Reflections on CoC through Lasse's Beautiful Framework

Marcin Kacperczyk (Imperial College)

December 5, 2025

Central Message

- Lasse's central calibration:
  - If the social cost of carbon is about 43 $/tCO$_2$, the sustainable discount rate should move by $(43 - \tau) \cdot X/v$
  - Yet empirically the greenium corresponds to only about 0–10 $/tCO$_2$ and roughly 4 $/tCO$_2$ in the ICC regressions

- My reflections:
  - What exactly is reflected in the cost of capital estimates?
  - What does this mean for quantities?
  - What does this mean for policy?

Reflection 1:

Objectives and Mechanisms

Remark 1a: Objectives

- **Do markets price intensity or levels of emissions?**
- From an asset-pricing perspective, thinking of intensity is natural
  - Investors price "how much CO_2 do I get per dollar invested?" → intensity
- But from a welfare or policy perspective, this may not be the case
  - Climate damages depend on tons (levels), not tons per dollar of market cap
- If I care about the planet (portfolio's cost of capital): regress ICC on levels (intensity)
  - Bolton & Kacperczyk show that expected returns respond to **levels** and to **percentage changes** in levels, even controlling for size (closer to what policy cares about)

- Can we show that penalizing intensity at the portfolio level translates into lower **aggregate emissions**, not just **better-looking portfolios**?

Remark 1b: Aggregation & Taste

- Aggregate Green Score as unique ESG factor with the largest greenium and explanatory power
  - A powerful positive result, but normatively slippery
  - GS is an **equilibrium** object reflecting investor mandates, political views, and marketing constraints. There is no reason why that particular combination of E, S, and G should coincide with the social-cost vector (Sj)
  - From a CoC perspective, we may be loading the discount-rate wedge on the wrong combination of E, S, and G relative to the social-cost vector Sj.
- Should not infer society's welfare weights from the factor that capital markets currently price
  - Risk that capital is steered toward whatever happens to be salient to today's ESG investors (executive diversity, etc.) rather than toward the highest-damage externalities

Remark 1c: What is $4/t actually measuring?

- Lasse's work relies on ICC as a proxy for expected returns
  - ICC inherits analyst biases and coverage patterns: the implicit carbon tax is identified on large, covered, relatively clean firms
  - Data lags and backfills mean the information set in our regression is *not* the information set that priced the assets at the time
- All of the estimates inevitably bundle together three ingredients
  - Risk compensation (climate cash-flow risk, policy risk)
  - Mispricing
  - Tastes
- The 4$/t$CO_2$ is a serious but highly conditional estimate

Reflection 2:

Does this move the real world?

Remark 2a: Investment elasticities

- How much do spreads actually move real investment?
- Standard Q-theory evidence suggests investment is only mildly elastic to the cost of capital (technology and product demand often bind more tightly than WACC)
- What elasticity of brown investment to a 30-50 bp WACC wedge do we need for the cost-of-capital channel to matter for the climate? That's largely unknown.

Remark 2b: Boundary of capital markets

- Lasse's framework stated at the firm level, but estimates come from listed equities and bonds
- A big share of global emissions comes from:
  - state-owned enterprises; privately held firms; agriculture, buildings, and small emitters
- Even if sustainable finance *perfectly* implemented the 43 $/tCO$_2$ wedge in public markets, that may only cover a subset of global emissions (the larger mismatch between assets and activities that emit CO_2, the weaker the aggregate impact of the cost-of-capital channel)

- The right object is not just the size of the greenium, but something like **greenium × coverage × investment elasticity**. Current work pins down only the first term

Remark 2c: Dynamic responses

- Empirical specification regresses expected returns on X/v. That's a static cross-section. In practice, firms can react in very different ways to a greenium:
  - **Decarbonize:** invest in abatement technology and transition plans
  - **Divest:** sell the high-emitting assets to less constrained owners (PE, state-owned enterprises)
  - **Run down:** harvest cash, cut investment, and let assets depreciate
- These all lower X/v mechanically but have very different implications for aggregate emissions

- We need dynamic tests of firm-level responses to greenium exposure: asset sales, private equity ownership, and real decarbonization trajectories, not just WACC estimates

Reflection 3:

What about the policy?

Remark 3: Complementarity with policy

- Lasse: "society should **either** tax externalities **or** use investment decisions based on untaxed externalities weighted by social costs"
- View: CoC and policy are **complements** not substitutes
  - Small greenium helping change lobbying balance / constituency
  - Carbon pricing **amplifying** greenium via climate-risk channel

Cost-of-capital as a lever that interacts with taxation, regulation, and innovation

Conclusion

- For the research community, future work should jointly model:
  - the mapping from characteristics to priced factors (aggregation)
  - the decomposition of the greenium
  - the mapping from WACC to real investment and emissions

- Lasse's framework gives us a precise price wedge. Our task is to understand **what it really measures** and **what it does to quantities**