The Translation Problem, Refined

Why climate data does not move capital, what the three distinct risk methodologies are, and where the insurance industry uniquely holds the key

A stress tested update following Ecosperity Week 2026 | Anthony Hobley

Why this version

My article on the eve of Ecosperity Week 2026 on this topic argued that we do not have a climate data problem, we have a translation problem. The data exists; the form in which it arrives is not the form in which capital allocation decisions are actually made. That central thesis holds, and the evidence for it has strengthened over the past year, from the BIS, the Bank of England, FCLTGlobal, and the 4th HLEG Report.

Feedback from readers of earlier articles has identified several specific points where the thesis can be further developed. Sophisticated readers from portfolio management, structured finance, and risk teams have provided valuable insights, and this version of the argument develops those points. It distinguishes more precisely between forecasts and scenarios; it concedes the current limits of forward curves at long horizons; it separates mark to market from mark to model; and it disentangles three risk types, physical, transition, and nature, that each require a different methodological response. It also addresses head on the discount rate question, which is where the broader argument about climate risk accumulating in the system actually plays out.

The thesis survives the stress test. What follows is a more nuanced and sophisticated argument.

This is not an argument that no one is working on the problem. Quite the opposite. The architecture is beginning to emerge in fragments: IPR for transition-policy forecasting, EIB/Howden-style cost curves for physical agricultural risk, bank-led work on climate-adjusted PD/LGD/EAD, and early nature-risk platforms for sovereign and ecosystem finance. The problem is that these remain fragmented pockets of innovation rather than a standardised valuation layer embedded in mainstream capital allocation.

The Paradox

We have the most comprehensive, institutionally validated body of risk data in human history. The IPCC has published six assessment reports synthesising the work of thousands of scientists. The TCFD framework has been adopted by regulators in over forty jurisdictions. CDP collects climate risk disclosures from more than twenty thousand companies. ESG ratings now cover the majority of global listed equity. Whilst climate analytics stand in contrast to mainstream financial analytics, the Inevitable Policy Response does produce a structured policy forecast covering 21 countries and 15 sectors.

And yet the people who allocate the world's capital, the finance ministers approving national budgets, the credit committees pricing twenty-year infrastructure loans, the asset managers running pension portfolios, are, on a day-to-day basis, making decisions as if that data does not exist.

This is not ignorance. The data has never been translated into the form they need to act on it.

The problem is not that investors are ignoring climate risk. It is that the signal arrives in the wrong analytical language.

The Language Capital Actually Speaks

Investment decisions are not made in degrees of warming, scenario narratives, or relative rankings. They are made in a precise and limited vocabulary built around four core elements.

1. Forecasts AND scenarios, not scenarios alone

This is the point at the heart of the mis-framing. Mainstream financial analysis uses both forecasts and scenarios. Much sustainable and climate finance analysis still relies primarily on scenarios, scores and disclosure outputs rather than forecast-based valuation inputs. That is the gap.

A forecast is a probability weighted central projection of a future variable: revenue, free cash flow, commodity price, interest rate, over a defined horizon, usually with explicit upside and downside ranges. It is a single most likely path of the world, with confidence intervals around it. Consensus estimates aggregated by I/B/E/S, FactSet, and Bloomberg are forecast averages.

A scenario, by contrast, is a coherent narrative of an alternative state of the world, used to stress test the central forecast. The standard convention is to run 1 in 20, 1 in 100, and 1 in 200 year stress scenarios alongside the base case. NGFS scenarios, IEA NZE/STEPS/APS, and IPCC SSPs are all scenario frameworks.

The asymmetry is decisive, though not absolute: capital is usually allocated around a base case or central forecast. Scenarios shape the risk envelope around that forecast, including covenants, capital buffers, downside cases and sometimes investment approval itself. They rarely, by themselves, price the deal. This is why TCFD scenario disclosure, however sophisticated, has not moved capital at the scale required. Too often, it sits in the wrong analytical box.

2. Forward price curves, where they exist

A forward price curve is the market implied price of a commodity, currency pair, or interest rate at specified future dates, derived from observable futures contracts or constructed from spot prices, storage costs, and supply demand fundamentals. The NYMEX WTI crude curve, the EEX German baseload power curve, the ICE EUA carbon curve, and the SOFR swap curve are canonical examples.

Forward curves do what scenarios cannot: they translate uncertainty into a single monetised number per period that can be inserted directly into a discounted cash flow. An infrastructure investment committee pricing a 15-year offtake does not ask which power price scenario to use; it asks what the forward curve shows for 2031, 2032, 2033, and what the implied volatility is around it.

A nuance the original article understated: forward curves are liquid only out to a limited horizon. Oil futures are liquid to roughly five years; power curves three to seven; carbon credit curves shorter still. Beyond those horizons, even mainstream finance relies on fundamental long term price forecasts produced by the IEA, Wood Mackenzie, BloombergNEF, or in house teams, not on observable market quotes. This means the contrast between mainstream finance and climate risk is not "they have forward curves, we have nothing." It is "they have a hybrid of observable curves out to five years and structured fundamental forecasts beyond that. We have neither for physical climate risk." That is the gap to fill.

A parallel point applies to transition risk. What we currently have for transition risk pricing is flimsy at best. Where transition risk is assessed at all, it is typically priced by reference to a carbon price that has itself been derived from a scenario. The resulting number carries the analytical weight of the scenario assumptions baked into it and is therefore unlikely to be taken seriously by mainstream analysts and investors who are used to forecast based pricing. Transition risk needs a much more robust forecast-based approach to price both the risk and the opportunity, of the kind the IPR has begun to provide and which is discussed later in this article.

3. Mark to model is more important than mark to market

Mark to market values an asset at the price observable in a liquid market on the valuation date: listed equities, government bonds, exchange traded commodities. Mark to model values an asset using a DCF or option pricing model calibrated to comparables, because no liquid market price exists. Private equity, infrastructure equity, real estate, illiquid bonds, private credit, carbon and biodiversity credits are all mark to model.

Roughly thirty to forty per cent of institutional capital sits in mark to model assets, and it is precisely these asset classes, long duration, illiquid, real economy, where climate physical risk matters most. The methodological gap is therefore not that climate risk is uniquely model based. It is that climate risk is missing from the inputs into models that already exist and are run thousands of times a day.

4. The valuation toolkit

The standard toolkit is consistent across asset classes. A discounted cash flow projects unlevered free cash flows over an explicit forecast horizon (typically five to ten years), applies a terminal value, and discounts everything back at the weighted average cost of capital (WACC). The cost of equity within WACC is derived from the Capital Asset Pricing Model: risk free rate plus beta multiplied by the equity risk premium. Beta measures the asset's sensitivity to systematic market risk.

Relative valuation uses multiples: EV/EBITDA, P/E, P/B, EV/Sales for industrials; P/NAV for resources and real estate; LCOE for energy assets; cap rates for property, drawn from comparable transactions or trading peers.

Credit analysis runs on three numbers: probability of default (PD), loss given default (LGD), and exposure at default (EAD). Expected loss equals PD multiplied by LGD multiplied by EAD. This is formally the same probabilistic framework insurers use for catastrophe risk. Credit risk and cat risk are methodologically siblings. Climate physical risk is the conspicuously absent peril in this framework.

Where ESG Analysis Actually Sits in the Stack

Once the toolkit above is laid out, it becomes clear what ESG and sustainable finance analysis as conventionally practiced does and does not contribute to it.

- **Negative or exclusionary screening:** removing sin stocks, fossil fuel producers, controversial weapons. This is a portfolio constraint, not a valuation method.
- **Positive or best in class screening:** ranking issuers within sectors by ESG ratings from MSCI, Sustainalytics, or S&P Global. These ratings are relative scores against sector peers, not monetised risk estimates.
- **Thematic investing:** allocating to themes (clean energy, water, transition metals) on the basis of macro conviction.

None of these go to fundamentals or valuation in the technical sense set out above. They sit alongside the DCF and multiples machinery as overlays or constraints. They do not determine cash flow forecasts, discount rates, or terminal values in any robust monetised way.

There is a fourth category, ESG integration, which claims to incorporate material ESG factors into the fundamentals themselves. In best practice cases at certain sovereign credit shops, some real estate teams, and a handful of asset managers, this is real and material. It deserves to be acknowledged rather than dismissed. But it remains the exception, and even where it is done well it relies on the same inputs (TCFD scenarios, ESG scores, disclosure quality) that the rest of this article identifies as the wrong analytical form. Even sophisticated ESG integration cannot produce a forward physical risk price curve out of inputs that were never designed to be one.

The honest version of the critique is therefore not that ESG fails by its own standards. It is that ESG was never built to do what fundamental valuation does. If we want the data resulting from climate and nature disclosure to move the dial, it needs to be redesigned to do so. Or perhaps not even redesigned but aligned with mainstream financial analysis, which is focused on investment fundamentals and valuation.

Because the truth is what we now have is a systems failure resulting in a failure to price the very real financial risks from climate change being incurred now and building up in the system through a subprime-like accumulation mechanism, albeit with different risk characteristics, timescales and transmission channels. The counterpoint is that the real investment opportunities which would address these risks are also not being priced correctly, and so are erroneously not seen as commercially viable.

Three Risk Types, Three Methodologies

Feedback has also focused on the fact that "climate risk" is not one undifferentiated problem. It is three distinct problems, each of which needs a different methodological response. Conflating them invites dismissal. Distinguishing them lets each claim its proper tool.

Physical risk: this is where the insurance methodology is strongest

Physical climate risk is fundamentally a stochastic hazard problem, exactly the kind of problem the catastrophe modelling industry solved between 1992 and 2003, under existential commercial pressure following Hurricane Andrew. The tools developed in that period (catastrophe risk models, stress tests at 1 in 100 and 1 in 200 year return periods, Average Annual Loss (AAL), Probable Maximum Loss (PML), and Exceedance Probability (EP) curves) are uniquely suited to translating physical hazard science into monetised probability distributions over time.

THE INSURANCE METHODOLOGY IN ONE LINE

AAL = the expected loss; sets premium and budget.

PML = the worst plausible loss at a return period; sets capital requirement.

EP curve = the distribution that links the two.

Together: a forward price for risk that a finance ministry, bond market, or credit committee can directly use.

The EIB and European Commission joint study on insurance and risk management tools for agriculture in the EU, published in May 2025 with support from Howden, is the cleanest contemporary proof of concept. The study found that the EU agricultural sector loses on average more than 28 billion euros annually to adverse weather, with climate induced losses projected to increase by between 42 and 66 per cent by 2050. Only 20 to 30 per cent of those losses are currently insured. Insurance backed by public funding is consistently more effective than government compensation paid after the fact.

That is a forward-looking climate risk cost curve and risk-layering proof point: a monetised, time bounded, sector specific estimate of how a defined category of climate risk will evolve, expressed in terms a finance minister or agricultural lender can directly insert into planning. It is what the broader financial system needs across other physical risk exposed sectors: coastal property, water utilities, energy infrastructure, ports, and agriculture beyond the EU.

Transition risk: forecast-based policy and technology analysis, not cat modelling

Transition risk is a different animal. It is not stochastic in the cat modelling sense; it is a structured forecasting problem about policy, technology cost curves, and demand shifts. The right methodology is therefore structured expert elicitation, not catastrophe modelling.

The Inevitable Policy Response, commissioned by the Principles for Responsible Investment and led by Energy Transition Advisers and Theia Finance Labs, surveys around 250 climate transition experts annually to produce median forecasts for policy and regulatory change across 21 countries and 15 sectors. The 2025 Transition Forecast generates forward carbon price curves (ranging from 30 USD per tonne in Saudi Arabia to 100 to 146 USD per tonne in Canada by 2030), timelines for coal phase outs, clean power adoption rates, and industrial decarbonisation.

This is the right shape of output for transition risk: forward, monetised, geographically and sectorally disaggregated, and updated regularly. It is the transition risk analogue of what catastrophe modelling does for physical risk. IPR is one important methodology rather than the only possible methodology. Insurance has limited natural advantage here. The insurance industry should be a sophisticated consumer of this type of transition forecast, not necessarily its producer.

Nature and biodiversity risk: an emerging hybrid

Nature risk lies between the two. Some elements (ecosystem collapse, fishery failure, pollinator decline) are stochastic in ways that lend themselves to insurance style probabilistic modelling. Others (policy on land use, market formation in biodiversity credits) are more like transition risk. The platforms now emerging in this space, including Natural Capital Reserve's sovereign marine biodiversity investment ratings, are explicitly hybrid: they combine cat model style risk adjusted indices with rating agency style governance assessments and integrate with sovereign debt instruments such as blue bonds.

Treating these three risk types as one is a recipe for muddled methodology. Treating them as distinct lets each claim its proper tool and lets the insurance industry stake a precise claim on physical risk without overreaching on transition or nature.

Why the Insurance Industry Has Not Yet Filled the Gap

If the insurance industry holds the methodology for physical risk forward pricing, why has it not produced and published the curves that the broader financial system needs?

The answer is structural, not technical. The sector's business model is built on twelve-month renewal cycles, with annual repricing and the option to non-renew. There is no commercial incentive to produce a 20 year forward price curve when next year's premium can simply be raised, restricted, or withdrawn entirely. State Farm, Allstate, and others have exited or severely curtailed California residential business. That is the market signalling that the twelve-month model is no longer adequate, and it sends precisely the wrong signal to investors, who read withdrawal as evidence of un insurability rather than as evidence of a broken pricing architecture.

Compounding the commercial disincentive, capital regulation can penalise insurers who attempt to offer long duration physical risk products. Solvency II and Basel III were not designed for long duration physical risk commitments and can treat them as capital intensive in ways that make them uneconomic for many carriers. But prudential regulation is only part of the structural incentive problem. It sits alongside demand-side unwillingness to pay, product complexity, model uncertainty, basis risk, political constraints on insurance pricing, and the lack of standardised climate-conditioned models.

Regulatory adaptation of Solvency II and Basel III to address this structural short-termism is one important frontier identified in the 4th HLEG Report: an area where the multilateral agenda has named the problem but has not yet delivered the solution.

Discount Rates and the Accumulation Problem

There is a deeper problem lurking beneath the translation gap, and it deserves to be addressed directly because it is the technical mechanism through which climate damages are kept out of decision relevant time horizons. That mechanism is the discount rate.

The mechanics are straightforward and brutal. A capital project evaluated at a 10 per cent discount rate values a one-hundred-pound cost in 2056 at roughly £5.73 today. At 7 per cent, it becomes £13.14. At 15 per cent, the corporate hurdle rate common in emerging market settings, it falls to £1.51. Whatever the climate science says about losses in 2050, those losses, once converted into a 2026 net present value calculation, disappear into the rounding error.

This is the technical mechanism behind what has been called "accidental time bias", and it is the reason why even sophisticated, well-intentioned investors who fully accept the climate science still find it rational to defer climate related capital expenditure. It is not denial. It is arithmetic.

The Stern Review's controversial use of a near zero pure time preference rate was an attempt to address this normatively, arguing that future generations' welfare should not be discounted simply because they are future. That ethical argument remains contested, and at any rate cannot be operationalised in private capital allocation, where investors are accountable to current shareholders and policyholders rather than to 2080. The normative discount rate debate, two decades old and unresolved, will not unlock climate finance.

But here is the point that the translation framework makes possible. The discount rate problem is not actually about discount rates. It is about where in time climate damages are modelled to occur.

When climate risk is modelled as a 2050 scenario, it sits in a time zone where any reasonable discount rate washes it away. When acute physical climate risk is modelled through insurance methodology, with expected annual loss, probable maximum loss, premiums, cover withdrawal, and probability distributions over current period losses, it can be brought into a time zone where discounting has minimal effect. A premium increase next year is a real cost next year. Withdrawal of cover now is a real reduction in current asset value now. A 1 in 100-year loss has an expected annualised cost that can hit the books every year.

The insurance methodology does not require investors to adopt a non-standard discount rate. For acute physical risks, it relocates climate risk from the deep future into the near present, where mainstream discounting practice is already fit for purpose. It does not solve every climate valuation problem, particularly slow-onset chronic risks, systemic tipping points and long-horizon nature degradation. But it is one of the clearest ways to make climate risk financially visible now.

The accumulation problem

From this follows the deeper concern, which connects this analytical question directly to the systemic risk argument. Every year of failure to invest in mitigation, adaptation, and resilience is not a year of preserved capital. It is a year in which an additional layer of unpriced climate risk is added to the financial system, just as every year of subprime mortgage origination from 2003 to 2007 added an additional layer of unpriced credit risk.

The subprime analogy is useful at the mechanism level, though the underlying risk, timescale and transmission channels differ. A risk that exists in reality but is not priced in market quotations accumulates until either it is priced in (a sharp repricing event, a market dislocation) or it is realised (the loss event itself, with all its non-linear, contagion driven consequences). The longer the accumulation period, the more violent the eventual correction can be.

The translation problem and the discount rate problem are therefore the same problem expressed two ways. The data exists. The methodology exists. What is missing is the translation layer that brings climate risk forward into time periods where it can be priced, where it shows up in next year's premium, in this year's covenant breach, in this quarter's revaluation. Once it does, no discount rate adjustment is needed. Standard finance does the rest.

The other side of the balance sheet

There is a corollary that follows directly, and it is the part of the argument that finally makes mitigation, adaptation, and resilience investment financially visible.

If the failure to invest builds up risk that is currently unpriced, then the act of investing builds up value that is currently unpriced. Investments in mitigation reduce a premium that has not yet been quoted. Investments in resilience reduce a probable maximum loss that has not yet been disclosed. Investments in nature-based adaptation reduce a sovereign risk premium that has not yet been calculated. The forward pricing of climate risk creates the missing other side of the balance sheet for the investments that today struggle to demonstrate financial return.

This is where the climate finance gap actually closes. Not by lowering the discount rate. By moving the timing of the priced cost forward, and by revealing the value of the avoided cost. Both effects happen in the near term, where the discount rate has no power to wash them away.

Anticipating the Counterarguments

It is worth engaging with the strongest potential challenges to this argument explicitly rather than leaving them to be raised. Five are worth taking head on.

"If forward climate risk pricing were truly decision useful, the market would already produce it."

This is the strongest counter, and the response is twofold. First, on the supply side, the Solvency II and Basel III treatment of long duration commitments and the twelve-month renewal model are part of a wider structural incentive problem that suppresses the supply of public forward pricing. Insurers run many of these calculations internally every day but have limited commercial reason to publish them. Second, on the demand side, the FCLTGlobal 2025 working group of major asset owners and managers found that CEOs were acutely concerned about how low carbon prices translate into weak investment signals, not that they disbelieved the underlying risk, but that the risk was not arriving in a form they could use. The demand is documented. The supply is structurally suppressed. That is a market failure that warrants intervention, not evidence that the data is unwanted.

"Long duration climate forecasts have wide uncertainty bands. A precise but wrong number is worse than an honest scenario range."

This is the next strongest counter, and again the response has historical precedent. The same critique was made of catastrophe modelling in the mid 1990s. The response was not to abandon the methodology but to build it iteratively and transparently, to disclose model assumptions and confidence intervals, and to triangulate across competing models (AIR, RMS, EQECAT). The same approach applies to long duration climate risk: imperfect forward pricing with disclosed uncertainty bands is strictly better than the current default of zero. The honest framing is that uncertainty in the forward price is itself decision useful. A price with a 50 per cent confidence interval still tells a credit committee something it cannot otherwise know.

"ESG integration already does this in best practice shops."

Some best practice ESG integration teams genuinely do incorporate material climate factors into fundamentals, particularly in sovereign credit, certain real estate franchises, and a handful of asset managers. The point is not that this work does not exist. It is that it remains the exception, that even best practice relies on the same TCFD scenarios and ESG scores as inputs, and that no team, however sophisticated, currently has access to a public, standardised, multi decade forward physical risk price curve at the granularity its models actually need. The market for the input does not yet exist. Building it benefits the integration practitioners as much as everyone else.

"Crude uniform discount rates are not how sophisticated investors work."

Sophisticated investors do indeed unbundle WACC by phase of the project, vary it by geography, and use term structure-based discounting. But two qualifications apply. First, the FCLTGlobal evidence suggests that even sophisticated investors collapse long duration risks into discount rate adjustments where they cannot model them directly, climate risk being the leading example. Second, corporate capital budgeting, which governs trillions in real economy investment decisions, routinely uses a single hurdle rate of 10 to 15 per cent across decades. The critique that all future risks are reduced to a single equity risk premium, what has been described by one commentator who provided feedback as "accidental time bias", is most acute at the corporate, sub investment grade end of the market, which is precisely where adaptation and resilience investment is most needed. The deeper response is given in the discount rate section above: the answer is not to argue with mainstream discount rate practice but to relocate the timing of climate cost into the near present, where standard discount rates do not wash it away.

"You are betting on a regulatory adaptation that may not come."

The strongest version of the argument is that the architecture is already being built by market participants, with regulatory adaptation as the catalyst that will accelerate rather than create it. The EIB Agriculture study, the IPR, Natural Capital Reserve's sovereign biodiversity ratings, the embedded natural disaster clauses in sovereign debt instruments, and the 4th HLEG report itself all demonstrate the architecture forming through market driven and policy driven channels in parallel. Regulatory adaptation of Solvency II and Basel III, when it comes, will codify what already exists. Where formal regulation has stalled, as Douglas's 2014 roadmap stalled in NGFS channels, the market has begun to deliver anyway. That is a stronger position than dependence on the regulatory cycle.

Three Shifts Required

The translation problem will not be solved by more disclosure frameworks or more sophisticated scenario analysis. It requires three structural shifts, all of which have begun.

First, change the insurance incentive structure. Longer duration insurance products (five, ten, twenty-five year contracts) would incentivise insurers to price climate risk forward, because they could no longer simply reprice at annual renewal. This requires regulatory adaptation of Solvency II and Basel III to reduce the structural penalty against long duration physical risk commitments, alongside action on demand, data standards, product design and model transparency. This is the regulatory frontier the 4th HLEG Report names but has not yet resolved.

Second, treat forward climate risk pricing as public infrastructure. The underlying inputs (hazard maps, long duration climate projections, catastrophe model outputs, expert transition forecasts) should not remain exclusively proprietary. Just as the reinsurance sector's transformation in the 1990s was enabled by progressive opening of its modelling infrastructure, a new generation of open-source catastrophe risk platforms and collaborative forecasting initiatives would allow the methodology to be applied across asset classes, geographies, and institutional contexts that individual market participants cannot reach alone.

Third, require forward climate risk disclosure in standardised monetary form. The TCFD framework established that climate risk is a material financial risk requiring disclosure. The next step is specificity of form: not scenario narratives but monetised forward exposure at standardised return periods, comparable across companies, sectors, and sovereigns. This would give rating agencies, credit committees, and sovereign bond markets the common language they currently lack and create the market signal that incentivises insurers to produce the forward pricing data in the first place.

The Bottom Line

The climate crisis is, among other things, a market failure, and like all market failures it is sustained by a missing information channel. In financial markets, that missing channel shows up as an absent or weak price signal. The data exists, in volumes never before assembled. The methodologies for translating different categories of risk into decision-relevant financial form also exist, though unevenly and across different communities. The early prototypes show the architecture beginning to form.

What is missing is systematisation. The pieces exist, but they have not yet been assembled into a common valuation layer that translates climate and nature risk into the operating units of mainstream finance. The thesis that survives the stress test is this:

THE THESIS, HARDENED

The climate finance gap is sustained not by absent data but by a missing translation layer.

The inputs institutional capital allocation already requires, monetised, forward, decision-relevant risk signals, are systematically produced for energy, rates, and credit, but not yet for physical climate risk.

For acute physical risk, much of the relevant methodology exists inside the insurance industry.

The structural disincentive against producing these signals publicly is partly regulatory, and partly commercial, data, product-design and political economy related.

The discount rate problem is not solved by changing discount rates but by relocating climate cost into the near present where standard finance can price it.

The early prototypes (EIB Agriculture, IPR, NCR and related nature-risk platforms) show the architecture beginning to form.

The pieces exist, but they have not yet been assembled into a common valuation layer.

The next phase of climate and nature finance is not disclosure. It is valuation infrastructure.

The transition will not be financed unless it is first priced. The pricing has begun.

Anthony Hobley is the former Deputy Chair, Climate Risk & Resilience at Howden Group and Co-Lead Insurance, High Level Climate Champions. This article draws on joint research conducted with Rowan Douglas CBE, Craig Pettengell, Sagarika Chatterjee, and the BCG team. The author thanks Nigel Topping for the constructive challenge that led to this hardened version of the argument.

***Climate Market-Maker | Strategic Change Broker*: Aligning capital, profit, risk, opportunity & impact by design**

For 35 years Anthony has worked at the intersection of capital, policy and impact: across environmental law, carbon markets, financial markets, heavy industry, philanthropy and most recently insurance. He is recognised for spotting tipping points and brokering the strategic change needed to capture them. He is an expert at leveraging organisations' institutional strengths to open doors and create trusted relationships that conventional outreach cannot deliver. He combines scientific training, the regulatory fluency of a senior corporate lawyer, and nearly two decades of front-line transactional experience (M&A, project finance, brownfield, commodities) with three decades of executive and convening leadership across the private, public and not-for- profit sectors.

Described as an *"Ambassador for carbon markets"* by Chambers Legal Directory and named in the Trillion Fund's top 25 most influential people in renewable energy.

Supporting Documents and Reports

Core regulatory and analytical support

1. Bank of England, Measuring Climate-Related Financial Risks Using Scenario Analysis, Quarterly Bulletin, 2024. Supports the distinction between scenarios and forecasts and the need to connect climate pathways to granular financial-risk modelling.
2. NGFS Climate Scenarios Technical Documentation and related scenario guidance. Supports the point that NGFS scenarios are not forecasts or predictions, but plausible pathways used for analysis and supervision.
3. BIS Papers No. 130, Pricing of Climate Risks in Financial Markets. Supports the claim that climate risks are only partially and inconsistently reflected in market pricing, with gaps in data, methodologies and comparability.
4. NGFS, Leveraging Physical Climate Risk Data, 2025. Supports the translation of physical climate information into expected-loss, PD/LGD/EAD and other credit and market risk use cases.
5. UNEP FI and Oliver Wyman, Bridging Climate and Credit Risk, 2025. Supports the practical integration of climate risk into bank credit modelling and probability of default analysis.

Physical-risk and insurance methodology

1. Rowan Douglas, Integrating Natural Disaster Risks and Resilience into the Financial System, Willis Research Network Concept Note, June 2014. Early blueprint for translating natural disaster risk into financial-system decision tools.
2. WEF Private Workshop Summary, Integrating Natural Disaster Risk into the Financial System, Geneva, 16 June 2014. Documents early institutional

appetite for applying insurance and catastrophe-modelling methods beyond insurance markets.

3. CISL / ClimateWise, Physical Risk Framework. Strong precedent for applying catastrophe modelling tools to investors, lenders and asset owners.
4. Howden and BCG, The Bigger Picture: The $10 Trillion Role of Insurance in Mobilising the Climate Transition, June 2024. Supports the role of insurance as enabling infrastructure for climate finance.
5. Howden and BCG, The Great Enabler, November 2024. Provides insurance use cases and sectoral examples for mobilising climate finance.
6. Howden and BCG, The Insurability Imperative, 2025. Frames insurance as a strategic tool for national and global climate transition planning.
7. EIB and European Commission, Insurance and Risk Management Tools for Agriculture in the EU, May 2025, with support from Howden. The strongest contemporary proof point for translating physical climate risk into a forward-looking cost curve, insurance gap and risk-layering architecture.

Transition-risk forecasting

1. Inevitable Policy Response, 2025 Transition Forecast, February 2025. Strong support for forecast-based transition-risk analysis across countries, sectors, carbon prices, policy timelines and technology pathways.
2. FCLTGlobal, Ahead of the Curve: Factoring the Cost of Carbon into Long-Term Decision-Making, 2025. Closely aligned with the thesis that future carbon costs need to enter long-term financial models.

Nature, planetary boundaries and natural capital

1. NGFS, Nature-related Financial Risks: A Conceptual Framework. Supports the treatment of nature degradation and nature-positive action as macroprudential, macroeconomic and microprudential financial risks.
2. Stockholm Resilience Centre planetary boundaries materials and Earth Commission safe and just Earth system boundaries work. Supports the Earth-system resilience and safe operating space framing.
3. UK Government, Enabling a Natural Capital Approach (ENCA) guidance. Supports the use of natural-capital valuation for decision-making.
4. Natural Capital Reserve, Sovereign Marine Biodiversity Investment Ratings. Useful as an emerging example of market infrastructure attempting to translate sovereign biodiversity and marine natural-capital performance into investment-relevant signals. It should be framed cautiously as an emerging market example rather than established mainstream practice.

Climate finance architecture and policy

1. HLEG 4th Report, Delivering an Integrated Climate Finance Agenda in Support of the Baku to Belem Roadmap to $1.3T, November 2025. Supports the strategic case for insurance, risk transfer and financial architecture in climate finance.
2. Circle of Finance Ministers Report, November 2025. Provides operational case studies and policy pathways for insurance and risk management tools in the climate finance architecture.
3. The Economics of Climate Change: The Stern Review, Nicholas Stern, Cambridge University Press, 2007. Foundational reference for discounting, intergenerational welfare and the economics of climate risk.

Annex: Landscape Review - Who Else Is Building the Valuation Layer?

The argument in this paper is not that nobody else has identified the problem. A growing body of regulatory, banking, investor, insurance and natural-capital work is moving in the same direction. The point is that these efforts remain fragmented. They have not yet been assembled into a common valuation layer that translates climate and nature risk into the operating units of mainstream finance.

1. Regulators and central banks

Central banks and financial regulators are making adjacent arguments. The Bank of England has focused on translating climate scenarios into granular financial-risk analysis across sovereign bonds, corporate bonds, mortgages and other asset classes. The NGFS has built the supervisory scenario infrastructure and increasingly recognises the need for physical climate-risk data that can feed expected-loss analysis. The BIS literature also recognises that climate risks are being priced only partially and inconsistently, with major gaps in data, comparability and methodology. These bodies are not yet using the language of a valuation layer, but they are building much of the supervisory plumbing that such a layer would require.

2. Banks and credit-risk teams

Banks are beginning to connect climate risk to credit models. UNEP FI and Oliver Wyman's 2025 work on bridging climate and credit risk is directly relevant because it addresses how climate risk can be incorporated into probability of default and credit-risk modelling. NGFS work on leveraging physical climate risk data similarly shows how climate information can be connected to expected loss, probability of default, loss given default and exposure at default. This is very close to the translation move described in this paper, but it remains an emerging practice rather than a mainstream lending standard.

3. Asset owners and internal carbon pricing

Long-term investors are moving in this direction through internal carbon pricing, transition assumptions and long-term cost-of-carbon analysis. FCLTGlobal's Ahead of the Curve is especially relevant because it focuses on how companies and investors can embed future carbon costs into long-term decision-making. This supports the transition-risk element of the thesis, but it is narrower than the full valuation-layer argument, which also covers physical risk, nature risk, insurance pricing, discount rates and mark-to-model assets.

4. Transition-risk forecasting

The Inevitable Policy Response is one of the most important existing examples of forecast-style transition-risk analysis. It is not simply another scenario exercise. It uses structured expert elicitation to produce forward-looking policy, carbon-price, technology and sectoral forecasts across multiple countries and sectors. That makes it one of the closest existing analogues to the kind of forecast-based transition-risk input that mainstream valuation models need.

5. Insurance and catastrophe modelling

The insurance sector already has the strongest methodological base for acute physical risk: catastrophe models, annual average loss, probable maximum loss and exceedance probability curves. The CISL/ClimateWise physical risk framework, Rowan Douglas's 2014 work on integrating natural disaster risk into the financial system, and the EIB/European Commission/Howden agriculture study all point in the same direction. The methodology exists. What has not yet happened is its systematic translation into mainstream capital allocation across banks, asset managers, finance ministries, sovereign bond markets and project finance.

6. Nature finance and natural capital

Nature finance is moving from disclosure toward risk and valuation, but it is much less mature. TNFD has accelerated disclosure. The NGFS nature-related financial risk framework connects nature degradation to macroprudential and microprudential risk. The UK ENCA guidance supports natural-capital valuation for decision-making. Early market-infrastructure examples, including sovereign biodiversity and marine natural-capital platforms, are beginning to connect nature performance, ecosystem risk and sovereign or project finance. But nature valuation remains fragmented across public natural-capital accounting, biodiversity credits, sovereign risk, insurance, project finance, ecosystem services and disclosure.

Is this already mainstream?

No. It is not yet mainstream. The precise answer is that it exists in pockets. Climate scenario analysis is now relatively mainstream for disclosure and supervision. Transition-risk forecasting is emerging through IPR, internal carbon pricing and long-term investor work. Physical-risk modelling is mature in insurance and only partially embedded in finance. Climate risk in credit models is emerging through banks, supervisors and UNEP FI work. Nature-risk finance is early-stage. What does not yet exist is a standardised, valuation-ready input layer for DCF, credit, insurance, sovereign-risk and factor-model analysis. That is the gap this paper identifies.