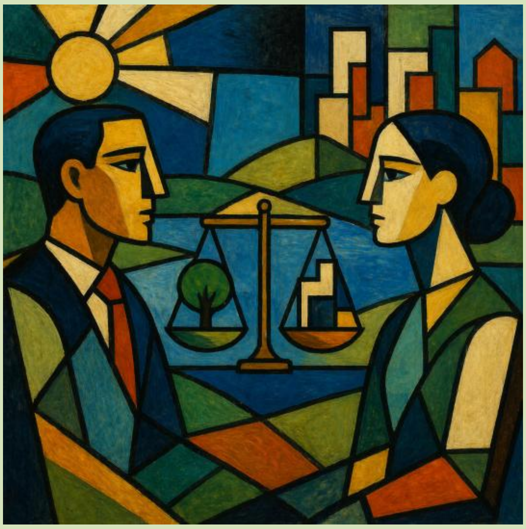

The Cooperation of Adversaries

Why Axelrod Belongs in the ESG Conversation, and Why the Plain Vanilla Version Does not Quite Fit

By MARTIN STUDER — April 2026

Abstract

Robert Axelrod's 1984 tournament is widely cited in sustainability writing as evidence that cooperation will emerge once players recognise the long game. This essay argues that the plain-vanilla reading misfits the ESG conversation in a specific way: the system is iterated, but the actors inside it are not. Politicians play to the next election, civil servants to the end of their mandate, executives to the length of their tenure. The gap between the system's horizon and the actor's horizon is the principal-agent problem at civilisational scale, compounded by overlapping prisoner's dilemmas that create epistemic cover for defection. Three forces determine whether a leader pushes past the truncation or settles into it: the upstream architecture of control, incentive alignment, and rationalisation. Together they form culture, understood as what people do when no one is looking, and culture is shaped by what the constituency is willing to tolerate in silence. Leadership of the repricing of nature cannot be handed uncritically to the long-horizon institutions, because these too are led by agents. Their role is narrower: they are pricing agents, not leaders. Common goods resist private capture. The ethical capitalist frame requires private actors who draw value from the commons to engage at least at replacement value. The leadership of the broader project belongs to a functioning, direct-democratic civic culture. Long-horizon institutions price. Citizens decide. Markets allocate within the boundaries citizens draw.

The tournament that should have been a footnote

In 1984, a political scientist at the University of Michigan ran a computer tournament that became a foundational text in applied ethics. Robert Axelrod invited game theorists from around the world to submit strategies for the iterated prisoner's dilemma. The winning entry was submitted by Anatol Rapoport. It was four lines of code and it was called Tit for Tat. Cooperate on the first move. Then do whatever your opponent did on the previous move.

Axelrod's finding, often summarised and almost as often oversimplified, was that nice, forgiving, provocable, and clear strategies tend to win against virtually any alternative, provided the game is iterated rather than one-shot, and provided the players expect to meet again.
I have come to think this is the most important insight in applied ethics that most executives have never read. I also think it is almost always applied too cleanly to the ESG conversation, in a way that obscures rather than illuminates what is actually going on.

Where the plain vanilla version stops working

The conventional use of Axelrod in sustainability writing goes something like this. Corporations, regulators, NGOs, and investors are locked in a prisoner's dilemma.

If they would recognise that the game is iterated, they would cooperate, and nature would get priced. If only they would read Axelrod.

The framing is appealing and not quite right. It assumes that the individual actors are themselves playing the iterated game. They are not. The system is iterated. The actors are playing much shorter games inside it. And the gap between the system's horizon and the actor's horizon is precisely the territory the economic literature has, for decades, called principal-agent theory.

The principal-agent problem at civilisational scale

Principal-agent theory describes what happens when the person acting, the agent, has different time horizons, different information, and different incentives than the person on whose behalf she or he is acting, the principal. The canonical example is the CEO and the shareholder. This is not a scandal. It is the foundational problem of modern corporate governance.

The ecological case is principal-agent at civilisational scale. The CEO is an agent. The current shareholder is one principal. Future generations are another principal, without a seat at the table. The politician is an agent. The current voter is the principal. The future citizen is, again, a principal without representation. Every actor in the system is acting on behalf of principals whose horizons are longer than the agent's tenure, and in many cases longer than the agent's life.
This is why the plain vanilla tit for tat does not quite fit. Axelrod's tournament assumed players who expected to keep meeting. The actors in the ecological-corporate system expect to play for as long as their role lasts, and then the game, for them, ends. The iteration is real at the system level. It is truncated at the level of every individual agent inside it.

The three actors and their truncated horizons

Consider them fairly.
A politician is playing a game whose primary horizon is the next election, rarely more than five years. What matters for her survival is whether the voter, or the voter she believes to be median, will reward or punish her at the ballot box. She may care privately about her grandchildren. She probably does. Altruism, however genuinely felt, will not get her re-elected.

A senior civil servant is playing a game whose horizon is the length of his mandate, which is longer than the politician's but still finite. He reports to ministers whose horizons are shorter than his, which caps his practical time horizon further.

A chief executive is playing a game whose horizon is the length of her tenure, typically four to eleven years, with median values closer to the lower end. Her LTIP vests over three to five years. Her share price is priced on a multiple of forward earnings that rarely looks past five. She wants to stay in role, which means she must optimise profit and valuation within the tenure she expects to have.

None is villainous. Each is playing, rationally, the game that defines his or her professional survival. And none will personally experience the consequences of the decisions they make today. The consequences arrive after their tenure ends.

Overlapping games and the space they create for hiding

There is a further wrinkle, and it is the one that makes the problem genuinely hard.

A leader is not playing one prisoner's dilemma. He is playing several simultaneously. Shareholders against competitors. Board against activists. Regulators against peer executives. Employees against the next restructuring cycle. Customers against the next margin review. The games overlap and interfere with each other.

This overlap is what makes hiding possible. A leader can defect in one game while cooperating in another, and then tell a narrative in which the cooperation is foregrounded and the defection disappears into the footnotes. He can invest in a headline sustainability programme while quietly cutting the maintenance budget on the watershed his supply chain depends on. Each observer sees a different, coherent version. Overlapping games create epistemic cover. And epistemic cover is what allows the truncated horizon to persist even when every actor privately knows it is not sustainable.

The three forces that together form culture

Culture is what people do when no one is looking. Not what they say when the camera is on. Not what appears in the annual report. What happens in the meeting that is not minuted, in the decision that is not sanctioned, in the scope creep that is tolerated rather than corrected.

Three forces, working together, produce that culture.

Control. By control I mean the upstream frame that defines what behaviour is even possible. The policies, procedures, risk tolerances, board charters, reporting lines, strategic planning horizons. The question is not whether a leader will be punished for stepping outside the frame. That is downstream. The question is how the frame is drawn in the first place. Are risk tolerances calibrated to the forty-year supply chain exposure, or to the five-year strategic plan? Is the board charter built to ask long-horizon questions, or to manage short-horizon reputation? The shape of the cage determines the shape of the leader's possible movements, and most cages today are drawn for the short horizon.

Incentive alignment. Does what the leader is rewarded for match the time horizon over which the consequences unfold? For most corporate leaders, the honest answer is no. Annual bonuses. Three-to-five-year LTIP vesting. Strategic goals designed to run to the end of a tenure. Personal life stages matter too. A leader with fifteen years to retirement behaves differently from one with three. Neither is wrong. Both are rational inside the incentive structure they inhabit.

Rationalisation. As long as a leader can demonstrate due compliance and move roughly in line with peers, he can credibly explain, to himself and to his stakeholders, that he is behaving responsibly. The herd provides the alibi. The competitive turn argument actually breaks this force, because it suggests that the middle of the herd is exactly where the loss is compounding. But until the competitive logic bites, rationalisation is what holds the structure together.

These three forces compound. A cage drawn for the short horizon makes incentive misalignment invisible. Invisible misalignment makes rationalisation easy. Easy rationalisation makes the constituency stop expecting anything different. The loop is self-reinforcing, which is why culture, in the sense I mean it, is so resistant to change.

Why the constituency, not the leader, produces culture

Culture is not made by the leader. Culture is what the constituency is willing to tolerate in silence. If voters accept scope creep without sanction, the scope creep becomes the norm. If a board tolerates quarterly framing in every leadership conversation, quarterly framing becomes the leadership default. If, in a democracy, free speech is not actively defended by those who benefit from it, free speech quietly erodes. Silence is cultural. Acceptance is cultural. What we do not sanction, we produce.

Who can actually lead, and who cannot

So when I say that the repricing of nature will not be led by the average politician, civil servant, or listed executive, I do not mean these actors are bad. I mean the principal-agent structure, the overlapping games, and the three cultural forces make it structurally unlikely that they will lead. They can follow, once price signals are clear enough. Leading requires something they cannot produce on their own.

It is tempting at this point to hand leadership to the long-horizon institutions. Reinsurers. Pension funds. Sovereign wealth vehicles. Patient private capital. Family offices. Their fiduciary horizons do extend past the individual iterated game. A pension fund's duty runs to a twenty-five-year-old worker who will retire in 2070. A reinsurance actuary models climate loss curves to the end of the century.

I want to resist that temptation, because these institutions are themselves led by agents. The pension fund manager still operates inside a truncated incentive. The sovereign fund is staffed by people with careers and bonuses. The family office principal has generational horizons but is also one person with blind spots. Handing leadership to them uncritically is a category error.

Their role is narrower and more precise. They are pricing agents, not leaders. They participate in the market's discovery of the correct price for ecosystem services, through their underwriting, their risk models, their capital allocation. They do not own the underlying systems, and should not. Common goods resist private capture. Watersheds, pollinator populations, coastal systems, soil carbon: these are priceable in the services they provide, but not ownable in any honest sense. The ethical capitalist frame, which I argue from, is clear about the boundary. Private ownership and risk-taking lead the way, within the limits that protect the common good, and at prices that reflect at least the replacement value of what is drawn from the commons.

Where leadership actually has to come from

The leadership of the larger project belongs to something else. It belongs to a functioning, direct-democratic civic culture, in which the people are the sovereign rather than the executive, in which public discourse is encouraged, fostered, trained, and practised, and in which free speech is actively defended rather than passively assumed.

This is the Swiss instinct, and it is the Popperian instinct, and it is the one that holds the whole system honest across generations. Long-horizon institutions price. Citizens decide. Markets allocate within the boundaries the citizens have drawn. This is not a call for more regulation. It is a call for more civic muscle, the kind that sets the boundaries within which free markets can operate without quietly liquidating the commons.

Closing

Axelrod did not set out to write a treatise on ecological leadership. He set out to understand cooperation. What he found was that cooperation is rational for anyone who expects to keep playing.

The question for the leader is whether to position the firm to cooperate, iteratively and credibly, with the small set of long-horizon counterparties who are genuinely playing at the time scale nature operates on. The reinsurer pricing supply chain risk. The pension fund whose fiduciary horizon matches its youngest member. The sovereign investor looking for assets still productive in 2070. The patient family capital that does not need to exit by 2030.

The question for the rest of us, the constituents who are not the leader, is harder and more important. It is whether we are willing to become the kind of sovereign public that draws the boundaries within which private ambition can safely operate. A leadership of agents alone will not solve this. A leadership of sovereigns might.

That is a softer argument than the usual Axelrod invocation. It is also, I think, a fitting one for the ESG conversation.

***

Sources and Further Reading

On cooperation and game theory. Robert Axelrod, *The Evolution of Cooperation* (Basic Books, 1984), remains the foundational text, including the full account of the Tit for Tat tournament and its implications. Anatol Rapoport's earlier work on conflict and cooperation frames the intellectual background.

On principal-agent theory. Michael C. Jensen and William H. Meckling, "Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure," *Journal of Financial Economics* 3 (1976), is the canonical starting point. Eugene Fama's work on agency and efficient markets provides useful extensions. For a contemporary application to corporate governance and time horizons, Lucian Bebchuk's scholarship at Harvard Law School is worth engaging.

On the commons and the limits of private ownership. Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action* (Cambridge University Press, 1990), is indispensable. It challenges the simplistic tragedy-of-the-commons framing and provides the empirical basis for the ethical capitalist argument that private activity must engage with the commons at replacement value. Garrett Hardin's 1968 *Science* essay, "The Tragedy of the Commons," remains useful as the foil Ostrom overturns.

On the open society and democratic discourse. Karl Popper, *The Open Society and Its Enemies* (Routledge, 1945), is the philosophical anchor for the civic conclusion of this essay. Popper's insistence that open societies depend on the active defence of free speech and public discourse is directly relevant to the argument about culture and constituency.

On direct democracy. The Swiss constitutional tradition, particularly the mechanisms of popular initiative and referendum, is the practical institutional expression of the civic argument made here. Wolf Linder, *Swiss Democracy: Possible Solutions to Conflict in Multicultural Societies* (Palgrave Macmillan, various editions), offers the most accessible English-language account.

On long-horizon institutional investment and nature. The Network for Greening the Financial System (NGFS) working papers, the Taskforce on Nature-related Financial Disclosures (TNFD) recommendations, and BioInt's *Biodiversity Briefs* series provide the technical infrastructure for the pricing argument. The 2023 Dasgupta Review, *The Economics of Biodiversity*, commissioned by HM Treasury, is the most rigorous public-policy treatment of ecosystem valuation to date.

On corporate time horizons and incentive design. Dominic Barton and Mark Wiseman's work on long-term capitalism, particularly their *Harvard Business Review* essays from 2014 onwards, is a useful bridge between the principal-agent critique and the practical reform of LTIP and board-level governance.

About Bluerain *Coherence* Series

This essay is part of a broader collection exploring how valuation, healthspan, collaboration, longevity, and biodiversity can be re-aligned with the rhythms of life — at a pace that enables regeneration rather than exhaustion.

The purpose of this compendium is to:

- Contribute meaningfully to public discourse on systems change and intergenerational strategy
- Offer fresh language and framing for leaders navigating complexity
- Share what we are learning, unlearning, and building in real time

We welcome dialogue, respectful critique, and collaboration — to deepen a shared understanding of what it means to steward coherence in a time of disruption.

For inquiries or conversation, please reach the author via linkedin.com/in/martinstuder.

About the Author

Martin Studer is an entrepreneur and impact investor focused on sustainability and longevity. Trained as an economist and certified public accountant, he brings decades of leadership experience in finance, governance, and executive management. As founder of Bluerain Partners, a Swiss family office, he works globally on projects spanning climate resilience, healthspan innovation, and intergenerational stewardship.

About Bluerain Partners Group

Bluerain Partners Group AG is a privately held, multi-generation family office headquartered in Switzerland. We do not offer commercial services or raise external capital, while some of our subsidiaries may do.

Our work is anchored in long-term purpose, driven by the belief that meaningful progress requires curiosity, collaboration, and courage.

We engage globally in projects that span climate resilience, health span and longevity, and the evolving nature of work and technology. We do so through investing, advising, governance, and venture development — always with an emphasis on coherence across generations, disciplines, and geographies.

These essays are authored by Martin Studer in a personal capacity. They reflect his experience, reflections, and aspirations — and may not necessarily represent the formal positions of Bluerain Partners Group AG, its subsidiaries, affiliated projects, or partners.

info@bluerainpartners.com.