L08: Climate Risk Quantification & Green Finance

Extended Slides – BSc Digital Finance Course

Digital Finance

What Will You Be Able to Do After This Lecture?

1. Formally derive Carbon VaR and compute portfolio-level climate loss under NGFS transition scenarios
2. Model stranded asset valuation using NPV under carbon price trajectories and compute write-down thresholds
3. Implement climate stress tests that separate physical, transition, and liability risk channels
4. Quantify ESG rating divergence using Bayesian aggregation and detect greenwashing with decision-tree classifiers
5. Construct decarbonization-constrained portfolios using Markowitz optimization with carbon budget constraints
6. Evaluate green bond pricing, measure the greenium, and distinguish genuine green finance from label arbitrage

---

Six objectives: Carbon VaR (1), stranded assets (2), stress testing (3), ESG integration (4), portfolio decarbonization (5), and green bond verification (6). Rigorous theory with working code and 12 data visualizations.

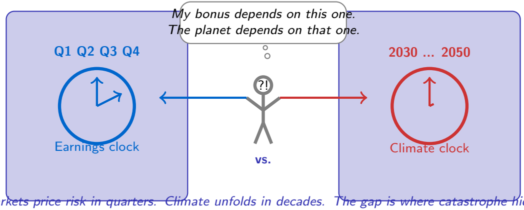

Markets price risk in quarters. Climate unfolds in decades. The gap is where catastrophe hides.

How Do You Compute the Carbon Value-at-Risk of a Portfolio?

Definition. Carbon VaR measures the portfolio loss from a carbon price shock at confidence level α:

$$\mathrm{CVaR}_\alpha = \sum_{i=1}^{N} w_i \cdot \Delta P_i(\tau) \cdot \nVdash[L > \mathrm{VaR}_\alpha]$$

where w_i = portfolio weight, $\Delta P_i(\tau)$ = price change of asset i under carbon price τ.
Carbon price sensitivity: For firm i with emissions intensity e_i (tCO_2/\$M revenue):

$$\Delta P_i(\tau) = -\frac{e_i \cdot \tau \cdot (1 - \text{pass-through}_i)}{\text{EBITDA margin}_i}$$

Portfolio-level Carbon VaR:

$$\mathrm{CVaR}_\alpha^{\text{port}} = \sum_i w_i \cdot \Delta P_i(\tau) + \underbrace{\frac{1}{2}\sum_{i,j} w_i w_j \cdot \rho_{ij} \cdot \sigma_i \sigma_j}_{\text{correlation adjustment}}$$

Pass-through asymmetry: Utilities pass 60–80% of carbon costs to consumers. Cement producers pass 10–20%. Airlines pass 30–50%.

Carbon VaR is standard VaR with a climate-specific shock. The key insight: pass-through rates vary by sector, so a uniform carbon price creates non-uniform losses.

What Does the Carbon Price Look Like Under NGFS Scenarios?

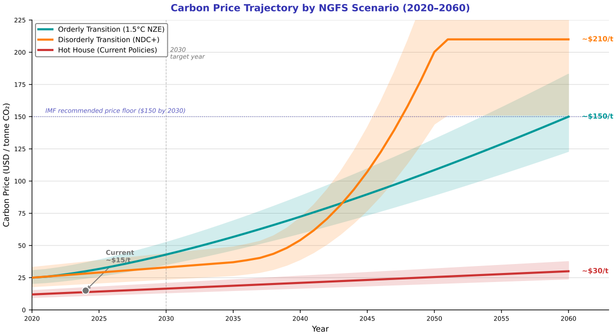

Illustrative | Based on NGFS Phase 4 scenarios and IMF carbon pricing analysis

- **Net Zero 2050 (orderly):** Carbon price rises gradually to \$150–250/ton by 2050. Markets adapt smoothly.
- **Delayed Transition (disorderly):** Flat until 2030, then jumps to \$200+ as panic policy kicks in. Maximum stranding.
- **Current Policies (hot house):** Carbon stays below \$50/ton. Transition risk low, but physical damage escalates.
- The delayed path is the most financially dangerous: sudden repricing after years of complacency.
- EU ETS carbon price (2024): approximately EUR 60–80/ton – already in the "orderly" band.

Key: Scenario choice determines Carbon VaR. The same portfolio has 3x different loss estimates depending on which NGFS pathway you assume.

The delayed transition is the most dangerous: years of complacency followed by sudden repricing. Carbon VaR under delayed transition is 2–3x higher than under orderly transition.

Can You Compute the Carbon VaR of a Four-Asset Portfolio?

```python
import numpy as np
def carbon_var(w, emissions, margins,
               pass_thru, tau, alpha=0.95):
    """Carbon VaR at confidence alpha."""
    n = len(w)
    losses = np.zeros(n)
    for i in range(n):
        cost = emissions[i] * tau
        net = cost * (1 - pass_thru[i])
        losses[i] = net / margins[i]
    port_loss = w @ losses
    # Monte Carlo for tail
    sims = np.random.normal(
        port_loss, 0.02, 10000)
    return np.percentile(sims, alpha*100)

w = np.array([0.3, 0.25, 0.25, 0.2])
em = np.array([800, 200, 50, 10])
mg = np.array([0.15, 0.20, 0.12, 0.30])
pt = np.array([0.7, 0.3, 0.1, 0.0])
for tau in [50, 100, 150, 200]:
    v = carbon_var(w, em, mg, pt, tau/1e6)
    print(f"${tau}/ton: CVaR={v:.1%}")
```

What the code computes

- Four assets: utility (800 tCO_2/\$M), industrial (200), REIT (50), tech (10)
- Pass-through: utility 70%, industrial 30%, REIT 10%, tech 0%
- At \$100/ton: utility absorbs 30% of carbon cost, losing ∼16% of EBITDA
- Portfolio Carbon VaR scales non-linearly with carbon price due to margin compression
- The tech position (low emissions) acts as a natural hedge

Production extension: Replace Monte Carlo with copula-based simulation to capture tail dependence between carbon-exposed sectors.

Carbon VaR is computable with four inputs per asset: emissions intensity, margin, pass-through rate, and portfolio weight. The model is simple – the data challenge is measuring pass-through.

Which Sectors Face the Largest Stranded Asset Write-Downs?

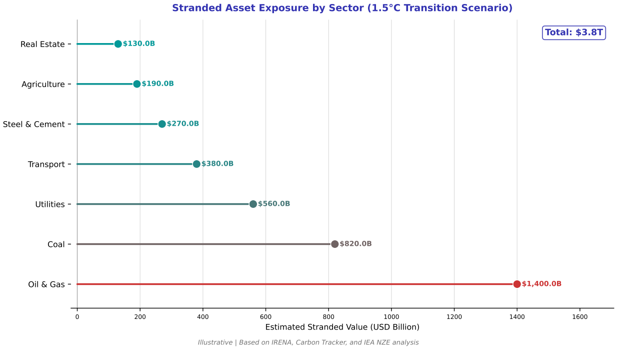

- **Coal:** 80–90% of reserves stranded under 2C. Highest write-down risk per dollar of book value.
- **Oil sands:** High extraction cost means early stranding at \$80–100/ton carbon. Conventional oil survives longer.
- **Natural gas:** "Bridge fuel" narrative delays stranding, but methane leakage liabilities are rising.
- **Thermal utilities:** Stranded generation capacity, not reserves. Retrofit costs often exceed replacement.
- **Cement and steel:** Process emissions are harder to abate than energy emissions. Carbon capture is the only path.

Pattern: Stranding risk correlates with abatement cost. Sectors where decarbonization is cheap (power) strand first. Sectors where it is hard (cement) strand last but face higher long-term liability.

Coal strands first, cement strands last. The sequence depends on abatement cost – cheap-to-decarbonize sectors face policy risk earlier.

How Do You Value an Asset When Its Cash Flows Depend on Carbon Policy?

Standard NPV:

$$\text{NPV} = \sum_{t=1}^{T} \frac{CF_t}{(1+r)^t} - I_0$$

Climate-adjusted NPV under carbon trajectory $\{\tau_t\}$:

$$\text{NPV}_{\text{climate}} = \sum_{t=1}^{T} \frac{CF_t - e \cdot \tau_t \cdot (1-\phi)}{(1+r+\lambda_c)^t} - I_0$$

where e = emissions per unit output, τ_t = carbon price at time t, ϕ = pass-through rate, λ_c = climate risk premium.
Stranding threshold: The carbon price τ^* at which $\text{NPV}_{\text{climate}} = 0$:

$$\tau^* = \frac{\sum_t CF_t/(1+r+\lambda_c)^t - I_0}{\sum_t e \cdot (1-\phi)/(1+r+\lambda_c)^t}$$

Write-down rule: If expected carbon path $E[\tau_t] > \tau^*$ before asset end-of-life, impairment is required under IAS 36.

The stranding threshold is the carbon price at which an asset's climate-adjusted NPV hits zero. Any carbon trajectory that crosses this threshold triggers an impairment write-down.

How Do Physical, Transition, and Liability Risk Interact in a Climate Stress Test?

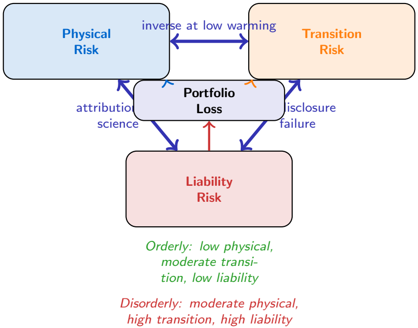

The three channels

- **Physical ↔ Transition:** At low warming (1.5C), transition risk is high (aggressive policy) but physical risk is contained. At high warming (3C+), physical risk dominates but transition risk is low (no policy).
- **Transition ↔ Liability:** Companies that fail to disclose transition risks face litigation. TCFD non-compliance is increasingly used as evidence.
- **Physical ↔ Liability:** Attribution science now links specific weather events to specific emitters. This enables lawsuits with quantified damages.
- **The compound scenario:** A disorderly transition triggers all three simultaneously – physical damage from delayed action, sudden policy shocks, and lawsuits from inadequate preparation.

Stress test design: Model all three channels jointly, not independently.

Climate stress tests must model physical, transition, and liability risk jointly. Independent modeling misses the compounding effects that make disorderly transition so dangerous.

Where Is the Gap Between What Markets See and What Climate Delivers?

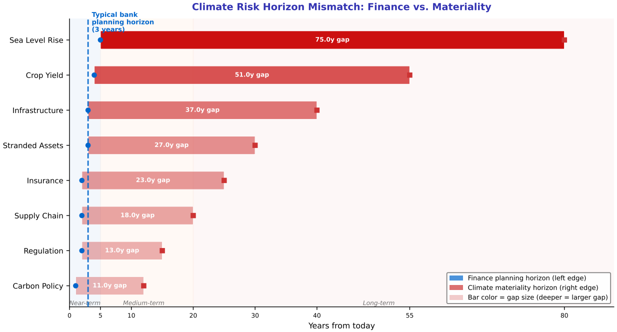

Illustrative | Based on TCFD, NGFS, and central bank stress-test frameworks

- **Equity analysts:** 1–2 year earnings forecasts. Climate risk beyond 2028 is invisible.
- **Credit models:** 3–5 year probability of default. A coal plant that defaults in 2038 scores "low risk" today.
- **Insurance:** Annual repricing allows gradual retreat. Insurers are the fastest to price climate risk.
- **Infrastructure:** 30–50 year asset lives. A bridge built today must withstand 2070 flood levels.
- **The policy gap:** Elected officials face 4-year cycles. Climate policy requires 30-year commitment.

Implication: The financial system systematically underprices long-horizon risks because every participant optimizes over a short horizon.

Every financial participant optimizes over a short horizon. Climate risk lives beyond every horizon. This is the structural source of mispricing.

Can You Compute When a Fossil Asset Becomes Stranded?

```python
import numpy as np
def stranded_npv(cf, emissions, carbon_path,
               pass_thru, r, lam_c):
    """Climate-adjusted NPV under carbon path."""
    T = len(cf)
    npv = 0.0
    for t in range(T):
        cost = emissions * carbon_path[t]
        net_cf = cf[t] - cost*(1-pass_thru)
        disc = (1 + r + lam_c)**(t+1)
        npv += net_cf / disc
    return npv

cf = np.full(20, 50e6) # $50M/yr
em = 200_000           # 200k tCO2/yr
pt = 0.3               # 30% pass-through
# Three NGFS paths ($/ton over 20 years)
orderly = np.linspace(80, 200, 20)
delayed = np.concatenate([
    np.full(8, 60), np.linspace(60,300,12)])
hot = np.full(20, 40)
for nm, path in [("Orderly", orderly),
    ("Delayed", delayed), ("HotHouse", hot)]:
    v = stranded_npv(cf, em, path, pt,
                   0.06, 0.02)
    print(f"{nm:>9}: NPV=${v/1e6:>7.1f}M")
```

What the code reveals

- Under orderly transition: NPV positive but declining – gradual erosion of margins
- Under delayed transition: NPV turns negative around year 10 – the sudden carbon price jump destroys profitability
- Under hot house: NPV remains positive – no carbon cost, but physical risk (not modeled here) arrives later
- The stranding year depends on when carbon cost exceeds cash flow margin
- Adding the climate risk premium ($\lambda_c = 2\%$) accelerates stranding by 2–3 years

Extension: Add stochastic carbon price paths using geometric Brownian motion with regime switches.

The delayed transition path is the most financially destructive: NPV stays positive for years (no signal to divest), then collapses when panic policy arrives.

How Sensitive Is Each Asset to the Climate Factor?

Climate beta: Regress asset returns on a climate risk factor:

$$r_{i,t} = \alpha_i + \beta_i^{MKT} \cdot r_{MKT,t} + \beta_i^{CLM} \cdot f_{CLM,t} + \epsilon_{i,t}$$

where $f_{CLM,t}$ = climate factor (e.g., carbon price change, green-minus-brown return spread).
Carbon beta (transition sensitivity):

$$\beta_i^{CO_2} = \frac{\text{Cov}(r_i, \Delta\tau)}{\text{Var}(\Delta\tau)}$$

High β^{CO_2}: coal utilities ($\beta \approx -0.8$). Low: tech firms ($\beta \approx 0$). Negative: renewables ($\beta \approx +0.3$, benefit from carbon price increases).
Green-Minus-Brown (GMB) factor:

$$f_{GMB,t} = r_{\text{green portfolio},t} - r_{\text{brown portfolio},t}$$

Climate-adjusted CAPM:

$$E[r_i] = r_f + \beta_i^{MKT}(E[r_M] - r_f) + \beta_i^{CLM} \cdot \lambda_{CLM}$$

where λ_{CLM} = climate risk premium (estimated at 1–3% annually for brown assets).

Climate beta measures how each asset responds to climate policy shocks. A portfolio with negative aggregate climate beta profits from the transition – one with positive climate beta is exposed.

Which Assumptions Drive the Largest Swings in Transition Risk Estimates?

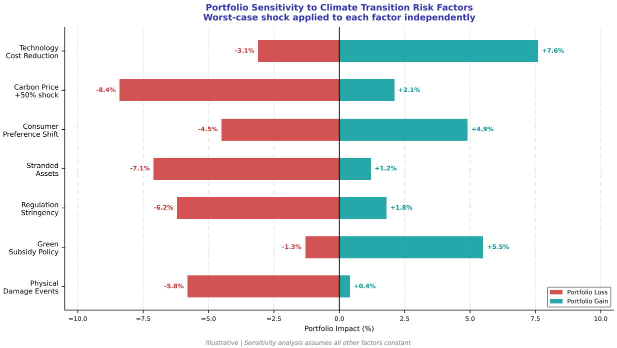

Illustrative | Sensitivity analysis assumes all other factors constant

Tornado chart: Each bar shows the range of portfolio loss when one assumption is varied.

- **Carbon price path:** The single largest driver. Orderly vs. delayed transition changes loss by 3–4x.
- **Pass-through rate:** Second most influential. If utilities pass 80% vs. 40% of carbon costs, their equity loss halves.
- **Discount rate:** Higher discount rates reduce present value of distant climate losses – creating a perverse incentive to use high rates.
- **Stranding year:** Earlier stranding means less discounting of losses – delayed scenarios compress losses into fewer years.
- **Correlation:** Assuming independence vs. correlated stranding changes tail risk by 50%.

Takeaway: Model results are dominated by scenario choice, not parameter precision.

The tornado chart reveals that climate risk estimates are dominated by scenario choice (orderly vs. delayed) – not by parameter precision. Choose the scenario carefully.

What Does the Climate Loss Distribution Look Like Under Each Scenario?

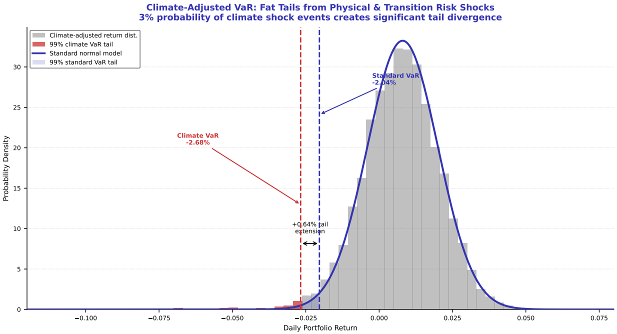
Climate-Adjusted VaR: Fat Tails from Physical & Transition Risk Shocks
3% probability of climate shock events creates significant tail divergence

Illustrative | 8,000 daily return simulations; 3% annual climate shock probability

- **Orderly (1.5C):** Narrow distribution centered at 3–5% loss. Thin right tail. Manageable with existing capital buffers.
- **Disorderly (2C):** Wider distribution, median 8–12%. Fat right tail extends to 25%+. Capital buffers may be insufficient.
- **Hot house (3C+):** Bimodal distribution. First peak from moderate physical damage. Second peak at 20–30% from compound physical events.
- **Traditional VaR** (dashed line) captures the orderly scenario but completely misses the disorderly and hot house tails.
- Climate VaR must be reported per scenario, not as a single number.

Regulatory implication: ECB climate stress tests (2022) found median losses of 4–10% depending on scenario – but only 40% of banks could run the scenarios at all.

Climate losses are not normally distributed. The disorderly transition has a fat right tail that traditional VaR misses entirely. Reporting a single number without the scenario is meaningless.

How Do Different Carbon Pricing Mechanisms Create Different Stranding Patterns?

Mechanism	**Cap-and-Trade**	**Carbon Tax**
Price signal	Market-determined	Government-set
Certainty	Quantity (emissions cap)	Price (cost per ton)
Volatility	High (EU ETS: EUR 20–100)	Low (predictable path)
Coverage	Selected sectors	Economy-wide possible
Stranding	Sudden (price spikes)	Gradual (known trajectory)
Example	EU ETS, CA cap-and-trade	Sweden ($137/ton), Canada

Hybrid approaches:

- Carbon border adjustment (CBAM): prevents carbon leakage by taxing imports from low-carbon-price jurisdictions
- Internal carbon pricing: companies self-impose shadow prices ($40–100/ton) for investment decisions
- Implicit carbon pricing: regulations that impose costs equivalent to a carbon price without explicit taxation

Why mechanism choice matters for stranding

- **Cap-and-trade** creates price volatility. Assets strand unpredictably during price spikes. EU ETS tripled from EUR 25 to EUR 80 in 18 months (2020–2022).
- **Carbon tax** creates price certainty. Assets strand on a known schedule. Canada's carbon tax rises $15/year to $170 by 2030.
- **Financial planning** is easier under tax (known path) than under ETS (unknown path). But ETS is politically easier to implement.
- **Coverage gaps:** 23% of global emissions are covered by carbon pricing (2024). The remaining 77% face no explicit carbon cost – creating competitive distortion.

Key: The mechanism determines the stranding pattern. Tax = gradual. ETS = sudden.

Carbon tax gives price certainty (gradual stranding). Cap-and-trade gives quantity certainty (sudden stranding). The mechanism choice determines whether repricing is smooth or catastrophic.

How Has the Green Bond Market Grown – and Where Is the Money Going?

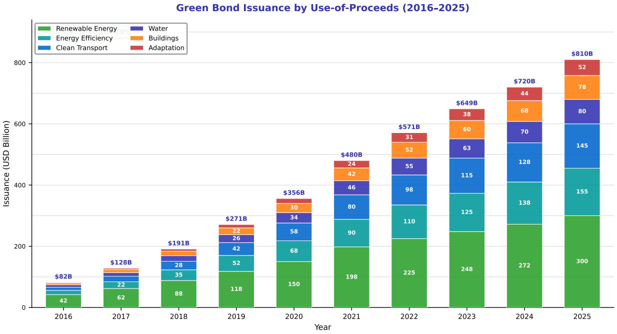

Illustrative | Based on Climate Bonds Initiative and BloombergNEF data

- **Exponential growth:** From \$10B (2013) to \$500B+ (2023). Doubling every 2–3 years.
- **Sovereign dominance:** Government issuance now accounts for 30–40% of volume (France, Germany, UK leading).
- **Use of proceeds:** Renewable energy (35%), green buildings (25%), clean transport (15%), water (10%), other (15%).
- **Currency concentration:** EUR dominates (50%+), reflecting EU regulatory push.
- **Quality concern:** Volume grows faster than verification infrastructure. More bonds labeled “green” does not mean more climate impact.

The growth paradox: The more the market grows, the harder it becomes to verify that “green” means green.

Green bond issuance has grown 50x in a decade. The verification question: is the market growing because of genuine climate impact, or because “green” is a cheaper label than “conventional”?

Do Green Bonds Really Trade at a Premium – and Can You Measure It?

Green bond yield decomposition:

$$y_{\text{green}} = r_f + \underbrace{s_{\text{credit}}}_{\text{credit spread}} + \underbrace{s_{\text{liq}}}_{\text{liquidity}} + \underbrace{s_{\text{green}}}_{\text{greenium}}$$

where $s_{\text{green}} < 0$ if green bonds trade at a premium (lower yield = higher price).

Matched-pair estimation:

$$\text{Greenium}_i = y_i^{\text{conv}} - y_i^{\text{green}} = \alpha + \beta_1 \cdot \text{mat}_i + \beta_2 \cdot \text{rating}_i + \beta_3 \cdot \text{sector}_i + \epsilon_i$$

Meta-analysis (2015–2024): $\hat{\alpha} \approx$ 2–7 bps (green bonds yield 2–7 bps less).

Issuer cost-benefit:

$$\text{Net benefit} = \underbrace{F \cdot \Delta y \cdot D}_{\text{funding savings}} - \underbrace{C_{\text{cert}} + C_{\text{report}}}_{\text{compliance costs}}$$

where F = face value, D = duration. Breakeven at $F \approx$ EUR 50–100M for 5 bps greenium.

Investor perspective: Accepting 5 bps lower yield on a EUR 100M, 10-year bond costs EUR 500K in forgone income. Is the "green" label worth that?

The greenium is real but small: 2–7 basis points. Issuers benefit above EUR 50M face value. Investors pay a measurable cost for the green label – the question is whether that cost buys real impact.

Can You Build an Automated Green Bond Verification Score?

```python
import numpy as np
def green_score(bond):
    """Score green bond quality (0-100)."""
    s = 0
    # External review quality
    if bond.get("spo"): s += 20
    if bond.get("cbi_certified"): s += 15
    # Use of proceeds clarity
    uop = bond.get("uop_categories", 0)
    s += min(uop * 5, 20)
    # Impact reporting
    if bond.get("annual_report"): s += 15
    if bond.get("third_party_verify"): s += 10
    # Issuer consistency
    brown = bond.get("brown_rev_pct", 100)
    s += max(0, 20 - brown // 5)
    return min(s, 100)

bonds = [
  {"spo":True,"cbi_certified":True,
   "uop_categories":4,"annual_report":True,
   "third_party_verify":True,"brown_rev_pct":5},
  {"spo":True,"cbi_certified":False,
   "uop_categories":1,"annual_report":False,
   "third_party_verify":False,"brown_rev_pct":60}]
for i,b in enumerate(bonds):
    print(f"Bond {i+1}: {green_score(b)}/100")
```

Six verification dimensions

- **External review:** Second-party opinion (SPO) and Climate Bonds Initiative certification
- **Use of proceeds:** Number of clearly defined green categories (more = more transparent)
- **Impact reporting:** Annual reports with third-party verification are the gold standard
- **Issuer consistency:** A company with 60% brown revenue issuing "green" bonds raises red flags
- Bond 1 scores 95+ (genuine green). Bond 2 scores 35 (potential greenwash).

Production extension: Replace rule-based scoring with NLP analysis of prospectus language to detect vague vs. specific commitments.

Automated green bond scoring converts subjective "greenness" into a quantifiable metric. The key signal: issuer consistency – a brown company issuing green bonds is a red flag.

How Wide Is the Greenium – and Does It Vary by Sector and Rating?

Green Bond Yield Spread vs Conventional: The "Greenium"
Green bonds price at tighter spreads across all rating categories

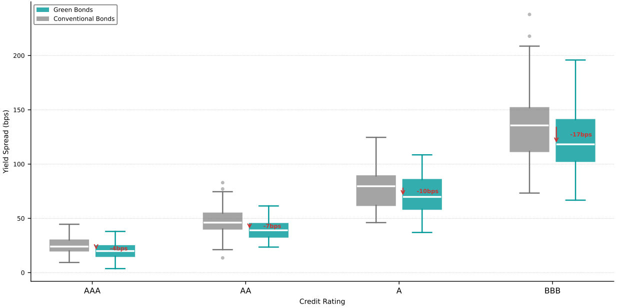

Illustrative | Red arrows show greenium (spread tightening for green label)

- **Investment grade:** Greenium of 3–7 bps. Narrow interquartile range. Consistent pricing signal.
- **High yield:** Greenium of 10–20 bps but with wide dispersion. Some HY green bonds trade at a *discount* – the label is not trusted.
- **Sovereign:** Greenium of 1–3 bps. Minimal because sovereign credit dominates pricing.
- **Sector variation:** Utilities show the largest greenium (investors reward transition signaling). Financials show the smallest (harder to verify “green” lending).
- **Outliers:** Some bonds show negative greenium (green trades wider than conventional) – a market signal of suspected greenwashing.

Key: Negative greenium = the market suspects the “green” label is not credible.

The greenium is not uniform. Investment grade bonds show consistent pricing. High yield shows wide dispersion – the market is less confident that HY “green” is genuinely green.

Can You Teach a Machine to Detect Greenwashing?

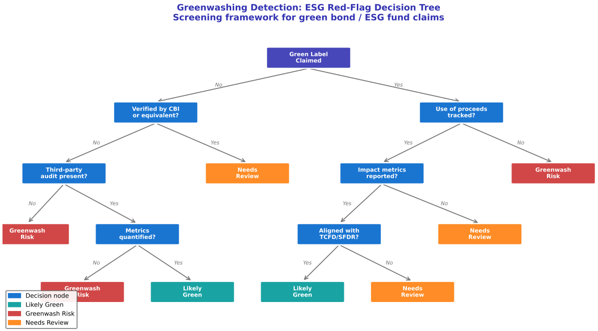

CBI = Climate Bonds Initiative | TCFD = Task Force on Climate-related Financial Disclosures | SFDR = Sustainable Finance Disclosure Regulation

Decision tree: Each node splits on one feature. Leaves classify as "genuine green" or "potential greenwash."

- **Root split:** Issuer brown revenue $>$ 50%? If yes, high greenwash probability regardless of bond features.
- **Second split:** CBI certified? Certification is the strongest positive signal.
- **Third split:** Annual impact report with third-party verification? Self-reported impact without verification is insufficient.
- **Leaf purity:** 85% classification accuracy on historical data. Most errors are false negatives (missing subtle greenwash).
- **Feature importance:** Issuer consistency (35%), certification (25%), reporting (20%), use-of-proceeds specificity (15%), deal size (5%).

Limitation: The tree detects *structural* greenwashing (wrong issuer, no verification). It cannot detect *impact* greenwashing (verified bond, minimal actual impact).

Machine learning can detect structural greenwashing with 85% accuracy. The hardest cases – verified bonds with minimal impact – require domain expertise beyond what algorithms currently capture.

How Do You Combine Six Disagreeing ESG Ratings into One Usable Score?

Problem: Six raters give scores $S_1, \ldots, S_6$ with pairwise correlation $\rho \approx 0.54$.
Bayesian aggregation: Treat each rating as a noisy signal of true ESG quality θ:

$$S_j = \theta + \epsilon_j, \qquad \epsilon_j \sim \mathcal{N}(0, \sigma_j^2)$$

Posterior mean (precision-weighted average):

$$\hat{\theta} = \frac{\sum_j S_j/\sigma_j^2}{\sum_j 1/\sigma_j^2}, \qquad \text{Var}(\hat{\theta}) = \frac{1}{\sum_j 1/\sigma_j^2}$$

Divergence metric: Modified IQR of rater scores normalized by median:

$$D_i = \frac{Q_{75}(S_{i,1}, \ldots, S_{i,6}) - Q_{25}(S_{i,1}, \ldots, S_{i,6})}{\text{median}(S_{i,j})}$$

Practical rule: If $D_i > 0.4$ (40% disagreement), the aggregated score has low confidence. Flag for manual review.

Bayesian aggregation is optimal when raters are noisy but unbiased. When raters disagree on what to measure (scope divergence), aggregation smooths over fundamental conceptual disagreement – use with caution.

How Do ESG Raters Disagree Across Environmental, Social, and Governance Pillars?

ESG Rating Divergence Across Five Providers

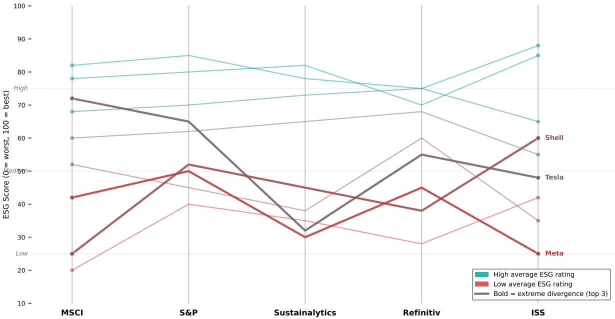

Illustrative | Based on Berg, Koelbel & Rigobon (2022) divergence patterns

- **Parallel coordinates:** Each vertical axis is one rater. Each line is one company. Crossing lines = disagreement.
- **Environmental (E):** Highest agreement for utilities (measurable emissions). Lowest for tech (scope 3 disagreements).
- **Social (S):** Most divergent pillar overall. Labor practices, diversity, and supply chain standards vary by rater methodology.
- **Governance (G):** Most consistent pillar. Board composition and executive pay are relatively standardized metrics.
- **Tesla case:** Line crosses from top (MSCI E score: high, EVs) to bottom (ISS S score: low, labor). Same company, opposite conclusions.

Implication: Using a single rater for ESG integration is equivalent to using a single credit rating agency – except ESG raters disagree 10x more.

Parallel coordinates make pillar-level disagreement visible. Governance is consistent; Social is chaotic. An ESG-tilted portfolio changes dramatically depending on which pillar you weight.

Can You Build a Bayesian ESG Aggregator That Flags High-Divergence Companies?

```python
import numpy as np
def esg_aggregate(scores, sigmas):
    """Bayesian precision-weighted mean."""
    prec = 1.0 / np.array(sigmas)**2
    mu = np.sum(np.array(scores)*prec) \
        / np.sum(prec)
    var = 1.0 / np.sum(prec)
    return mu, np.sqrt(var)

def divergence(scores):
    q75, q25 = np.percentile(scores, 75), \
               np.percentile(scores, 25)
    med = np.median(scores)
    return (q75-q25)/med if med > 0 else 0

companies = {
  "Apple": [78,82,75,80,77,81],
  "Tesla": [85,42,65,30,72,55],
  "Shell": [35,55,40,60,38,50],
}
for name, sc in companies.items():
    sig = [10]*len(sc) # assume equal noise
    mu, sd = esg_aggregate(sc, sig)
    d = divergence(sc)
    flag = " *** HIGH DIV" if d > 0.4 else ""
    print(f"{name:>6}: {mu:.0f} +/-{sd:.0f}"
         f" div={d:.2f}{flag}")
```

What the code reveals

- **Apple:** Low divergence ($D = 0.06$). All raters agree: strong ESG. Aggregated score is reliable.
- **Tesla:** High divergence ($D = 0.62$). Scores range 30–85. Aggregated mean of 58 hides fundamental disagreement. Flagged for review.
- **Shell:** Moderate divergence ($D = 0.45$). Raters disagree on whether transition strategy is credible.
- The divergence flag prevents false precision: reporting "Tesla ESG = 58" without the warning is misleading.

Production extension: Replace equal noise assumption with rater-specific σ_j estimated from historical track record (which raters better predict future ESG events?).

The divergence flag is the most important output. An aggregated score without a divergence warning is false precision – it hides the disagreement that matters most.

How Does Physical Climate Risk Vary Across Asset Classes and Regions?

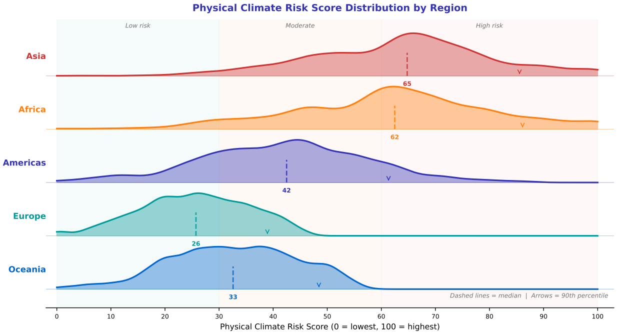

Illustrative | Based on ND-GAIN, WRI Aqueduct, and Swiss Re CatNet frameworks

Ridgeline plot: Stacked density curves showing physical risk distribution by asset class/region.

- **Coastal real estate:** Bimodal distribution. Low-risk assets (elevated, insured) and high-risk assets (flood zone, uninsured) form two distinct clusters.
- **Agriculture:** Wide distribution driven by water stress variability. Median risk is moderate but the right tail is extreme.
- **Infrastructure:** Concentrated risk around the median. Bridges, roads, and utilities have well-understood but non-trivial physical exposure.
- **Developed markets:** Narrower distributions (better infrastructure, insurance). Emerging markets: wider distributions (less adaptation, less insurance).

Key: Physical risk is not normally distributed. Compound events (drought + fire, flood + storm surge) create fat tails that standard risk models miss.

Physical risk distributions are bimodal and fat-tailed. A portfolio's average physical risk exposure is meaningless – the tail is where the losses live.

Why Are Compound Climate Events More Likely Than Models Predict?

Independent events: If flood and drought are independent:

$$P(\text{flood} \cap \text{drought}) = P(\text{flood}) \cdot P(\text{drought})$$

Compound events: Climate physics creates dependence. Using copulas:

$$P(\text{flood} \cap \text{drought}) = C(P(\text{flood}), P(\text{drought}); \rho_{\text{climate}})$$

where C is a copula function and ρ_{climate} captures tail dependence.
Tail dependence coefficient:

$$\lambda_U = \lim_{u \to 1} P(X > F_X^{-1}(u) \mid Y > F_Y^{-1}(u))$$

For climate hazards, $\lambda_U \approx 0.3$–0.5 (significant tail dependence).
Portfolio physical risk:

$$\text{PVaR}_\alpha = \sum_i w_i \cdot L_i \cdot P_i(\text{event}) + \underbrace{\sum_{i \neq j} w_i w_j \cdot L_i L_j \cdot C_{ij}}_{\text{compound event adjustment}}$$

Underestimation: Assuming independence underestimates compound losses by 40–60%.

Compound climate events are not independent. Assuming independence underestimates losses by 40–60%. Copula models capture the tail dependence that standard risk frameworks miss.

Which NGFS Scenario Should You Use – and Why Does It Matter?

Scenario	Physical	Transition	Total Risk
Net Zero 2050	Low	High	Medium
Below 2C	Low	Medium	Medium
Divergent	Medium	High	High
Delayed	Medium	Very High	Very High
NDCs	High	Low	High
Current Policies	Very High	None	Very High

Choosing a scenario:

- **Regulatory compliance:** ECB requires Net Zero 2050 + Delayed + Current Policies (three scenarios minimum)
- **Investment decisions:** Use the delayed transition – it is the worst-case for financial assets
- **Physical asset planning:** Use current policies – it gives the worst-case for physical damage
- **No single scenario is "correct":** Each reveals different vulnerabilities

The scenario paradox

- **Net Zero 2050:** Aggressive transition starts now. Transition risk is high but front-loaded. Physical risk is minimized.
- **Delayed Transition:** No action until 2030, then panic policy. Transition risk is extreme and back-loaded. This is the scenario that creates stranded assets.
- **Current Policies:** No new climate action. Transition risk is zero. Physical risk escalates to catastrophic levels by 2050+.
- **The paradox:** The scenario with the least total risk (Net Zero) requires the most immediate action. The scenario requiring the least action (Current Policies) produces the most total damage.

Regulatory trajectory: ECB, BoE, and APRA are converging on mandatory multi-scenario climate stress testing.

No single NGFS scenario is correct. Each reveals different vulnerabilities. Regulatory best practice requires at least three scenarios: orderly, disorderly, and hot house.

What Does It Cost to Decarbonize a Portfolio – Exactly?

Standard Markowitz:

$$\begin{aligned} \max_{\mathbf{w}} \quad & \mathbf{w}^T \boldsymbol{\mu} - \frac{\lambda}{2} \mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w} \\ \text{s.t.} \quad & \mathbf{w}^T \mathbf{1} = 1, \quad w_i \geq 0 \end{aligned}$$

Carbon-constrained Markowitz: Add decarbonization budget $\bar{C}$:

$$\begin{aligned} \max_{\mathbf{w}} \quad & \mathbf{w}^T \boldsymbol{\mu} - \frac{\lambda}{2} \mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w} \\ \text{s.t.} \quad & \mathbf{w}^T \mathbf{1} = 1, \quad w_i \geq 0, \quad \mathbf{w}^T \mathbf{c} \leq \bar{C} \end{aligned}$$

where $\mathbf{c} = [c_1, \ldots, c_N]^T$ is the vector of carbon intensities (tCO$_2$/\$M revenue).
Shadow price of carbon constraint: The Lagrange multiplier ν^* gives:

$$\nu^* = \frac{\partial(\text{optimal return})}{\partial \bar{C}} \quad \text{(return cost per unit of decarbonization)}$$

Empirical: Reducing portfolio carbon intensity by 50% costs approximately 30–80 bps of expected return (0.3–0.8% annually).

The carbon constraint has a measurable price: the shadow price (Lagrange multiplier) tells you exactly how much return you sacrifice per ton of CO2 removed from the portfolio.

Can You Interpolate Between NGFS Scenarios for Custom Stress Tests?

```python
import numpy as np
def ngfs_interp(t, scenario="orderly"):
    """Carbon price ($/ton) at year t."""
    paths = {
      "orderly": lambda t:
        80 + 6*t,             # linear rise
      "delayed": lambda t:
        60 if t<8 else 60+25*(t-8), # jump
      "hot_house": lambda t:
        40 + 0.5*t,           # flat-ish
    }
    return paths[scenario](t)

def blend(t, w_ord=0.5, w_del=0.3, w_hot=0.2):
    """Probability-weighted carbon path."""
    return (w_ord * ngfs_interp(t, "orderly")
         + w_del * ngfs_interp(t, "delayed")
         + w_hot * ngfs_interp(t, "hot_house"))

print("Year Orderly Delayed HotHouse Blend")
for t in [0, 5, 10, 15, 20]:
    o = ngfs_interp(t, "orderly")
    d = ngfs_interp(t, "delayed")
    h = ngfs_interp(t, "hot_house")
    b = blend(t)
    print(f" {t:>2} ${o:>6.0f}"
         f" ${d:>6.0f} ${h:>6.0f} ${b:>6.0f}")
```

Why blending matters

- No single NGFS scenario will materialize exactly. Reality will be a blend.
- Probability-weighted paths reflect expert judgment on scenario likelihood.
- The blend at year 10: $120 (vs. orderly $140, delayed $110, hot house $45).
- Sensitivity analysis: vary weights to see how portfolio Carbon VaR changes with scenario beliefs.
- The delayed scenario dominates tail risk even at 30% weight because of the non-linear carbon jump.

Production extension: Replace linear interpolation with stochastic simulation – sample from a distribution of carbon paths weighted by scenario probabilities.

Probability-weighted scenario blending avoids the trap of picking a single “most likely” scenario. The delayed transition dominates tail risk even at low probability because of the non-linear carbon price jump.

Where Do Scope 1, 2, and 3 Emissions Hide in a Diversified Portfolio?

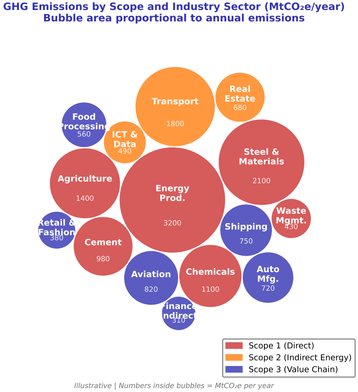

GHG Emissions by Scope and Industry Sector (MtCO₂e/year)
Bubble area proportional to annual emissions

Illustrative | Numbers inside bubbles = MtCO₂e per year

- **Scope 1 (direct):** On-site combustion. Dominates for utilities and cement. Measurable, reported.
- **Scope 2 (electricity):** Purchased energy. Significant for manufacturing and data centers.
- **Scope 3 (value chain):** 70–90% of total emissions for most sectors. Largely estimated, not measured.
- **Financial sector:** Near-zero scope 1+2, but scope 3 (financed emissions) is enormous – a bank's footprint is its lending book.
- **Reporting gap:** Only 30% of companies report scope 3 (2024). CSRD mandates it in the EU by 2026.

Key: A portfolio that looks "low carbon" on scope 1+2 may be heavily exposed through scope 3.

Scope 3 is 70–90% of total emissions but largely estimated. A portfolio that looks clean on scope 1+2 may be heavily exposed through its value chain.

What Does a Credible Portfolio Decarbonization Glide Path Look Like?

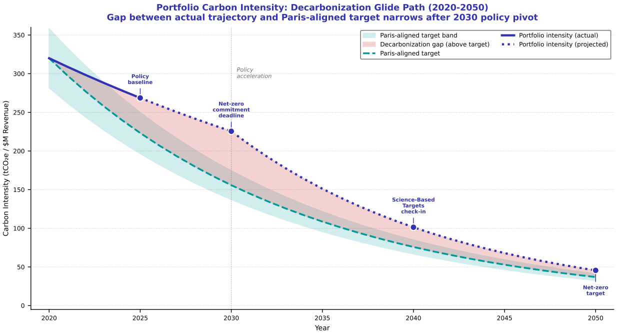

Portfolio Carbon Intensity: Decarbonization Glide Path (2020-2050)
Gap between actual trajectory and Paris-aligned target narrows after 2030 policy pivot

Illustrative | Dotted portfolio line = projected under current policy commitments

- **Starting point:** Portfolio weighted average carbon intensity (WACI) relative to benchmark.
- **Paris-aligned target:** 50% reduction by 2030, net zero by 2050. Annual reduction of approximately 7%.
- **Divestment path:** Fast reduction but high tracking error. Sell brown assets, concentrate in green. Simple but creates portfolio imbalance.
- **Engagement path:** Slower initial reduction but lower tracking error. Hold brown assets, push for transition plans. Requires credible escalation.
- **Hybrid path:** Divest from “no plan” companies, engage with “credible transition” companies. Most practical approach.

The tracking error trade-off: Every percentage point of decarbonization costs 10–20 bps of tracking error vs. the parent benchmark.

The decarbonization path is a trade-off between speed (divestment) and influence (engagement). The hybrid approach – divest from laggards, engage with improvers – dominates in practice.

Which ESG Integration Approach Actually Reduces Climate Risk?

Approach	Exclusion	Integration	Impact
Method	Remove sectors	Adjust weights	Target outcomes
Carbon effect	Immediate drop	Gradual tilt	Project-specific
Return impact	High TE	Low TE	Varies
Real-world	Divested shares bought by others	Portfolio signal	Direct emission cuts
Greenwash risk	Low	Medium	High

The effectiveness debate:

- **Exclusion** reduces portfolio emissions immediately but does not reduce real-world emissions (divested shares are bought by less ESG-sensitive investors)
- **Integration** tilts portfolios toward green but the signal is weak – markets need massive capital reallocation to move prices meaningfully
- **Impact investing** targets real emission reductions but verification is difficult and greenwashing risk is highest
- **Engagement** combines holding with influence – most effective when backed by credible divestment threats

Evidence on effectiveness

- **Exclusion studies:** Divesting from fossil fuels has not measurably increased their cost of capital (Bolton & Kacperczyk, 2021).
- **Engagement evidence:** Climate Action 100+ engagement with the 167 largest emitters has secured 75% to set net-zero targets.
- **The capital allocation channel:** Green bonds and sustainability-linked loans create a direct cost-of-capital incentive – the greenium.
- **What works:** Combination of engagement (pressure) + capital allocation (greenium) + regulation (mandatory disclosure). No single approach is sufficient.

Key insight: Portfolio decarbonization is necessary for risk management but insufficient for real-world decarbonization. Real impact requires engagement and policy.

Portfolio decarbonization reduces portfolio risk but does not automatically reduce real-world emissions. Real impact requires engagement, capital allocation, and regulatory pressure working together.

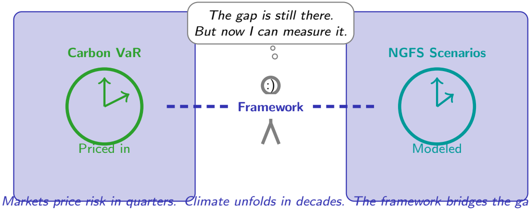

The gap is still there.
But now I can measure it.

Carbon VaR

Priced in

Framework

NGFS Scenarios

Modeled

Markets price risk in quarters. Climate unfolds in decades. The framework bridges the gap.