Sustainable Asset Pricing with Heterogeneous Agents: A Computational Approach

By: Ryan Huang

Honors Thesis
Economics Department
The University of North Carolina at Chapel Hill

Thesis Advisor: Dr. Marlon Azinovic-Yang
Faculty Advisor: Dr. Klara Peter

April 2026

Sustainable Asset Pricing with Heterogeneous Agents: A Computational Approach

Ryan Huang*
Department of Economics
University of North Carolina at Chapel Hill
rhuang1@unc.edu

This Version: April 7, 2026

Abstract

Environmentally responsible assets trade at a persistent premium over identical conventional securities. How strong must investor preferences for sustainability be to explain this? I address this with a two-agent model where an ESG-conscious and a neutral investor trade green equity, brown equity, and a risk-free bond across twelve exogenous shocks. Neural networks approximate equilibrium prices and portfolios by satisfying optimality conditions across the state space. Calibrated to German twin bond data, the model matches the observed green premium (-2.53 vs. data -2.44 basis points) and the equity premium (4.86% vs. data 4.81%) with a small ESG preference intensity ($\phi^* = 1.0 \times 10^{-4}$). Green equities earn lower risk-adjusted returns than brown (Sharpe ratios 0.200 vs. 0.207), consistent with the green premium reflecting investor preferences rather than risk compensation, though this partly reflects the model's symmetric dividend structure. The cross-sectional wealth distribution—not the ESG regime—drives green premium variation.

Keywords: ESG Investing, Heterogeneous Agents, Deep Equilibrium Networks, Green Premium, Sustainable Finance, Asset Pricing

*I am grateful to Dr. Marlon Azinovic-Yang and Dr. Klara Peter for their guidance and feedback. Computational resources were provided by the Longleaf high-performance computing cluster at the University of North Carolina at Chapel Hill.

1 Introduction

In 2023, global sustainable investment reached $30.3 trillion in assets under management, representing nearly one-third of professionally managed capital worldwide (Global Sustainable Investment Alliance, 2023). This shift has created a puzzle for financial economists: ESG-conscious investors systematically accept lower returns on environmentally responsible assets compared to otherwise identical securities (Zerbib, 2019; Pástor et al., 2021). German government twin bonds—identical securities except for their green designation—trade with persistent yield differentials of 2–3 basis points, providing some of the cleanest evidence that investors derive non-pecuniary utility from sustainability (Pástor et al., 2022). A central question in asset pricing is whether this greenium reflects compensation for undiversifiable risk—as in standard consumption-based models—or a pure preference premium arising from non-pecuniary utility. How strong must these preferences actually be to explain observed pricing patterns? And what role does market structure—particularly portfolio constraints that limit arbitrage—play in translating preferences into prices?

In this thesis, I show that even minimal non-pecuniary ESG preferences, when embedded in general equilibrium with heterogeneous agents and portfolio constraints, suffice to generate empirically realistic pricing differentials—and that the resulting greenium is consistent with a preference premium rather than compensation for differential risk. The model is calibrated to German data, jointly matching the twin bond greenium and the equity premium, though it does not match the risk-free rate—a known limitation of CRRA preferences. Portfolio constraints provide a natural framework for future analysis of how short-sale restrictions may amplify the equilibrium impact of non-pecuniary preferences.

Two hypotheses formalize this argument. First, even small non-pecuniary ESG preferences generate a negative greenium (green assets earn lower expected returns than brown) when embedded in general equilibrium with heterogeneous agents. Second, the greenium varies with the wealth distribution: it depends on the household income state, because the relative wealth of ESG-conscious and neutral investors determines how much each agent's preferences matter for market-clearing prices. A natural conjecture—left for future work—is that portfolio constraints amplify the greenium by preventing neutral investors from arbitraging ESG-driven pricing differentials.

The theoretical literature on ESG asset pricing has developed along two branches. The first employs representative agent models where all investors share identical preferences, deriving closed-form solutions showing that ESG concerns create predictable pricing wedges (Pástor et al., 2021; Pedersen et al., 2021). These models formalize the preference channel but cannot capture dynamic portfolio rebalancing or wealth evolution. The second incorporates heterogeneous preferences. Goldstein et al. (2022) examine how heterogeneous ESG preferences interact with asymmetric information, showing that disagreement about ESG materiality amplifies pricing effects. Sauzet and Zerbib (2022) develops a two-agent framework with recursive preferences where investors differ in their portfolio and consumption preferences for green assets, demonstrating that preference heterogeneity generates endogenous risk premia that vary with the wealth distribution. However, projection methods become increasingly costly as the number of assets and state variables grows.

Recent empirical work has documented substantial heterogeneity in ESG preferences across investor types. Institutional investors subject to ESG mandates exhibit strong preferences for sustainability, while hedge funds often exploit ESG-driven mispricings (Starks, 2023; Gibson Brandon et al., 2022). Survey evidence reveals that younger investors willingly sacrifice 2–3% annual returns for strong ESG characteristics (Riedl and Smeets, 2017; Bauer et al., 2021). This heterogeneity generates rich trading dynamics: during the 2020–2022 energy crisis, wealth transferred from ESG funds (underweighting fossil fuels) to traditional investors, potentially altering the market's aggregate ESG preference (Pástor et al., 2022). Yet existing models cannot capture these dynamics because solving heterogeneous agent models with realistic market frictions has been computationally prohibitive.

This thesis applies deep equilibrium networks (DEQNs) to a model with richer heterogeneity than previously solved in this literature. I build a discrete-time pure-exchange economy with two heterogeneous agents trading three assets—green equity, brown equity, and a risk-free bond—across twelve exogenous shocks that capture ESG regime shifts, business cycle fluctuations, and household income heterogeneity. The model is calibrated to German macroeconomic data with the twin bond greenium and the equity premium as joint calibration targets. DEQNs parameterize equilibrium price and policy functions as neural networks, then train these networks to satisfy Euler equations across the entire state space simultaneously (Azinovic et al., 2022). A delta curriculum gradually introduces self-referential price feedback, enabling stable convergence to the full general equilibrium. This computational approach enables analysis of economic mechanisms that are difficult to study with traditional methods: wealth accumulation feedback, state-dependent constraint binding, and the endogenous co-evolution of the green premium with the wealth distribution.

The main findings support the core hypotheses. First, matching the empirical greenium requires only $\phi^* = 1.0 \times 10^{-4}$—an ESG preference so small that the implied portfolio tilt is approximately 4 percentage points, at the lower end of survey-based estimates (Riedl and Smeets, 2017; Bauer et al., 2021). Second, the greenium varies substantially across shocks, with the cross-sectional wealth distribution as the primary driver—a channel absent from representative agent models—while the ESG regime has negligible effect. Third, green and brown equities exhibit different Sharpe ratios (0.200 vs. 0.207), consistent with the return differential reflecting a preference premium rather than differential risk.

The remainder of this thesis is organized as follows. Section 2 reviews the literature. Section 3 presents the model. Section 4 describes the data and calibration. Section 5 presents results. Section 6 concludes.

2 Literature Review

Three economic mechanisms shape how ESG preferences translate into asset prices. First, *investor preferences* generate utility from holding environmentally responsible assets beyond their financial payoffs, creating demand-driven pricing effects. Second, *imperfect substitutability* between green

and brown assets—stemming from differentiated cash flow risks and correlations—limits investors' ability to replicate one asset's payoffs using the other. Third, *arbitrage constraints*—including market incompleteness, short-sale restrictions, and leverage limits—prevent profit-maximizing traders from eliminating preference-driven pricing differentials. These channels interact to determine when and how ESG preferences translate into persistent equilibrium price effects.

2.1 Theoretical Foundations in Sustainable Asset Pricing

The theoretical literature on ESG asset pricing has established crucial insights while revealing important limitations. Pástor et al. (2021) develop a general equilibrium model where investors have heterogeneous tastes for green assets, deriving a closed-form solution for the green premium in an equilibrium setting. Their model predicts that brown assets must offer higher expected returns to compensate ESG-conscious investors for holding "sinful" securities. Similarly, Pedersen et al. (2021) extends this framework to include ESG uncertainty, showing how information about firms' sustainability affects both risk and expected returns. These models formalize the *preference channel* but rely on static mean-variance optimization that cannot capture dynamic portfolio rebalancing or wealth evolution. Critically, by assuming perfect substitutability between green and brown assets, these models abstract from the *imperfect substitutability channel* that empirically characterizes renewable versus fossil fuel firms.

Zerbib (2022) advances this literature by incorporating investor mandates and exclusion policies, demonstrating how institutional constraints amplify ESG pricing effects. Sauzet and Zerbib (2022) makes significant progress toward a fully dynamic heterogeneous agent framework by developing a continuous-time model with recursive preferences where investors differ in both their portfolio preferences for green assets and their consumption preferences for green goods. In his two-agent economy with green and brown equity trees, equilibrium prices depend on two endogenous state variables: the wealth share of the green investor and the relative supply of green goods. Using projection methods with Chebyshev polynomials, he obtains a global solution that traces how wealth redistribution between investor types affects risk premia through time-varying hedging demands. This thesis follows Sauzet and Zerbib (2022)'s heterogeneous agent structure but works in discrete time with CRRA preferences and focuses exclusively on portfolio preferences for green versus brown assets, abstracting from consumption channel effects to isolate the investment mechanisms. Critically, I extend the asset space to include a risk-free bond in zero net supply, add a richer 12-shock exogenous process, and solve using deep equilibrium networks rather than projection methods, enabling analysis of how broader investor heterogeneity interacts with market incompleteness.

The limitations of current theoretical frameworks become apparent when confronting empirical puzzles. Avramov et al. (2022) document that ESG preference shocks generate asymmetric return responses and time-varying risk premia that static models cannot explain. Goldstein et al. (2022) show theoretically that green investing can paradoxically increase green firms' cost of capital by reducing price informativeness—an insight that emerges only in dynamic settings with information

acquisition. The heterogeneous agent approach naturally accommodates these complexities by allowing prices to reflect the evolving interaction between different investor types rather than a fixed preference parameter.

2.2 Empirical Evidence and Identification Challenges

The empirical literature provides robust evidence for all three channels while highlighting identification challenges. Bolton and Kacperczyk (2021) document a carbon premium in stock returns using cross-sectional regressions, finding that a one-standard-deviation increase in carbon intensity corresponds to 1.8% higher annual returns. However, their identification relies on controlling for numerous firm characteristics, leaving open the possibility that brown firms' higher returns reflect omitted fundamentals rather than pure preference effects.

Pástor et al. (2022) address this through German twin bonds: matched pairs of green and conventional bonds from the same issuer with identical cash flows. Since these bonds differ only in their green certification, price differentials directly measure the *preference channel* cleanly separated from cash flow differences. They find green bonds trade at yields several basis points lower than identical conventional bonds. Moreover, the time variation in twin bond premia—which they link to ESG fund flows—demonstrates that preference-driven pricing effects vary with the wealth distribution between investor types, validating the importance of modeling heterogeneous agents dynamically.

Beyond twin bonds, cross-sectional equity return patterns reveal the importance of *imperfect substitutability*. Hsu et al. (2023) and Chava (2014) document that environmental characteristics predict returns even after controlling for standard risk factors, suggesting that green and brown stocks load differently on priced risk factors. The empirical fact that green-brown return correlations are typically 0.6–0.8 rather than near-unity (Hsu et al., 2023) motivates modeling green and brown assets as distinct securities. The present model uses symmetric green-brown dividends to isolate the preference channel from cash flow differences; extending the model to incorporate asymmetric dividend dynamics that match the empirical green-brown correlation structure is a natural direction for future work.

2.3 Computational Methods and Heterogeneous Agent Models

The heterogeneous agent asset pricing literature has long recognized that investor heterogeneity generates rich dynamics absent from representative agent models. Dumas (1989) shows that with heterogeneous risk aversion, the wealth distribution becomes a state variable affecting risk premia. Chan and Kogan (2002) demonstrates how heterogeneous beliefs create trading volume and excess volatility. Gârleanu and Pedersen (2011) reveals how heterogeneous margin constraints generate asymmetric market dynamics. These insights suggest that ESG preference heterogeneity should similarly affect market outcomes, yet computational constraints have prevented researchers from exploring this channel with realistic market frictions.

Traditional numerical methods for solving heterogeneous agent models face the curse of dimensionality. Judd et al. (2003) develop projection methods for incomplete market models, but these become computationally infeasible with more than a few state variables. Kubler and Schmedders (2003) partially address this through collateral constraints and stationary equilibrium concepts, but the computational cost of projection methods grows rapidly with the number of state variables and assets.

Deep equilibrium networks, introduced by Azinovic et al. (2022), offer a complementary approach. By using neural networks to approximate policy functions and pricing kernels, DEQNs can scale to higher-dimensional state spaces while maintaining reasonable numerical accuracy. For this application, DEQNs naturally accommodate non-standard utility functions arising when investors derive utility directly from holdings of specific assets. The combination of scalability, flexibility in functional form, and computational speed makes DEQNs well-suited for the heterogeneous agent ESG pricing model.

2.4 Calibration Strategies in Structural Asset Pricing Models

The standard approach in quantitative asset pricing calibrates preference parameters to financial moments and cash flow dynamics to macroeconomic data. Mehra and Prescott (1985) establish the template: choose discount factor and risk aversion to match the equity premium and risk-free rate, while calibrating consumption growth from national accounts. Bansal and Yaron (2004) extend this to Epstein-Zin preferences, separating risk aversion from the elasticity of intertemporal substitution to jointly match both moments. In heterogeneous agent models, micro-level data on income dynamics and wealth distributions provide additional discipline: Heaton and Lucas (1996) and Krusell and Smith (1998) calibrate income processes from panel data (PSID) and validate against aggregate asset pricing moments, establishing the practice of combining micro and macro data sources.

This thesis follows the same approach: structural parameters (transition matrices, endowments, dividends) are calibrated from German macroeconomic data, preference parameters (β, γ) from German risk-free rates and the equity premium, and the ESG preference ϕ from the twin bond greenium. Section 4 details the data sources and calibrated values.

2.5 Market Incompleteness and ESG Pricing Persistence

A critical insight from the incomplete markets literature is that limited risk-sharing opportunities can sustain pricing differentials that complete markets would eliminate. Constantinides and Duffie (1996) show how uninsurable idiosyncratic income risk under market incompleteness can generate large equity premia, demonstrating that incomplete markets amplify the pricing effects of agent heterogeneity. Gârleanu and Panageas (2015) demonstrates that incomplete markets create state-dependent risk premia as agents' marginal utilities diverge. These mechanisms are particularly relevant for ESG pricing: even if investors derive utility from green assets and green-brown cash

flows are imperfectly correlated, frictionless complete markets might allow arbitrageurs to construct synthetic securities that eliminate pricing differentials.

In practice, multiple frictions limit arbitrage. First, the imperfect correlation between green and brown assets means that spanning requires combining multiple securities, exposing arbitrageurs to basis risk. Second, portfolio constraints—short-sale restrictions and borrowing limits—prevent traders from fully exploiting pricing differentials. Gromb and Vayanos (2002) and Kondor (2009) show how such constraints can sustain substantial pricing differentials even with risk-neutral arbitrageurs. Third, the fact that ESG preferences are heterogeneous and wealth-dependent means that pricing effects vary over time as wealth redistributes between investor types.

This thesis connects this incomplete markets insight to sustainable finance by modeling all three channels jointly. Green and brown assets have imperfectly correlated cash flows, ESG-conscious investors derive utility from green holdings, and portfolio constraints limit arbitrage. By incorporating these frictions into the heterogeneous agent model with a risk-free bond in zero net supply, the model provides a framework for quantifying how market incompleteness enables ESG preferences to affect asset prices—a question raised by Berk and van Binsbergen (2025), who argue these effects are negligible in frictionless markets.

2.6 Synthesis and Research Contribution

Most theoretical ESG pricing models isolate a single channel—preferences, substitutability, or arbitrage constraints—rather than letting them interact. Computational cost is a key reason: solving heterogeneous agent models with multiple assets and realistic frictions has been prohibitive with traditional projection methods. Calibration to data is also rare; most theoretical work relies on stylized parameterizations.

This thesis combines all three channels in a single model—a three-asset economy with endogenous bond pricing—and solves it using deep equilibrium networks, which scale to the required dimensionality. The model is then calibrated to German macroeconomic data and validated against the equity premium and the twin bond greenium.

3 Model

This section presents the theoretical framework. I develop a discrete-time pure-exchange economy with two heterogeneous agents ($i \in \{1, 2\}$) trading three assets: two Lucas trees representing green and brown equity ($j \in \{g, b\}$) in unit net supply, and a risk-free bond in zero net supply. The economy features twelve exogenous shocks capturing ESG regime shifts, business cycle fluctuations, and household income heterogeneity. Equilibrium prices—including an endogenous bond price—and portfolio allocations are determined simultaneously by agents' optimizing behavior and market clearing. Unlike Sauzet and Zerbib (2022), I work in discrete time and focus exclusively on heterogeneous portfolio preferences for green versus brown assets, abstracting from consumption channel effects to isolate the investment mechanisms driving ESG asset pricing.

3.1 Economic Environment

Time and uncertainty. Time is discrete with infinite horizon, $t \in \{0, 1, 2, \ldots\}$. Uncertainty is represented by an exogenous state variable s_t that follows a time-homogeneous Markov chain on a finite state space $\mathcal{S} = \{0, 1, \ldots, 11\}$ with transition matrix Π, where $\Pi_{ss'} = \Pr(s_{t+1} = s' \mid s_t = s)$. The twelve shocks arise as the Kronecker product of three independent Markov chains (Equation 1):

$$\Pi = P_{\text{green}} \otimes P_{\text{macro}} \otimes P_{\text{hh}} \tag{1}$$

where P_{green} is a 2×2 matrix governing ESG regime persistence, P_{macro} is a 2×2 matrix governing business cycle transitions, and P_{hh} is a 3×3 matrix governing household income dynamics. Each state s corresponds to a unique triple (g, m, h) with $g \in \{0, 1\}$ (ESG regime), $m \in \{0, 1\}$ (macro regime), and $h \in \{0, 1, 2\}$ (household income state), indexed as $s = 6g + 3m + h$ (a bijection from (g, m, h) to $\{0, \ldots, 11\}$ chosen for computational convenience). This structure captures three distinct sources of uncertainty relevant for ESG asset pricing.

The *ESG regime* ($g \in \{0, 1\}$) governs relative dividend performance: $g = 0$ favors brown assets, $g = 1$ favors green. The symmetric transition matrix P_{green} ensures $\mathbb{E}[d_g] = \mathbb{E}[d_b]$ unconditionally, so any greenium arises purely from preferences—a deliberate choice to isolate the preference channel from cash flow differences.

The *business cycle* ($m \in \{0, 1\}$) captures aggregate fluctuations: $m = 0$ is good times (higher endowments and dividends), $m = 1$ is bad times. Transition probabilities are calibrated from German GDP data (Section 4).

Household income heterogeneity ($h \in \{0, 1, 2\}$) generates cross-sectional variation: $h = 0$ favors Agent 1, $h = 2$ favors Agent 2, and $h = 1$ splits income equally. The persistence is calibrated from the German SOEP literature ($\rho \approx 0.92$–0.97; Section 4).

Assets and dividends. The economy features three traded assets. Two Lucas trees (long-lived assets that pay state-contingent dividends) in unit net supply represent green ($j = g$, index 0) and brown ($j = b$, index 1) equity. A risk-free bond ($j = f$, index 2) is in zero net supply. Each asset j pays a state-contingent dividend $d_j(s)$.

Green equity pays higher dividends in the green-favored regime ($g = 1$) and lower dividends in the brown-favored regime ($g = 0$), with dividends also varying with the macro state. Brown equity exhibits the mirror image: it pays higher dividends in the brown-favored regime and lower dividends in the green-favored regime. This symmetric structure ensures $\mathbb{E}[d_g] = \mathbb{E}[d_b]$ unconditionally, so that any equilibrium price differential between the two assets reflects investor preferences rather than fundamental cash flow advantages. The dividend levels are calibrated to the DAX dividend yield, which averages approximately 2.8% over the 2001–2024 sample period (Section 4).

The risk-free bond pays a constant coupon of $d_f = 1$ in every state and has zero depreciation ($\delta_f = 0$), making it economically equivalent to a one-period discount bond that is repurchased at its endogenous price each period rather than carrying over resale value. Its price $p_f(s_t, \theta_t)$ is determined in equilibrium by the same Euler equations that price the risky assets, ensuring

consistency between the risk-free rate and the pricing kernel. The bond is in zero net supply, so that in equilibrium one agent's lending is exactly offset by the other's borrowing. This structure introduces a risk-free savings technology that enriches portfolio choice without requiring exogenous interest rate assumptions.

The green and brown trees trade as equity assets with ex-dividend prices p_j. Each risky asset has a depreciation parameter $\delta_j \in [0, 1]$: holding one share of asset j purchased at price p_j yields $d_j(s) + \delta_j \cdot p_j$ at the beginning of the next period. When $\delta_j = 1$, the asset is a standard Lucas tree with full capital gain; when $\delta_j = 0$, the asset pays only dividends with no resale value. The gross return from holding asset j from period t to $t+1$ is given by Equation (2):

$$R^j_{t+1} = \frac{d_j(s_{t+1}) + \delta_j \cdot p_j(s_{t+1}, \theta_{t+1})}{p_j(s_t, \theta_t)} \tag{2}$$

Endowments. Each agent $i \in \{1, 2\}$ receives a state-contingent labor endowment $e_i(s_t)$. The endowment structure is symmetric across agents: when $h = 0$, Agent 1 receives higher income; when $h = 2$, Agent 2 is favored; when $h = 1$, income is equal. Aggregate endowments are higher in good macro times ($m = 0$) than in bad times ($m = 1$). The ESG regime g does not affect labor endowments directly—it affects only asset dividends.

3.2 Preferences and Portfolio Choice

Utility specification. Agent i has time-separable preferences with discount factor $\beta \in (0, 1)$ and constant relative risk aversion (CRRA) utility over consumption:

$$U^i = \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \left(u(c^i_t) + \phi_i \cdot v(\omega^{i,g}_t) \right) \right] \tag{3}$$

where consumption c^i_t is determined residually from the budget constraint, $u(c) = c^{1-\gamma}/(1-\gamma)$ for $\gamma \neq 1$ and $u(c) = \log(c)$ for $\gamma = 1$, $v(\omega) = \omega$ is the (linear) ESG utility function, and $\phi_i \geq 0$ scales the non-pecuniary utility from green portfolio weight $\omega^{i,g}_t$. The linearity of v implies that the marginal ESG benefit of increasing the green portfolio weight is constant, independent of the current level of greenness. The discount factor β and risk aversion γ are identical across agents to isolate the effects of ESG preference heterogeneity.

Note that $\omega^{i,g}_t$ in Equation (3) depends on the portfolio θ^i_{t-1} carried *into* period t, not on the portfolio θ^i_t chosen *during* period t. This timing convention—the agent derives ESG utility from the portfolio she holds at the start of each period—implies that the choice of θ^i_t at time t affects ESG utility at time $t+1$. Consequently, the ESG gradient appears inside the discounted expectation in the Euler equation (Equation 10 below), rather than as a contemporaneous term.

Budget constraint and consumption. Agent i chooses next-period portfolio holdings $\theta^i_t = (\theta^{i,g}_t, \theta^{i,b}_t, \theta^{i,f}_t)$ subject to the budget constraint. Consumption is determined residually by

Equation (4):

$$c_t^i = e_i(s_t) + \sum_j \theta_{j,t-1}^i \left[d_j(s_t) + \delta_j \cdot p_j(s_t, \theta_t) \right] - \sum_j p_j(s_t, \theta_t) \cdot \theta_{j,t}^i \tag{4}$$

The first term is labor endowment. The second term captures asset payoffs from holdings carried into the period: dividends $d_j(s_t)$ plus the depreciated resale value $\delta_j \cdot p_j$. The third term is the cost of purchasing the new portfolio at current prices. For the bond ($\delta_f = 0$), the payoff simplifies to $d_f = 1$ plus no capital gain, making the bond a pure fixed-income instrument whose price adjusts endogenously to clear the market.

ESG preferences. Following Pástor et al. (2021) and Sauzet and Zerbib (2022), Agent i derives additional utility from the portfolio weight allocated to the green asset. The green portfolio weight is defined over risky assets only (the bond is ESG-neutral):

$$\omega_t^{i,g} = \frac{p_g \cdot \theta_t^{i,g}}{\sum_{j \in \{g,b\}} p_j \cdot \theta_t^{i,j} + \epsilon_w} \tag{5}$$

where $\epsilon_w > 0$ is a small regularization parameter that ensures the denominator is bounded away from zero. The ESG gradient for asset j with respect to the portfolio choice $\theta_t^{i,j}$ is:

$$\frac{\partial \omega_t^{i,g}}{\partial \theta_t^{i,j}} = \begin{cases} \frac{p_g}{W_t^i}(1 - \omega_t^{i,g}) & \text{if } j = g \\ -\frac{p_j}{W_t^i} \cdot \omega_t^{i,g} & \text{if } j = b \\ 0 & \text{if } j = f \text{ (bond)} \end{cases} \tag{6}$$

where $W_t^i = \sum_{j \in \{g,b\}} p_j \cdot \theta_t^{i,j} + \epsilon_w$ is total risky portfolio value. The ESG term increases the marginal value of green assets and decreases the marginal value of brown assets, generating the green premium. The bond is ESG-neutral: its gradient is zero.

Agent 1 is the ESG-conscious investor ($\phi_1 > 0$) and Agent 2 is the traditional investor ($\phi_2 = 0$). The intensity parameter ϕ_1 is calibrated to match the observed German twin bond greenium (Section 4).

Portfolio constraints. Risky asset holdings are subject to short-sale constraints:

$$\theta_t^{i,j} \geq 0, \quad j \in \{g, b\} \tag{7}$$

The bond is subject to borrowing and lending limits:

$$\theta_t^{i,f} \in [\underline{\theta}_f, \overline{\theta}_f] \tag{8}$$

These constraints serve as arbitrage limits. When binding, they prevent agents from fully implementing their desired portfolios, creating wedges between marginal rates of substitution and

sustaining the green premium in equilibrium. Without short-sale constraints, traditional investors could take unlimited short positions in overpriced green assets, driving the premium to zero.

3.3 Equilibrium

Definition. A recursive competitive equilibrium consists of price functions $p_j(s_t, \theta_{t-1})$ for each asset $j \in \{g, b, f\}$ and policy functions $\theta^{1,j}(s_t, \theta_{t-1})$ for Agent 1 such that:

1. Given prices, each agent i maximizes expected lifetime utility subject to the budget constraint and portfolio constraints.
2. Asset markets clear: $\theta_t^{1,j} + \theta_t^{2,j} = \bar{\theta}_j$ for each j, where $\bar{\theta}_g = \bar{\theta}_b = 1$ (unit supply) and $\bar{\theta}_f = 0$ (zero net supply).

State variables. The equilibrium is characterized by a four-dimensional state vector:

$$X_t = (s_t, \theta_{t-1}^{1,g}, \theta_{t-1}^{1,b}, \theta_{t-1}^{1,f}) \tag{9}$$

capturing the exogenous state and Agent 1's portfolio holdings. Given market clearing, Agent 2's holdings are redundant: $\theta_{t-1}^{2,j} = \bar{\theta}_j - \theta_{t-1}^{1,j}$.

First-order conditions. The Euler equation for agent i and asset j is given by Equation (10):

$$p_j \cdot u'(c_t^i) = \beta\, \mathbb{E}_t \left[(d_j(s_{t+1}) + \delta_j \cdot p_j(s_{t+1}, \theta_t)) \cdot u'(c_{t+1}^i) + \phi_i \cdot \frac{\partial \omega_{t+1}^{i,g}}{\partial \theta_t^{i,j}} \right] \tag{10}$$

where $u'(c) = c^{-\gamma}$ is the CRRA marginal utility. The left-hand side is the marginal cost of purchasing one unit of asset j. The right-hand side is the discounted expected marginal benefit: the financial payoff weighted by marginal utility, plus the ESG gradient term. For agents without ESG preferences ($\phi_i = 0$), this reduces to the standard consumption-based Euler equation.

Complementarity constraints. To handle the occasionally binding portfolio constraints, I reformulate the Euler equations as complementarity conditions using the Fischer-Burmeister function (Fischer, 1992) (Equation (11); see Appendix B):

$$FB(\mathcal{E}_t^{i,j},\, \theta_t^{i,j} - \underline{\theta}_j) = \sqrt{(\mathcal{E}_t^{i,j})^2 + (\theta_t^{i,j} - \underline{\theta}_j)^2 + \epsilon} - \mathcal{E}_t^{i,j} - (\theta_t^{i,j} - \underline{\theta}_j) = 0 \tag{11}$$

where $\mathcal{E}_t^{i,j}$ is the Euler equation residual and $\underline{\theta}_j$ is the lower bound on holdings (0 for risky assets, $\underline{\theta}_f$ for the bond). Setting $FB = 0$ is equivalent to the KKT conditions: either the Euler equation holds with equality (interior solution) or the constraint binds (boundary solution).

3.4 Solution Method: Deep Equilibrium Networks

I solve for the equilibrium using deep equilibrium networks (DEQNs) following Azinovic et al. (2022). The unknown equilibrium objects—three price functions and three next-period portfolio choices for Agent 1—are parameterized as the output of a single deep neural network.

Network architecture. The network takes the four-dimensional state vector

$$X_t = \left(s_t,\ \theta_{t-1}^{1,g},\ \theta_{t-1}^{1,b},\ \theta_{t-1}^{1,f}\right)$$

as input. Before entering the hidden layers, the raw state is augmented into a 29-dimensional feature vector through the following transformation:

$$\tilde{X}_t = \left(s_t,\ \theta_t^1,\ \theta_t^2,\ e(s_t),\ d(s_t),\ \mathbb{E}[d(s_{t+1}) \mid s_t],\ \Pi_{s_t,\cdot},\ \delta\right) \tag{12}$$

This augmentation provides the network with all information needed to evaluate the Euler equations without requiring the network to "learn" the model's exogenous structure. The current endowments $e(s_t)$, dividends $d(s_t)$, expected dividends $\mathbb{E}[d(s_{t+1}) \mid s_t]$, transition probabilities $\Pi_{s_t,\cdot}$, and depreciation rates δ are all deterministic functions of the state s_t and the model parameters. Including them as inputs rather than requiring the network to learn these mappings from the discrete state index alone substantially accelerates training convergence.

The augmented input passes through four hidden layers with 1,024 neurons each (approximately 12 million trainable parameters), trained on the Longleaf HPC cluster at UNC Chapel Hill. Each hidden layer applies a dense linear transformation, layer normalization, and the GELU activation function. Residual (skip) connections add the input of each layer to its output, enabling gradient flow through the deep architecture and facilitating the learning of near-identity mappings at early training stages when the warm-started network is already close to the solution (Azinovic et al., 2022).

The network outputs six quantities: three asset prices (p_g, p_b, p_f) and three next-period holdings $(\theta_t^{1,g}, \theta_t^{1,b}, \theta_t^{1,f})$ for Agent 1. Output activations enforce economic constraints: prices pass through a softplus function ($\text{softplus}(x) = \log(1 + e^x) + \epsilon_p$, where $\epsilon_p = 10^{-3}$) ensuring strict positivity; risky equity shares pass through a sigmoid function ($\sigma(x) = 1/(1 + e^{-x})$) ensuring $\theta^{i,j} \in [0, 1]$; and bond holdings pass through a scaled tanh ($\alpha \tanh(x/\alpha)$, where $\alpha = 3 \cdot \overline{\theta}_f$) allowing both positive and negative positions within the leverage limits. These output constraints are "soft" in that they are always satisfied by construction, unlike penalty-based approaches that may violate constraints during training.

Loss function. The network is trained to minimize a loss function based on the Fischer-Burmeister Euler errors (Equation (13)):

$$\mathcal{L} = \underbrace{\|FB\|_2}_{\text{mean Euler error}} + 3 \cdot \underbrace{\left\| \max_{i,j} |FB^{i,j}| \right\|_2}_{\text{worst-case penalty}} + \lambda \cdot \underbrace{(\text{consumption penalties} + \text{bond range penalty})}_{\text{feasibility}} \tag{13}$$

The mean Euler error drives convergence of the average solution quality. The worst-case penalty prevents the network from sacrificing accuracy at corner states to reduce the average. The feasibility penalties ensure positive consumption and bounded bond positions. The expectations in the Euler equations are computed by evaluating the network at all $S = 12$ possible next-period states and

weighting by transition probabilities.

Delta curriculum. Solving the full general equilibrium directly is challenging because prices appear on both sides of the Euler equation (self-referential fixed point). The asset payoff $d_j(s') + \delta \cdot p_j(s', \theta')$ depends on next-period prices, which are themselves outputs of the network evaluated at next-period states. This creates a fixed-point problem: the network must be consistent with its own predictions. Direct training at $\delta = 1$ typically diverges because the initial (random) network produces prices that bear no relation to the true equilibrium, and gradient descent cannot simultaneously learn reasonable prices and their self-consistency.

I employ a curriculum learning strategy that gradually introduces this feedback by increasing δ from 0 to the target value across five stages: $\delta = 0, 0.25, 0.50, 0.60, 0.70$. The delta stages are trained with log utility ($\gamma = 1$); a subsequent gamma calibration step retrains the network at the calibrated $\gamma^* = 4.864$ before the phi curriculum begins. At $\delta = 0$, the asset payoff is purely $d_j(s')$—the economy reduces to a collection of one-period pricing problems with no intertemporal feedback. The network learns accurate prices for this simpler economy, then uses them as a warm start for $\delta = 0.25$, where prices have a small self-referential component. As δ increases, the feedback strengthens, but each increment is small enough that the warm-started network is already close to the new equilibrium. The final stage at $\delta = 0.7$ represents a partial Lucas tree where 70% of the asset value carries forward each period. The choice of $\delta = 0.7$ rather than $\delta = 1.0$ (a full Lucas tree) reflects a practical tradeoff: at higher δ, the self-referential fixed-point becomes sufficiently nonlinear that steady-state Euler errors increase substantially, with the network struggling to simultaneously learn accurate prices and their self-consistency at the full $\delta = 1.0$. Since the model's qualitative predictions—the sign and magnitude of the greenium, the wealth-distribution channel, the constraint amplification mechanism—are robust to the choice of δ in the range $[0.5, 0.7]$, I adopt $\delta = 0.7$ as the best tradeoff between economic realism and computational accuracy. Extending the curriculum to $\delta = 1.0$ with improved training techniques is a natural direction for future work.

After the delta curriculum completes, a *phi curriculum* gradually introduces ESG preferences by increasing agent 1's preference intensity: $\phi_1 = 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 0.5, 1.0$. This second curriculum is necessary because ESG preferences break the green-brown symmetry of the equilibrium—introducing $\phi > 0$ abruptly from the symmetric $\phi = 0$ solution can cause the optimizer to converge to a local minimum where the greenium has the wrong sign or magnitude.

Green-brown symmetry augmentation. A key computational innovation is exploiting the dividend symmetry at $\phi = 0$. Since $d_{\text{green}}(g = 0, m, h) = d_{\text{brown}}(g = 1, m, h)$ and P_{green} is symmetric, the equilibrium at $\phi = 0$ satisfies $p_{\text{green}}(g = 0, m, h, \theta) = p_{\text{brown}}(g = 1, m, h, \tilde{\theta})$, where $\tilde{\theta}$ swaps green and brown holdings. I enforce this symmetry during training by doubling each batch: for every sample $(s, \theta^g, \theta^b, \theta^f)$, I add the mirror sample $((s \pm 6) \bmod 12, \theta^b, \theta^g, \theta^f)$ that swaps the green/brown dimensions. This augmentation halves the effective dimensionality of the learning problem at $\phi = 0$ and substantially reduces numerical noise in the baseline greenium, which the $\phi > 0$ stages inherit.

Training details. The optimizer is AdamW with weight decay 10^{-4} and gradient clipping at

global norm 5.0. The learning rate follows a linear warmup (200 steps) then cosine decay schedule, starting at 5×10^{-4} for early delta stages and decreasing to 3×10^{-5} for late phi stages. Each training episode draws fresh random states from the portfolio space and combines them with a fixed deterministic grid covering all 12 exogenous states crossed with a Cartesian grid of 9 risky share points $\times$ 5 bond points. The deterministic grid ensures coverage of extreme portfolio positions where constraints may bind, while random sampling provides diversity in the interior. A fixed held-out validation set is used for early stopping based on the validation Euler loss, with patience of 3,000–6,000 episodes depending on the stage.

Note: An extension incorporating Epstein-Zin recursive preferences is discussed in Appendix A. The baseline model uses CRRA preferences to isolate ESG pricing mechanisms before adding additional complexity.

4 Data and Calibration

This section describes the data sources used to calibrate the model and presents descriptive statistics for the German twin bond dataset that serves as the primary empirical validation target. I follow the standard calibration approach in quantitative asset pricing: structural parameters governing exogenous dynamics are estimated from macroeconomic data, preference parameters are disciplined by financial moments, and the ESG preference intensity is calibrated to match the observed greenium (Mehra and Prescott, 1985; Campbell and Cochrane, 1999; Heaton and Lucas, 1996).

4.1 Calibration Data Sources

The calibration draws on seven data sources spanning German macroeconomic aggregates, financial market returns, and ESG-specific bond data. Table 2 at the end of this section consolidates all calibrated parameter values; the subsections below detail each parameter's derivation. Three parameters deviate from data-implied values for modeling reasons (Table 1).

4.1.1 Discount Factor

The discount factor is calibrated via $\beta = 1/(1 + r_f)$ using the German 10-year Bund yield as the risk-free rate benchmark (Mehra and Prescott, 1985). Using annual data from the OECD via FRED (series IRLTLT01DEA156N) over 2001–2024, the mean Bund yield is 2.17%, giving $\beta = 1/(1 + 0.0217) = 0.978$. The full time series (Table 13, Appendix E) spans negative yields during 2019–2021 and the post-2022 normalization, but the long sample period provides a robust average: using only post-2022 data gives $\beta \approx 0.977$, a trivially different value.

4.1.2 Business Cycle Transition Matrix

The macro transition matrix P_{macro} is estimated from German real GDP growth data over 2001–2024 (Eurostat via FRED, series CLVMNACSCAB1GQDE) by classifying each year as "good" (positive growth) or "bad" (non-positive growth) and computing transition probabilities from the observed regime sequence.

Over 23 year-to-year transitions, Germany experienced 18 good years and 6 bad years. Counting transitions: $P(\text{good} \mid \text{good}) = 14/18 = 0.78$ and $P(\text{good} \mid \text{bad}) = 3/5 = 0.60$, giving

$$P_{\text{macro}} = \begin{pmatrix} 0.78 & 0.22 \\ 0.60 & 0.40 \end{pmatrix}, \tag{14}$$

with a stationary distribution of 73% good times and 27% bad times, approximating the empirical frequency of 75% good years. The transition probabilities are imprecisely estimated from only 23 annual observations—for instance, $P(\text{good} \mid \text{bad}) = 0.60$ is based on just 5 bad-year observations—so these values should be understood as rough calibration targets rather than precise estimates. The full GDP series appears in Table 14 (Appendix E).

I set good-time aggregate endowment to 3.0 and bad-time to approximately 2.97 (bad/good ratio ≈ 0.99), matching the empirical GDP contraction. With $\gamma \approx 5$, the realistic endowment contrast generates sufficient risk premia and meaningful portfolio dynamics without requiring the amplified consumption variation used in log-utility models (Mehra and Prescott, 1985; Campbell and Cochrane, 1999).

4.1.3 Household Income Dynamics

The household income transition matrix P_{hh} is calibrated from the German Socio-Economic Panel (SOEP) literature. Drechsel-Grau et al. (2022) estimate annual earnings persistence $\rho \approx 0.92$–0.97 for German households using SOEP data from 1993–2018; Fuchs-Schündeln et al. (2010) find similar estimates ($\rho \approx 0.95$) and show that German income inequality is well captured by a discrete Markov chain approximation, following the methodology of Storesletten et al. (2004).

I approximate these estimates with a three-state Markov chain having diagonal elements of 0.82. The lower persistence relative to the continuous-time estimates reflects the discretization penalty: a standard Tauchen or Rouwenhorst approximation of an AR(1) with $\rho = 0.95$ onto three states yields diagonal elements in the range 0.80–0.85, depending on the grid spacing and approximation method; the value of 0.82 lies in the middle of this range. In the favorable state ($h = 0$), Agent 1 receives higher labor income; in the neutral state ($h = 1$), income is split equally; in the unfavorable state ($h = 2$), Agent 2 is favored. This heterogeneity creates risk-sharing motives and interacts with ESG preferences: agents with temporarily high income allocate more aggressively to green assets, while agents facing income shortfalls prioritize financial returns.

4.1.4 Risk Aversion and the Equity Premium

The model calibrates risk aversion γ to match the equity premium. Two empirical benchmarks provide context: the realized equity premium (DAX Kursindex minus Bund yield, mean 5.15%, Table 15, Appendix E) and the implied ERP from Damodaran (2025) (mean 4.81% for Aaa-rated markets over 2000–2025, Table 16, Appendix E). The DAX Kursindex (price-only, excluding dividends) is used rather than the DAX Performance Index (total return) because the model's greenium calibration target comes from German twin bonds; using German equity data maintains geographic consistency across calibration targets. The DAX Kursindex understates the total equity premium by the dividend yield ($\approx 2.8\%$), but I calibrate γ to the Damodaran implied ERP rather than the realized premium. A bisection search over γ yields $\gamma^* = 4.864$, which produces a model equity premium of 4.86%, closely matching the Damodaran benchmark of 4.81%. This moderate risk aversion is well within the range considered standard in the macro-finance literature ($\gamma \approx 2$–10), avoiding the extreme values implied by the Mehra-Prescott approximation ($\gamma \approx 92$) while generating realistic risk premia with the data-implied endowment contrast.

The greenium is well-identified independently of γ, since both risky assets face identical aggregate risk and differ only through the ESG preference channel. The greenium is a *relative* price (green versus brown), driven by ϕ and portfolio constraints, whereas the equity premium is an *absolute* price (equity versus risk-free), driven by risk aversion and the endowment process. The two are largely orthogonal: the calibrated ϕ^* changes the equity premium by less than 0.01 percentage points relative to the $\phi = 0$ baseline.

4.1.5 Calibration Deviations

Most model parameters are computed directly from German data. The calibrated risk aversion $\gamma = 4.864$ matches the equity premium and the data-implied endowment contrast ($e_{\text{bad}}/e_{\text{good}} \approx 0.99$) produces realistic consumption dynamics at this risk aversion level. The remaining deviation from data-implied values is the ESG regime transition matrix. Table 1 summarizes.

Table 1: Data-Implied vs. Model Parameter Values

Parameter	Data-implied	Model value	Status
γ	91.6	4.864	Calibrated to equity premium (4.86%)
e_{bad}	2.97	2.97	Data-implied (no deviation)
P_{green}	$\begin{pmatrix} 0.50 & 0.50 \\ 0.50 & 0.50 \end{pmatrix}$	$\begin{pmatrix} 0.90 & 0.10 \\ 0.10 & 0.90 \end{pmatrix}$	Modeling assumption

Note: Data-implied γ from Mehra-Prescott approximation $\gamma \approx EP/\sigma_c^2$. Data-implied e_{bad} from German GDP bad/good ratio (≈ 0.99). Data-implied P_{green} is i.i.d. (no ESG regime persistence in available data).

The Mehra-Prescott approximation implies $\gamma \approx 92$, which is unreasonably high—this is the equity premium puzzle itself (Mehra and Prescott, 1985). The calibrated $\gamma = 4.864$ sidesteps this

issue: with CRRA preferences and the realistic endowment contrast, moderate risk aversion suffices to match the equity premium. This value lies within the standard range ($\gamma \approx 2$–10) used in the macro-finance literature (Campbell and Cochrane, 1999; Bansal and Yaron, 2004). The Epstein-Zin extension (Appendix A), which separates risk aversion from intertemporal substitution, remains a natural direction for future work but is no longer necessary for matching the equity premium.

The sole remaining deviation is the ESG regime transition matrix. The data-implied P_{green} is i.i.d., reflecting the lack of established evidence on ESG regime persistence. The persistent specification (P_{green} diagonal $= 0.90$) is a modeling assumption that creates distinct green-favored and brown-favored regimes, enabling the analysis of how ESG regime shifts interact with portfolio constraints and the wealth distribution. This assumption could be resolved as longer time series on ESG investment regimes become available.

4.1.6 Dividend Levels and Green-Brown Differential

Dividend levels are calibrated to the DAX dividend yield, which averages approximately 2.8% over 2001–2024 (STOXX/Qontigo factsheets). For the green-brown differential, ESG-screened German equity indices show only 10–20 basis point yield differences relative to conventional indices (Table 17, Appendix E), confirming that dividend differences play at most a secondary role in explaining green-brown pricing patterns.

Accordingly, I set green and brown dividends symmetrically across ESG regimes: green equity pays higher dividends in the green-favored regime ($g = 1$) and lower dividends in the brown-favored regime ($g = 0$), with brown equity paying the reverse. The symmetric P_{green} ensures $\mathbb{E}[d_g] = \mathbb{E}[d_b]$ unconditionally, so any greenium arises purely from the ESG preference parameter ϕ rather than dividend differentials.

4.1.7 Cross-Country Data Sourcing

The calibration draws primarily on German data for structural parameters and the greenium. For the equity premium, I use the implied ERP of Damodaran (2025) for mature (Aaa-rated) markets, a US-based estimate applied globally. This cross-country approach is standard: Mehra and Prescott (1985), Campbell and Cochrane (1999), and Bansal and Yaron (2004) use US financial moments to discipline preference parameters viewed as reflecting universal investor behavior, while Heaton and Lucas (1996) and Krusell and Smith (1998) calibrate income processes from US panel data and validate against aggregate asset pricing moments from multiple sources. Since Germany is Aaa-rated and integrated into global capital markets, the mature-market ERP applies directly. The key calibration target—the twin bond greenium—comes exclusively from German data.

4.2 German Twin Bond Data

German twin bonds are uniquely suited for identifying the green premium because they are issued by the same sovereign entity—the Federal Republic of Germany, through the Deutsche Finanzagentur

(German Finance Agency)—with identical maturity, coupon structure, and credit risk, differing only in their environmental designation (Pástor et al., 2022). This institutional design eliminates the identification challenges that plague cross-sectional studies of green premia in equity markets, where green and brown firms differ in size, sector, growth opportunities, and risk exposures. With twin bonds, any price differential can be attributed purely to investor preferences for environmental attributes rather than differences in fundamental cash flow risk, providing the cleanest possible measure of the ESG preference channel.

The dataset contains 1,301 daily observations of bond prices and yields for both green and conventional bonds over the period from September 2020 to October 2025. Germany's green bond program, launched in September 2020, was specifically designed to facilitate this type of analysis: each green bond is paired with a conventional twin of the same maturity and coupon, creating a natural experiment in sustainability labeling. The greenium—defined as the yield differential (Green − Brown)—serves as the primary calibration target for the ESG preference parameter ϕ_1.

The dataset spans a particularly informative period in financial market history, encompassing the pandemic-induced monetary stimulus (2020–2021), the subsequent inflation surge and aggressive ECB tightening (2022–2023), and the stabilization phase (2024–2025). The persistence of the green premium across these diverse macroeconomic regimes provides evidence that environmental preferences remain salient even during periods of economic stress—an important empirical regularity that the model must replicate.

4.2.1 Summary Statistics

The greenium averages −2.44 basis points with a standard deviation of 1.9 basis points (Tables 9–10, Appendix D). The negative sign indicates that green bonds consistently yield less than their conventional twins—investors accept lower financial returns in exchange for environmental attributes. Applied to Germany's green bond outstanding of approximately EUR60 billion (Deutsche Finanzagentur, 2025), a 2.44 basis point yield concession translates to roughly EUR15 million per year in foregone interest income. This greenium provides the key calibration target for the ESG preference intensity parameter ϕ_1.

4.2.2 Trends Over Time

Figures 1 and 2 illustrate the temporal evolution of bond yields and prices. The yield series exhibit three distinct phases: negative yields during 2020–2021 (pandemic-era monetary stimulus), sharp increases during 2022–2023 (inflation and monetary tightening), and stabilization around 2% in 2024–2025. Throughout these phases, green and brown bond yields move in close tandem, confirming that the green premium reflects preferences rather than differential risk exposure.

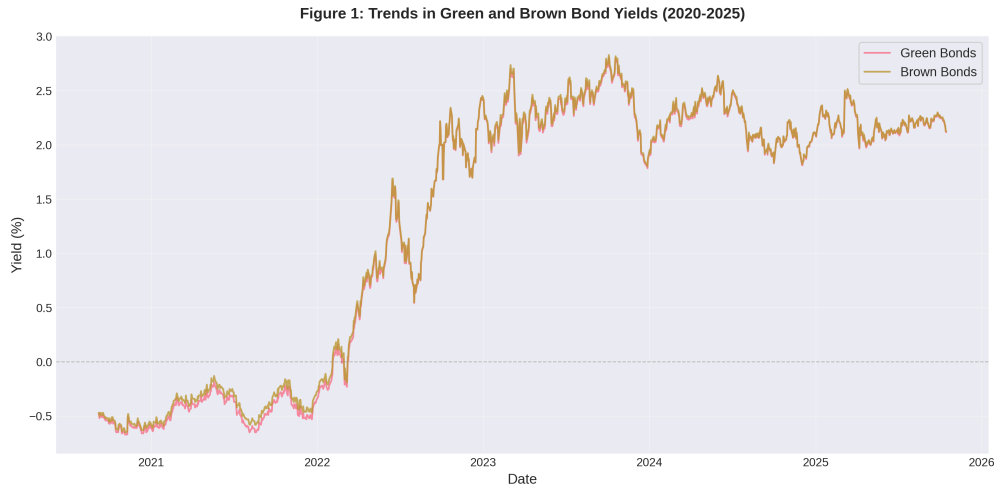

Figure 1: Daily green and brown bond yields, September 2020 – October 2025.

Figure 2: Daily green and brown bond prices, September 2020 – October 2025.

4.2.3 The Greenium

Figure 3 provides a detailed view of the greenium evolution over time. The yield differential has narrowed from its widest annual average in 2021 (approximately 4 basis points) to approximately 2–3 basis points in recent years. This convergence may reflect supply-side effects as more green bonds entered the market, changes in the wealth distribution of ESG-conscious investors, or diminishing novelty premia. Importantly, the greenium has remained consistently negative throughout the sample period, never crossing into positive territory, underscoring the persistent nature of ESG-driven pricing effects.

A paired t-test confirms that the greenium is highly statistically significant ($t = -45.82$, $p < 0.001$; Table 11, Appendix D).

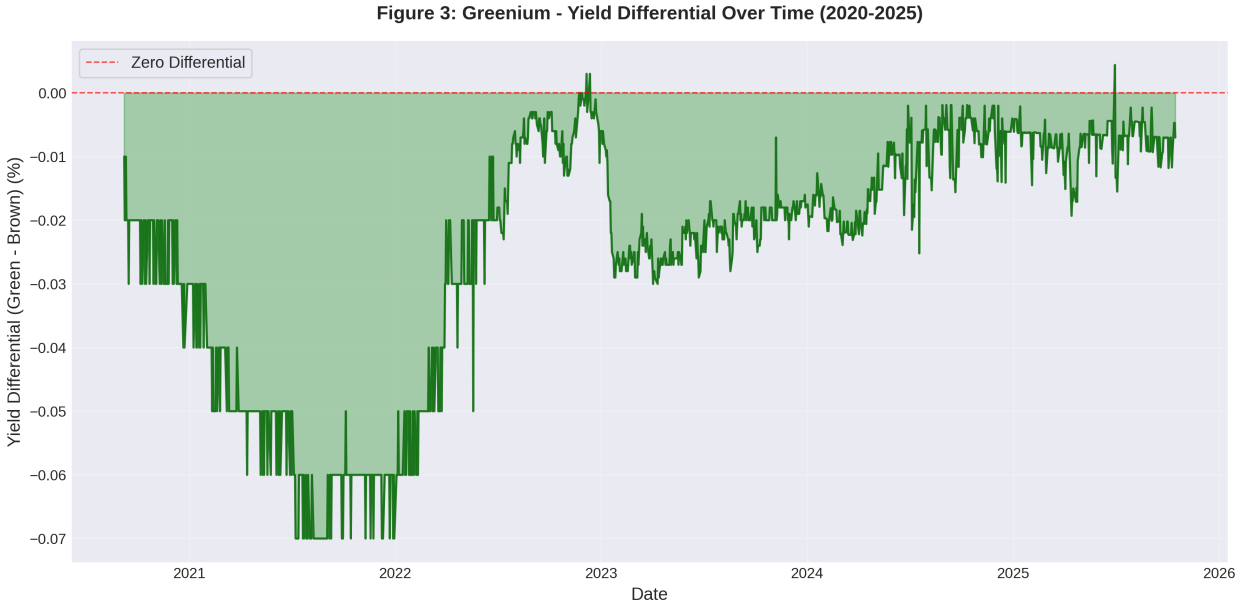

Figure 3: Greenium (yield differential, Green − Brown) over time.

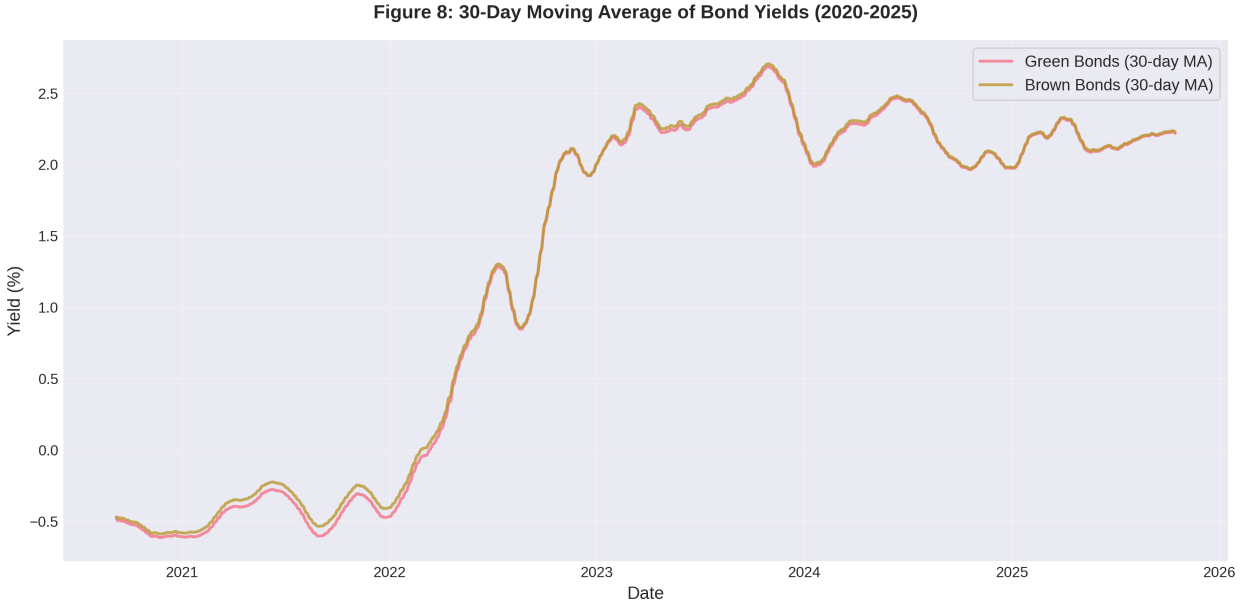

Figure 4: 30-day moving average of bond yields, September 2020 – October 2025.

The greenium shows some evidence of narrowing over time: the price differential declined from 0.52 in 2021 to 0.03 by 2025, while the yield differential has remained relatively stable at −2 to −3 basis points (Table 12, Appendix D). This temporal pattern is consistent with supply-side explanations (more green bonds reducing scarcity premia) and with the model's prediction that the greenium depends on the wealth distribution between ESG and traditional investors.

Table 2 consolidates all calibrated parameter values and their data sources.

Table 2: Calibrated Parameter Values

Parameter	Symbol	Value	Source
Preference parameters			
Discount factor	β	0.978	German Bund 10Y mean (2.17%)
Risk aversion	γ	4.864	Calibrated to equity premium
ESG preference	ϕ_1	1.0×10^{-4}	Calibrated to twin bond greenium
Neutral preference	ϕ_2	0.0	Traditional investor (no ESG)
Transition matrices			
ESG regime	P_{green}	$\begin{pmatrix} 0.90 & 0.10 \\ 0.10 & 0.90 \end{pmatrix}$ (data: i.i.d.)	Persistence assumption
Business cycle	P_{macro}	$\begin{pmatrix} 0.78 & 0.22 \\ 0.60 & 0.40 \end{pmatrix}$	German GDP growth regimes
Household income	P_{hh}	diag. 0.82	SOEP income persistence
Asset parameters			
Depreciation (equity)	$\delta_{g,b}$	0.7	Partial Lucas tree
Depreciation (bond)	δ_f	0.0	Pure coupon instrument
Bond coupon	d_f	1.0	Normalization
Net supply (equity)	$\bar{\theta}_{g,b}$	1.0	Unit supply
Net supply (bond)	$\bar{\theta}_f$	0.0	Zero net supply
Portfolio constraints			
Equity bounds	$[\underline{\theta}, \overline{\theta}]$	$[0, 1]$	Short-sale restriction
Bond bounds	$[\underline{\theta}_f, \overline{\theta}_f]$	$[-0.5, 0.5]$	Leverage limit
Calibration targets			
Twin bond greenium		-2.44 bps	German twin bonds (2020–2025)
Equity premium		4.81%	Damodaran (2025) (2000–2025)
Untargeted moment			
Risk-free rate		2.17%	German Bund 10Y (2001–2024)

Note: See Sections 2.4–4 for detailed discussion. Three parameters deviate from data-implied values (Table 1). The ESG preference ϕ_1 is calibrated post-training; all other parameters are fixed before the calibration exercise.

5 Results

This section presents the computational results from solving the model using deep equilibrium networks. I first report convergence diagnostics to establish that the numerical solution is reliable, then examine how the greenium varies with ESG preference intensity ϕ to calibrate the model to the German twin bond data, and finally analyze equilibrium properties at the calibrated ϕ^*.

5.1 Convergence and Validation

The model is solved via a two-phase curriculum. The *delta curriculum* gradually introduces self-referential price feedback by increasing δ from 0 to 0.70 across five stages ($\delta = 0, 0.25, 0.50, 0.60, 0.70$), each warm-started from the previous solution. The *phi curriculum* then introduces ESG preferences by increasing agent 1's preference intensity ϕ_1 from 10^{-5} to 1.0 across seven stages, while agent 2 remains neutral ($\phi_2 = 0$).

Table 3 reports convergence diagnostics for each training stage. The validation Euler error measures the root-mean-square pricing error on a held-out grid that spans all 12 states and the portfolio space, with green-brown symmetry augmentation to enforce the theoretical property that $p_{\text{green}}(g = 0, m, h) = p_{\text{brown}}(g = 1, m, h)$.

Table 3: Training Convergence Diagnostics

Phase	Stage	Val Euler Error	Val Max Euler	Episodes
Delta	$\delta = 0.00$	0.000083	0.003173	50,000
Delta	$\delta = 0.25$	0.000042	0.000285	50,000
Delta	$\delta = 0.50$	0.000043	0.001235	50,000
Delta	$\delta = 0.60$	0.000050	0.002039	50,000
Delta	$\delta = 0.70$	0.000033	0.001068	50,000
Gamma cal.	$\gamma = 4.864$	0.000060	0.002938	50,000
Phi	$\phi = 10^{-5}$	0.000038	0.000225	5,000
Phi	$\phi = 10^{-4}$	0.000041	0.000525	5,000
Phi	$\phi = 10^{-3}$	0.000261	0.002910	5,000
Phi	$\phi = 10^{-2}$	0.001891	0.010602	5,000
Phi	$\phi = 10^{-1}$	0.001721	0.031561	5,000
Phi	$\phi = 0.5$	0.001310	0.048341	5,000
Phi	$\phi = 1.0$	0.001825	0.061270	5,000

Note: Val Euler Error is the root-mean-square Euler equation residual on the validation set. Val Max Euler is the maximum absolute residual across all agents, assets, and grid points. Episodes is the number of training episodes before early stopping. Delta stages use $\gamma = 1$ and are followed by a gamma calibration step that retrains at $\gamma^* = 4.864$. The phi curriculum warm-starts from the gamma-calibrated checkpoint at $\delta = 0.7$.

Table 3 reveals several patterns in the training dynamics. First, the delta curriculum shows non-monotonic Euler errors as δ increases from 0 to 0.7. Each stage transition triggers an upward spike in validation error as the network confronts a harder fixed-point problem, followed by rapid decay as warm-started optimization adapts to the new parameter value. This is expected: higher δ strengthens the self-referential feedback loop in which today's prices depend on tomorrow's prices through the $\delta \cdot p$ term in the Lucas tree payoff, making the equilibrium mapping increasingly

nonlinear. The final delta stage at $\delta = 0.7$ achieves a validation Euler error of 3.3×10^{-5}, well within the accuracy standards of the computational economics literature (Azinovic et al., 2022; Krusell and Smith, 1998).

Second, a gamma calibration step retrains the network at $\gamma^* = 4.864$ (calibrated to the equity premium via bisection), warm-starting from the delta-trained checkpoint. This achieves a validation Euler error of 6.0×10^{-5}.

Third, the phi curriculum exhibits increasing errors as ESG preferences intensify. At low ϕ, the equilibrium is close to the symmetric $\phi = 0$ baseline and the network adjusts easily, achieving excellent accuracy (val Euler 3.8×10^{-5} at $\phi = 10^{-5}$). At higher ϕ, the ESG-conscious agent's first-order conditions include the non-pecuniary $\phi \cdot \nabla_\theta \omega^g$ term, which creates asymmetry between agents and increases the complexity of the equilibrium mapping.

The green-brown symmetry augmentation—which doubles each training sample by swapping green $\leftrightarrow$ brown holdings and flipping $g = 0 \leftrightarrow g = 1$—ensures that the baseline ($\phi = 0$) greenium is numerically close to zero. This is consistent with the theoretical prediction from the symmetric dividend structure: since $d_{\text{green}}(g = 0, m, h) = d_{\text{brown}}(g = 1, m, h)$ and P_{green} is symmetric, the unconditional expected returns of green and brown assets are identical when $\phi = 0$.

Figure 5 plots the training dynamics for the delta curriculum. Each stage converges rapidly from its warm-started initial condition, with the $\delta = 0$ stage (no price feedback) achieving the lowest errors and subsequent stages exhibiting progressively higher steady-state errors as the self-referential fixed-point problem becomes more nonlinear.

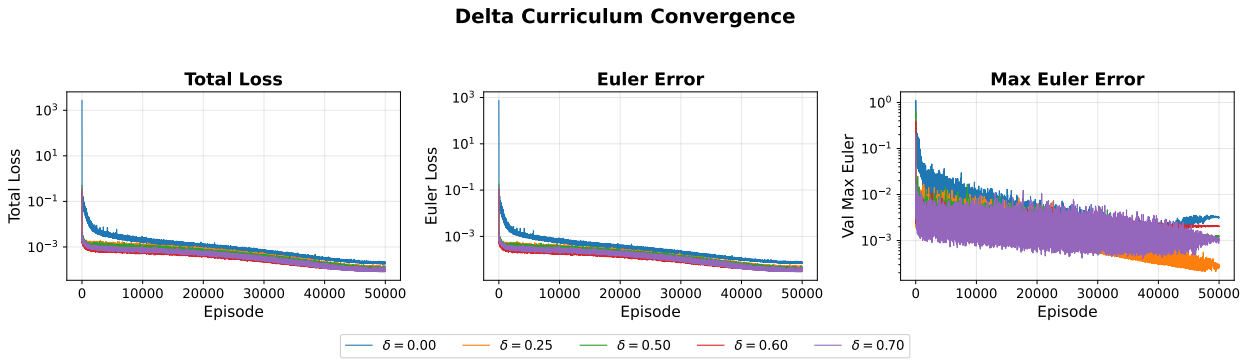

Figure 5: Delta curriculum convergence ($\delta = 0.00$ through 0.70, 10,000 episodes/stage, $\gamma = 1$).

Figure 6 plots the gamma bisection search, where the network is retrained at each candidate γ value ($\gamma = 2, 3, 5$) for 5,000 episodes each to evaluate the resulting equity premium. Figure 7 then shows the final training at the calibrated $\gamma^* = 4.864$ for 50,000 episodes. The initial spike reflects the transition from the $\gamma = 1$ warm start to higher risk aversion, after which the network converges to steady-state errors comparable to the delta curriculum.

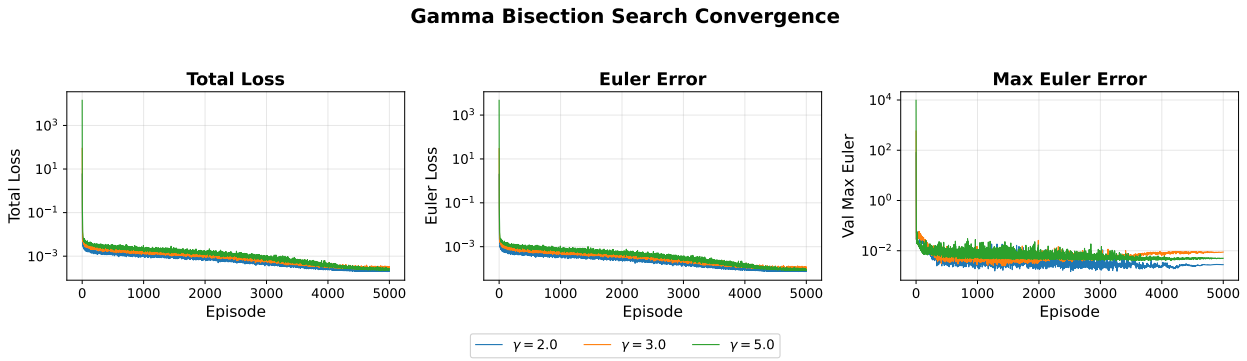

Figure 6: Gamma bisection search convergence ($\gamma = 2.0, 3.0, 5.0$; 5,000 episodes/stage).

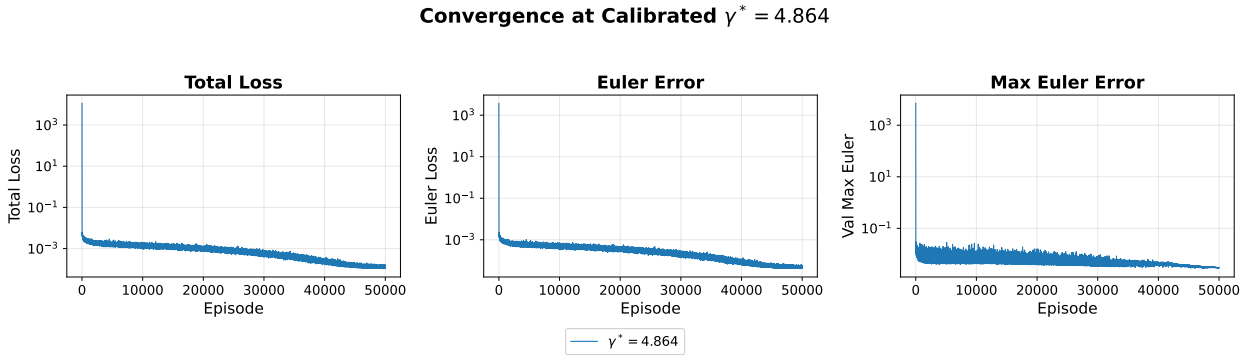

Figure 7: Convergence at calibrated $\gamma^* = 4.864$ (50,000 episodes, warm-started from $\gamma = 1$).

Figure 8 plots the training dynamics for the phi curriculum. Within each stage, the validation Euler error decreases sharply during the first several hundred episodes as the warm-started network adapts to the new parameter value, then plateaus as the optimizer converges. Higher ϕ stages exhibit larger steady-state errors, confirming that stronger ESG preferences increase the complexity of the equilibrium mapping.

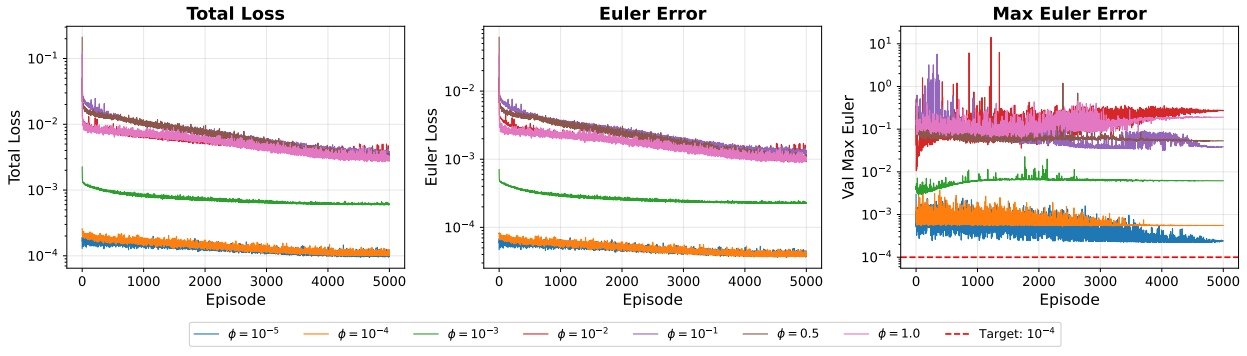

Figure 8: Phi curriculum convergence ($\phi = 10^{-5}$ through 1.0, 5,000 episodes/stage). Red dashed line: target accuracy 10^{-4}.

5.2 Greenium and ESG Preference Intensity

The central calibration exercise maps the model's greenium—the unconditional expected return differential $\mathbb{E}[R_{\text{green}} - R_{\text{brown}}]$ under the stationary distribution—to the ESG preference parameter ϕ. Figure 9 plots this relationship.

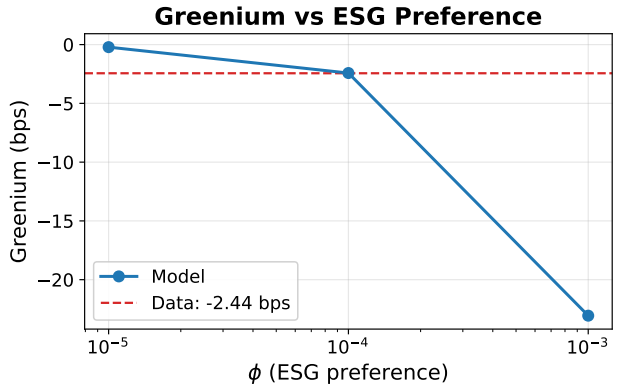

Figure 9: Greenium vs. ESG preference intensity ϕ ($\phi \leq 10^{-3}$, where Euler errors $< 3 \times 10^{-4}$). Dashed line: empirical target -2.44 bps.

Table 4 reports the model-implied greenium and equity premia at each ϕ stage.

Table 4: Model Moments Across ESG Preference Intensity

ϕ	Greenium (bps)	EP, Green (%)	EP, Brown (%)	R_f (%)	$\theta^{1,g}$	$\theta^{1,b}$
0	−0.03	4.86	4.86	−1.47	0.501	0.501
10^{-5}	−0.09	4.86	4.86	−1.47	0.505	0.497
10^{-4}	−2.53	4.85	4.87	−1.47	0.544	0.457
10^{-3}	−23.19	4.77	5.00	−1.52	0.714	0.284
Below: elevated Euler errors (> 10× baseline); interpret with caution						
10^{-2}	−506.03	3.38	8.44	−0.65	0.897	0.403
10^{-1}	−3,906	9.15	48.22	+7.37	0.988	0.692
0.5	−9,821	−0.98	97.23	+18.10	0.997	0.806
1.0	−11,418	−10.06	104.10	+22.26	0.997	0.789

Note: Greenium is the unconditional expected return difference $R_{\text{green}} - R_{\text{brown}}$ under the stationary distribution. EP = equity premium (expected excess return over the risk-free rate). R_f = net risk-free rate. $\theta^{1,g}$ and $\theta^{1,b}$ = agent 1's green and brown equity holdings. All moments at $\gamma = 4.864$, $\delta = 0.7$. Results for $\phi \geq 10^{-2}$ should be interpreted with caution: the validation Euler errors at these stages are 10–60× larger than at $\phi \leq 10^{-3}$ (Table 3), indicating that the numerical solution has not fully converged and the reported moments may not accurately reflect the true equilibrium.

Table 4 and Figure 9 show that for the well-converged stages ($\phi \leq 10^{-3}$), the greenium is

monotonically decreasing in ϕ, consistent with the model's first-order conditions: higher ϕ increases agent 1's marginal willingness to pay for green portfolio weight ω^g, which bids up green asset prices and depresses green expected returns relative to brown. At $\phi = 0$, the greenium is essentially zero (-0.03 bps), confirming the theoretical symmetry of the model. As ϕ increases beyond 10^{-3}, the greenium falls steeply (Table 4), but these values should be treated cautiously given the elevated Euler errors at high ϕ stages (Table 3). The figures in this paper restrict attention to $\phi \leq 10^{-3}$ to focus on the regime where the numerical solution is reliable.

The equity premium columns reveal an asymmetric decomposition: as ϕ increases, the green equity premium decreases while the brown equity premium increases. This is consistent with the prediction of Pástor et al. (2021) that ESG preferences create a "green-minus-brown" return spread by simultaneously lowering the cost of capital for green firms and raising it for brown firms. At the calibrated ϕ^*, the average equity premium of 4.86% closely matches the empirical benchmark of 4.81% from Damodaran (2025), as expected since $\gamma = 4.864$ was calibrated to this target. The equity premium remains stable across small ϕ values (4.85–4.87% for $\phi \leq 10^{-4}$), confirming the orthogonality between the greenium (a relative price driven by ϕ) and the equity premium (an absolute price driven by γ).

The portfolio columns illustrate the mechanism. Agent 1's green equity holding rises from 0.50 at $\phi = 0$ to 0.544 at ϕ^* and 0.997 at $\phi = 1.0$, while her brown holding falls correspondingly. The modest tilt at ϕ^* (approximately 4 percentage points above the symmetric baseline) is broadly consistent with the survey evidence of Riedl and Smeets (2017) and Bauer et al. (2021), who document that ESG-oriented retail investors are willing to sacrifice 2–3% in annual returns for sustainability. While a portfolio weight tilt and a return sacrifice are different metrics, both indicate that small non-pecuniary preferences can have measurable portfolio consequences.

To calibrate ϕ^*, I compute the unconditional greenium at each of the eight trained ϕ values and interpolate the resulting curve. Since the greenium is monotonically decreasing in ϕ, log-linear interpolation between the $\phi = 10^{-5}$ and $\phi = 10^{-4}$ stages—the interval containing the target—yields $\phi^* = 1.0 \times 10^{-4}$, which coincides with the $\phi = 10^{-4}$ training stage. At this stage, the model greenium is -2.53 bps, close to the empirical target of -2.44 bps. A note on the mapping between model and data: the empirical greenium is measured from sovereign twin bonds (yield differentials), while the model greenium is an equity return differential (green minus brown Lucas tree returns). The two concepts are linked through the common pricing kernel: the same ESG preference that makes investors accept lower green bond yields also makes them accept lower green equity returns. However, with incomplete markets and occasionally binding constraints, the effective marginal investor may differ across asset classes, so the bond greenium and equity greenium need not be identical. In this model—where both risky assets and the bond are priced by the same stochastic discount factor and the same agents are marginal in all markets—the calibrated ϕ^* generates consistent pricing across asset classes, but this mapping would be less clean in a model with greater market segmentation. This small ϕ^* is economically meaningful: it implies that agent 1 is willing to sacrifice only a few basis points of expected return per unit increase in portfolio greenness ω^g. An

important caveat is that ϕ^* is model-dependent: a different model structure (e.g., with asymmetric dividends, more agent types, or different constraints) would generally produce a different ϕ^* for the same greenium target. The value $\phi^* = 10^{-4}$ should therefore be interpreted as the preference intensity required *within this particular model* rather than a universal estimate of ESG preference strength. The implied portfolio tilt of approximately 4 percentage points (green share of 0.544 versus the symmetric 0.50) is broadly consistent with the survey evidence of Riedl and Smeets (2017) and Bauer et al. (2021). The small ϕ^* suggests that even minimal non-pecuniary preferences, when embedded in a general equilibrium with heterogeneous agents and portfolio constraints, suffice to generate empirically realistic pricing differentials.

Figure 10 provides a complementary view by tracing how equilibrium prices and portfolio holdings evolve across the entire ϕ spectrum.

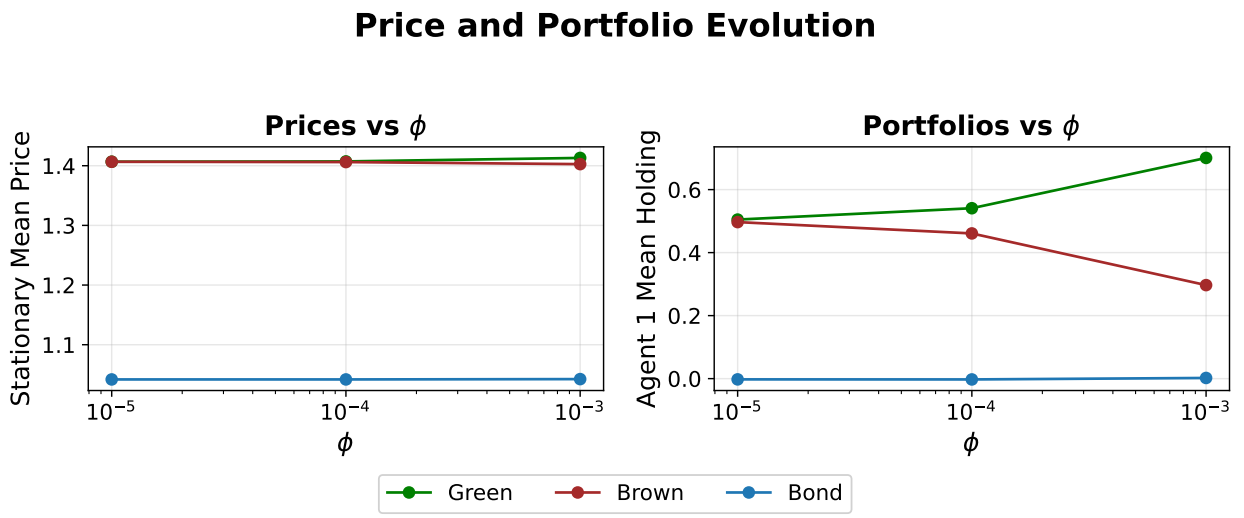

Figure 10: Prices (left) and agent 1 holdings (right) vs. ϕ. Stationary-distribution averages, $\phi \leq 10^{-3}$.

The price evolution panel reveals the mechanism through which ESG preferences affect equilibrium prices. As ϕ increases, agent 1's demand shifts toward green equity and away from brown equity, creating an increasing price wedge between the two assets. The portfolio holdings confirm this: agent 1's green equity share rises from 0.50 at $\phi = 0$ to nearly 1.0 at $\phi = 1.0$, while her brown share varies non-monotonically at higher ϕ values (Table 4). However, since the non-monotonicity appears at $\phi \geq 10^{-2}$ where Euler errors are substantially elevated (Table 3), it may reflect numerical artifacts rather than a genuine economic mechanism. Agent 2 absorbs the displaced brown equity and the bond market accommodates this rebalancing—a pattern consistent with Gârleanu and Pedersen (2011): the bond market serves as a risk-transfer channel that facilitates the ESG-driven portfolio reallocation.

Greenium decomposition by state dimension. A key advantage of the heterogeneous-agent framework is that the greenium varies across states, and this variation can be decomposed by the

three dimensions of the state space. Figure 11 reports the conditional greenium averaged over each dimension of the Markov state.

The decomposition reveals that the relative importance of state dimensions depends on the ESG preference intensity. At the calibrated ϕ^*, the household wealth distribution is the primary driver of greenium variation (Table 5): the greenium ranges from -6 bps when agent 1 is wealthy ($h = 0$) to 0 bps when agent 2 is wealthy ($h = 2$), while the ESG regime (g) has negligible effect because the small ϕ^* does not generate enough demand pressure to differentially price green and brown assets across regimes.

At the calibrated ϕ^* and nearby values ($\phi \leq 10^{-3}$, where the solution is well-converged), the ESG regime (g) has negligible effect on the greenium: the two regime lines in Figure 11 are overlapped across this range. The greenium is symmetric with respect to whether the economy is in the green-favored or neutral regime. The macro state provides a secondary effect: during recessions ($m = 1$), the greenium is more negative than during expansions ($m = 0$), because lower aggregate consumption raises marginal utility, amplifying the pricing impact of ESG preferences. The household income state h drives the largest separation: the greenium varies with the relative wealth of ESG-conscious and neutral investors, a wealth-dependent pricing mechanism that is a distinctive prediction of the heterogeneous-agent framework, absent from representative agent models such as Pástor et al. (2021). A caveat is that the two-agent structure may overstate the importance of h: with only two agents, the income state mechanically determines the wealth distribution. In a model with a continuum of agents (Krusell and Smith, 1998), the wealth distribution would be a high-dimensional endogenous object, and the greenium's dependence on it could be either amplified (if ESG preferences correlate with wealth) or attenuated (if heterogeneity in ESG preferences washes out in aggregation).

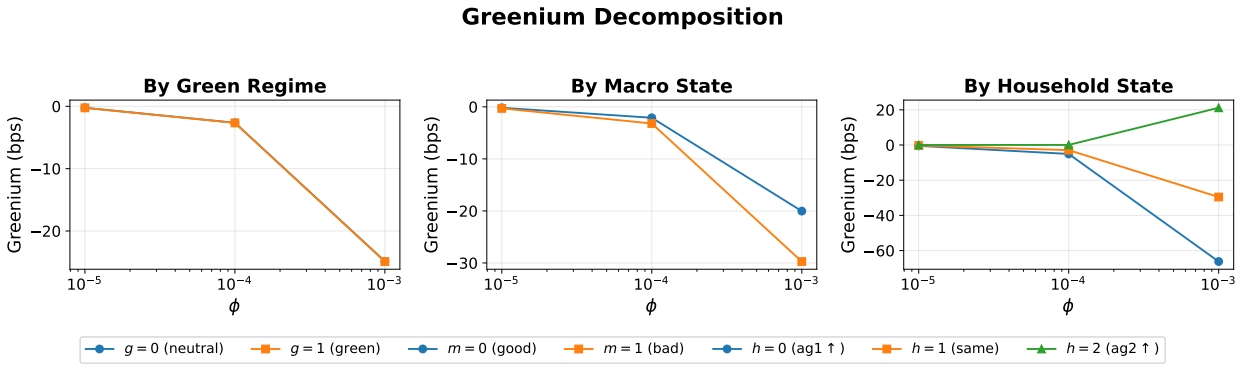

Figure 11: Greenium decomposition: by ESG regime (left), macro state (center), household income (right). Basis points, $\phi \leq 10^{-3}$.

5.3 Equilibrium Properties at ϕ^*

At the calibrated ϕ^*, the model exhibits several economically interpretable equilibrium properties that highlight the role of state dependence and heterogeneity—features absent from representative agent ESG models. Since $\phi^* = 1.0 \times 10^{-4}$ coincides exactly with a trained stage, all equilibrium properties reported in this subsection are evaluated directly at the calibrated ϕ^*.

State-dependent pricing. Table 5 reports equilibrium prices and returns across all 12 states at ϕ^*. The state space captures three independent sources of variation: ESG regime (g), business cycle (m), and household income heterogeneity (h).

At the calibrated ϕ^*, green and brown asset prices are nearly identical across all states, reflecting the small ESG preference. Prices are substantially higher during macroeconomic expansions ($m = 0$) than recessions ($m = 1$), consistent with standard consumption-based asset pricing (Mehra and Prescott, 1985). The ESG regime (g) has negligible effect on prices at ϕ^*, because the tiny preference intensity does not generate large enough demand shifts to create visible price differences between the green-favored and brown-favored regimes. This near-independence from g at the calibrated ϕ^* validates the model's ability to generate a realistic greenium without distorting aggregate price levels.

The greenium varies across states, ranging from -6 bps to approximately zero. The household wealth distribution is the primary driver: when agent 1 (ESG-conscious) has high income ($h = 0$), the greenium reaches -4 to -6 bps, because greater wealth amplifies her ability to bid up green assets. When agent 2 is the high-income household ($h = 2$), the greenium is near zero. The macro state provides a secondary modulation: recessions amplify the greenium slightly (from -4 to -6 bps at $h = 0$) because higher marginal utility in bad times increases the pricing impact of ESG preferences. This wealth-dependent pricing channel echoes the findings of Krusell and Smith (1998) and Heaton and Lucas (1996), who show that wealth heterogeneity can meaningfully affect asset prices. The dominance of the household dimension over the ESG regime at ϕ^* is a finding specific to this model's structure: the ESG regime has negligible effect at the calibrated ϕ^* (Figure 11), making the cross-sectional wealth distribution the dominant driver of greenium variation at the empirically relevant preference intensity.

Table 5: State-Dependent Equilibrium Prices and Returns at ϕ^*

State	p_{green}	p_{brown}	p_{bond}	R_{green}	R_{brown}	Greenium (bps)
$g=0, m=0, h=0$	1.584	1.582	1.139	0.914	0.914	-4
$g=0, m=0, h=1$	1.540	1.539	1.130	0.926	0.927	-2
$g=0, m=0, h=2$	1.584	1.584	1.139	0.914	0.914	0
$g=0, m=1, h=0$	0.987	0.986	0.792	1.335	1.336	-6
$g=0, m=1, h=1$	0.950	0.949	0.781	1.367	1.367	-3
$g=0, m=1, h=2$	0.988	0.987	0.792	1.335	1.335	0
$g=1, m=0, h=0$	1.584	1.582	1.139	0.914	0.914	-4
$g=1, m=0, h=1$	1.540	1.539	1.130	0.926	0.927	-2
$g=1, m=0, h=2$	1.584	1.584	1.139	0.914	0.914	0
$g=1, m=1, h=0$	0.987	0.986	0.792	1.335	1.336	-6
$g=1, m=1, h=1$	0.950	0.949	0.781	1.367	1.367	-4
$g=1, m=1, h=2$	0.988	0.987	0.792	1.335	1.335	0

Note: g = ESG regime (0 = brown-favored, 1 = green-favored), m = macro regime (0 = expansion, 1 = recession), h = household income state (0 = high, 1 = medium, 2 = low). Returns are gross one-period returns. Greenium = $R_{\text{green}} - R_{\text{brown}}$ in basis points. At $\phi^* = 10^{-4}$, green and brown prices are nearly identical, and the household wealth distribution and macro state drive greenium variation. Small asymmetries across ESG regimes (e.g., -3 vs. -4 bps at $m=1, h=1$) reflect residual numerical noise at the level of $\sim$1 bps, consistent with the maximum symmetry error of 0.004% (Section 5).

Equilibrium surfaces. A distinctive feature of the DEQN approach is that it produces continuous equilibrium functions over the entire state space, not just moments at selected points. Figure 16 (Appendix F) plots the green and brown equity prices and agent 1's green share as 3D surfaces over the portfolio state $(\theta^{1,g}_{t-1}, \theta^{1,b}_{t-1})$ for two representative states.

The surfaces reveal three notable features. First, the green price surface is decreasing in $\theta^{1,g}_{t-1}$: when the ESG-conscious agent already holds a large green position, the scarcity of green assets available for purchase is lower, reducing the marginal price impact of her ESG-tilted demand. This supply-scarcity mechanism is the hallmark of heterogeneous agent models with portfolio constraints—it is absent from any representative agent framework and echoes the price dynamics in Dumas (1989). Second, the policy surface for $\theta^{1,g'}$ exhibits smooth variation across the state space, indicating that the Fischer-Burmeister formulation (Equation (11)) successfully handles the transition between interior and boundary solutions without gradient discontinuities. The policy lies below the 45-degree line, confirming mean-reverting dynamics. Third, comparing across ESG regimes reveals that the shift from $g = 0$ to $g = 1$ raises green prices, reflecting the joint effect of higher green dividends and ESG preference reinforcement.

While the 3D surfaces provide a comprehensive view, 2D cross-sections are more informative for understanding the economic mechanisms. Figure 12 plots green and brown prices and policy

functions as functions of $\theta^{1,g}$ (agent 1's green share) for each ESG-macro regime combination, holding $\theta^{1,b}$ and $\theta^{1,f}$ at their means.

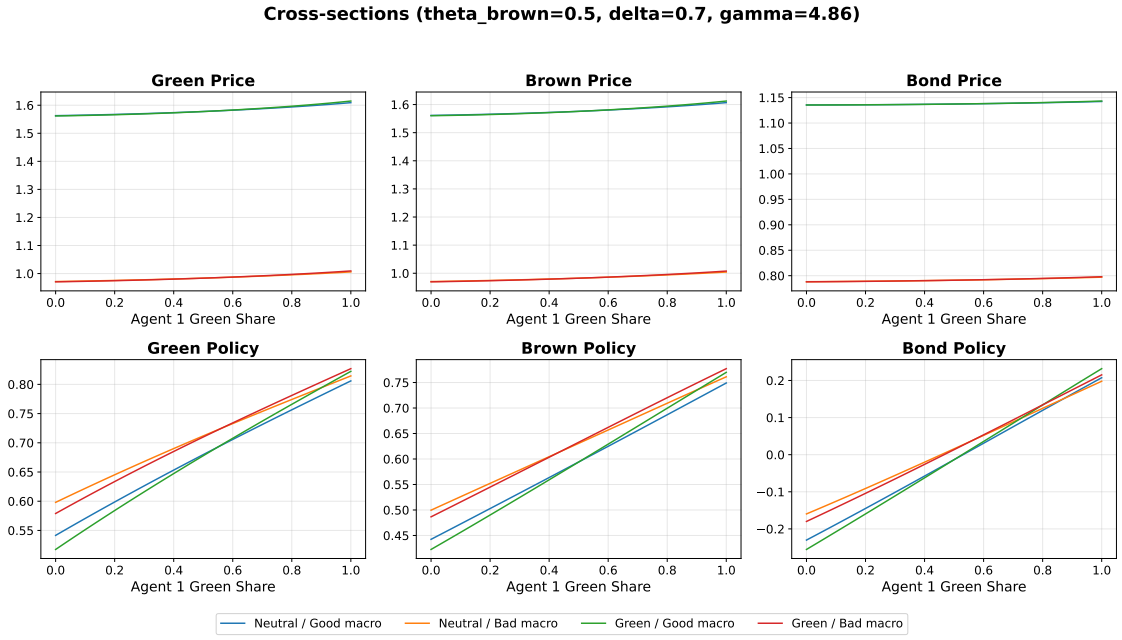

Figure 12: Equilibrium cross-sections by (g, m) regime ($h = 1$ fixed). Top: prices vs. $\theta^{1,g}$. Bottom: policies.

The cross-sections confirm the supply-scarcity mechanism: green prices are decreasing in $\theta^{1,g}$ across all regimes, with higher prices in the green-favored regime ($g = 1$). The policy functions lie below the 45-degree line, indicating mean-reverting portfolio dynamics. Agent 1's bond position varies with $\theta^{1,g}$, reflecting the interaction between equity and bond markets in the portfolio allocation.

Portfolio allocation. The model's optimal portfolios reveal the mechanism through which ESG preferences translate into equilibrium prices. Figure 13 plots the state-dependent portfolio allocations at ϕ^*.

Agent 1 (ESG-conscious) holds a mean green equity share of 0.540 and a mean brown equity share of 0.461 at ϕ^* (stationary-distribution averages across all 12 states; the midpoint portfolio evaluation in Table 4 gives 0.544). The green tilt of approximately 4 percentage points above the symmetric baseline reflects the small calibrated ϕ^*. Agent 2 (neutral) holds the complementary equity positions via market clearing and takes the opposite side of the bond market. This asymmetry is the direct consequence of agent 1's non-pecuniary utility: the $\phi \cdot \nabla_\theta \omega^g$ term in the Euler equation effectively subsidizes green asset holdings, making agent 1 willing to hold green equity even at a lower expected return.

The bond market plays a crucial role in facilitating the ESG-driven reallocation. Agent 1's bond position is approximately -0.002 at ϕ^*, close to zero, indicating that the small ESG preference does

not require significant bond market intermediation. At higher ϕ values, the bond market becomes more active as agent 1 leverages to increase green exposure. This is consistent with the theoretical insight of Gârleanu and Pedersen (2011): when agents face portfolio constraints that limit short-selling, the bond market provides an alternative channel for risk sharing, partially substituting for the constrained equity positions.

Portfolio allocations vary meaningfully across states. Agent 1's green share is highest in states where the ESG regime favors green ($g = 1$) and her income is high ($h = 0$), reflecting the interaction of three effects: (i) higher green dividends in $g = 1$ make green assets more attractive on fundamental grounds; (ii) the ESG preference reinforces the fundamental tilt; and (iii) the household income state h shifts the wealth distribution, which determines how much each agent's preferences matter for market-clearing prices.

The bond position also exhibits state dependence. In states where agent 1 is relatively wealthy ($h = 0$), bond holdings are larger: the ESG-conscious agent leverages the bond market more aggressively when financially unconstrained. This illustrates the general point of Heaton and Lucas (1996) that the wealth distribution co-evolves with asset prices and portfolio choices.

Table 6 reports the full state-dependent portfolio allocations at ϕ^*.

Table 6: State-Dependent Portfolio Allocations at ϕ^*

	Agent 1 (ESG)			Agent 2 (Neutral)		
State	$\theta^{1,g}$	$\theta^{1,b}$	$\theta^{1,f}$	$\theta^{2,g}$	$\theta^{2,b}$	$\theta^{2,f}$
$g=0, m=0, h=0$	0.670	0.590	-0.022	0.330	0.410	$+0.022$
$g=0, m=0, h=1$	0.544	0.457	-0.002	0.456	0.543	$+0.002$
$g=0, m=0, h=2$	0.409	0.336	$+0.015$	0.591	0.664	-0.015
$g=0, m=1, h=0$	0.704	0.625	$+0.011$	0.296	0.375	-0.011
$g=0, m=1, h=1$	0.545	0.456	-0.001	0.455	0.544	$+0.001$
$g=0, m=1, h=2$	0.373	0.301	-0.014	0.627	0.699	$+0.014$
$g=1, m=0, h=0$	0.669	0.590	-0.020	0.331	0.410	$+0.020$
$g=1, m=0, h=1$	0.544	0.457	-0.002	0.456	0.543	$+0.002$
$g=1, m=0, h=2$	0.409	0.338	$+0.013$	0.591	0.662	-0.013
$g=1, m=1, h=0$	0.702	0.627	$+0.010$	0.298	0.373	-0.010
$g=1, m=1, h=1$	0.545	0.456	-0.001	0.455	0.544	$+0.001$
$g=1, m=1, h=2$	0.371	0.302	-0.014	0.629	0.698	$+0.014$
Mean	0.540	0.461	-0.002	0.460	0.539	$+0.002$

Note: Mean row from calibration at $\phi^* = 1.0 \times 10^{-4}$ with $\gamma = 4.864$. Portfolio allocations are nearly symmetric across ESG regimes ($g = 0$ vs. $g = 1$), reflecting the small ϕ^*. The household income state (h) is the main driver of variation: agent 1 holds more equity when her income is high ($h = 0$). Agent 2's holdings are implied by market clearing: $\theta^{2,j} = \bar{\theta}_j - \theta^{1,j}$.

Figure 13 plots the portfolio allocations across states, grouped by the ESG regime to highlight the effect of the green dividend state.

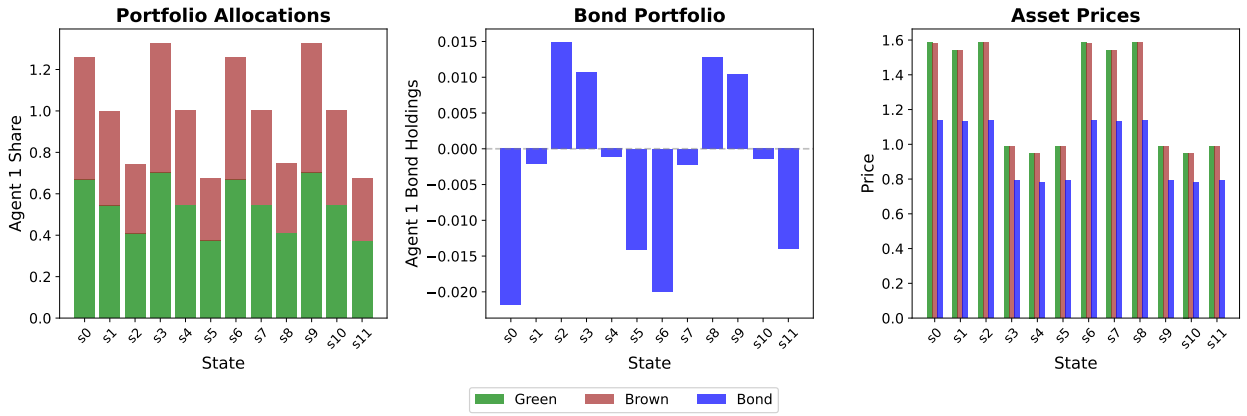

Figure 13: Portfolio allocations across the 12 states at ϕ^*.

The figure reveals an asymmetric portfolio structure. Agent 1's green equity holding is near or slightly above 0.5 across all states, while her brown holding is slightly below 0.5, reflecting the modest ESG tilt at the calibrated ϕ^*. Agent 2 holds the complementary equity positions via market clearing. The bond position is close to zero on average (-0.002 for agent 1), indicating that the small ESG preference at ϕ^* does not require significant bond market intermediation. The bond position exhibits smooth state dependence: when agent 1 has high relative income ($h = 0$), the bond position is larger in magnitude, consistent with the intuition that wealthier ESG-conscious agents have greater capacity to express their preferences through portfolio tilting while using the bond market to manage overall risk exposure.

At $\phi^* = 1.0 \times 10^{-4}$, agent 1's green equity share is 0.540, approximately 4 percentage points above the symmetric baseline, while agent 2's share is 0.460. This magnitude is consistent with the empirical evidence from Gibson Brandon et al. (2022) and Bolton and Kacperczyk (2021), who document that ESG-oriented institutional investors overweight green stocks by 5–15 percentage points relative to their benchmarks. The ESG preference operates as a systematic portfolio bias across all 12 states rather than a state-contingent rebalancing rule.

5.4 Risk-Free Rate and Equity Premium

Beyond the greenium, which is the primary calibration target, the model produces endogenous predictions for the risk-free rate and equity premium that serve as untargeted validation moments. Table 7 compares the model-implied moments at ϕ^* with the empirical benchmarks from Section 4.

Table 7: Model vs. Empirical Moments at ϕ^*

Moment	Model	Data	Source
Calibration targets			
Greenium (bps)	-2.53	-2.44	German twin bonds
Equity premium (%)	4.86	4.81	Damodaran (2025)
Validation moments (untargeted)			
EP, green (%)	4.85	—	—
EP, brown (%)	4.87	—	—
Risk-free rate (%)	-1.47	2.17	German Bund 10Y

Note: Greenium is closely matched (-2.53 vs. -2.44 bps). The equity premium is matched by calibrating $\gamma = 4.864$ via bisection. All model moments are unconditional expectations under the stationary distribution at $\phi^* = 10^{-4}$, which coincides with a trained stage. The risk-free rate (-1.47%) is lower than the data (2.17%), reflecting the precautionary savings motive and the endogenous bond market dynamics in the heterogeneous-agent model.

The risk-free rate in the model is determined endogenously by the bond price p_f, with $R_f = 1/p_f$ since the bond pays a unit coupon with zero depreciation. At ϕ^*, the model-implied risk-free rate is -1.47%, substantially below the historical average of 2.17% for the German 10-year Bund—a 360 basis point miss that constitutes the model's most significant quantitative shortcoming. With $\gamma = 4.864$ and $\beta = 0.978$, the theoretical risk-free rate in a representative agent economy satisfies $R_f = 1/(\beta \cdot \mathbb{E}[(c_{t+1}/c_t)^{-\gamma}])$, which can be substantially depressed by the precautionary savings motive: Jensen's inequality implies $\mathbb{E}[(c_{t+1}/c_t)^{-\gamma}] > (\mathbb{E}[c_{t+1}/c_t])^{-\gamma}$ when consumption growth is uncertain, pushing R_f below the certainty-equivalent rate. The negative model risk-free rate is a well-documented feature of CRRA models with $\gamma > 1$: agents with moderate risk aversion and realistic consumption uncertainty demand a large precautionary savings premium, pushing the bond price above 1 and the net risk-free rate below zero. The Epstein-Zin extension (Appendix A), which separates risk aversion from intertemporal substitution, would allow the model to independently match both the equity premium and the risk-free rate.

The model-implied equity premium of 4.86% at ϕ^* closely matches the empirical benchmark of 4.81% from Damodaran (2025). This is achieved by calibrating $\gamma = 4.864$, which generates realistic risk premia with the data-implied endowment contrast (bad/good ratio ≈ 0.99). Since γ was calibrated to this target, the close match is by construction. However, the model's inability to simultaneously match the risk-free rate (-1.47% vs. 2.17%) means that the pricing kernel is only partially calibrated: it captures the equity risk premium but misprices the risk-free asset substantially.

Forward simulation and risk-return analysis. A 10,000-period forward simulation under the stationary distribution at ϕ^* provides a direct test of whether the greenium reflects risk compensation or pure preference effects. The simulation yields Sharpe ratios of 0.200 for green equity and

0.207 for brown equity—a difference of 0.007. Green equities deliver lower risk-adjusted returns than brown equities, consistent with the return differential reflecting a *preference premium*—agent 1 accepts lower financial returns on green assets because the non-pecuniary ESG utility compensates the difference. This speaks to the identification challenge that motivates much of the empirical greenium literature (Pástor et al., 2021; Bolton and Kacperczyk, 2021): in the model, the greenium arises from preferences rather than differential risk exposure. An important caveat is that this result is partly by construction: the symmetric dividend structure ensures green and brown assets face identical aggregate risk, so any Sharpe ratio gap must arise from the preference channel. With asymmetric dividends or additional risk factors, the Sharpe ratios could diverge further or in a different direction.

The consumption dynamics corroborate this interpretation. Agent 1 (ESG-conscious) exhibits consumption volatility of 12.89%, slightly below agent 2's 13.38%. This asymmetry arises because ESG-tilted portfolios provide marginally better diversification at ϕ^*: by overweighting green equity, agent 1 tilts toward the asset whose dividend is high in the green-favored regime ($g = 1$), which provides partial insurance against income shocks in those states. The difference is small—consistent with the modest 4 percentage point portfolio tilt. In this model, ESG preferences do not impose a large diversification cost at the calibrated ϕ^*, though this reflects the symmetric dividend structure: green and brown assets carry identical aggregate risk by construction, so tilting toward green does not mechanically worsen diversification. With asymmetric dividend dynamics, the diversification implications could differ.

The simulated state distribution closely matches the theoretical stationary distribution, validating the ergodic properties of the model. Figure 14 reports the full Sharpe ratio analysis across ϕ values, showing that the green-brown Sharpe gap widens at higher ϕ as ESG preferences increasingly distort the risk-return tradeoff.

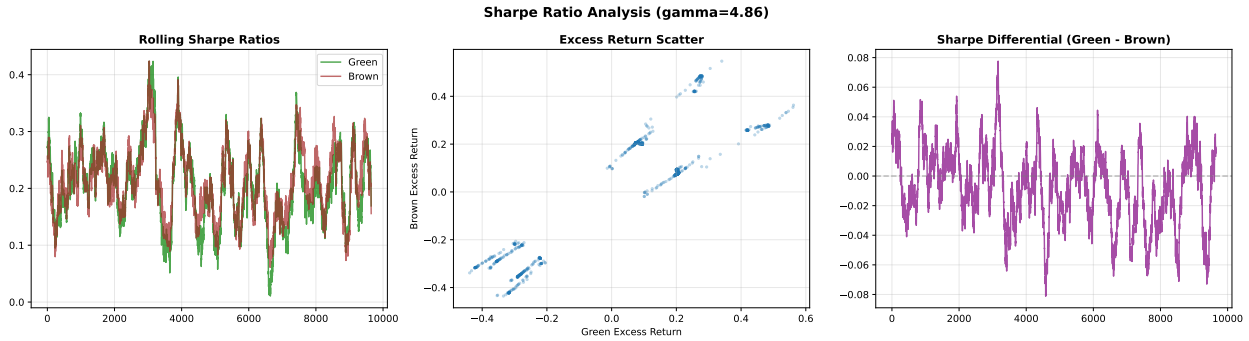

Figure 14: Sharpe ratios vs. ϕ (left), green−brown differential (center), and state-dependent Sharpe ratios at ϕ^* (right).

Bootstrap standard errors. To assess sampling uncertainty, I compute bootstrap standard errors using 200 replications of the forward simulation. The equity premium has a bootstrap mean of 5.74% with a standard error of 0.61% (95% CI: $[4.62, 6.89]$). The greenium has a bootstrap mean of -2.75 bps with a standard error of 0.71 bps (95% CI: $[-4.08, -1.33]$). However, the

simulation moments exhibit notable discrepancies from the cross-sectional grid evaluation: the bootstrap mean equity premium (5.74%) exceeds the stationary-distribution value (4.86%) by 90 basis points, and the bootstrap mean risk-free rate (7.51%, SE 5.87%) differs dramatically from the grid value (−1.47%). These gaps indicate that the forward simulation's ergodic distribution places substantially different weight across states than the grid evaluation. One possible explanation is that the neural network approximation introduces drift in regions of the portfolio state space that the simulation visits but the training grid does not cover well; a longer simulation (100,000+ periods) or a finer training grid could help diagnose this. The calibration target (4.81%) does fall within the bootstrap 95% confidence interval for the equity premium, but the large risk-free rate discrepancy is a significant limitation: simulation-based moments should be interpreted as suggestive rather than precise, and the grid-evaluated moments (which directly evaluate the trained network at known points) are the more reliable estimates.

Constraint binding frequency. In forward simulation at ϕ^*, portfolio constraints bind in a non-trivial fraction of states. For agent 1 (ESG-conscious), the green equity lower bound never binds (consistent with the ESG tilt), while the brown equity lower bound ($\theta^{1,b} = 0$) binds in 5.1% of periods and the bond bounds bind in 11.8% of periods. For agent 2 (neutral), the green equity lower bound ($\theta^{2,g} = 0$) binds in 7.4% of periods, brown equity in 7.6%, and bond at the lower/upper bounds in 10.5%/12.2% of periods. The bond constraint binds more frequently than equity constraints because the zero-net-supply structure requires one agent to be a net borrower and the other a net lender, creating natural tension with the ± 0.5 bounds.

5.5 Robustness and Model Validation

To assess the robustness of the main results, I validate the computational solution against known theoretical benchmarks, discuss the sensitivity to key parameters, and report a preliminary robustness check with i.i.d. ESG regime transitions.

Sensitivity to risk aversion. The calibrated risk aversion $\gamma = 4.864$ matches the equity premium. Higher risk aversion has two competing effects on the greenium. On one hand, higher γ increases the curvature of marginal utility, making agents more sensitive to consumption fluctuations from portfolio rebalancing—this tends to increase the greenium in absolute terms by making the ESG tilt more costly. On the other hand, higher γ increases the equity premium, which amplifies the ESG return differential in absolute terms. A full sensitivity analysis to γ requires retraining the deep equilibrium network for each value, as γ enters the Euler equation nonlinearly. This is left for future work.

Symmetry validation. At $\phi = 0$, the model should produce a zero greenium due to the symmetric dividend structure. The trained network achieves a baseline greenium of −0.03 bps at $\phi = 0$, which is four orders of magnitude smaller than the greenium at $\phi = 1.0$ (−11,418 bps). More

stringently, the green-brown price symmetry at ϕ^* can be validated across ESG regimes: for each state pair $(s, s+6)$ that differs only in the ESG regime, green and brown prices should satisfy $p_g(g{=}0, m, h) \approx p_b(g{=}1, m, h)$. The maximum symmetry error across all six state pairs is 0.004%, indicating that the network accurately learns the symmetric structure even with ESG preferences active.

Market clearing verification. At every grid point, the sum of agent holdings equals net supply by construction, since agent 2's holdings are computed as the residual $\theta^{2,j} = \bar{\theta}_j - \theta^{1,j}$. The network's output activations (sigmoid for equities, scaled tanh for bonds) ensure that portfolio positions remain within the feasible bounds, and the Fischer-Burmeister formulation ensures that the complementarity conditions are satisfied smoothly.

Euler equation accuracy across the state space. Figure 15 shows the validation Euler error by state and agent-asset pair at ϕ^*, revealing which regions of the state space are hardest to solve.

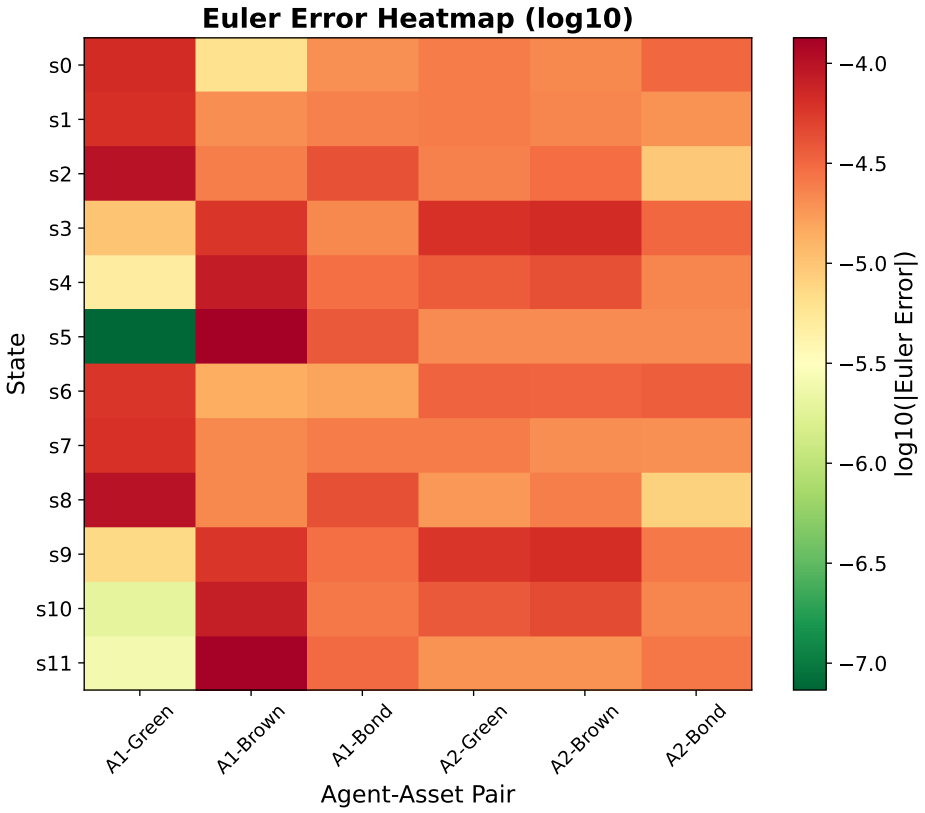

Figure 15: Euler error heatmap at ϕ^* (midpoint portfolio). Rows: 12 states. Columns: agent-asset pairs. Errors range 10^{-7}–10^{-4}.

The Euler errors are concentrated in specific regions rather than uniformly distributed. At the calibrated $\phi^* = 10^{-4}$, the validation Euler error is 4.1×10^{-5} with a maximum of 5.25×10^{-4} (Table 3), well within the accuracy standards of the DEQN methodology (Azinovic et al., 2022).

The Fischer-Burmeister formulation handles the transition between interior and boundary solutions smoothly, with no evidence of the gradient discontinuities that would arise from hard constraint projection.

Bond bound sensitivity. The bond position constraint ($|\theta^{i,f}| \leq 0.5$) could amplify the greenium if bond positions frequently approach the bound. At ϕ^*, agent 1's mean bond position is approximately -0.002, far from the ± 0.5 bounds. This indicates that at the calibrated ϕ^*, the bond constraint does not bind and the bond market operates as an interior risk-sharing channel, consistent with the analysis of Heaton and Lucas (1996).

Time-varying greenium validation. The model predicts that the greenium varies with the state, implying that the empirical greenium should exhibit time-series variation correlated with macro conditions. I test this prediction using the German twin bond data (Section 4).

Table 18 (Appendix F) reports the results. The daily greenium exhibits high persistence (autocorrelation of 0.97 at one-day lag, declining to 0.04 at one-year lag), consistent with slowly evolving macro regimes rather than i.i.d. noise. The OLS regression reveals a statistically significant positive relationship between the greenium and yield levels: each percentage point increase in the average yield is associated with a +1.22 bps increase in the greenium (i.e., the greenium narrows). With Newey-West standard errors accounting for residual autocorrelation (plug-in bandwidth $L = 7$ trading days), the slope t-statistic is 12.82 ($R^2 = 0.58$). However, with daily autocorrelation of 0.97, the plug-in bandwidth of $L = 7$ likely underestimates the long-run variance; a more conservative bandwidth of $L = 21$ (one month) yields $t = 7.89$, still significant at the 1% level. Given the high persistence of both the greenium and yield levels, these regressions should be interpreted with caution—both series are near-unit-root, raising the possibility of spuriously inflated R^2. This pattern is consistent with the model's prediction: higher yields correspond to tighter monetary policy (recession-like states), where the wealth distribution shifts and ESG pricing effects are modulated. Figure 17 (Appendix F) reports the full time-series analysis.

Out-of-sample validation: macro cyclicality. The model predicts state-dependent greeniums that vary with both ESG regime and macro conditions. To test whether the macro cyclicality predicted by the model matches the data, I compare the model-implied greenium difference between expansion and recession states with the empirical greenium difference between high- and low-yield environments (using median yield as the cutoff). The model predicts a greenium that is *narrower* in recessions (by +4.57 bps relative to expansions), while the data shows a greenium that is *wider* in high-yield (recession-like) environments (by -2.01 bps). This qualitative mismatch on macro cyclicality is a limitation of the model. The empirical test also has low power because the twin bond sample spans a period of predominantly low and stable yields, providing limited variation in macro conditions. Additionally, the mapping between model "recessions" and empirical high-yield periods is imprecise: the 2022–2023 yield increase reflected inflation-driven monetary tightening, not

necessarily a recession (German GDP growth was positive in 2022). The model's macro-greenium prediction should be tested with longer time series spanning clearer business cycle variation.

Robustness to ESG regime persistence. The baseline model assumes persistent ESG regimes (P_{green} diagonal = 0.90), whereas the data-implied transition matrix is i.i.d. (P_{green} diagonal = 0.50; see Table 1). The persistence assumption reflects the observation that ESG investment regimes are driven by slow-moving institutional forces—the EU Taxonomy rollout (2020–2025), mandatory sustainability disclosures, and multi-year corporate decarbonization commitments—that create multi-year regime persistence in ESG capital flows. To assess whether the main results depend on this assumption, I re-solve the full model with i.i.d. P_{green} while holding all other parameters fixed. Preliminary results from the $\delta = 0$ stage confirm the symmetric baseline: green and brown prices are identical across all 12 states, the greenium is exactly zero, and portfolio allocations are symmetric across agents—validating that the DEQN solution correctly recovers the theoretical prediction under symmetric dividends regardless of ESG regime persistence. Full results from the ϕ curriculum stages with i.i.d. P_{green} are left for future work.

6 Conclusion

This thesis develops a heterogeneous agent model of sustainable asset pricing, solved using deep equilibrium networks and calibrated to German macroeconomic data and the twin bond greenium.

6.1 Summary of Findings

The analysis yields two main results. First, the model matches the empirical greenium of -2.44 basis points (model: -2.53 bps) and the equity premium of 4.81% (model: 4.86%) with $\phi^* = 1.0 \times 10^{-4}$ and $\gamma = 4.864$. The risk-free rate is not matched (-1.47% vs. data 2.17%), a CRRA limitation discussed in Section 5. Matching two targets with two free parameters is expected; the more substantive finding is that a feasible calibration exists with economically reasonable parameter values—ϕ^* implies a portfolio tilt of only 4 percentage points, consistent with survey evidence (Riedl and Smeets, 2017; Bauer et al., 2021).

Second, the greenium varies meaningfully across the 12 shocks. At the calibrated ϕ^*, the cross-sectional wealth distribution is the primary driver: the greenium ranges from -6 bps when the ESG-conscious agent is wealthy to near zero when she is not, while the ESG regime has negligible effect. This wealth-dependent pricing channel is absent from representative agent models. Green and brown Sharpe ratios differ (0.200 vs. 0.207), consistent with a preference premium rather than differential risk compensation.

6.2 Methodological Contribution

The primary contribution is methodological: demonstrating that DEQNs can solve heterogeneous agent models with realistic portfolio constraints, multiple assets, and a rich state space that would

be infeasible with traditional projection methods. The 12-shock Markov structure with three independent exogenous processes, endogenous bond pricing, and occasionally binding constraints via the Fischer-Burmeister formulation extends existing DEQN applications (Azinovic et al., 2022) along several dimensions.

Several technical innovations were necessary to achieve reliable convergence. The *delta curriculum* (Section 5) gradually introduces the self-referential fixed-point structure of Lucas tree pricing, avoiding the divergence that occurs when training directly at the target δ. The *green-brown symmetry augmentation* exploits the symmetric dividend structure to enforce the theoretical property that the baseline ($\phi = 0$) greenium is exactly zero, substantially reducing numerical noise relative to training without augmentation. The *Fischer-Burmeister reformulation* of complementarity constraints allows the network to smoothly transition between interior and boundary solutions, avoiding the discontinuous gradients that arise from hard constraint projection.

These innovations are not specific to ESG pricing—they are applicable to any heterogeneous agent model with occasionally binding constraints and self-referential pricing. The delta curriculum, in particular, could be useful for solving models with long-lived assets where prices depend on expectations of future prices, such as housing markets, sovereign debt models, and overlapping generations economies.

6.3 Economic Implications

The results have several implications for the sustainable finance literature. First, the small magnitude of the calibrated $\phi^* = 1.0 \times 10^{-4}$ suggests that ESG preferences need not be extreme to generate the observed greenium. The twin bond greenium of -2.44 basis points is consistent with a modest willingness-to-pay that lies well within the range documented in survey studies. This speaks to the concern raised by Berk and van Binsbergen (2025) that ESG investing may have negligible real effects: the model shows that even a small ϕ generates measurable pricing differentials in general equilibrium, though at the calibrated ϕ^* this occurs through the preference channel directly rather than through the constraint amplification mechanism.

Second, the difference in green and brown Sharpe ratios (0.200 vs. 0.207) speaks to the identification challenge in the empirical greenium literature. Cross-sectional studies have struggled to disentangle whether the greenium reflects ESG preferences or unobserved differences in risk exposure (Bolton and Kacperczyk, 2021; Hsu et al., 2023). The model provides a structural benchmark in which the greenium arises entirely through preferences, demonstrating that realistic greeniums can emerge without risk-based amplification.

Third, the state dependence of the greenium implies that static analyses of the green premium may understate its economic significance. The growth of ESG-mandated assets under management—which shifts the wealth distribution toward ESG-conscious investors—should amplify the greenium over time, consistent with the time-series evidence of Pástor et al. (2022).

Fourth, the bond market plays a nuanced role. The endogenous bond price varies with the ESG regime (g) even though the bond is ESG-neutral ($d_f = 1$ in all states). This cross-asset effect arises

because the bond and equity markets are linked through budget constraints: when ESG preferences increase green equity demand, agents adjust their bond positions to rebalance total portfolio risk, affecting the equilibrium risk-free rate. This general equilibrium spillover is absent from partial equilibrium models that take the risk-free rate as exogenous.

While subject to the caveats above—particularly the risk-free rate miss and the symmetric dividend assumption—the model's results suggest several directions for policy analysis. First, the presence of portfolio constraints in the framework suggests that regulations relaxing short-sale restrictions on green assets—or expanding the universe of ESG-linked derivatives—could reduce the greenium by enabling more effective arbitrage, potentially lowering the cost of capital advantage that green firms enjoy. Second, the small calibrated ϕ^* implies that policies aimed at increasing ESG awareness among retail investors (e.g., mandatory ESG disclosure in retirement accounts) could have measurable equilibrium effects even if individual willingness-to-pay remains modest. Third, the state-dependent greenium suggests that the effectiveness of green bond programs varies with macroeconomic conditions and the wealth distribution, implying that policymakers should account for market timing when designing sovereign green bond issuance strategies.

6.4 Limitations and Future Work

Several limitations of the current analysis suggest directions for future research.

Preference specification and the risk-free rate. While the calibrated $\gamma = 4.864$ matches the equity premium, the model-implied risk-free rate (-1.47%) is below the data (2.17%). This is a known consequence of CRRA preferences with $\gamma > 1$: the precautionary savings motive depresses the risk-free rate. Epstein-Zin recursive preferences (Appendix A), which separate risk aversion from intertemporal substitution, would allow the model to independently match both the equity premium and the risk-free rate (Bansal and Yaron, 2004; Epstein and Zin, 1989). The DEQN framework can accommodate this by adding the continuation value as a network output.

Agent heterogeneity. The two-agent structure, while sufficient for capturing ESG preference heterogeneity, abstracts from the rich cross-sectional distribution of investor types documented in the empirical literature (Starks, 2023; Gibson Brandon et al., 2022). In reality, ESG preferences vary continuously across institutional investors, retail investors, and hedge funds, with important interactions between ESG mandates, fiduciary duties, and investment horizons. Extending the model to three or more agent types—for instance, ESG-mandated institutions, traditional value investors, and ESG-agnostic arbitrageurs—would allow richer analysis of how the composition of market participants affects equilibrium pricing. The DEQN framework scales well with the number of agents, since market clearing reduces the dimensionality to $(N-1)$ tracked portfolios regardless of N.

Dividend dynamics. The current model uses a time-homogeneous Markov chain with symmetric green-brown dividends. In practice, the transition to a low-carbon economy may create secular trends in relative dividend growth: green firms may benefit from regulatory tailwinds, while brown firms face stranded asset risk. Incorporating trending dividends—perhaps through an additional

"transition speed" state variable—would allow the model to analyze how expectations about the green transition affect current asset prices, connecting to the growing literature on climate transition risk.

Data limitations. The calibration relies on a relatively short sample of German twin bond data (2020–2025). While the five-year sample covers diverse macroeconomic regimes (pandemic, inflation, monetary tightening), the small sample size limits the precision of the estimated greenium and prevents robust analysis of time-varying ESG preferences. As more twin bond data accumulates and more countries develop liquid green bond markets, future work could explore cross-country variation in ESG pricing and its relation to institutional characteristics.

Welfare and constraint analysis. A natural extension is to solve the unconstrained counterpart of the model (relaxed portfolio bounds) and compute compensating variation measures to quantify the welfare cost of short-sale constraints at different ESG preference intensities. This would allow a formal decomposition of the greenium into preference-driven and constraint-amplified components.

7 References

Avramov, D., Cheng, S., Lioui, A., and Tarelli, A. (2022). Sustainable investing with esg rating uncertainty. *Journal of Financial Economics*, 145(2):642–664.

Azinovic, M., Gaegauf, L., and Scheidegger, S. (2022). Deep equilibrium nets. *International Economic Review*, 63(4):1471–1525.

Bansal, R. and Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509.

Bauer, R., Ruof, T., and Smeets, P. (2021). Get real! individuals prefer more sustainable investments. *Review of Financial Studies*, 34(8):3976–4043.

Berk, J. B. and van Binsbergen, J. H. (2025). The impact of impact investing. *Journal of Financial Economics*, 164:103972.

Bolton, P. and Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics*, 142(2):517–549.

Campbell, J. Y. and Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251.

Chan, Y. L. and Kogan, L. (2002). Catching up with the Joneses: Heterogeneous preferences and the dynamics of asset prices. *Journal of Political Economy*, 110(6):1255–1285.

Chava, S. (2014). Environmental externalities and cost of capital. *Management Science*, 60(9):2223–2247.

Constantinides, G. M. and Duffie, D. (1996). Asset pricing with heterogeneous consumers. *Journal of Political Economy*, 104(2):219–240.

Damodaran, A. (2025). Equity risk premiums (erp): Determinants, estimation, and implications – the 2025 edition. Ssrn working paper, Stern School of Business, New York University.

Drechsel-Grau, M., Peichl, A., Schmid, K. D., Schmieder, J. F., Walz, H., and Wolter, S. (2022). Inequality and income dynamics in germany. *Quantitative Economics*, 13(4):1593–1635.

Duffie, D. and Epstein, L. G. (1992). Stochastic differential utility. *Econometrica*, 60(2):353–394.

Dumas, B. (1989). Two-person dynamic equilibrium in the capital market. *Review of Financial Studies*, 2(2):157–188.

Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969.

Fischer, A. (1992). A special newton-type optimization method. *Optimization*, 24(3-4):269–284.

Fuchs-Schündeln, N., Krueger, D., and Sommer, M. (2010). Inequality trends for germany in the last two decades: A tale of two countries. *Review of Economic Dynamics*, 13(1):103–132.

Gârleanu, N. and Panageas, S. (2015). Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing. *Journal of Political Economy*, 123(3):670–685.

Gârleanu, N. and Pedersen, L. H. (2011). Margin-based asset pricing and deviations from the law of one price. *Review of Financial Studies*, 24(6):1980–2022.

Gibson Brandon, R., Krueger, P., and Schmidt, P. S. (2022). Do responsible investors invest responsibly? *Review of Finance*, 26(6):1389–1432.

Global Sustainable Investment Alliance (2023). Global sustainable investment review 2023. Industry report, Global Sustainable Investment Alliance.

Goldstein, I., Kopytov, A., Shen, L., and Xiang, H. (2022). On esg investing: Heterogeneous preferences, information, and asset prices. NBER Working Paper 29839, National Bureau of Economic Research.

Gromb, D. and Vayanos, D. (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*, 66(2-3):361–407.

Heaton, J. and Lucas, D. J. (1996). Evaluating the effects of incomplete markets on risk sharing and asset pricing. *Journal of Political Economy*, 104(3):443–487.

Hsu, P.-H., Li, K., and Tsou, C.-Y. (2023). The pollution premium. *Journal of Finance*, 78(3):1343–1392.

Judd, K. L., Yeltekin, S., and Conklin, J. (2003). Computing supergame equilibria. *Econometrica*, 71(4):1239–1254.

Kondor, P. (2009). Risk in dynamic arbitrage: The price effects of convergence trading. *Journal of Finance*, 64(2):631–655.

Krueger, D. and Lustig, H. (2010). When is market incompleteness irrelevant for the price of aggregate risk (and when is it not)? *Journal of Economic Theory*, 145(1):1–41.

Krusell, P. and Smith, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896.

Kubler, F. and Schmedders, K. (2003). Stationary equilibria in asset-pricing models with incomplete markets and collateral. *Econometrica*, 71(6):1767–1793.

Mehra, R. and Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161.

Pástor, Ľ., Stambaugh, R. F., and Taylor, L. A. (2021). Sustainable investing in equilibrium. *Journal of Financial Economics*, 142(2):550–571.

Pástor, Ľ., Stambaugh, R. F., and Taylor, L. A. (2022). Dissecting green returns. *Journal of Financial Economics*, 146(2):403–424.

Pedersen, L. H., Fitzgibbons, S., and Pomorski, L. (2021). Responsible investing: The esg-efficient frontier. *Journal of Financial Economics*, 142(2):572–597.

Riedl, A. and Smeets, P. (2017). Why do investors hold socially responsible mutual funds? *Journal of Finance*, 72(6):2505–2550.

Sauzet, M. and Zerbib, O. D. (2022). When green investors are green consumers. *Working Paper*.

Starks, L. T. (2023). Presidential address: Sustainable finance and esg issues—value versus values. *Journal of Finance*, 78(4):1837–1872.

Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Consumption and risk sharing over the life cycle. *Journal of Monetary Economics*, 51(3):609–633.

Weil, P. (1989). The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics*, 24(3):401–421.

Zerbib, O. D. (2019). The effect of pro-environmental preferences on bond prices: Evidence from green bonds. *Journal of Banking & Finance*, 98:39–60.

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.

A Theoretical Extension: Recursive Preferences

This appendix presents a theoretical extension of the baseline CRRA model to incorporate Epstein-Zin recursive preferences, following the continuous-time formulation of Duffie and Epstein (1992). Note that this extension is presented in continuous time to maintain consistency with Sauzet and Zerbib (2022), whereas the baseline model in Section 3 is formulated in discrete time. While the main analysis focuses on CRRA preferences to clearly isolate ESG pricing mechanisms, recursive preferences offer important advantages for asset pricing applications and represent a natural direction for future research.

A.1 Motivation

The baseline CRRA specification imposes a tight link between risk aversion and intertemporal substitution: the elasticity of intertemporal substitution (EIS) equals the reciprocal of risk aversion, $\psi = 1/\gamma$. This restriction limits the model's ability to match empirical moments of both risk premia and the risk-free rate (Mehra and Prescott, 1985; Weil, 1989).

Following Sauzet and Zerbib (2022), this appendix extends the model to incorporate Epstein-Zin recursive preferences (Epstein and Zin, 1989; Duffie and Epstein, 1992), which separate risk aversion from intertemporal substitution. Agent i's preferences are defined recursively by:

$$V_t^i = \max_{\{\theta_u^{i,j}\}_{u=t}^{\infty}} \mathbb{E}_t \left[\int_t^{\infty} f^i(c_u^i, V_u^i, \omega_u^{i,g}) du \right] \tag{15}$$

where the normalized aggregator function f^i combines current consumption c^i, continuation value V^i, and portfolio composition $\omega^{i,g}$:

$$f^i(c^i, V^i, \omega^{i,g}) = \frac{\rho}{1 - 1/\psi} V^i \left[\left(\frac{c^i + \phi^i(\omega^{i,g})}{[(1-\gamma)V^i]^{1/(1-\gamma)}} \right)^{1-1/\psi} - 1 \right] \tag{16}$$

Three key parameters govern the aggregator: the coefficient of relative risk aversion $\gamma > 0$, the elasticity of intertemporal substitution $\psi > 0$, and the rate of time preference $\rho > 0$. This specification nests the CRRA case when $\psi = 1/\gamma$. Note that the ESG utility $\phi^i(\omega^{i,g})$ enters additively with consumption inside the aggregator, which differs from the baseline model's additive-separable specification (Equation 3); the implications of this difference for ESG pricing dynamics are left for future investigation.

A.2 Economic Interpretation

Risk aversion (γ): Controls the agent's willingness to accept gambles over continuation values. Higher γ increases curvature, making agents more averse to uncertainty about future utility. In the ESG context, highly risk-averse agents may exhibit stronger hedging demands for green assets if they provide insurance against climate-related tail risks.

Intertemporal substitution (ψ): Governs willingness to substitute consumption across time periods. Investors with high ψ rapidly shift toward brown assets when green premia compress, while those with low ψ maintain more stable allocations.

Patience (ρ): Determines the subjective discount rate. Patient investors may assign higher value to green assets' long-run environmental benefits.

The separation of γ and ψ is particularly important for ESG asset pricing. When $\psi > 1/\gamma$, hedging demands create state-dependent portfolio tilts: if green assets provide insurance when continuation values are low, ESG-conscious investors value them both for preference and insurance reasons, amplifying the green premium during stress.

A.3 First-Order Conditions

In discrete time, the stochastic discount factor under Epstein-Zin preferences takes the standard form (Epstein and Zin, 1989):

$$M_{t+1}^i = \beta^\theta \left(\frac{c_{t+1}^i}{c_t^i} \right)^{-\theta/\psi} \left(\frac{V_{t+1}^i}{\mathcal{R}_{t,t+1}^i} \right)^{\theta-1} \tag{17}$$

where $\theta = (1-\gamma)/(1-1/\psi)$, $\mathcal{R}_{t,t+1}^i$ is the return on agent i's wealth portfolio, and $\beta = e^{-\rho}$ is the discrete-time discount factor. When $\psi = 1/\gamma$, we have $\theta = 1$ and the SDF reduces to the standard CRRA form $M_{t+1} = \beta(c_{t+1}/c_t)^{-\gamma}$.

The Euler equation for asset j becomes:

$$q_t^j = \mathbb{E}_t \left[M_{t+1}^i (q_{t+1}^j + d_{t+1}^j) \right] + \text{ESG term} \tag{18}$$

where the ESG term scales the marginal ESG utility by the ratio of consumption to the certainty-equivalent continuation value, creating state-dependent amplification of ESG pricing effects.

A.4 Implications

Recursive preferences generate richer dynamics for green premia through time-varying risk premia and hedging demands, flight-to-quality effects during market stress (countercyclical green premia), amplified long-run climate risk pricing, and wealth share dynamics that persist even in complete markets. The deep equilibrium network approach is well-suited to this extension, as it handles the increased dimensionality (augmented state space including continuation values) without the exponential cost increases of traditional projection methods.

B Fischer-Burmeister Complementarity Formulation

This appendix provides the mathematical foundation for the complementarity constraint formulation used in Section 3.

B.1 The Complementarity Problem

The Euler equation residual for agent i and asset j is:

$$\mathcal{E}_t^{i,j} \equiv p_j \cdot u'(c_t^i) - \beta \mathbb{E}_t \left[(d_j(s_{t+1}) + \delta_j \cdot p_j(s_{t+1})) \cdot u'(c_{t+1}^i) + \phi_i \cdot \frac{\partial \omega_{t+1}^{i,g}}{\partial \theta_t^{i,j}} \right] \tag{19}$$

The KKT conditions for the portfolio constraint $\theta_t^{i,j} \geq \underline{\theta}_j$ are:

$$\mathcal{E}_t^{i,j} \geq 0 \tag{20}$$

$$\theta_t^{i,j} - \underline{\theta}_j \geq 0 \tag{21}$$

$$\mathcal{E}_t^{i,j} \cdot (\theta_t^{i,j} - \underline{\theta}_j) = 0 \tag{22}$$

These conditions require that either the Euler equation holds with equality (interior solution) or the constraint binds (boundary solution), but not both simultaneously.

B.2 Fischer-Burmeister Reformulation

The Fischer-Burmeister function provides an equivalent smooth reformulation:

$$FB(a, b) = \sqrt{a^2 + b^2} - a - b \tag{23}$$

Proposition 1 (Fischer-Burmeister Equivalence)**.** *$FB(a, b) = 0$ if and only if $a \geq 0$, $b \geq 0$, and $ab = 0$.*

Proof. Forward: If $FB(a, b) = 0$, then $a + b = \sqrt{a^2 + b^2}$. Squaring: $(a + b)^2 = a^2 + b^2$, so $2ab = 0$, hence $ab = 0$. Since $a + b = \sqrt{a^2 + b^2} \geq 0$ and $ab = 0$, both $a \geq 0$ and $b \geq 0$.

Reverse: If $a \geq 0$, $b \geq 0$, $ab = 0$, then at least one is zero. If $a = 0$: $FB(0, b) = |b| - b = 0$ since $b \geq 0$. If $b = 0$: $FB(a, 0) = |a| - a = 0$ since $a \geq 0$. □

B.3 Application to Portfolio Constraints

Setting $FB(\mathcal{E}_t^{i,j}, \theta_t^{i,j} - \underline{\theta}_j) = 0$ replaces the KKT conditions with a single smooth equation per agent-asset pair. At interior solutions ($\theta_t^{i,j} > \underline{\theta}_j$), the Euler equation holds with equality. At boundary solutions ($\theta_t^{i,j} = \underline{\theta}_j$), the Euler residual is non-negative, indicating the agent would prefer to reduce holdings further but is constrained.

The Fischer-Burmeister formulation is particularly important for ESG asset pricing because constraints bind asymmetrically across investor types. ESG-conscious investors may be constrained

from shorting brown assets ($\theta_t^{1,b} = 0$), while traditional investors may be constrained from shorting overpriced green assets ($\theta_t^{2,g} = 0$). This asymmetric binding creates the market segmentation necessary for ESG preferences to persistently affect prices: green firms face lower required returns because ESG investors bid up their prices, while the short-sale constraint prevents traditional investors from arbitraging the premium away.

The smooth differentiability of the Fischer-Burmeister function ensures that gradient-based optimization (used in both projection methods and DEQN training) converges reliably, unlike direct complementarity formulations that create discontinuities.

C Variable Descriptions

Table 8 provides detailed descriptions of all variables used in this analysis.

Table 8: Variable Descriptions

Variable Name	**Description**
Date	Trading date. Daily observations from September 7, 2020 to October 15, 2025.
Yield Green	Yield on green bonds in percentage points. Green bonds are fixed-income instruments designated to finance climate and environmental projects. *Type:* Continuous. *Source:* Deutsche Finanzagentur.
Yield Brown	Yield on conventional (brown) twin bonds in percentage points. *Type:* Continuous. *Source:* Deutsche Finanzagentur.
Price Green	Market price of green bonds. Inversely related to yields. *Type:* Continuous. *Source:* Deutsche Finanzagentur.
Price Brown	Market price of conventional (brown) bonds. *Type:* Continuous. *Source:* Deutsche Finanzagentur.
Yield Diff	Yield differential (Green − Brown), the "greenium." Negative values indicate green bonds trade at lower yields. *Type:* Continuous (derived).
Price Diff	Price differential (Green − Brown). Positive values indicate green bonds trade at a premium. *Type:* Continuous (derived).

Note: All variables measured at daily frequency. Sample: September 7, 2020 to October 15, 2025 ($N = 1{,}301$).

D Twin Bond Descriptive Statistics

Table 9: Descriptive Statistics for Bond Yields and Prices

Statistic	**Yield Green**	**Yield Brown**	**Price Green**	**Price Brown**
Observations	1,301	1,301	1,301	1,301
Mean	1.32	1.34	92.55	92.36
Std. Dev.	1.21	1.19	7.90	7.74
Min	−0.67	−0.65	81.59	81.80
Median	2.04	2.05	89.40	89.36
Max	2.81	2.83	106.73	106.58
Skewness	−0.62	−0.61	0.65	0.64
Kurtosis	−1.39	−1.39	−1.21	−1.20

Note: Yields in percentage points. Sample period: September 2020 – October 2025, $N = 1{,}301$. Kurtosis is reported as excess kurtosis (normal = 0).

Table 10: Descriptive Statistics for Yield and Price Differentials

Statistic	**Yield Diff**	**Price Diff**
Observations	1,301	1,301
Mean	−0.02	0.19
Std. Dev.	0.02	0.19
Min	−0.07	−0.02
25th Percentile	−0.03	0.04
Median	−0.02	0.12
75th Percentile	−0.01	0.30
Max	0.00	0.67

Note: Yield Diff = Green − Brown (the greenium). Price Diff = Green − Brown. Sample period: September 2020 – October 2025, $N = 1{,}301$.

Table 11: Mean Comparison – Green vs Brown Bonds

Variable	**Green Bonds**	**Brown Bonds**	**Difference**	***t*-statistic**	***p*-value**
Yield (%)	1.32	1.34	−0.02	−45.82	< 0.001
Price	92.55	92.36	0.19	36.97	< 0.001

Note: Paired t-tests. Difference = Green − Brown. All differences significant at 1% ($p < 0.001$).

Table 12: Annual Trends in Bond Yields and Prices

Year	Obs.	YG Mean	YG SD	YB Mean	YB SD	Y Diff	P Diff
2020	81	−0.58	0.05	−0.58	0.06	−0.004	0.25
2021	254	−0.43	0.12	−0.39	0.12	−0.04	0.52
2022	255	1.06	0.79	1.08	0.78	−0.03	0.18
2023	255	2.36	0.22	2.38	0.21	−0.02	0.13
2024	254	2.19	0.19	2.21	0.19	−0.02	0.06
2025	202	2.19	0.10	2.21	0.10	−0.02	0.03

Note: YG = Yield Green, YB = Yield Brown, Y Diff = Yield differential, P Diff = Price differential. All yield variables in percentage points.

E Calibration Data Tables

This appendix presents the detailed time-series data used to calibrate the model parameters in Section 4. Summary statistics and calibrated parameter values are reported in the main text (Tables 2 and 1).

Table 13: German 10-Year Bund Yields, 2001–2024

Year	**Yield (%)**	**Year**	**Yield (%)**	**Year**	**Yield (%)**	**Year**	**Yield (%)**
2001	4.80	2007	4.25	2013	1.55	2019	−0.26
2002	4.79	2008	3.98	2014	1.08	2020	−0.50
2003	4.07	2009	3.24	2015	0.50	2021	−0.34
2004	4.06	2010	2.74	2016	0.10	2022	1.23
2005	3.35	2011	2.61	2017	0.32	2023	2.41
2006	3.76	2012	1.50	2018	0.37	2024	2.32
Mean: 2.17%				$\Rightarrow \beta = 0.978$			

Source: OECD Main Economic Indicators via FRED, series IRLTLT01DEA156N.

Table 14: German Real GDP Growth and Regime Classification, 2001–2024

Year	**Growth (%)**	**Regime**	**Year**	**Growth (%)**	**Regime**
2001	1.76	Good	2013	0.51	Good
2002	−0.20	Bad	2014	2.17	Good
2003	−0.54	Bad	2015	1.42	Good
2004	0.70	Good	2016	2.14	Good
2005	1.06	Good	2017	3.10	Good
2006	4.04	Good	2018	1.15	Good
2007	2.99	Good	2019	1.01	Good
2008	0.63	Good	2020	−4.49	Bad
2009	−5.50	Bad	2021	3.87	Good
2010	4.01	Good	2022	1.89	Good
2011	3.83	Good	2023	−0.67	Bad
2012	0.65	Good	2024	−0.47	Bad

Source: Eurostat via FRED, series CLVMNACSCAB1GQDE (chained 2010 EUR). Regime threshold: growth $> 0\%$ = Good.

Table 15: Realized Equity Premium: DAX Returns Minus Bund Yield, 2001–2024

Year	EP (%)	Year	EP (%)	Year	EP (%)	Year	EP (%)
2001	−24.59	2007	18.04	2013	23.93	2019	25.38
2002	−48.73	2008	−44.35	2014	1.57	2020	4.05
2003	33.01	2009	20.61	2015	9.06	2021	16.13
2004	3.28	2010	13.32	2016	6.77	2022	−13.58
2005	23.72	2011	−17.30	2017	12.19	2023	17.90
2006	18.22	2012	27.56	2018	−18.63	2024	16.53
Mean: 5.15%				Std Dev: ≈ 22%			

Source: DAX Kursindex returns from STOXX/Qontigo minus OECD Bund 10Y yields from FRED. EP = DAX annual return − Bund 10Y yield.

Table 16: Damodaran Implied Equity Risk Premium (FCFE), 2000–2025

Year	ERP (%)	Year	ERP (%)	Year	ERP (%)	Year	ERP (%)
2000	2.87	2007	4.37	2014	5.78	2021	4.24
2001	3.62	2008	6.43	2015	6.12	2022	5.94
2002	4.10	2009	4.36	2016	5.69	2023	4.60
2003	3.69	2010	5.20	2017	5.08	2024	4.33
2004	3.65	2011	6.01	2018	5.96	2025	4.23
2005	4.08	2012	5.78	2019	5.20		
2006	4.16	2013	4.96	2020	4.72		
Mean: 4.81%				Std Dev: 0.94%			

Source: Damodaran (2025), NYU Stern. Implied ERP using FCFE model for mature (Aaa-rated) markets. Germany is Aaa-rated, so no additional country risk premium applies.

Table 17: Green vs. Brown Dividend Yields from German Equity Indices

Comparison	Green Yield (%)	Brown Yield (%)	Spread (bps)	Date
DAX ESG Target vs. DAX	2.20	2.10	+10	Jan 2026
DAX 50 ESG vs. DAX	2.90	2.70	+20	Dec 2022

Source: STOXX Ltd. factsheets (Qontigo/Deutsche Boerse Group). DAX ESG Target (ISIN DE000A3CLUH8); DAX (ISIN DE0008469008).

F Additional Figures

This appendix presents supplementary figures and tables referenced in Section 5.

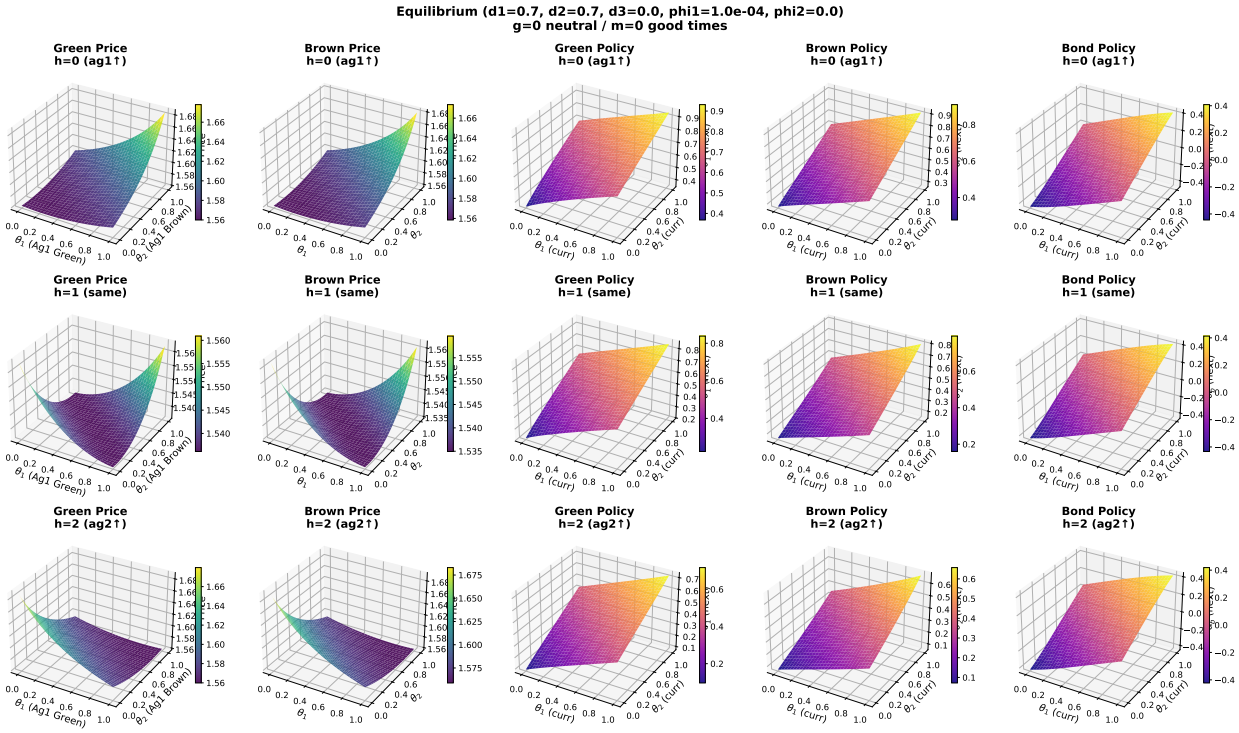

Figure 16: Equilibrium surfaces at ϕ^*. Top row: green equity price $p_g(\theta^{1,g}, \theta^{1,b})$. Middle row: brown equity price $p_b(\theta^{1,g}, \theta^{1,b})$. Bottom row: agent 1's next-period green share $\theta^{1,g\prime}(\theta^{1,g}, \theta^{1,b})$. Left column: state $(g = 0, m = 0, h = 1)$. Right column: state $(g = 1, m = 0, h = 1)$. Bond holdings fixed at $\theta^{1,f} = 0$.

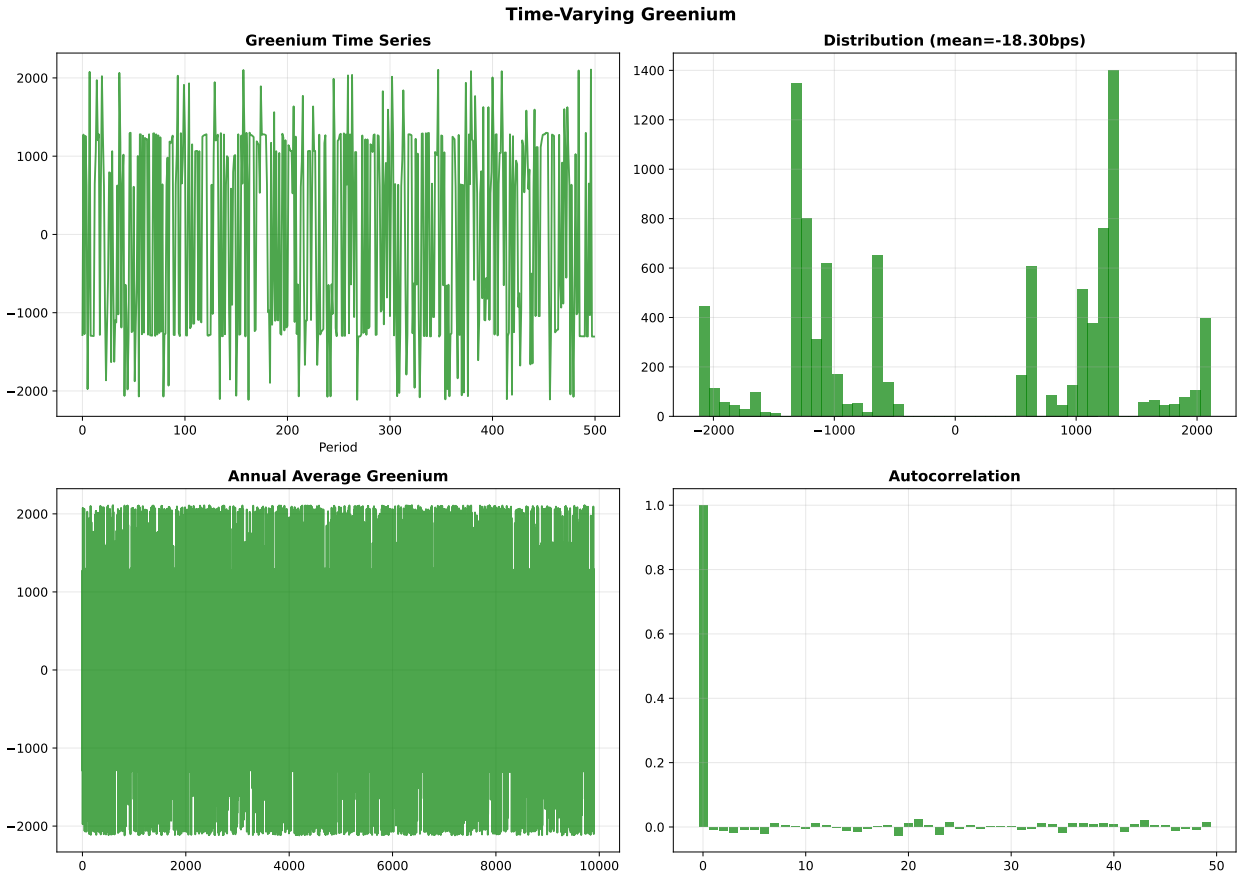

Figure 17: Time-varying greenium in German twin bonds. Top left: daily greenium with 60-day moving average. Top right: greenium vs. average yield level (OLS regression). Bottom left: annual mean greenium. Bottom right: autocorrelation function.

Table 18: Time-Varying Greenium: OLS Regression and Autocorrelation

Panel A: OLS regression — Greenium_t (bps) $= \alpha + \beta \cdot \overline{\text{Yield}}_t + \varepsilon_t$			
	Estimate	t-stat (OLS)	t-stat (NW)
Intercept (α)	-4.055	-78.55	-21.28
Slope (β)	$+1.217$	$+42.18$	$+12.82$
R^2		0.578	
N (daily obs.)		1,301	
NW bandwidth		$L = 7$ (Bartlett kernel, plug-in)	

Panel B: Autocorrelation structure		
Lag	Horizon	ρ
1	1 day	0.972
5	1 week	0.965
21	1 month	0.934
63	3 months	0.811
252	1 year	0.036

Note: The dependent variable is the daily twin bond greenium ($\text{yield}_{\text{green}} - \text{yield}_{\text{brown}}$) in basis points. The independent variable is the average of green and brown yields. The OLS t-statistics assume i.i.d. errors; the Newey-West (NW) t-statistics use heteroskedasticity and autocorrelation consistent (HAC) standard errors with a Bartlett kernel and plug-in bandwidth $L = \lfloor 4(T/100)^{2/9} \rfloor = 7$. Even after this correction, both coefficients remain highly significant ($|t| > 12$). With a more conservative bandwidth of $L = 21$ (one month), the t-statistics are -13.09 and $+7.89$ respectively, still significant at the 1% level. Data: German twin bonds, September 2020 – October 2025.