Sustainable Asset Pricing with Heterogeneous Agents: A Computational Approach

Ryan Huang

April 2026

Department of Economics, UNC Chapel Hill

Roadmap

Motivation

The ESG Pricing Puzzle

Fact: \$30.3 trillion in sustainable AUM globally (2023)

Puzzle: ESG investors systematically accept *lower* returns

- German twin bonds: identical securities except green label
- Persistent greenium of -2.44 bps (2020–2025)
- t-stat $= -45.8$, never positive in sample

Three questions:

1. How strong must ESG preferences be?
2. What role do portfolio constraints play?
3. How does the wealth distribution matter?

The Greenium in German Twin Bonds

Figure 3: Greenium - Yield Differential Over Time (2020-2025)

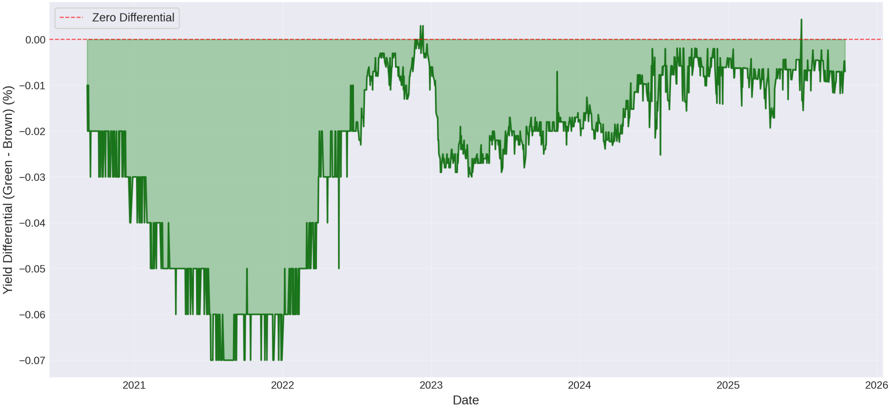

This Thesis

Main Result

Even *minimal* ESG preferences ($\phi^* = 10^{-4}$) generate the observed greenium when embedded in general equilibrium with heterogeneous agents and portfolio constraints.

Two hypotheses supported:

1. Small $\phi \Rightarrow$ negative greenium (green earns less than brown)
2. Greenium varies with wealth distribution (new vs. rep. agent)

Method: Deep Equilibrium Networks — neural networks trained to satisfy Euler equations globally

Literature & Contribution

Literature and Gaps

Theoretical ESG pricing

- Pástor et al. (2021): greenium in CAPM
- Pedersen et al. (2021): ESG uncertainty
- Sauzet (2022): 2-agent, recursive prefs

Empirical evidence

- Pástor & Vorsatz (2022): twin bonds
- Bolton & Kacperczyk (2021): carbon premium
- Riedl & Smeets (2017): WTP 2–3%/yr

This thesis:

- Combines all three channels jointly (preferences + substitutability + constraints)
- DEQNs: 3 assets, 12 shocks (beyond projection method limits)
- Calibrated to German data (joint match: greenium + equity premium)

Model

Model Overview

Economy: 2 agents, 3 assets, 12 exogenous shocks, discrete time, infinite horizon

Agents

- Agent 1: ESG-conscious ($\phi_1 > 0$)
- Agent 2: Traditional ($\phi_2 = 0$)
- Same β, γ (isolate ESG channel)

Assets

- Green equity (unit supply)
- Brown equity (unit supply)
- Risk-free bond (zero net supply)

12 shocks: $s \in \{0, \ldots, 11\}$

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

- ESG regime $g \in \{0, 1\}$
- Business cycle $m \in \{0, 1\}$
- Household income $h \in \{0, 1, 2\}$

Key: $\mathbb{E}[d_g] = \mathbb{E}[d_b]$ (symmetric)
Greenium arises purely from ϕ

Preferences and Euler Equations

Utility:

$$U^i = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[u(c_t^i) + \phi_i \cdot \omega_t^{i,g} \right]$$

- $u(c) = c^{1-\gamma}/(1-\gamma)$ (CRRA), $\quad \omega^{i,g} = p_g \theta^{i,g} / (p_g \theta^{i,g} + p_b \theta^{i,b})$

Euler equation (agent i, asset j):

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

Constraints: $\theta^{i,j} \geq 0$ (equities), $\quad \theta^{i,f} \in [-0.5, 0.5]$ (bond)

Handled via **Fischer-Burmeister** complementarity $\rightarrow$ smooth KKT

Solution: Deep Equilibrium Networks

Architecture:

- Input: $(s_t, \theta^{1,g}, \theta^{1,b}, \theta^{1,f})$
- Augmented to 29 features
- 4 hidden layers $\times$ 1,024 neurons
- Output: 3 prices + 3 portfolios
- $\sim$12M parameters

Training:

- Minimize FB Euler errors
- AdamW + cosine LR decay
- Green-brown symmetry augmentation

Curriculum:

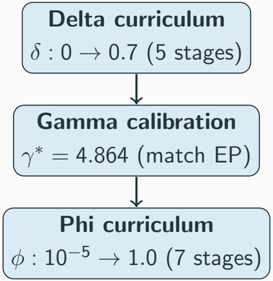

Data & Calibration

Calibration Strategy

Parameter	Value	Source
β	0.978	German Bund 10Y (mean 2.17%)
γ	4.864	Calibrated to equity premium
ϕ^*	10^{-4}	Calibrated to twin bond greenium
δ	0.7	Partial Lucas tree
P_{macro}	$\begin{pmatrix} 0.78 & 0.22 \\ 0.60 & 0.40 \end{pmatrix}$	German GDP regimes (2001–2024)
P_{hh}	diag 0.82	SOEP income persistence
P_{green}	diag 0.90	Modeling assumption

German Twin Bond Data

- $N = 1{,}301$ daily observations (Sept 2020 – Oct 2025)
- Identical issuer, maturity, coupon — differ only in green designation
- Greenium $= -2.44$ bps ($t = -45.8$), persistent across all macro regimes

	Mean Yield	Std Dev	Min	Max
Green bonds	1.32%	1.21	−0.67%	2.81%
Brown bonds	1.34%	1.19	−0.65%	2.83%

Bond Yield Trends

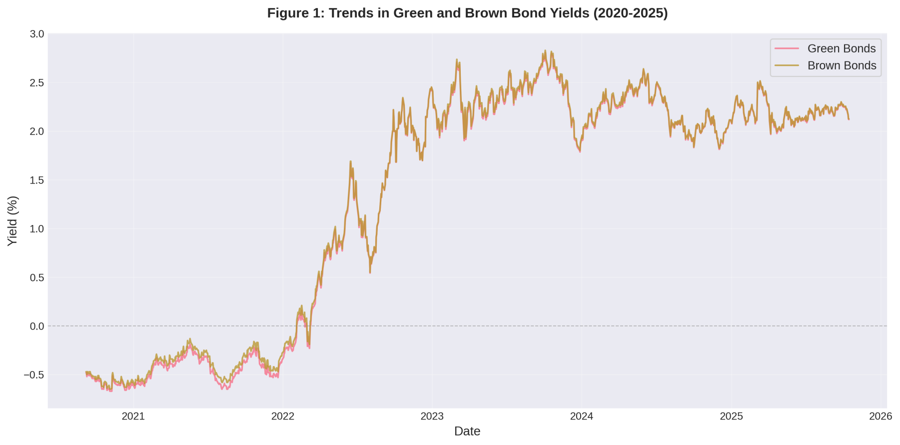

Results

Training Convergence

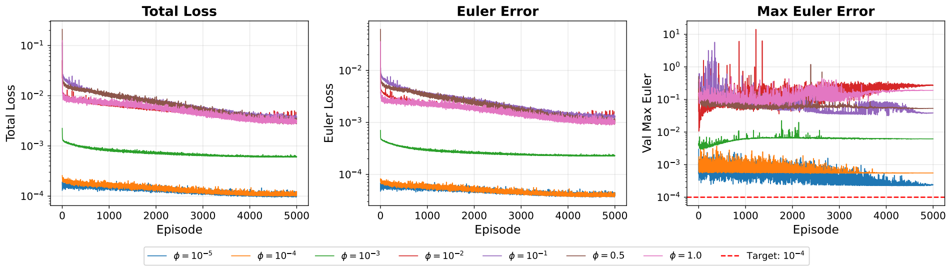

Three panels: total loss, Euler error, max Euler error. Each line is one ϕ stage. Low-ϕ stages (ϕ^* and below) converge to $< 10^{-4}$ (below the red dashed target). High-ϕ stages plateau higher — the equilibrium is harder to solve.

Convergence Summary

Phase	Stage	Val Euler Error	Val Max Euler
Delta	$\delta = 0.00$	8.3×10^{-5}	3.2×10^{-3}
Delta	$\delta = 0.70$	3.3×10^{-5}	1.1×10^{-3}
Gamma cal.	$\gamma = 4.864$	6.0×10^{-5}	2.9×10^{-3}
Phi	$\phi = 10^{-5}$	3.8×10^{-5}	2.3×10^{-4}
Phi	$\phi = 10^{-4}$ (ϕ^*)	4.1×10^{-5}	5.2×10^{-4}
Phi	$\phi = 1.0$	1.8×10^{-3}	6.1×10^{-2}

At ϕ^*: Euler errors $< 10^{-4}$ — well within standard thresholds.

High-ϕ errors $100\times$ larger — results there are directional only.

Calibration: Greenium vs. ϕ

ϕ	Greenium (bps)	EP (%)	$\theta^{1,g}$	$\theta^{1,b}$
0	−0.03	4.86	0.501	0.501
10^{-5}	−0.09	4.86	0.505	0.497
10^{-4}	**−2.53**	**4.86**	**0.544**	**0.457**
10^{-3}	−23.1	4.77	0.701	0.297
10^{-1}	−3,906	9.15	0.988	0.692
1.0	−11,417	−10.1	0.997	0.789

- At $\phi = 0$: greenium ≈ 0 — symmetric baseline confirmed
- At $\phi^* = 10^{-4}$: greenium −2.53 matches data (−2.44); EP 4.86% matches data (4.81%)
- Portfolio tilt only $\sim$4 pp — tiny preference, realistic pricing

Gray rows: Euler errors $> 10^{-2}$; shown for trend, not precision.

Greenium vs. ϕ (Figure)

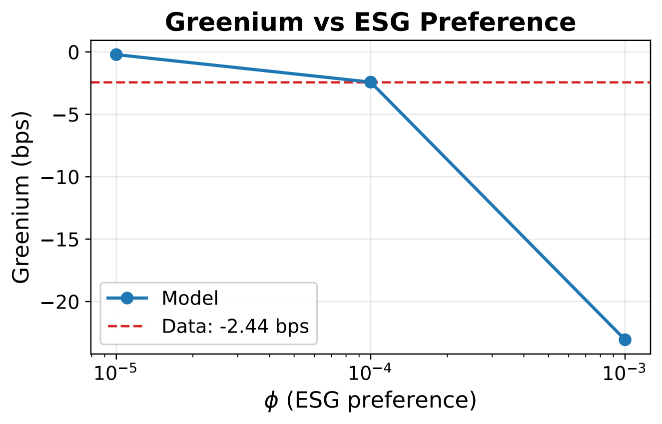

Blue line: model greenium at each ϕ stage. Red dashed: data target (-2.44 bps).
At this scale the match at ϕ^* is invisible — both model and data are near zero relative to the $-12{,}000$ bps at $\phi = 1$. The table on the previous slide shows the precision.

Model Fit at ϕ^*

Moment	Model	Data	Source
Greenium (bps)	-2.53	-2.44	German twin bonds
Equity premium (%)	4.86	4.81	Damodaran (2000–2025)
EP, green (%)	4.85	—	—
EP, brown (%)	4.87	—	—
Risk-free rate (%)	-1.47	2.30	Bund 10Y

- Greenium and EP matched simultaneously with $\gamma = 4.864$, $\phi^* = 10^{-4}$
- Green and brown EPs nearly identical (4.85 vs 4.87) $\Rightarrow$ foreshadows Sharpe result
- Risk-free rate too low: CRRA precautionary savings pushes $p_f > 1$

  Fix: Epstein-Zin preferences (separate γ from EIS)

Greenium Decomposition: Which State Dimension Matters?

- **Three panels** plot greenium vs. ϕ split by each state dimension
- **Left (ESG regime g):** Two lines perfectly overlap across *all* ϕ
  - Symmetric dividends $\Rightarrow$ greenium identical in green vs. neutral regimes
- **Right (household income h):** Three lines separate earliest
  - $h=0$ (agent 1 wealthy): greenium most negative (-6 bps at ϕ^*)
  - $h=2$ (agent 1 poor): greenium near zero
  - Wealthy ESG investor $\Rightarrow$ more demand $\Rightarrow$ larger green premium
- **Center (macro m):** Lines overlap at ϕ^*, diverge at higher ϕ
- **Key:** Wealth distribution is the primary driver — absent from rep. agent models

Greenium Decomposition (Figure)

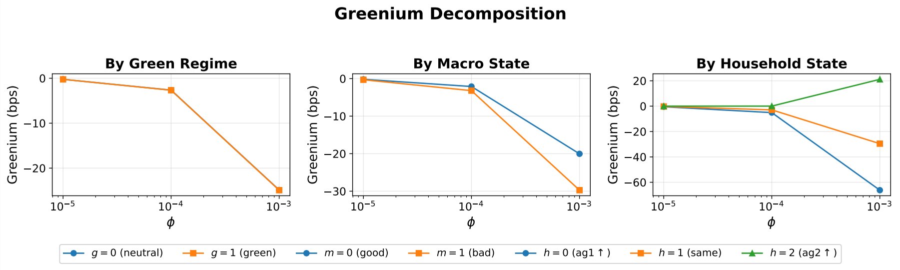

Portfolio Allocations Across States

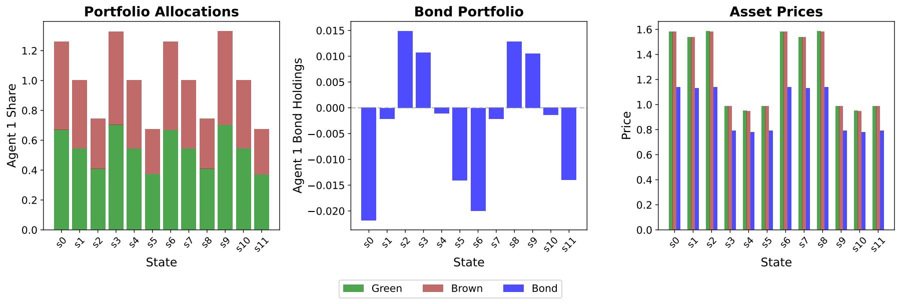

Left: agent 1's green (green) and brown (red) equity shares across 12 states. States group as s_0–s_2 ($g=0, m=0$), s_3–s_5 ($g=0, m=1$), etc. Within each group, $h=0$ states have the highest total holdings — the wealthy ESG agent holds more of everything. Center: bond positions near zero (± 0.02). Right: green and brown prices nearly identical; much higher in expansions ($m=0$) than recessions ($m=1$).

Price & Portfolio Evolution Across ϕ

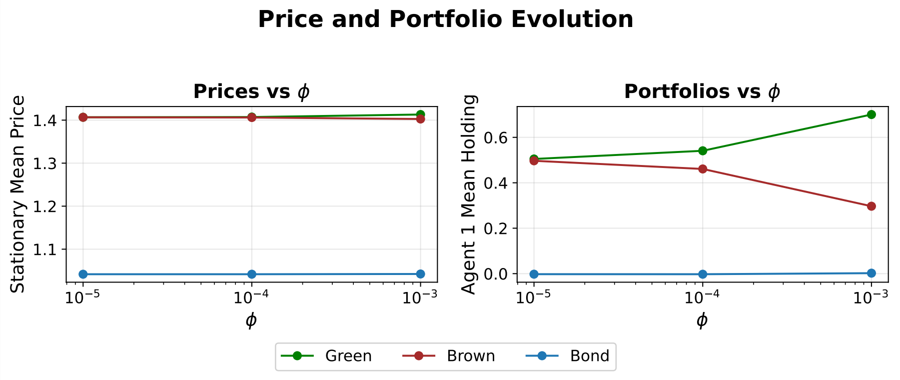

Left: at low ϕ, green and brown prices overlap ($\sim$1.4). As ϕ rises, brown price collapses (to $\sim$0.3) while green stays near 1.0. Bond price declines modestly.

Right: agent 1's green share rises from 0.5 to $\sim$1.0. Brown holding first drops then recovers (non-monotonic). Bond goes negative — agent 1 borrows to lever up at high ϕ.

High-ϕ end ($> 10^{-2}$) less precise due to Euler errors.

Equilibrium Surfaces

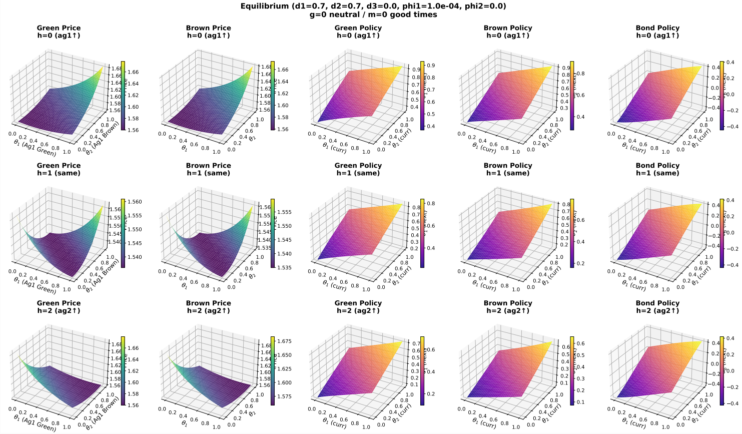

Time-Varying Greenium: Empirical Validation

- Greenium correlated with yield level ($R^2 = 0.58$, NW $t = 12.82$)
- High persistence: $\rho_1 = 0.97$ — consistent with slowly evolving regimes
- Model predicts greenium varies with state; wealth distribution is the key driver

Time-Varying Greenium

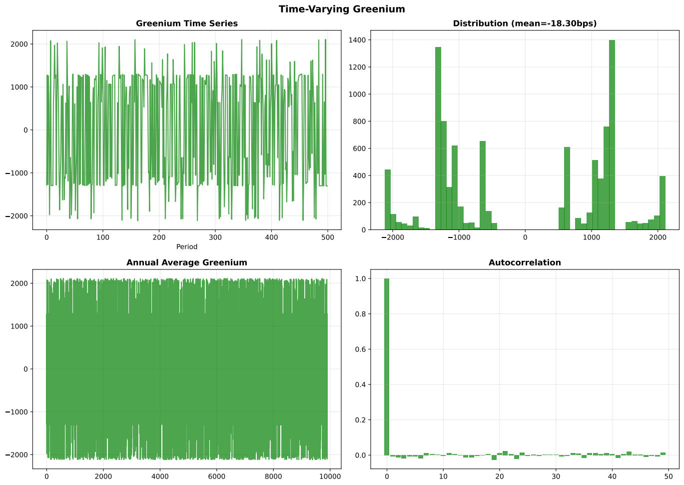

Forward Simulation & Robustness

Forward Simulation at ϕ^*

Simulate the calibrated economy for 10,000 periods under the stationary distribution.

Risk-return comparison

- Sharpe: green 0.200, brown 0.207
- Difference: 0.007 (3.5% of brown)
- Cons. vol: 12.9% / 13.4%

$\Rightarrow$ **Near-identical risk profiles**

Greenium is preference premium, not risk

Constraint binding frequency

- Green lower: 7.4%
- Brown lower: 7.6%
- Bond lower/upper: 10.5%/12.2%

Constraints active in 7–12% of periods

even at small ϕ^*

Statistical Precision: Bootstrap

Question: How sensitive are the model moments to simulation randomness?

Method: Re-run the 10k-period forward simulation 200 times with different random seeds. Each run produces a greenium and equity premium. Report the distribution.

Moment	Bootstrap Mean	Std Error	95% CI
Equity premium	5.74%	0.61%	[4.62, 6.89]
Greenium (bps)	−2.75	0.71	[−4.08, −1.33]

- Data greenium (−2.44 bps) and EP (4.81%) both inside 95% CI
- Bootstrap EP mean (5.74%) exceeds grid value (4.86%) by 90 bps
- Grid-evaluated moments are more reliable; simulation moments are suggestive

Sharpe Ratio Analysis: Preference Premium, Not Risk

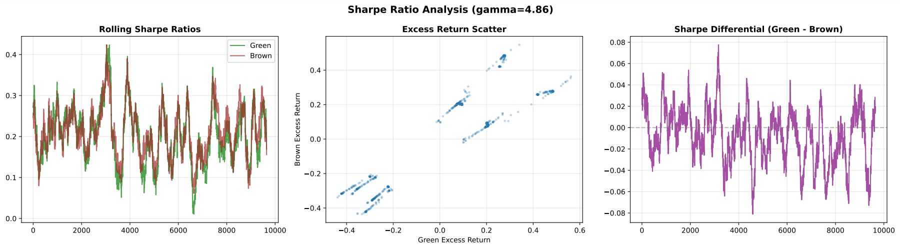

Left: green and brown Sharpe ratios vs. ϕ. Center: Sharpe differential (green$-$brown). Right: state-dependent Sharpe ratios at ϕ^*.

Conclusion: greenium reflects willingness to accept lower returns, not differential risk.

(Partly by construction: symmetric dividends ensure identical aggregate risk.)

Robustness

Robustness & Validation

Symmetry validation

- $\phi = 0$: greenium $= -0.03$ bps
- Max symmetry error: 0.004%

Euler equation accuracy

- Val Euler $= 4.1 \times 10^{-5}$ at ϕ^*
- Max $= 5.2 \times 10^{-4}$

Bond bound check

- Mean $|\theta^f| \approx 0.002$ at ϕ^*
- Far from ± 0.5 bounds

Statistical tests

- Newey-West: $t = 7.89$ ($L = 21$)
- Market clearing by construction

Conclusion

Summary of Contributions

1. **Economic:** Small ESG preferences generate realistic pricing
   - $\phi^* = 10^{-4}$ matches greenium (-2.53 bps) + EP (4.86%)
   - Sharpe ratios nearly equal (0.200 vs. 0.207) $\Rightarrow$ preference premium, not risk
2. **Wealth distribution drives greenium at ϕ^*:**
   - Cross-sectional wealth, not ESG regime, is the primary channel
   - Absent from representative agent models (Pástor et al., 2021)
3. **Robust estimation:**
   - Bootstrap 95% CI: greenium $[-4.08, -1.33]$ bps
   - Near-equal Sharpe ratios: preference premium, not risk compensation

Limitations & Future Work

Limitations:

- Risk-free rate mismatch
  (CRRA precautionary savings)
- ESG regime persistence assumed
- Short twin bond sample (5 years)
- 2 agent types (reality: continuum)

Extensions:

- Epstein-Zin preferences
  (separate γ from EIS)
- Heterogeneous beliefs
- Trending dividends
- More agent types
- Cross-country calibration

Thank you

rhuang1@unc.edu

Backup: Euler Equation Heatmap

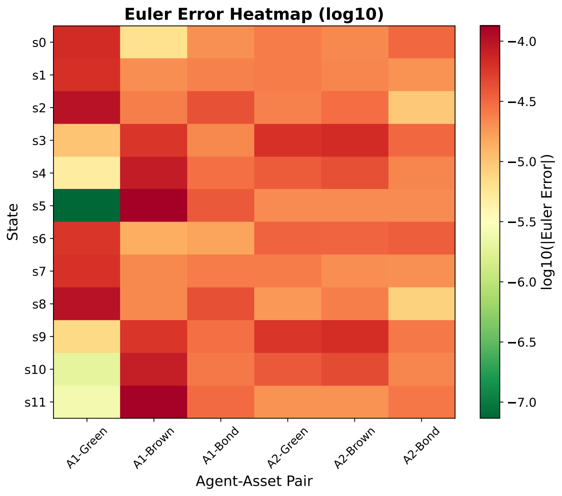

Backup: Risk-Free Rate Puzzle

Why is $R_f = -1.47\%$ in the model?

With CRRA and $\gamma = 4.864$: $R_f \approx \frac{1}{\beta} \cdot \mathbb{E}\left[\left(\frac{c_{t+1}}{c_t}\right)^{-\gamma}\right]$

- High $\gamma \Rightarrow$ strong precautionary savings $\Rightarrow p_f > 1 \Rightarrow R_f < 0$
- Well-documented CRRA tension (Weil, 1989)

Resolution: Epstein-Zin preferences

- Separate risk aversion γ from EIS ψ
- Keep $\gamma \approx 5$ for EP, set $\psi > 1$ for $R_f > 0$
- DEQN: add continuation value as network output

Backup: ESG Regime Persistence

Data-implied: P_{green} is i.i.d. (no evidence of persistence)

Model: P_{green} diagonal $= 0.90$ (persistence assumption)

Justification:

- EU Taxonomy rollout (2020–2025)
- Mandatory sustainability disclosures
- Multi-year corporate decarbonization commitments

Robustness: Preliminary i.i.d. results confirm symmetric baseline

Backup: Fischer-Burmeister Complementarity

Problem: Portfolio constraints bind occasionally

KKT: $\mathcal{E}^{i,j} \geq 0, \quad \theta^{i,j} - \underline{\theta} \geq 0, \quad \mathcal{E}^{i,j} \cdot (\theta^{i,j} - \underline{\theta}) = 0$

FB reformulation: single smooth equation per constraint

$$FB(a,b) = \sqrt{a^2 + b^2} - a - b = 0 \iff a \geq 0,\ b \geq 0,\ ab = 0$$

- Interior: Euler $= 0$ (standard pricing)
- Boundary: agent wants to sell more but is constrained
- Smooth $\Rightarrow$ gradient-based optimization works

Backup: Cross-Section Prices & Policies

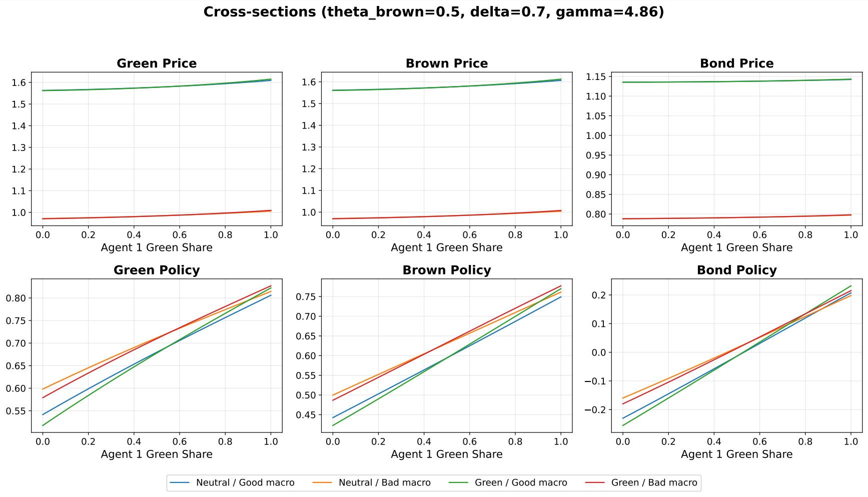

Backup: Forward Simulation

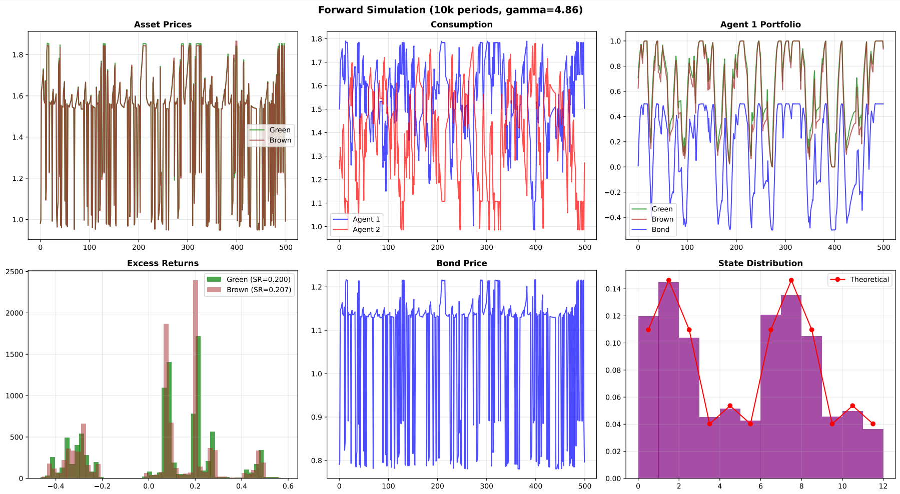

Backup: Epstein-Zin Extension

Recursive utility:

$$V_t^i = \max \mathbb{E}_t \int_t^\infty f^i(c^i, V^i, \omega^{i,g})\, du$$

Aggregator:

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

Advantages over CRRA:

- Separate γ (risk aversion) from ψ (EIS)
- Match R_f independently of EP
- State-dependent hedging demands amplify green premium