period and constrained portfolios in terms of Sharpe ratios, alphas, and information ratios, highlighting the value of the dynamic ESG framework. For example, when $(d, \gamma) = (0.1, 1)$,

[13] The portfolio alpha is the intercept from the Fama-French five-factor regression, and the tracking error is the standard deviation of the regression residuals. The information ratio is defined as the ratio of alpha to the tracking error.

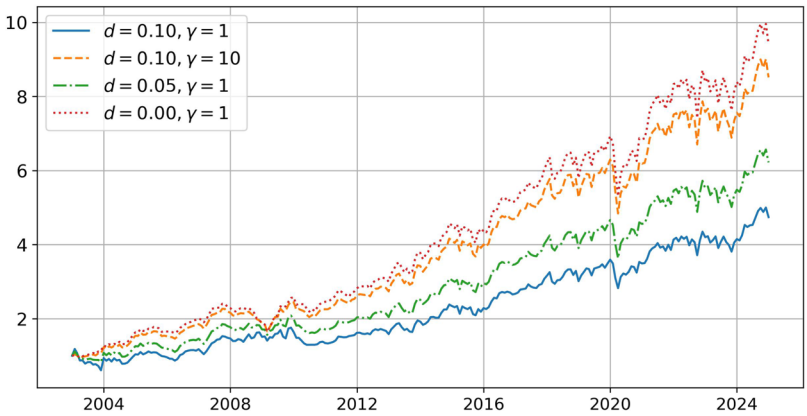
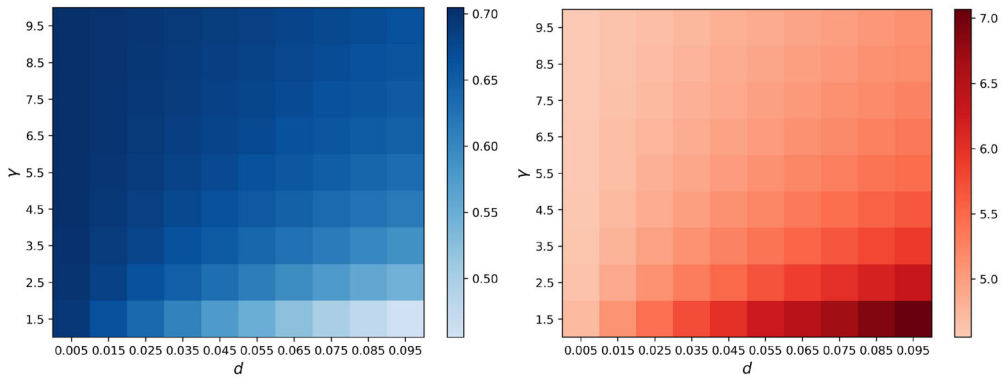


Figure 5. Cumulative wealth of the optimal portfolios.

the multi-period portfolio earns a Sharpe ratio of 0.3652, compared with 0.2929 for the single-period portfolio and 0.2333 for the constrained portfolio. Second, the comparison between the multi-period and single-period portfolios shows the benefit of the multi-period optimization mechanism, since both portfolios are based on the same model but differ in the investment horizon. Third, comparing the first two setups, we find that a higher pecuniary preference ($\gamma = 10$) leads to better financial performance but a lower ESG score than when $\gamma = 1$. In particular, when $\gamma = 10$ ($\gamma = 1$), the multi-period portfolio earns a Sharpe ratio of 0.6652 (0.3652) and an average ESG score of 4.9883 (7.6250). Finally, examining the first ($d = 0.1$), third ($d = 0.05$), and fourth ($d = 0$) setups, we observe that a higher ESG taste results in higher ESG scores but weaker financial performance. For example, the multi-period portfolio with $d = 0$ ($d = 0.1$) earns a Sharpe ratio of 0.7069 (0.3652) but an average ESG score of 4.4051 (7.6250). Figure 5 shows the cumulative wealth of the four multi-period portfolios, with initial wealth normalized to 1. Consistent with the performance metrics in Table 2, portfolios with higher pecuniary preference (γ) or lower ESG taste (d) deliver stronger financial performance.

Figure 6 illustrates how the Sharpe ratio and the average ESG score of the multi-period portfolio change with the model parameters. The horizontal and vertical axes represent the ESG taste (d) and the pecuniary preference (γ), respectively, while the color scale corresponds to the Sharpe ratio in (Figure 6(a)) and the average ESG score in (Figure 6(b)) of the multi-period portfolio. We observe that as γ increases and d decreases, the Sharpe ratio increases while the average ESG score decreases. These results are consistent with those in Table 2 and Figure 5, highlighting the important trade-off between financial performance and ESG performance. As a robustness check, we repeat the portfolio construction procedure using a longer estimation window: instead of estimating the model parameters from the previous three years of monthly data, we use the previous five years. Table 3 reports the corresponding performance results. The main conclusions remain unchanged: multi-period portfolios continue to outperform the corresponding single-period and constrained portfolios in terms of risk-adjusted performance, and the comparative statics with respect to γ and d are qualitatively consistent with those reported in Table 2.



(a) Sharpe ratio. (b) Average ESG score.

Figure 6. The Sharpe ratio and average ESG score of the multi-period portfolio under different parameter setups. (a) Sharpe ratio. (b) Average ESG score.

Table 3. Performance measures of the multi-period portfolio, the single-period portfolio, and the constrained portfolio under different parameter setups using five-year data to estimate model parameters.

setup	$d = 0.1, \gamma = 1$			$d = 0.1, \gamma = 10$		
portfolio	multi-period	single-period	constrained	multi-period	single-period	constrained
return	9.32%	9.24%	9.29%	9.70%	9.56%	9.67%
std	14.88%	15.51%	18.00%	14.21%	14.20%	16.88%
SR	0.5270	0.5011	0.4339	0.5780	0.5691	0.4851
MDD	24.51%	24.35%	36.66%	33.64%	31.12%	41.15%
α	0.69%	0.97%	−2.11%	0.03%	0.10%	−1.61%
TE	0.79%	0.88%	0.86%	0.56%	0.60%	0.66%
IR	0.0619	0.0773	−0.1727	0.0033	0.0120	−0.1715
ESG	5.3527	5.7412	5.3527	4.1648	4.3855	4.1648
setup	$d = 0.05, \gamma = 1$			$d = 0, \gamma = 1$		
return	9.45%	9.33%	9.44%	9.83%	9.73%	9.79%
std	14.41%	14.86%	17.45%	14.32%	14.26%	16.73%
SR	0.5527	0.5286	0.4564	0.5825	0.5781	0.4968
MDD	26.86%	24.47%	38.41%	36.59%	36.12%	42.33%
α	0.41%	0.66%	−1.90%	−0.12%	−0.17%	−1.46%
TE	0.69%	0.78%	0.78%	0.51%	0.52%	0.61%
IR	0.0413	0.0593	−0.1725	−0.0169	−0.0235	−0.1675
ESG	4.8821	5.2913	4.8821	3.8493	3.8794	3.8493

5. Conclusion

In this article, we develop a stochastic control framework for constructing optimal multi-period dynamic ESG portfolios that explicitly account for both financial returns and sustainability objectives. Our theoretical solution shows that in each period, the optimal weight can be decomposed into three components: a financial component reflecting expected returns, an ESG component reflecting current and future sustainability considerations, and a last-period-weight component reflecting past allocations to reduce transaction costs. This decomposition highlights the trade-off faced by sustainable investors between profitability, greenness, and cost efficiency.

The numerical results provide further insights into the structure of the optimal portfolios. We find that stronger pecuniary preferences increase the role of the financial and last-period-weight components, while stronger ESG preferences tilt the portfolio more toward greener

assets. Higher transaction costs increase the weight placed on the last-period-weight component, thereby reducing turnover. These results clarify how investors' preferences and market frictions shape the dynamic balance between financial and non-financial objectives.

Our empirical analysis using U.S. stock market data confirms the advantages of the dynamic approach. Portfolios constructed under our framework deliver higher Sharpe ratios and larger alphas than both their single-period counterparts and comparable ESG-constrained portfolios. The empirical evidence demonstrates that explicitly incorporating the dynamics of ESG scores and transaction costs allows investors to improve both financial performance and sustainability outcomes. We also show that the trade-off between financial performance and ESG performance can be systematically adjusted by tuning the preference parameters, offering a flexible tool for investors with heterogeneous objectives.

Beyond the quantitative results, our findings have several broader implications. First, dynamic management of ESG portfolios is essential because repeated static optimization may ignore important intertemporal links between ESG scores, financial returns, and trading costs. Moreover, our framework shows that the trade-off between financial performance and sustainability can be made explicit and adjustable, which provides a transparent tool for investors with heterogeneous objectives. In addition, the decomposition of the optimal strategy into financial, ESG, and persistence components offers a clear structure for reporting and governance, making it easier to communicate how portfolio choices align with both financial goals and sustainability commitments.

Overall, our results contribute to the literature on both sustainable investing and dynamic portfolio theory, and provide practical guidance for investors who seek to balance long-term financial goals with sustainability considerations. Our framework also suggests several directions for further exploration. An important extension would be to incorporate time-varying ESG preferences, reflecting shifts in ESG attitudes or regulatory requirements, and to analyze how such changes affect long-term allocation. It would also be useful to study how our framework interacts with different ESG rating methodologies, given the increasing diversity and inconsistency across data providers (Avramov et al., 2022; Christensen et al., 2022). These extensions would not only deepen the theoretical foundations of dynamic ESG investing, but also enhance its practical relevance in evolving markets.

Author contributions

CRediT: **Danping Li:** Conceptualization, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Supervision, Validation, Writing – original draft, Writing – review & editing; **Ruixun Zhang:** Conceptualization, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Supervision, Validation, Writing – original draft, Writing – review & editing; **Chaoyi Zhao:** Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by the National Key R & D Program of China [grant number 2022YFA1007900]; National Natural Science Foundation of China [grant numbers 12271013, 12271171, 12671572, 72342004]; Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China [grant number JYB2025XDXM904]; Fundamental Research Funds

for the Central Universities [grant numbers 2024QKT008, 2026QKT006, 2025ECNU-WLJC006]; Shanghai Sci-tech Co-research Program [grant number 25HB2702200]; Shanghai Philosophy Social Science Planning Office Project [2022ZJB005].

ORCID

Danping Li http://orcid.org/0000-0003-3882-0971
Ruixun Zhang http://orcid.org/0000-0002-7670-8393
Chaoyi Zhao http://orcid.org/0000-0002-5921-3612

References

Albuquerque, R., Koskinen, Y., & Zhang, C. (2019). Corporate social responsibility and firm risk: Theory and empirical evidence. *Management Science*, *65*(10), 4451–4469.
Alves, R., Krüger, P., & van Dijk, M. (2025). Drawing up the bill: Are ESG ratings related to stock returns around the world? *Journal of Corporate Finance*, *93*, 102768.
Avramov, D., Cheng, S., Lioui, A., & Tarelli, A. (2022). Sustainable investing with ESG rating uncertainty. *Journal of Financial Economics*, *145*(2), 642–664.
Bansal, R., Wu, D., & Yaron, A. (2022). Socially responsible investing in good and bad times. *The Review of Financial Studies*, *35*(4), 2067–2099.
Berg, F., Lo, A. W., Rigobon, R., Singh, M., & Zhang, R. (2024). Quantifying the returns of ESG investing: An empirical analysis with six ESG metrics. *The Journal of Portfolio Management*, *50*(8), 216–238.
Bertsekas, D. (2012). *Dynamic Programming and Optimal ontrol* (Vol. 1). Athena Scientific.
Blitz, D., & Fabozzi, F. J. (2017). Sin stocks revisited: Resolving the sin stock anomaly. *The Journal of Portfolio Management*, *44*(1), 105–111.
Bolton, P., Halem, Z., & Kacperczyk, M. (2022). The financial cost of carbon. *Journal of Applied Corporate Finance*, *34*(2), 17–29.
Bolton, P., & Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics*, *142*(2), 517–549.
Bolton, P., & Kacperczyk, M. (2023). Global pricing of carbon-transition risk. *The Journal of Finance*, *78*(6), 3677–3754.
Bolton, P., & Kacperczyk, M. (2026). Firm commitments. *Management Science*, *72*(3), 2134–2167.
Breen, W. J., Hodrick, L. S., & Korajczyk, R. A. (2002). Predicting equity liquidity. *Management Science*, *48*(4), 470–483.
Brown, D. B., & Smith, J. E. (2011). Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. *Management Science*, *57*(10), 1752–1770.
Campbell, J. Y., Chan, Y. L., & Viceira, L. M. (2003). A multivariate model of strategic asset allocation. *Journal of Financial Economics*, *67*(1), 41–80.
Chan, Y., Hogan, K., Schwaiger, K., & Ang, A. (2020). ESG in factors. *The Journal of Impact and ESG Investing*, *1*(1), 26–45.
Chen, A., Gerick, L., & Jin, Z. (2025). Optimizing portfolios under carbon risk constraints: Setting effective constraints to favor green investments. *Energy Economics*, *148*, 108634.
Chen, H. C., & Lin, Y. N. (2026). ESG migration: Evidence from the Taiwan stock market. *Review of Quantitative Finance and Accounting*, *67*(2), 617–648.
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*(1), 147–175.
Dybvig, P. H., & Pezzo, L. (2020). Mean-variance portfolio rebalancing with transaction costs. Available at SSRN 3373329.
Ehsani, S., & Linnainmaa, J. T. (2022). Factor momentum and the momentum factor. *The Journal of Finance*, *77*(3), 1877–1919.
Escobar-Saldívar, L. J., Villarreal-Samaniego, D., & Santillán-Salgado, R. J. (2025). The effects of ESG scores and ESG momentum on stock returns and volatility: Evidence from US markets. *Journal of Risk and Financial Management*, *18*(7), 1–21.

Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, *33*(1), 3–56.
Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, *116*(1), 1–22.
Flora, M., & Tankov, P. (2023). Green investment and asset stranding under transition scenario uncertainty. *Energy Economics*, *124*, 106773.
Galema, R., Plantinga, A., & Scholtens, B. (2008). The stocks at stake: Return and risk in socially responsible investment. *Journal of Banking and Finance*, *32*(12), 2646–2654.
Gârleanu, N., & Pedersen, L. H. (2013). Dynamic trading with predictable returns and transaction costs. *The Journal of Finance*, *68*(6), 2309–2340.
Glasserman, P., & Xu, X. (2013). Robust portfolio control with stochastic factor dynamics. *Operations Research*, *61*(4), 874–893.
Goldstein, I., Kopytov, A., Shen, L., & Xiang, H. (2025). On ESG investing: Heterogeneous preferences, information, and asset prices. Available at SSRN 3823042.
Idzorek, T. M., Kaplan, P. D., & Ibbotson, R. G. (2021). The popularity asset pricing model. Available at SSRN 3451554.
Jiao, Y., & Liu, H. (2025). Optimal portfolio choice with ESG considerations and asymmetric information. *Quantitative Finance*, *25*(7), 1163–1176.
Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, *53*(6), 1315–1335.
Li, D., & Ng, W. L. (2000). Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Mathematical Finance*, *10*(3), 387–406.
Li, X., Uysal, A. S., & Mulvey, J. M. (2022). Multi-period portfolio optimization using model predictive control with mean-variance and risk parity frameworks. *European Journal of Operational Research*, *299*(3), 1158–1176.
Lillo, F., Farmer, J. D., & Mantegna, R. N. (2003). Master curve for price-impact function. *Nature*, *421*(6919), 129–130.
Lindsey, L. A., Pruitt, S., & Schiller, C. (2025). ESG and the conditional pricing of risk. Available at SSRN 3975077.
Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*, *47*(1), 13–37.
Liu, X., Yang, C., & Chao, Y. (2022). The pricing of ESG: Evidence from overnight return and intraday return. *Frontiers in Environmental Science*, *10*, 927420.
Lo, A. W., Wu, L., Zhang, R., & Zhao, C. (2024). Optimal impact portfolios with general dependence and marginals. *Operations Research*, *72*(5), 1775–1789.
Lo, A. W., & Zhang, R. (2024). Quantifying the impact of impact investing. *Management Science*, *70*(10), 7161–7186.
Lo, A. W., & Zhang, R. (2025). Performance attribution for portfolio constraints. *Management Science*, *71*(9), 7537–7559.
Lo, A. W., Zhang, R., & Zhao, C. (2022). Measuring and optimizing the risk and reward of green portfolios. *The Journal of Impact and ESG Investing*, *3*(2), 55–93.
Madhavan, A., Sobczyk, A., & Ang, A. (2021). Toward ESG alpha: Analyzing ESG exposures through a factor lens. *Financial Analysts Journal*, *77*(1), 69–88.
Pástor, L., Stambaugh, R. F., & Taylor, L. A. (2021). Sustainable investing in equilibrium. *Journal of Financial Economics*, *142*(2), 550–571.
Pástor, L., Stambaugh, R. F., & Taylor, L. A. (2022). Dissecting green returns. *Journal of Financial Economics*, *146*(2), 403–424.
Pedersen, L. H., Fitzgibbons, S., & Pomorski, L. (2021). Responsible investing: The ESG-efficient frontier. *Journal of Financial Economics*, *142*(2), 572–597.
Petkova, R. (2006). Do the Fama–French factors proxy for innovations in predictive variables? *The Journal of Finance*, *61*(2), 581–612.
Pysarenko, S., Alexeev, V., & Tapon, F. (2019). Predictive blends: Fundamental indexing meets Markowitz. *Journal of Banking and Finance*, *100*, 28–42.
Renneboog, L., Ter Horst, J., & Zhang, C. (2008). Socially responsible investments: Institutional aspects, performance, and investor behavior. *Journal of Banking and Finance*, *32*(9), 1723–1742.

Ross, S. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, *13*(3), 341–360.
Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics*, *51*(3), 239–246.
Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, *19*(3), 425–442.
Sun, W., Fan, A., Chen, L. W., Schouwenaars, T., & Albota, M. A. (2006). Optimal rebalancing for institutional portfolios. *The Journal of Portfolio Management*, *32*(2), 33–43.
Tankov, P., & Zhang, R. (2025). *Handbook of Quantitative Sustainable Finance*. Chapman & Hall.
Wedajo, A. D., Salah, A. A., Bhat, M. A., Iqbal, R., & Khan, S. T. (2024). Analyzing the dynamic relationship between ESG scores and firm value in Chinese listed companies: Insights from generalized cross-lagged panel model. *Discover Sustainability*, *5*(1), 1–26.
Yang, R., Jimenez-Martin, J. A., & Caporin, M. (2025). The non-linear ESG premium. *Quantitative Finance*, *25*(5), 817–840.
Zerbib, O. D. (2022). A sustainable capital asset pricing model (S-CAPM): Evidence from environmental integration and sin stock exclusion. *Review of Finance*, *26*(6), 1345–1388.

Appendix

In this appendix, we provide the proof of Theorem 2.1 in our main paper.

Proof of Theorem 2.1: The optimization problem in Equation (5) can be solved using dynamic programming. Specifically, Bertsekas (2012, Exercise 1.7) shows that the optimal solution u_k can be obtained using a dynamic programming algorithm. The value function is denoted by $V_k(x_k, u_{k-1}, x_{k-1}, g_k, \widetilde{f}_k)$. It satisfies the boundary condition

$$V_T(x_T, u_{T-1}, x_{T-1}, g_T, \widetilde{f}_T) = -\gamma\, x_T \tag{A1}$$

and the recursive formula

$$\begin{aligned} &V_k(x_k, u_{k-1}, x_{k-1}, g_k, \widetilde{f}_k) \\ &\quad = \min_{u_k} \left\{ -d(1-\rho)^k x_k g_k^\top u_k + \ln \mathbb{E}_k \left[\exp\left(V_{k+1}(x_{k+1}, u_k, x_k, g_{k+1}, \widetilde{f}_{k+1}) \right) \right] \right\}, \end{aligned} \tag{A2}$$

with $\mathbb{E}_k(\cdot)$ denoting the conditional expectation given all information available up to and including time k.

We claim that the value function is given by

$$\begin{aligned} V_k(x_k, u_{k-1}, x_{k-1}, g_k, \widetilde{f}_k) = &-\gamma\, x_k (1 + r^f)^{T-k} + g_k^\top A_k g_k + \widetilde{f}_k^\top S_k \widetilde{f}_k + 2 g_k^\top E_k \widetilde{f}_k \\ &+ x_{k-1}^2 u_{k-1}^\top B_k u_{k-1} + 2 x_{k-1} u_{k-1}^\top C_k g_k + 2 x_{k-1} u_{k-1}^\top K_k \widetilde{f}_k \\ &+ 2 \mu^\top D_k^\top g_k + 2 \mu^\top L_k \widetilde{f}_k + 2 x_{k-1} \mu^\top F_k^\top u_{k-1} + h_k. \end{aligned} \tag{A3}$$

We prove that Equation (A3) satisfies both Equations (A1) and (A2) using backward induction. At time $k = T$, Equation (A3) reduces to Equation (A1). Now suppose that Equation (A3) holds at time $k+1$. We show that it is the solution to Equation (A2) at time k.

Let

$$z_k = x_k u_k, \quad y_k = x_{k-1} u_{k-1}, \quad a_k = \gamma\, (1 + r^f)^{T-k-1},$$

and write

$$\bar{g}_{k+1} = (I - \Phi) g_k, \quad \bar{f}_{k+1} = (I - \Psi) \widetilde{f}_k.$$

By Equations (2), (4), and (A3), the deterministic part of the exponent in $\mathbb{E}_k[\exp(V_{k+1})]$ is

$$\begin{aligned} &-\gamma\, x_k (1 + r^f)^{T-k} - a_k z_k^\top \left(\mu + B \bar{f}_{k+1} + c g_k \right) + \frac{1}{2} \gamma\, (1 + r^f)^{T-k} (z_k - y_k)^\top \Lambda (z_k - y_k) \\ &+ \bar{g}_{k+1}^\top A_{k+1} \bar{g}_{k+1} + \bar{f}_{k+1}^\top S_{k+1} \bar{f}_{k+1} + 2 \bar{g}_{k+1}^\top E_{k+1} \bar{f}_{k+1} + z_k^\top B_{k+1} z_k \\ &+ 2 z_k^\top C_{k+1} \bar{g}_{k+1} + 2 z_k^\top K_{k+1} \bar{f}_{k+1} + 2 \mu^\top D_{k+1}^\top \bar{g}_{k+1} + 2 \mu^\top L_{k+1} \bar{f}_{k+1} + 2 \mu^\top F_{k+1}^\top z_k + h_{k+1}. \end{aligned}$$

The residual return shock e_{k+1} gives

$$\mathbb{E}_k\left[\exp\left(-a_k e_{k+1}^\top z_k\right)\right] = \exp\left(\frac{1}{2}a_k^2 z_k^\top \Sigma_e z_k\right). \tag{A4}$$

It remains to evaluate the expectation with respect to η_{k+1} and ε_{k+1}. Using the notation in the theorem with the index $k+1$, define

$$\Delta_{k+1}^f = U_{z,k+1}z_k + U_{g,k+1}g_k + U_{f,k+1}\widetilde{f}_k + U_{\mu,k+1}\mu.$$

The terms involving η_{k+1} are

$$\eta_{k+1}^\top S_{k+1}\eta_{k+1} + 2\left(\Delta_{k+1}^f + E_{k+1}^\top \varepsilon_{k+1}\right)^\top \eta_{k+1}.$$

By the moment generating function of the quadratic form of a normally distributed random variable,

$$\begin{aligned}
&\mathbb{E}_k\left[\exp\left(\eta_{k+1}^\top S_{k+1}\eta_{k+1} + 2\left(\Delta_{k+1}^f + E_{k+1}^\top\varepsilon_{k+1}\right)^\top\eta_{k+1}\right) \mid \varepsilon_{k+1}\right] \\
&\quad = \left|I - 2\Sigma_F^{1/2}S_{k+1}\Sigma_F^{1/2}\right|^{-\frac{1}{2}} \exp\left(2\left(\Delta_{k+1}^f + E_{k+1}^\top\varepsilon_{k+1}\right)^\top N_{k+1}^f\left(\Delta_{k+1}^f + E_{k+1}^\top\varepsilon_{k+1}\right)\right).
\end{aligned} \tag{A5}$$

After this step, the remaining exponent is quadratic in ε_{k+1}. It can be written as

$$\varepsilon_{k+1}^\top \widetilde{A}_{k+1}\varepsilon_{k+1} + 2(\Delta_{k+1}^g)^\top\varepsilon_{k+1},$$

where

$$\Delta_{k+1}^g = W_{z,k+1}z_k + W_{g,k+1}g_k + W_{f,k+1}\widetilde{f}_k + W_{\mu,k+1}\mu.$$

Applying the same Gaussian formula to ε_{k+1} gives

$$\begin{aligned}
&\mathbb{E}_k\left[\exp\left(\varepsilon_{k+1}^\top\widetilde{A}_{k+1}\varepsilon_{k+1} + 2(\Delta_{k+1}^g)^\top\varepsilon_{k+1}\right)\right] \\
&\quad = \left|I - 2\Omega^{1/2}\widetilde{A}_{k+1}\Omega^{1/2}\right|^{-\frac{1}{2}} \exp\left(2(\Delta_{k+1}^g)^\top N_{k+1}^g\Delta_{k+1}^g\right).
\end{aligned} \tag{A6}$$

Combining Equations (A4)–(A6) with the deterministic terms above, adding the current-period ESG term in Equation (A2), and collecting the coefficients of $z_k = x_k u_k$, we obtain

$$\begin{aligned}
&V_k(x_k, u_{k-1}, x_{k-1}, g_k, \widetilde{f}_k) \\
&= \min_{u_k}\left\{\frac{1}{2}\left[z_k - Q_k^{-1}\left(P_k g_k + H_k\mu + J_k\widetilde{f}_k + y_k\gamma\left(1+r^f\right)^{T-k}\Lambda\right)\right]^\top Q_k\right. \\
&\quad\times\left[z_k - Q_k^{-1}\left(P_k g_k + H_k\mu + J_k\widetilde{f}_k + y_k\gamma\left(1+r^f\right)^{T-k}\Lambda\right)\right] \\
&\quad - \gamma x_k\left(1+r^f\right)^{T-k} + g_k^\top A_k g_k + \widetilde{f}_k^\top S_k\widetilde{f}_k + 2g_k^\top E_k\widetilde{f}_k \\
&\quad + x_{k-1}^2 u_{k-1}^\top B_k u_{k-1} + 2x_{k-1}u_{k-1}^\top C_k g_k + 2x_{k-1}u_{k-1}^\top K_k\widetilde{f}_k \\
&\quad \left.+2\mu^\top D_k^\top g_k + 2\mu^\top L_k\widetilde{f}_k + 2x_{k-1}\mu^\top F_k^\top u_{k-1} + h_k\right\},
\end{aligned}$$

where the matrices $Q_k, P_k, H_k, J_k, A_k, S_k, E_k, B_k, C_k, K_k, D_k, L_k, F_k$ and h_k are exactly those given by the backward recursions in the theorem. Since Q_k is positive definite, the minimum is attained at

$$z_k = Q_k^{-1}\left(P_k g_k + H_k\mu + J_k\widetilde{f}_k + x_{k-1}\gamma\left(1+r^f\right)^{T-k}\Lambda u_{k-1}\right).$$

Because $z_k = x_k u_k$, the optimal solution is Equation (6), and the value function is given by Equation (A3). This proves the result. ■