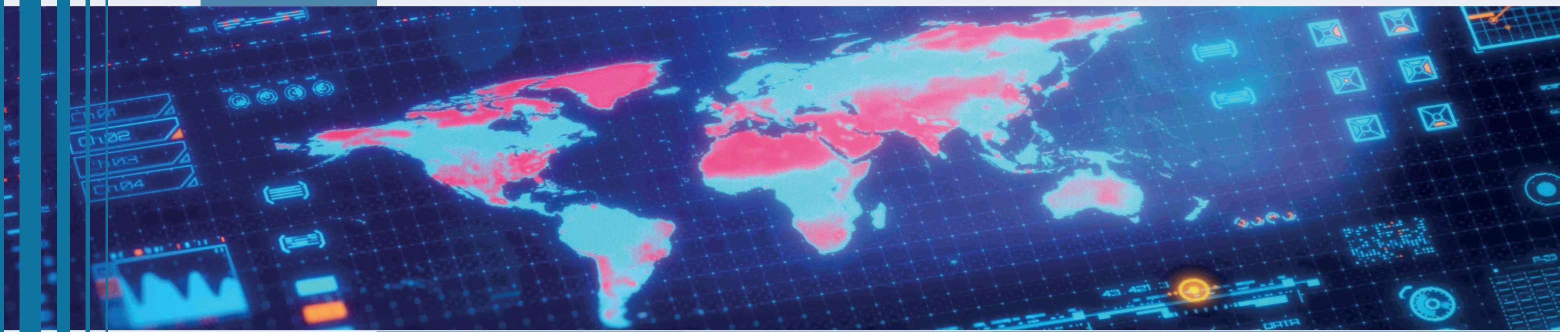

May 2025

An EDHEC Climate Institute Publication

How to Assign Probabilities to Climate Scenarios

EDHEC Climate Institute
London - Nice - Paris - Singapore

Table of Contents

Authors and acknowledgments

This publication was prepared by Riccardo Rebonato, Lionel Melin, and Fangyuan Zhang. It is a pleasure to acknowledge substantial contributions to this work offered by Dr D. Kainth (who is also specifically acknowledged in the text). Thanks are also due to Dr A. Schrapffer, Dr R. Estran and Mr F. Ducoulombier.

About the Authors

Lionel Melin, PhD is Associate Researcher at EDHEC Climate Institute. He works with the "Climate Scenarios" research team, focusing on the implementation of a novel Dynamic Integrated Climate-Economy (DICE) model. He is the CEO of MacroLucid Research, an independent economic consulting firm specialising in ESG finance. He began his financial career as an emerging markets economist and fixed income strategist for Deutsche Bank and later as a senior cross-asset strategist for Société Générale Lyxor Asset Management. He holds a PhD in economics from the University of Chicago and a Master's degree in statistics from the ENSAE in Paris. While a doctoral student, Lionel conducted research projects with the Federal Reserve of Chicago, the International Monetary Fund, Bank of America Merrill Lynch and BNP Paribas. He has held lectureships at the University of Chicago and has co-authored many academic papers published in the *Journal of Impact and ESG Investing*, *Journal of Asset Management* and *Journal of Portfolio Management*.

Riccardo Rebonato, PhD is Senior Advisor of EDHEC Climate Institute and Professor of Finance at EDHEC Business School, where he heads the "Climate Scenarios" research programme. He was previously a professorial visiting fellow at the University of Edinburgh (Political Economics and Sociology), a visiting lecturer at Oxford University (Mathematical Finance), an adjunct professor at Imperial College, London (Financial Economics) and a research fellow in physics at Corpus Christi College, Oxford. Riccardo has been Head of Derivatives Trading, Risk Management and Research and served on the boards of ISDA and GARP. Holding a PhD in nuclear engineering and condensed matter physics, he has more than 10 published books, including "How to Think About Climate Change" (Cambridge University Press), as well as 50 articles in refereed academic journals on the economics of climate change, derivative and asset pricing and risk management. He was named the 2022 PMR Quant Researcher of the Year by *The Journal of Portfolio Management*.

Fangyuan Zhang, PhD is a Senior Research Engineer at EDHEC Climate Institute where she contributes to the research programme on climate scenarios, seeking to understand and apply them while addressing their uncertainties. Fangyuan was previously a postdoctoral researcher for the Data Science Department at EURECOM, where she worked on the interdisciplinary research between economics and machine learning. She holds a PhD in economics, an MSc in finance and insurance from the University of Ulm in Germany and also Master's and bachelor's degrees in economics from Fudan University in China. She has published academic research on pension systems and portfolio optimisation and is a recipient of the Chris Daykin Prize (2023), the best paper award presented by the International Actuarial Association.

Executive Summary

The social and economic impacts of climate change are, literally, unprecedented. For this reason, forward-looking climate scenarios have played an unusually large role in the toolbox of policymakers, regulators and investors. For the present discussion, their most salient feature is that they almost never have probabilities attached to them. We argue that this strongly limits their applicability: policy makers need to understand where to focus their attention, what should be considered a 'tail event' and what a 'clear and present danger'; as for investors, from an asset pricing perspective prices are discounted expectations of future cashflows, but without probabilities the expectation cannot be carried out and the asset pricing attempt is stillborn.

We recognise that traditional financial scenarios are often provided without probabilities. With non-climate-related scenarios, the situation is, however, different, because in this case users can avail themselves of decades, and sometime hundreds, of years of macrofinancial information, which can provide either an 'intuitive' appreciation of the severity of a given scenario, or the basis for a proper probabilistic estimate. Since this is not the case with climate scenarios, scenario users requre a degree of probabilistic guidance for climate scenario that they typically do not need with 'traditional' financial scenarios.

Given this state of affairs, our goal is to make a contribution to the current body of work in two distinct directions. First, to help investors and policymakers, we propose two ways in which approximate probabilities can be assigned to scenarios created by a variety of scenario engines. Second, we recognise that 'narratives' (discursive descriptions of how the world may evolve, and of which abatement paths may be undertaken) have found widespread currency with the scenario community. We therefore show i) how and when probabilities can be associated to existing narrative; and ii) how new probability-endowed narratives can be created.

To fulfil the first goal, our strategy employs a two-pronged approach, following either what we call an informative (elicitation-based) route or a least-committal (maximum-entropy) method. With the informative approach, we make use of three pieces of information, namely:
1. fiscal, monetary and technological bounds on the feasible pace of abatement;
2. the distribution of views about optimal abatement policies elicited from professional economists; and
3. the observed disconnect between their recommendations on the desirable pace of abatement and actual abatement policy in order to obtain a probability distribution of abatement policy aggressiveness.

With the second (least-committal) approach, we do not assume that what economists recommend conveys useful information about future abatement policies, and we find the maximum-entropy distribution for the speed of emission abatement subject to the fiscal, monetary and technological bounds mentioned above, and to the observed social cost of carbon (or SCC, the 'carbon tax') today. This does not mean, however, that we know absolutely nothing about the likelihood of different abatement paths. With the maximum-entropy approach we are asserting that we should use this information – however limited – but no other.

This approach has solid theoretical and practical foundations, being rooted as it is in the maximum-entropy technique, which has widely and successfully been employed in fields as diverse as biological systems, natural language processing, statistical physics, and many others. It is particularly reassuring that the probability distributions for physical quantities, such as the end-of-century temperature anomaly, obtained from our two very different starting points, show strong similarities. Moreover, the small differences between them can be easily explained.

It is important to stress that the approach we present is, to a large extent, scenario-engine agnostic. By this we mean that in this study we present results that depend on our modelling choices for the non-policy quantities. However, the quantification of the policy uncertainty that is at the core of our work can be transported onto very different scenario engines. Indeed, we find that the DICE-model-inspired scenario model that we use is of sufficient flexibility that it can be 'calibrated' so as to recover the joint dynamics of the non-policy variables obtained by very different models – such as, but not limited to, the Oxford Economics or the Greening of the Financial System (NGFS) models.

Ultimately, the uncertainty in abatement policies is of great importance because it is a component of the uncertainty about abatement costs on the one hand (linked to what is usually called 'transition risk'), and about physical climate damages on the other. As for the latter, these are usually derived by applying a 'damage function' to a temperature distribution – a function, that is, that maps temperatures to damages. The temperature anomaly at a given horizon is therefore a powerful proxy not only for the whole climate, but also for its impact on macrofinancial quantities. By constraining our climate-economy scenario engine with the policy information gained by the methods above, we are able to attach probabilities to atmospheric CO_2 concentrations and therefore temperature anomalies at a chosen horizon.

When we do so, the first observation is that the likelihood of limiting end-of-century temperature increases to 1.5°C is very small: the exact value depends on the modelling choices, but these probabilities are never larger than a few percentage points. We stress that the goal is technologically achievable, but it would require a major and sudden alignment of actual abatement policy towards the consensus (median) recommendations of economists. Since economists have put forth these abatement recommendations for the best part of half a century, and their suggestions have gone largely unheeded, our method finds that the probability of an imminent correction of the politician/economist disconnect is very low.

Since the highest transition (abatement) costs are associated with the achievement of the 1.5°C target, this has direct asset valuation implications, as it gives low weight to the most 'costly' abatement paths.

Next, we find that the median 2100-temperature anomaly (around 2.75°C) is well above the 2.0°C end-of-century target, and that there is a significant probability (around 35-40%) that the temperature will exceed 3°C. To put these figures in perspective, the human species, let alone civilization, has never experienced temperature anomalies of 3°C or higher. Such high temperatures would push us into uncharted territory, increasing the likelihood of tipping points – sudden and potentially irreversible climate shifts triggered by crossing critical thresholds. These events, while difficult to predict, would severely challenge adaptation efforts and lead to significant physical damages. In any case, the large probability mass that we estimate for relatively high temperatures suggests that physical damages are likely to be greater than transition costs. Another way to look at our results is that our estimates assign low probabilities to high-transition-cost scenarios.

These are bold probabilistic statements, and a lot hinges on their credibility. Admittedly, in order to estimate what we call the 'informative' probability distribution (ie, the distribution derived from the economists' estimates of the optimal SCC) we have to make many and debatable assumptions. However, the least-committal (maximum-entropy) distribution we develop in parallel rests on much fewer, and much more easily defensible, assumptions: fundamentally, that we are very unlikely to divert to climate abatement a fraction of consumption greater than what we currently devote to health care (globally, 8-10% of Gross World Product); and that the current price of emission permits traded in regulated carbon markets is a reasonable proxy for the *actually implemented* SCC. The distinction between the recommended and the actually implemented cost of carbon is key to our approach, as we base on this difference the 'correction' of the economists' views about optimal carbon taxes. When we give 'reasonable' variations to our two key input quantities – where 'reasonable' can mean something as drastic as halving or almost doubling the cost of carbon permits – we find that our results change relatively little, and well within the range of uncertainty that we consider acceptable, given the nature of the problem.

1. Introduction

1.1 Why We Need Scenarios with Probabilities

Humanity is witnessing an unprecedented degree of anthropogenic change in climate.[1] Because of the historically unprecedented nature of the climate changes that may lie ahead of us, statistical tools, based on backward-looking information, are of little help. This is why forward-looking climate scenarios have become a major analytical tool for policymakers, regulators and investors.

It is in this context that an impressive body of work has been carried out, initially under the auspices of the Intergovernmental Panel on Climate Change (IPCC), and then by other organisations, such as the NGFS, to create a number of narrative-based reference climate scenarios. This work has been truly path-breaking, and has given policymakers, regulators and investors much-needed tools in their analysis of climate actions (such as emissions and abatements) and of their consequences. Our work firmly builds on these foundations.

For a variety of reasons, the original scenarios firmly eschewed any probabilistic statement. What we try to add to this well-established scenario framework is a probabilistic dimension. This, we believe, is sorely needed by all the climate-affected actors. Investors want to understand the effect of climate change on the valuation of their assets. But a valuing an asset means carrying out an expectation of its projected future cashflows. Carrying out an expectation, in turn, starts from assigning probabilities to what can happen. Without some notion of probability, the whole valuation project cannot even start. But it is not only investors who are keenly interested in valuation: regulators want to make sure that the value of the investment portfolios and collateral held by systemically important financial institutions may not be severely impaired by climate-induced losses, and they would like to know how likely the most severe impairments may be – Rebonato, Kainth, and Melin (2024) have recently argued, for instance, that in some scenarios the impairment of equity portfolios could be severe: but how likely are these scenarios? And as for policymakers, they certainly need guidance about what constitutes 'clear and present danger', and what can be considered a 'tail event'. For all these applications, some notion of likelihood is necessary. This probabilistic information can also inform the choice of scenarios used by corporates to guide strategic orientation and resilience planning by, as recommended by the Task Force on Climate-Related Financial Disclosures and integrated into law in certain jurisdictions such as the EU and the UK. Probabilities can add a layer of prioritisation to these initiatives.

Admittedly, financial scenarios are often provided without any probabilities attached to them. But in this case the scenario users can avail themselves of very long histories of financial data, on the basis of which formal or 'intuitive' probability assessments can be carried out. But when it comes to climate outcomes, it is again their unprecedented nature that calls for a more explicit assessment of their likelihood. In this paper we therefore try to enrich the existing scenario framework with approximate, but actionable, probabilistic information.

1.2 Narratives and the Current Scenario Architecture

Since our approach should be seen as complementary to the currently established scenario framework, it is useful to give a broad-brush description of the common features of the prevailing scenario architecture, and of the points of difference and contact with our proposed methodology. Examples of the 'prevailing scenario architecture' are the scenario approaches derived from the IPCC AR6 framework (made up of Shared Socioeconomic Pathways (SSPs) and Representative Carbon Pathways (RCP) – see IPCC (2019)) and the attending Process- Based Integrated Assessment Models. The NGFS scenarios also fall in this category.[2] We discuss in Section 5 how our approach can be integrated with this prevailing framework by looking in

1 - For a more technical presentation of the material to follow, see Rebonato, Melin, and Zhang (2025).

2 - There are many (overlapping) classes of scenario models. Broadly speaking, the SSP-RCP scenarios are designed to guide scientific research and policy discussion. Climate finance and risk management models (of which the NGFS approach could be an example) are designed to guide financial risk assessment and strategic planning for financial institutions, regulators, and policymakers. Economic and business impact models (such as the Oxford Economics scenarios) are designed to guide business strategy and economic impact assessment, providing actionable insights into how climate change and transition policies affect macroeconomic trends, industries, and regional economies. The International Energy Agency scenarios have been created to guide energy policy and transition planning. The boundaries between the different scenario classes are very porous: the NGFS model, for instance, is built on one specific SSP-RCP scenario, and many different scenario engines use in the background the same class of Process Based Integrated Assessment models (such as GCAM, which is used both by Oxford Economics and, in part, by the NGFS).

detail at the case of the Oxford Economics scenarios, but in this section we limit ourselves to a high-level description of the established conceptual framework, with a special focus on the concept of 'climate narrative'.

The SSPs, described in detail in Riahi et al. (2017) and van Vuuren et al. (2014), are narratives that describe a variety of social, polit-ical and economic global developments from now to the end of the century. SSPs describe alternative socioeconomic futures, focusing on factors like population growth, economic de-velopment, technological advancement, and social trends. They provide insights into how societal choices might affect vulnerability, resilience, and mitigation capacity in the face of climate change. Alongside SSPs, the IPCC framework specifies Representative Carbon Pathway (RCPs), which represent potential greenhouse gas concentrations at a given hori-zon and their associated radiative forcing levels.[3]

Each SSP is coupled with a Representative Carbon Pathway via a carbon tax: given a socioeconomic pathway, a given forcing (temperature) is achieved by charging an appropriate tax and using the proceeds to finance the abatement. Within the models, there is a one-to-one correspondence between terminal CO_2 concentrations, temperatures and forc ings, so an RCP specification effectively pins down a single end-of-century temperature, a single terminal CO_2 concentration, and all the emission schedules that add up to the same concentration.[4] The combination of an SSP with an RCP is a 'scenario'.

Given this temperature (emission/concentration) target (the chosen RCP), and given the economic pathway (the chosen SSP), a model (chosen from a set of Process-Based Integrated Assessment Models, PB IAM) finds the minimum-cost strategy to conjoin the two: in practice, this means that the PB IAM finds the mix of nuclear, fossil fuel, renewable, hydrogen energy, and possibly carbon removal technologies, that achieves the RCP target in the economically most efficient way. The financing for this abatement effort is assumed to be obtained through a carbon tax. The mix of technologies deployed to reach the target forcing varies a lot across PB IAMs (with some models giving, for instance, a big role to negative emission technologies, and others a negligible one), but aggregate quantities (such as temperature or cumulative emissions) are very similar across PB IAMs.

For our purposes, a few observations are important. The first is that the narratives are vivid and complex, but, ultimately, a 'narrative' is just an expressive label associated to a particular joint evolution of the macrofinancial quantities chosen to characterise the world. Whatever the name of the narrative, it is the combination of these paths that the underlying models really 'understand'. We find that, despite the sometimes very high number of variables that many scenario models set out to project, the effective model dimensionality is very low – between 2 or 3: this means that we can recover very effectively dozens of variables given the value of just two or three well chosen ('state') variables. We document in Section 5 that this is indeed the case for the Oxford Economics approach, but the observation is more general (it applies almost as strongly, for instance, to the NGFS scenarios). Speaking of paths of economic output and of emission schedules may not have what has been referred to as 'cognitive resonance', but models understand the narratives only in terms of well-specified paths for a small handful of macrofinancial quantities.

The second remark is that, by *starting* from narratives, it is easy to make some plausible worlds 'impossible'. For instance, in the Oxford Economics framework we discuss in Section 5, there are no paths of low economic growth associated with low end-of-century temperatures: this means that it is *impossible* to have what Acemoglou et al. (2012) call a 'Greenpeace world'.[5] Conversely, there is no world with high economic growth and high temperatures. The scenarios only cover the hand-picked trajec-tories for economic growth and

3 - Forcing is balance of energy (per unit time, per unit area) in minus energy (per unit time, per unit area) out.
4 - The last statement is only an approximation, because of natural reabsorption of CO_2, but, for reasonably smooth emission paths, the approximation is acceptable. See Appendix B for a discussion of this point.
5 - We note in passing that also some SSP/RCP pairings find no solution – ie, the parameters of the models that try to couple the GDP path and the final temperature cannot be adjusted to find a so lution. Implicitly, this means that these joint paths are 'impossible'. However, if a solution is found, perhaps by pushing the model parameters to their boundaries, the resulting SSP/RCP pair is treated as any other scenario, with no warning that it has been obtained 'under duress'.

emissions. Therefore, asking whether a set of narratives cover everything that can happen really means asking whether the combinations of GDP growth and emissions associated with these narratives really cover all that can happen. If they do not, then we know that something must have been left out.

1.3 Models, Uncertainty and Probabilities

Since the goal of our work is to associate probabilities with scenarios, it is important to understand how these probabilities can in principle be arrived at, and how they are linked to the various sources of uncertainty one encounters in the joint climate-economics system that we want to model. The considerations we present in this section (eg, the distinction between intrinsic stochasticity and model uncertainty or between conditional and unconditional probabilities) are general, and apply to any scenario framework. For the sake of concreteness, we present the discussion in terms of our scenario model. We stress, however, that the original offering of this work – the quantification of policy risk – can be transported to different scenario set-ups (as we show in Section 5).

The main *independent* sources of stochasticity/uncertainty that any scenario engine should capture pertain to:
1. economic growth;
2. the climate physics;
3. the damage function; and
4. the aggressiveness of the abatement policy.[6]

A scenario engine is a model that captures these different sources of uncertainty, and produces macrofinancial quantities its output. In our set-up the inputs are uncertain, hence the output macrofinancial quantities will be described by probability distributions. The fourth item is the key focus of this work, and the rest of the paper is devoted to its assessment. In this section we briefly discuss how we handle the uncertainty associated with items 1 to 3.

The first distinction is between (model) uncertainty, structural stochasticity, and policy uncertainty. Model uncertainty applies to a quantity that could in principle be known with high precision, but about which there still is scientific disagreement. An example is the uncertainty surrounding climate equilibrium sensitivity – a fundamental parameter in climate models that has a single 'true' value but remains uncertain due to limitations in current knowledge and modelling. Structural stochasticity applies instead to quantities (such as economic or population growth) for which we believe we have phenomenological model, but this model includes a stochastic component. And then we have policy uncertainty, for which we do not have a model, stochastic or otherwise, and that has to be quantified in a different manner. The reader is referred to Rebonato, Melin, and Zhang (2025) for adiscussion about how we model uncertainty in climate modelling, in economic growth and in the damage function. Important as these sources of stochasticity are, the quantification of abatement policy uncertainty is the focus of this paper. We show in what follows how we can use either bias-corrected opinion elicitation or maximum entropy to arrive at a probabilistic assessment of policy uncertainty. More precisely, we define as a *conditional* distribution for a focus quantity (say, temperature) a distribution that reflects the uncertainty in climate models and economic growth, but is conditioned on a particular emission path. An *unconditional* distribution for the same quantity is then obtained by assigning probability weights to the different conditional probabilities. We show in what follows that these weights are in turn obtained by characterising each emission policy by an effective abatement speed, and by assigning probabilities to these latter quantities.

6 - Abatement results from a combination of policy-driven, market-driven, and innovation-driven factors. The focus on abatement policy in discussions or modelling often reflects the importance of deliberate actions in accelerating and scaling emissions reductions where market forces alone might not suffice.

2.Data

Our strategy to obtain the informative unconditional distributions for the quantities of interest is made up of the following steps:

1. collect and 'curate' the economists' estimates of the optimal SCC;
2. associate to each SCC an effective abatement speed that fully describes the abatement policy;
3. using this relationship, from the curated distribution of SCC obtain the distribution of effective abatement speed;
4. correct this distribution for the economist-politician bias.

In this section we explain how the first step i s carried out.[7]

2.1 The Input Data

The SCC depends on the preferences (mainly, aversion to static risk, aversion to uneven consumption and preference for early resolution of uncertainty) of the agents in the economy and on the future, possibly state-dependent, consumption stream. More precisely, if we denote the today's welfare of the agents by $W(0)$, the SCC at time 0 of time-τ emissions, $SCC(0, t)$, can be obtained via the implicit function theorem as

$$SCC(0,t) = \frac{\partial W(0)/\partial e(t)}{\partial W(0)/\partial c(t)} \tag{1}$$

where $c(t)$ and $e(t)$ are the consumption and the emissions at time t, respectively. In words, the SCC expresses how much consumption changes for a change in emissions, everything else remaining the same. Because of the '*everything else remaining the same*' clause, this ratio is negative: higher emissions reduce consumption because of climate damages.[8] Since greenhouse gas emission is an externality, the SCC is the non-distorting tax that, in principle, would allow market mechanisms to obtain *a* (Pareto) efficient allocation of resources without the intervention of a 'benevolent dicta-tor'. The SCC is expressed either in \$/tonne of carbon emitted (t/C) or in\$/tonne of carbon dioxide emitted (t/CO_2).[9]

By surveying 207 papers, Tol (2023) has carried out a metastudy of the estimates (5,905 in total) of the SCC produced by professional economists.[10] In general the SCC is an (increasing) function of time, and the 5,905 quantities surveyed in the metastudy refer to the SCC at the time of the estimate, ie, of the quantity $SCC(0, t)$ in Equation 1. In addition, 94 of the 207 surveyed papers also provide an estimate of the growth rate of the SCC. Since the estimates collated and studied in Tol (2023) have been produced over the last 40 years, for comparability the dif-ferent estimates are grown or deflated to obtain 2010-equivalent values using the growth rate of the SCC. They have also been adjusted for inflation. To ensure consistency, we have reimplemented and reproduced Tol (2023)'s analysis. This time equalisation is the fist level of what we call 'data curation'. Tab 1 reports the descriptive statistics for the first-level-curated estimates of the SCC.

Table 1: Descriptive statistics of the first-round-curated SCC data (ie, for the raw data expressed in 2010 \$ equivalents per metric tonne of C).

Quantity	Mean	Std dev	Min	Max	5th perc	95th perc
SCC	106,881	2,165,172	−770.0	1.1 × 108	8.7	7,409.0

The first observation about the distribution of these SCC estimates is that they cover an extremely wide range: after rounding, from −\$1, 000 to +\$110 · 10^6 / tonne C.[11] To give a feel for the numbers, if we round the figure of the world GDP to \$100 Trl (\$ 10^{14}), and the figure of current world emissions to 40 Gt CO_2 ($4 \cdot 10^{10}$ tonnes) the estimate of an SCC of \$110 · 10^6/tonne C would imply a carbon tax far greater than the world GDP. Conversely, the lowest estimate suggests that the positive effects of carbon fertilisation are so

7 - A note about units. All the SCC data Tol (2023) are expressed in 2010-equivalent \$ per metric tonne of carbon (C), not of carbon dioxide. When discussing the SCC, for direct comparability we use consistently the same units. However, the prices of traded carbon permits are universally expressed in \$ per tonne of carbon *dioxide* (CO_2). Since these prices are well known and widely quoted, in that context, and only in that context, we also use \$ per tonne of carbon dioxide. When SCC for CO_2 is provided, the SCC estimate must be adjusted by the factor of 3.66, to reflects the approximate ratio of the 8 weight of a CO_2 molecule to a C atom.

8 - If everything else is not the same, the sign is indeterminate, because higher emissions are linked not only to higher damages, but also to higher economic output.

9 - When SCC for CO_2 is provided, the SCC estimate must be adjusted by the factor of 3.66, to reflect the approximate ratio of the weight of a CO_2 molecule to a C atom.

10 - Despite the fact that Tol (2023) denotes the SCC as \$/tC, the numbers reported in his metastudy refer to prices for tonne of CO_2 emitted. (see, Tol (2023), page 2: 'These are estimates of the social cost of carbon *for carbon dioxide emitted*', emphasis added.) Unless otherwise stated, all prices in the rest of the paper will refer to the cost per tonne of CO_2.

11 - The negative values are obtained on the basis of estimates of the positive effects of carbon fertilisation – see, eg, Chen et al. (2022).

large, and all the other effect so small, that we should currently *subsidise* emissions to the tune of several percentage points of GDP. The most charitable interpretation that can be given to these estimates is that they have been obtained without giving much thought to their implementability. However, they are clearly not just useless, but highly distorting for our purposes. We discuss in the next section how this problem can be solved.

2.2 Setting Lower and Upper Limits for the SCC Distribution

Because of the presence of such extreme values, parametric distributions do not fit these raw data well, and therefore Tol (2023) explores a number of non-parametric estimates. Whatever procedure is used to fit a distribution to the raw data, all the problems with the extremely wide range of values still remain. The 'date-equivalenced' data therefore needs to undergo a second round of 'curation' before any fitting (parametric or otherwise) can be attempted. We want to find a non-arbitrary procedure to cull the raw data.

We start from the low-value end of the data. We note first that the most negative estimates of the SCC is –\$770.8, but that the second-most-negative jumps to a value of – \$17.77. When we apply Extreme Value Theory to the raw data, we find that the probability of obtaining –\$770.8 given all the other negative values is more than ten orders of magnitude smaller than the probability for the second smallest estimate. We therefore exclude in further analysis the estimate of –\$770. We discuss later what to do with the remaining negative estimates.

For the upper cut-off, an extremely conservative value is obtained by capping the total carbon tax to the world GDP. Using 2010 values expressed in \$-equivalents (as in the analysis in Tol (2023)), we have a total GDP of $\$64.48 \cdot 10^{12}$, and global CO_2 emissions of $\approx 33.0 \cdot 10^9$ tonnes, or $\approx 50.2 \cdot 10^9$ tonnes of all greenhouse gases expressed in CO_2 equivalent terms. This would give an upper bound for the SCC of \$4,700/tonne C if all greenhouse-gas emissions are taken into account, or of \$7,200/tonne C if only carbon from CO_2 emissions are considered (as done in Tol (2023)).

Setting the cut-off for the total carbon tax equal to the world GDP is clearly unrealistic.

Tol argues that the total amount of tax that governments can raise is limited, and caps the additional tax at 15%. To test the sensitivity of the final results to this choice for the cutoff, we have also used an alternative procedure. Since global healthcare (which adds up to more than the combined expenditure for education and defence) is one of the largest ticket items for government expenditure, we have made the assumption that the global carbon tax should not exceed the global healthcare expenditure. With this choice, the lower limit for the SCC becomes \$709 or \$1080, depending on whether CO_2 only or CO_2-equivalent emissions are considered.

In sum: we shall assume in what follows that satisfactory estimates have been carried out
1. for the lower limit;
2. for a 'realistic upper limit', *RUL* (a flat 15% for Tol, or the healthcare-expenditurerelated value as we discussed);
3. for a 'theoretical upper limit', *TUL* (the whole GDP).

We want to give importance weights to the observations such that i) for values of the SCC greater than zero but smaller than the *RUL* all the weights should be 1; ii) for values of the SCC above the *TUL* the weights should identically be zero; iii) for values of the SCC above *RUL* and *TUL* the weights should decay monotonically, with continuous first derivatives at the two extremes; and iv) the weight function, $w(SCC)$, should have the minimum root-mean-squared curvature between *RUL* and *TUL*.[12] When we impose these conditions, we end up with cubic tapering (in essence a cubic spline – see Greville (1969) in this respect).

12 - We note that Tol uses linear tapering. Doing so, however, introduces discontinuities in the first derivative of the weight function to be interpolated.

Finally, we have to decide what to do with the negative estimates for the SCC. Once Tol's raw data have been date-equalised, weighted and fitted, there is an area of slightly less of 5% assigned to negative social costs of carbon. This raises an important modelling question. As we discuss in Section 3.4, there is a close positive link between the SCC and the emission intensity (emissions per unit of GDP). A negative SCC would mean that deliberate efforts are made to *increase*, rather than reduce, the amounts of emissions per unit of GDP. Alternatively, a negative SCC could be interpreted as a costly measure to inject *more* CO_2 in the atmosphere than the industrial processes require. Admittedly, there still are substantial subsidies for fossil fuels;[13] however, if we included these negative SCC values in our data set, our estimates of the expected temperature anomaly at all horizons would be even higher than what we find. To err on the side of conservatism, in what follows we therefore give zero probability to values of the SCC below zero.

Given this choice, Fig 2 shows that an excellent fit to the 'curated', weighted empirical distribution for values of SCC greater than 0 can be obtained as a mixture (convex combination) of a truncated Gaussian (TN) and a lognormal (LN) distribution:

$$\phi_{emp} = A \cdot [w \cdot TN(\mu_1, \sigma_1) + (1 - w) \cdot LN(\mu_2, \sigma_2)] \tag{2}$$

with μ_1 = 4.9638, σ_1 = 9.1582, μ_2 = 4.3797, σ_2 = 1.3206, w = 0.1940 and the constant A = 1.0602 takes care of the normalisation. Since the first moment of the distribution will play an important role in the analysis to follow, we determine the coefficients of the fitted density in Eq 2 in such a way to minimise both the squared differences b etween the fitted and empirical densities, and the squared difference b etween t he s ample mean a nd model expectation. The sample average is $157.29/tonne C and we obtain $157.19/tonne C for the corresponding expectation from the fitted distribution.

We now have a curated distribution that reflects the expectations of the SCC elicited from professional economists. We shall refer to this distribution as the 'informative' one. This distribution provides one of the two routes that we follow to generate a distribution of abatement speeds (the other route being the maximum-entropy approach). In either case, information about the actual cost of carbon today must be taken into account. The next section explains how this is done.

Figure 2: The fit to the empirical probability density (blue line) obtained using a truncated Gaussian and a lognormal distribution (red line). SCC in 2010 USD/tonne C on the x axis.

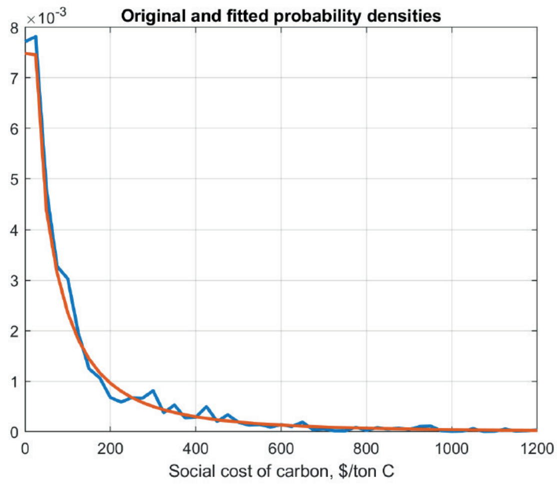

13 - In 2020 explicit fossil fuel subsidies were approximately 0.60% of world GDP (Black et al. (2023)), which corresponds to a (negative) carbon tax of approximately 15$/tonne CO_2.

3. Methodology

Our strategy is to build a distribution for the SCC that takes into account the information we have about the carbon tax actually implemented. We discuss how this can be done for the informative and the maximum-entropy distribution.

3.1 The Corrected Informative Distribution

As stated in the introduction, the distribution of social costs of carbon estimated by economist who do not face re-election concerns is likely to be biased (upwards) with respect to the politicians' distribution. Since both politicians' actions and economists' recommendations have been around for decades, we can estimate this bias by imposing that the first moment of the politicians distribution of SCC should match the current observed carbon cost. We therefore modify the curated economists' distribution of SCC in such a way that its mean coincides with the empirically observed carbon tax today.

We stress that making this adjustment is important: in a recent report (World Bank (2024)), the World Bank comments that '[p]rice levels continue to fall short of the ambition needed to achieve the Paris Agreement goals.' Since the average cost of carbon from the economists' distribution, \$154/tonne, is broadly in line with the achievement of at least the upper Paris Agreement target, there is a sizeable discrepancy between the actual total cost of carbon and what economists say this price should be. A substantial correction is therefore necessary.

To make this correction, we have assumed that the traded price of emission targets can be taken as a proxy of the actual SCC. From a recent comprehen-sive study of carbon pricing by the World Bank (2024), we estimate a global carbon tax of approximately \$30/tonne CO_2 (\$110.1/tonne C) and \$50/tonne CO_2 (\$183.5/tonne C). For comparability with the values in Tol (2023) (which are expressed in 2010 dollars per metric tonne of carbon, not of CO2), the latest values must first be reduced by a factor of 1.25 to account for inflation, and multiplied by the ratio of the weights of CO_2 and C, which is $44/12 \approxeq 3.667$. Furthermore, only a fraction of current emissions comes under some form of emission-trading-permit scheme, as the 36 carbon trading systems in opera-tion around the world by early 2024 cover approximately 20% of global emissions.[14] When these adjustments are applied to the prices of traded emission permits, the current SCC, expressed in 2010 dollars per tonne of carbon is in the range \$20/tonne C to \$40/tonne C. Since, as stated, we want to produce conservative temperature projections, we use the average of these estimates and take a forward-looking approach by assuming emission coverage of 30% or 40%, corresponding, after rounding, to effective carbon costs in the range of \$50/tonne C to \$70/tonne C.

At this stage we know by how much the first moment of the economists' distribution should be corrected. Formally, our problem can be stated as follows: starting from a distribution, $\phi'(x)$, about which we only know the normalization constraint, $\int_{\Omega'} \phi'(x) \cdot dx = 1$ (where Ω' denotes the finite support of $\phi'(x)$), we want to create a distribution, $\phi(x)$, such that, in addition to the normalization constraint, also the condition $\int_{\Omega'} \phi(x) \cdot x \cdot dx = \mu$ should be satisfied. (Here μ denotes the observed cost of traded carbon permits.) We show in Appendix C how this can be done by preserving as much as possible the shape of, and hence the information in, the unshifted distribution. Fig 3 shows the result, ie, the unshifted (labelled 'Original' ') and shifted (labelled 'KL' for Kullback-Leibler) distributions (the figure also displays the maximum-entropy distribution, discussed later).

14 - Source: International Carbon Action Partnership, 2024, available at https://icapcarbonaction.com/en/about-emissions-trading-systems

Figure 3: The original (blue line), Kullback-Leibler (red line) and maximum-entropy probability densities for the SCC for the case < SCC >= $70. SCC in 2010 USD per tonne C on the x axis.

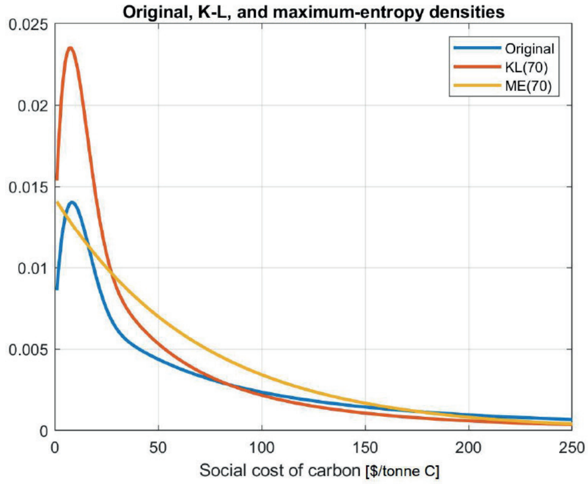

As we said, an alternative approach is to disregard whatever information there may be in the economists' distribution, and to take the least-committal approach – which amounts to assuming only that the expectation of the SCC should be equal to the observed price of the traded carbon permits. We show in Appendix A that the resulting (exponential) distribution for the SCC is given by

$$\phi(x) = K \exp(\lambda \cdot x)$$

with K and λ derived in Appendix A. The maximum-entropy distribution is also shown in Fig 3.

3.2 From the Distribution of SCC to the Abatement Distribution

Ultimately we want to assign probabilities to abatement policies characterised by different degrees of aggressiveness. To do this, first we must quantify the qualitative term 'aggressiveness' (possibly in terms of a single parameter); second, we want to establish a functional relationship between this 'aggressiveness parameter' and the SCC; last, we want to use this functional dependence in order to obtain the probability density for the aggressiveness parameter. We present these one step at a time.

3.3 Characterising the Aggressiveness of Abatement Policies

Abatement policies can, in principle, be very complex functions of time and state. It would therefore seem that characterising them in a parsimonious and tractable will be a very difficult task. We can simplify the task as described in Appendix B. The up shot of the analysis is that all abatement functions with same emission-weighted average produce (to a very good approximation) the same temperature at a given horizon. For reasons that will become clear in the following, without loss of generality we therefore choose a functional form for the abatement function $\mu(t)$ characterised by a single *free* parameter, κ, which describes the speed of the abatement policy:

$$\mu(t) = \mu(0) \cdot \exp(-\kappa \cdot t) + [1 - \exp(-\kappa \cdot t)] \tag{4}$$

The additional parameter of the distribution, μ_0, is calibrated so as to recover today's abatement, $\mu_0 = 0.05$. Fig 4 displays the single-parameter function $\mu(t; \kappa)$ for different values of the abatement speed, κ.

Apart from enjoying analytical tractability, the simple parametric form of Equation 4 allows the user to carry out some useful sanity checks: for instance, combining the information in the lines in Fig 4 with the information conveyed in Fig 3 one can see that a SCC of, say, \$250, corresponds to an abatement speed of approximately $\kappa = 0.1125\ y^{-1}$, which (from Fig 4) in turn implies that in just 10 years' time 70% of the economy will be decarbonised. Going back to Fig 3, this implausibly aggressive abatement policy would have a significant probability of occurrence for the unshifted (economists') distribution, but becomes very unlikely both for the shifted (politicians') distribution or for the maximum-entropy distribution. This stresses again the importance of adjusting the economists' distribution in order to obtain reasonable results.

Figure 4: The abatement function, μ(t), for different abatement speeds, κ. Time in years on the x axis.

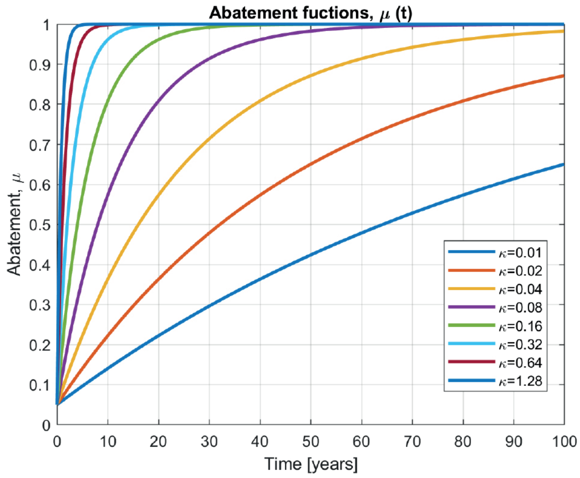

3.4 Establishing a Link Between the Abatement Speed and the SCC

In this section we explain why the choice of the functional form in Eq 4 for the abatement function is particularly useful for the task at hand. We show in fact that a close, smooth and monotonic relationship can be established between *the* one parameter of the abatement function μ and the SCC. To show this we proceed as follows.

First, we make use of a broad generalization of the DICE (Nordhaus and Sztorc (2013)) model described in Rebonato et al. (2024) to obtain the optimal abatement schedule, and we calculate the associated optimal SCC today. This generalisation nests many similar Integrated Assessment models as particular cases (in particular, it nests the Stern (2007) model for an appropriate choice of the preference parameters). We remind the reader that with DICE-like model savings and the abatement function are the control variables used to optimise the welfare of a utility-optimising agent who is exposed to climate damages to economic output. These damages can be mitigated through abatement, the cost of which increases exponentially with the degree of abatement, as captured by the function $\mu(t)$. We now input into the DICE-like models an abatement function, $\mu(t)$, of the form specified in Eq 4 : from the discussion in Section 3.3 we know that, as long as we are interested in a single projection horizon, T, this choice is not as restrictive as it seems. We repeat the optimiation exercise using:

- deterministic or stochastic (à la Jensen and Traeger (2014)) processes for the total factor of production;
- widely different assumptions about the best climate models;
- widely different assumptions about the damage function (ranging from the 'tame' Nordhaus and Sztorc (2013) quadratic formulation to the 'aggressive' Howard and Sterner (2017) with-tipping-point[15] parametrization);
- values for the utility discount rate ranging from 0.0010 (the value used in the Stern (2007) review) to 0.035 (more than double the value used in the DICE model);
- values for the elasticity of intertemporal substitution ranging from 0.5 (close to the value used in the DICE model) to 1.7 (close to the value used in the Bansal and Yaron (2004) approach);

15 - Howard and Sterner (2017) refer to this parametrization as the 'catastrophic' one.

- both time-separable and recursive (Epstein and Zin (1989)-type) utility functions;
- following Dietz and Stern (2015), damages to affect not just economic output, but also capital.

Each different choice for the preference, physics or economics parameters obviously produces very different estimates for the optimal value of the SCC 'today' ($SCC(0)$) and for the optimal abatement schedule, $\mu(t)$. However, one readily notices (see Fig 5) that higher values of the SCC are always associated with more aggressive abatement policies, and that there is a surprisingly tight relationship between the characteristic abatement speed, κ, and the optimal cost of carbon.

Figure 5: The abatement speed, κ_1 (years^{-1}, y axis) as a function of the optimal SCC (\$/tonne CO_2, x axis. The continuous curve is a LOWESS (Cleveland (1979)) quadratic smooth fit to the calculated points, shown as filled dots.

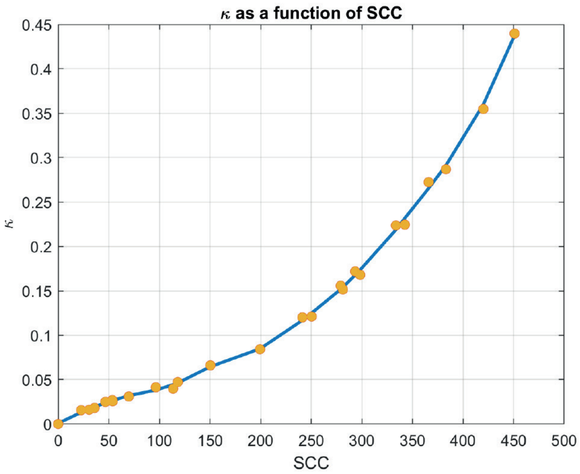

Empirically, we find that the relationship in Fig 5 can be excellently approximated by the following empirical fit (with SCC expressed in \$/tonne CO_2):

$$\kappa = 0.000428 \cdot SCC \text{ for } SCC \le 200 \tag{5}$$

$$\kappa = 0.1656 - 0.00115 \cdot SCC + 0.00000375 \cdot SCC^2 \text{ for } SCC > 200 \tag{6}$$

The reason for expressing the abatement speed, κ, as a monotonic function of the SCC, is that this will allow us to derive the distribution for κ from the SCC distri-bution. Since the equivalent abatement speed, κ, is a powerful proxy for the aggressiveness of the abatement policy, this would effectively enable us to tackle the most intractable of the tasks outlined in Section 1.3, ie, it would allow us to associate probabilities with different abatement policies. To carry out this last step, we proceed as follows.

3.5 Deriving a Distribution for the Abatement Speed

We have obtained a maximum-entropy or an informative distribution for the SCC. Let's generically denote the associate probability density by $\Psi(s)$, with s indicating the SCC. Since the relationship between the SCC and the equivalent abatement speed, κ, is monotonic, we can write

$$\phi(\kappa) = \frac{\Psi(s)}{\left|\frac{d\kappa}{ds}\right|} \tag{7}$$

The probability distributions for the abatement speed, κ, are shown in Fig 6 for the unshifted distribution (blue line, labelled 'Unshifted'), the shape-preserving shifted distribution (red line, labelled 'Shifted'), the minimum-divergence K-L shifted distribution (yellow line, labelled 'K-L'), and the Maximum-Entropy distribution, (purple line, labelled 'ME'). As in the case of the SCC distribution, the lack of assumed information in the maximum-entropy case about the likelihood of extremely slow abatement policies explains the differences between the two distributions.

Our methodological task is now complete, and we can move to the results.

4. Results

We now have all the elements in place to obtain unconditional distributions of several quantities of economic interest, such as the temperature anomaly at the end of the century. To give concrete results, we must combine (through our scenario engine) the quantification of policy uncertainty that is the focus of this paper with the other sources of uncertainty discussed in Section 1.3. Different scenario engine may describe these non-policy quantities in different ways (as, indeed, we show in Section 5), and arrive at somewhat different outcomes; our general procedure, however, remains valid.

Figure 6: The probability density of the effective abatement speed, κ, for the unshifted distribution (blue line, labelled 'Unshifted'), the shape-preserving shifted distribution (red line, labelled 'Shifted'), the minimum-divergence K-L shifted distribution (yellow line, labelled 'K-L'), and the Maximum-Entropy distribution, (blue line, labelled 'ME').

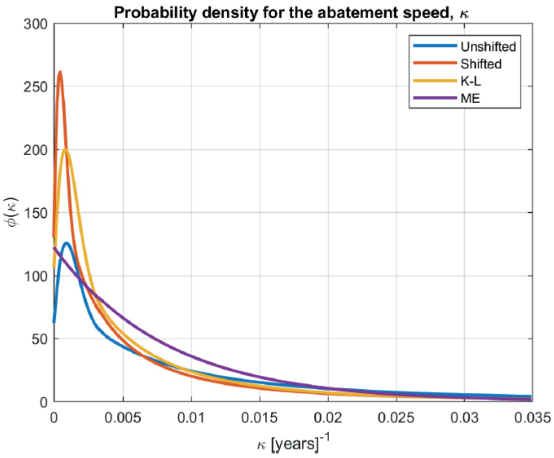

As explained in Section 1.3, we first obtain the *conditional* distributions for the quantities of interest – where the conditioning is on an abatement policy. We then use the probabilities of the abatement policies (ie, of the associated equivalent abatement speeds, κ) to weigh the conditional distributions. The unconditional distribution is therefore obtained as a mixture of the conditional distributions, each weighted by the abatement probability. We only discuss briefly how the conditional distributions are arrived at because different scenario users may prefer different modelling choices from the ones we made. Our original contribution lies in the estimation of the mixture weights.

The emissions associated with the different paths of economic output produced by the Jensen and Traeger (2014) model (traced from Bansal and Yaron (2004)) will be abated to different extents, depending on the abatement policy undertaken. We gauge the likelihood of different policies from the probability density shown in Fig 6 and, for every distribution, we run our simulations for each of the percentiles (from the 2.5th to the 97.5th, in increments of 5). We repeat this exercise using all the non-rejected reduced-form climate models calibrated from the Coupled Model Intercomparison Phase VI (CMIP6). To be clear: we run a full simulation of economic outcomes for each combination of abatement speed and climate model, thereby producing a different probability-weighted CO_2 concentrations at different horizons. After averaging over the climate models, this procedure therefore gives twenty (one for each percentile) distributions of terminal CO_2 concentrations, each conditioned on a particular abatement speed. The corresponding unconditional distribution is obtained as the equal-weighted average of the twenty conditional ones.

Figure 7: Percentiles of the distribution of the 2100 temperature anomaly for the unshifted distribution (blue line, labelled 'Unshifted'), the shape-preserving shifted distribution (red line, labelled 'Shifted'), the minimum-divergence K-L shifted distribution (yellow line, labelled 'K-L'), and the Maximum-Entropy distribution, (purple line, labelled 'ME'), for the case when the expectation of the SCC, is $30/tonne C, $50/tonne C and$70/tonne C (left, middle and right panels, respectively).

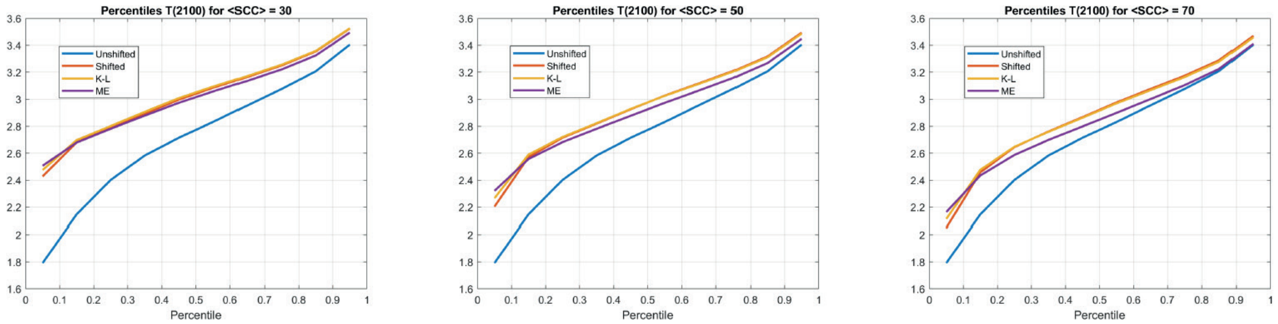

The results that we obtain using this approach are synthetically displayed in Figs 7 and 8: Fig 7 displays the percentiles of the distribution of the 2100 temperature anomaly for the unshifted distribution (blue line, labelled 'Unshifted'), the shape-preserving shifted distribution (red line, labelled 'Shifted'), the minimum-divergence K-L shifted distribution (yellow line, labelled 'K-L'), and the Maximum-Entropy distribution, (purple line, labelled 'ME'), for the case when the expectation of the *SCC* is $30/tonne C, $50/tonne C, and $70/tonne C, (top left, top right and bottom panels, respectively). Fig 8 displays the same information, but each panel shows the effect of imposing different expectations of the social cost of carbon for the shifted, K-L and Maximum-Entropy distributions (again, top left, top right and bottom panels, respectively).

Figure 8: Percentiles of the distribution of the 2100 temperature anomaly as a function of the expectation of the SCC (values of $30, $50 and $70) for the unshifted distribution (blue line, labelled 'Unshifted'), the shape-preserving shifted distribution (red line, labelled 'Shifted'), the minimum-divergence K-L shifted distribution (yellow line, labelled 'KL'), and the Maximum-Entropy distribution, (purple line, labelled 'ME'), displayed in the left, middle and right panels, respectively.

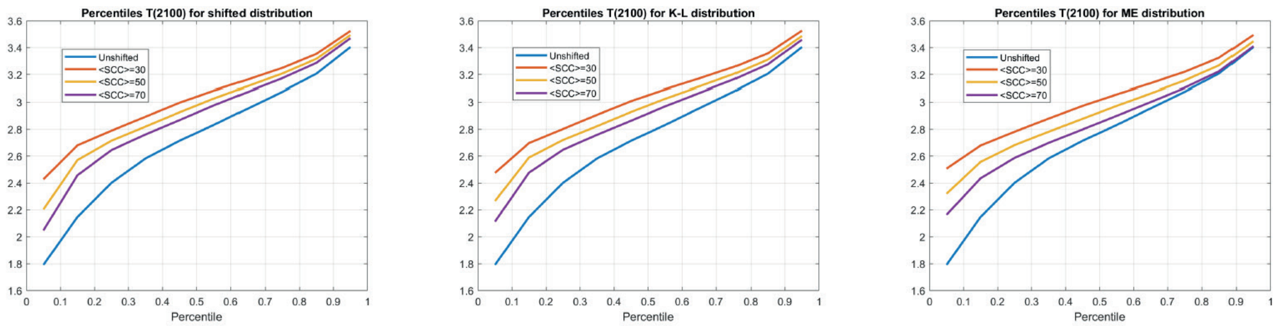

The first observation is that, while choosing different expected values for the expectation of the SCC, gives significantly different percentiles, for a given expectation the different ways to shift the economists' distribution yield remarkably similar values – see Fig 7. This suggests that the overall methodology we have proposed is very robust. This is all the more remarkable, given that the Maximum-Entropy procedure com-pletely by-passes the elicitation of the economists' views about the optimal cost of carbon, yet recovers very similar temperature distributions once the same first-moment constraint has been imposed.

Next, we observe that the greatest difference in percentiles between the economists' distribution and all the other distributions, howsoever obtained, is for the 'left tail' (lowtemperature) end of the distribution: the 5th percentile of the unshifted 2100 distribution is between 0.65 and 0.30 °C l ower than the shifted distributions, but the 95th percentiles only differ by about a tenth of degree centigrade: so, all approaches roughly agree on how high the temperatures can get, but disagree on how low they may be contained. Different expected carbon taxes, or *SCC*, affect the low percentiles much more than the high ones: see Fig 8.

As for the most likely levels of the temperature anomaly, we find values between 2.9 and 3.0 °C for the median value for the shifted distributions, and 2.8 °C for the unshifted distribution. If these estimates are correct, they should give food for serious thought, as the human species has never experienced temperatures 3.0 °C higher than pre-industrial levels. Given their importance, we critically discuss these results in Section Five but, at this stage we note that, for all distributions (including the unshifted one), the lower Paris-Agreement target of 1.5 °C is virtually unachievable, and even the achievement of the 2.0 °C target is a very-low-probability event. Since these are the end-of-century temperatures whose attainment would require the most significant transition costs, we find that the probability of the economy suffering large transition costs is low, but the associated physical damages can be very high (see, in this respect, the recent analyses in Burke, Hsiang, and Miguel (2015) and Kotz, Leverman, and Wenz (2024)).

16 - The description of the scenarios is provided in Appendix E.

5. Assigning Probabilities to the Oxford Economics Scenarios

In this section we want to show how our probability-based scenario approach can be applied to one common and widely used narrative-based scenario model, ie, the scenarios provided by Oxford Economics. We take the Oxford Economics approach (which uses the Process-Based Integrated Assessment Model GCAM to determine the joint paths for the quantities it projects) as a representative example of narrative-based scenario approaches. *Mutatis mutandis*, most of our considerations apply to similar scenario models.

Our goal is to use the approach we have sketched above to assign probabilities to climate outcomes in a way consistent with the Oxford Economics assumptions. We shall see that, *as a starting point*, the Oxford Economics assumptions and modelling choices are quite different from the default assumptions and choices we adopt in our scenario model. The purpose of the exercise that we present in the sections is to see whether our probabilitybased modelling approach is flexible enough to recover the Oxford Economics dynamics and to obtain probabilities consistent with the OE 'view of the world'. We stress than, when we highlight differences between the macroeconomic and financial assumptions of the Oxford Economics scenarios and the assumptions made in the default configuration of our model, we do not imply that the Oxford Economics approach is in any way wrong or deficient. We simply welcome and make a virtue of the significant difference in modelling choices to stress-test the flexibility of our modelling approach.

5.1 Features of the Oxford Economics Scenarios

Oxford Economics provide seven scenarios, that describe possible pathways for a large number of macrofinancial quantities from 2023 to 2050 (we focus on 17 variables, but there are several 'variants' in their dataset). In each scenario, one single pathway is assigned to every one of these financial quantities. The seven scenarios presented by Oxford Economics (and the labels we have used to identify them in tables and figures) are:[16]

1. Climate Catastrophe (Cl Cat)
2. Baseline (Baseline)
3. Climate Distress (Cl Distr)
4. Delayed Transition (Del Trans)
5. Low Demand (Low Dem)
6. Net Zero (NZ)
7. Net Zero Transformation (NZ Tr)

Table 9: The percentage of the total variability explained by the first four principal components for the seven Oxford Economics scenarios.

	Cl Cat	Base	Cl Distr	Del Trans	Low Dem	Net Zero	Net Zero Tr
1st eig	85.47%	86.38%	82.46%	93.92%	92.09%	92.09%	93.27%
2nd eig	96.63%	96.81%	95.22%	98.48%	97.25%	97.25%	97.98%
3rd eig	99.68%	99.65%	99.22%	99.41%	99.30%	99.30%	99.11%
4th eig	99.88%	99.89%	99.80%	99.79%	99.86%	99.86%	99.83%

Before discussing the nature of the different scenarios we note that, despite the very high number of variables that the Oxford Economics scenarios project, the effective dimensionality of their model is low. We show in Tab 9 the percentage of the variability explained by the first four principal components for all the seven Oxford Economics scenarios. As one can see, for the 17 macrofinancial variables we consider three principal components always explain more than 99% of the overall variability, and as few as two always account for more than 95%.[17]

16 - The description of the scenarios is provided in Appendix E.
17 - The principal component analysis was carried out on changes for the 17 variables. Results obtained using levels instead of changes show similar results.

A small number of principal components may explain a lot, but, in general, they are complex combinations of the original variables. Fortunately, we find by regressions that three key variables (economic output, emissions and temperatures) – variables that are also simulated by the EDHEC engine – afford almost the same explanatory power as the first three principal components. We therefore display in Tabs 10, 11 and 12 the paths (sampled at five years intervals) for these quantities, which are important in themselves, but also recover well most of the other variables.

The first observation is that three scenarios (*Net Zero, Net Zero Transformation* and *Delayed Transition*) display extremely aggressive decarbonization paths – in one, in 25 years' time we will actually be *subtracting* CO_2 from the atmosphere, which requires the deployment of large-scale negative emission technologies. Second, there is relatively little variability in the rate of growth of the economy, with the highest and lowest growth rates being 2.56% and 1.17%. This variability mainly arises from superimposing scenariodependent climate damages onto a relatively constant GDP path. (This observation will become important in Section 5.2.) The bulk of the variation across scenarios comes from the emission patterns.

Last, and perhaps most important, we note from from Tabs 11 and 12 that, in the Oxford Economics world, if the 2050 temperature is low, the GDP *must* have high growth: there is no possibility for low temperatures and low GDP (eg, no chance for a world in which temperatures are low *because* economic activity is reduced.) Conversely, the paths of lowest GDP growth are always associated with the paths with the highest temperatures. Overall, there is a strong (almost perfect!) negative correlation between GDP growth on the one hand, and temperatures and emissions on the other.[18]

In the terminology of Giglio, Kelly, and Stroebel (2021), the Oxford Economics scenarios therefore inhabit a 'Barro universe' – a universe, that is, in which climate damages are *so* large that, in and of themselves, cause GDP to be low. We stress that this 'Barro view of the world' is possible, but, as Giglio, Kelly, and Stroebel (2021) discuss, there can also be a different causal link: high economic activity causing high emissions, high concentrations, high temperatures and hence high damages (as, roughly speaking, in the Bansal, Kiku, and Ochoa (2016) model).

There is not enough empirical evidence at the moment to determine whether the causal link goes from high damages to low GDP growth, or from high GDP growth to high emissions, but, arguably, a scenario engine should leave the door open to both possibilities. The EDHEC simulation engine allows for the modelling of both worlds, and, in the context of this analysis, we modify its degrees of freedom so as to behave `a la Oxford Economics. The reason for doing so is clear: if the same relationships among variables enforced by the Oxford Economics model can be mimicked by the EDHEC scenario engine, then the scenario probabilities obtained by the EDHEC model can be associated to the Oxford Economics scenarios. We show in Section 5.2 how this can be done.

18 - This almost perfect negative correlation can be probably explained by the assumption of path of no-climate-change GDP applied for all scenarios with very little variability, onto which physical damages are superimposed. As damages detract from economic output, by construction the correlation between temperature and damages on the one hand, and economic output on the other, must be negative.

Table 10: Emissions paths for the seven Oxford Economics scenarios. The figures reflect emissions of CO_2 in Gtonne.

	Cl Cat	Low Dem	Net Zero	Net Zero T	Del Trans	Cl Distr	Baseline
2025	39.4	35.6	36.4	36.0	38.8	39.1	38.8
2030	43.6	22.6	26.3	26.4	37.9	40.9	38.4
2035	47.4	17.8	18.2	19.5	27.5	42.1	36.9
2040	50.3	14.5	10.9	11.7	18.3	43.1	35.4
2045	52.2	12.3	5.3	4.3	7.6	43.7	34.1
2050	53.5	10.9	1.1	0.3	-0.3	44.1	32.8

Table 11: Temperature anomalies for the seven Oxford Economics scenarios

	Cl Cat	Low Dem	Net Zero	Net Zero T	Del Trans	Cl Distr	Baseline
2025	1.39	1.39	1.39	1.39	1.39	1.39	1.39
2030	1.51	1.45	1.46	1.46	1.49	1.5	1.49
2035	1.67	1.51	1.53	1.53	1.59	1.64	1.6
2040	1.86	1.56	1.58	1.58	1.66	1.79	1.72
2045	2.06	1.6	1.62	1.61	1.7	1.95	1.83
2050	2.28	1.64	1.65	1.64	1.72	2.12	1.94

Table 12: Percentage increase in GDP from 2025 for the seven Oxford Economics scenarios.

	Cl Cat	Low Dem	Net Zero	Net Zero T	Del Trans	Cl Distr	Baseline
2025	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
2030	112.8%	106.4%	111.5%	110.0%	113.5%	113.3%	113.8%
2035	124.0%	120.9%	126.1%	124.3%	121.8%	125.7%	127.6%
2040	131.7%	136.3%	143.9%	140.3%	136.1%	136.6%	141.7%
2045	135.1%	151.4%	164.5%	158.4%	153.2%	145.3%	156.1%
2050	134.1%	167.0%	189.9%	180.0%	172.5%	151.2%	170.7%

5.2 Assigning Probabilities to the Oxford Economics Scenarios

In order to assign probabilities to the Oxford Economics scenarios we proceed as follows. Recall that emissions, e_t, are linked to the no-control emission intensity, σ_t, the abatement function, μ_t, and the GDP, y_t, by the relationship

$$e_t = \sigma \cdot (1 - \mu_t) \cdot y_t \tag{8}$$

In the case of no abatement at all, emissions are just given by $e_t = \sigma \cdot y_t$. We therefore make the very reasonable assumption that the *Climate Catastrophe* scenario corresponds to little-to-no abatement. Since both emissions and GDP are given in the Oxford Economics scenarios, from the *Climate Catastrophe* Oxford Economics scenario we can impute the no-control emission intensity function, σ_t. This function is 'universal', ie, it is not scenario dependent.[19] Therefore we can use it for *all* the Oxford Economics scenarios to obtain from Eq 8 the scenario-dependent abatement schedules, μ_t^j, where the superscript *j* denotes the scenario.

Fig 17 show how similar these Oxford Economic abatement functions so obtained are to the stylised abatement functions in Eq 4. But each of these stylised functions is characterised by its abatement speed, κ, we have shown in the introductory section how to obtain a probability distribution for this parameter. This means that, by associating to each Oxford Economics scenario one abatement schedule, we can find the abatement speed, κ, associated with it. And, given the association between scenarios and the effective abatement speeds, κ, we can estimate the associated probabilities.

19 - We note in passing that the OE implied function, σ_t, is virtually constant as a function of time – this is contrary to the 'natural' decarbonization of rich economies away from goods towards services.

In sum: if these three explanatory variables that account for 99% of the variability among Oxford Economics variables are made to behave in our model similarly to how they behave in the Oxford Economics world, we have an equivalent *probability-based* description for all the Oxford Economics variables. The remaining step in the procedure is therefore to show how the EDHEC model can be 'calibrated' as to reproduce the Oxford Economics joint dynamics (along the paths provided) for the key explanatory variables. To do so, we proceed as follows.

First we recalibrate our model to recover the Oxford Economics no-controls emission intensity, ie, the function σ in Eq 8; the rate of growth of the population; and the rate of growth of the total factor productivity. Next, we switch off the stochasticity in economic growth and the uncertainty in the damage exponent. Finally, we replace the Howard and Sterner (2017) damage function with the Oxford Economics damage function.

When we do this, our model is set up to behave as much as possible as the Oxford Economics model. What remains to do is to assign probabilities to the Oxford Economics scenarios. To do so we proceed as follows. Recall that the three quantities that account for most of the observed variability are economic output, emissions and temperature. From our recalibrated simulations, we can estimate at each point in time the expectations of, and the covariance between, these variables. These are, of course, a function of the abatement distribution we have previously determined. The Oxford Economics scenarios provide paths for these same variables. If the scenarios were independent random samples from their underlying population, equally-weighted sample expectations and covariances would provide the best estimate of the corresponding statistics for the population. However, the Oxford Economics scenarios have not been randomly and independently selected, but 'hand picked'. We can therefore ask the questions: what weights do we have to assign to the Oxford Economics scenarios in order to match the expectations and covariances produced by the Oxford Economics engine and by our model? And, how close are the estimates of quantities produced by our model and by the Oxford Economics approach? Figs 13 and 14 show that a set probabilities can indeed be found such that a close match between the Oxford Economics and the EDHEC expectations and covariances can be achieved. What do these probabilities look like?

Figure 13: The time-dependent expectations of the normalised GDP, emissions and temperature anomaly.

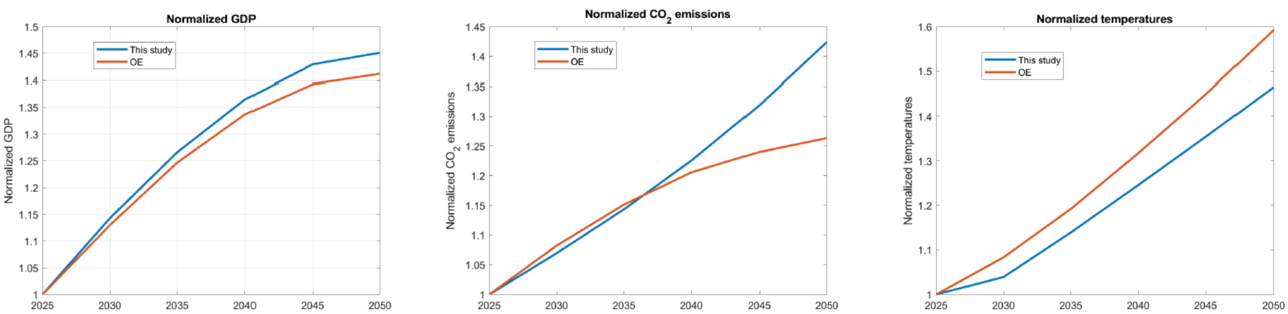

Figure 14: The time-dependent covariances between the normalised GDP, emissions and temperature anomaly.

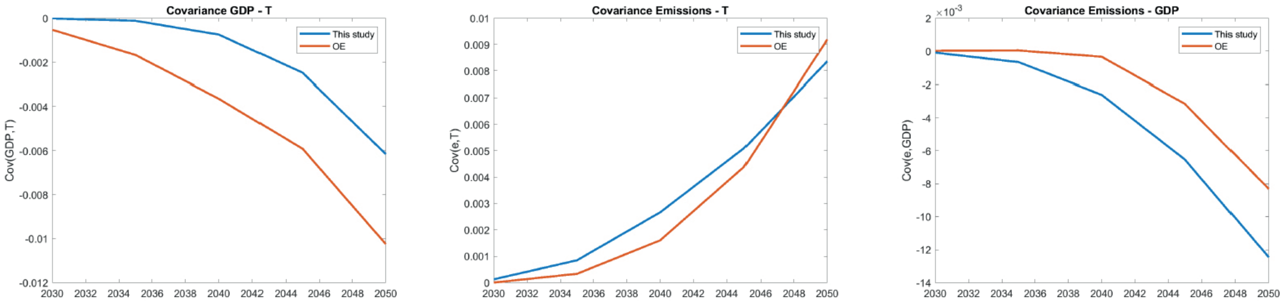

The best-fit probabilities are shown in Tab 15. What is apparent is a strong 'polarization' of the distribution: taken together, the two scenarios associated with the slowest pace of emissions account for more than 90% of the total probability, with the more severe *Climate Catastrophe* scenario taking the lion's share of the probability mass (57.4%). At the opposite end of the probability spectrum, we have very 'optimistic' scenarios – scenarios, that is, that have such high abatement speeds that our calculations assign them very low probabilities.

Table 15: Probabilities assigned to the various Oxford Economic scenarios that best recover the expectations of, and covariances among, the rescaled variables, GDP, emissions and temperature.

Cl Cat	Low Dem	NZ Tr	NZ	Del Tr	Cl Distr	Base
57.5%	0.1%	1.0%	1.0%	0.4%	35.0%	5.0%

Is this very pronounced concentration of probability masses on the slowest-abatement scenarios justifiable? We can answer this question (and explain *why* the probability density is so concentrated) along two different lines. First, from the Oxford Economic data, we can reconstruct the carbon tax as a percentage of GDP for the various scenarios. The results we obtain are displayed in Fig 16.

We note that the *Low Demand* scenario implies a steady increase in the carbon tax from less than 1% of GDP to approximately 11% of GDP in little more than five years. This would mean that, globally, the carbon tax would exceed what is currently spent on healthcare. The *Net Zero* and *Net Zero Transformation* scenarios also imply implausibly rapid, and implausibly high, increases in carbon tax, that would climb in five years' time at level equal to the sum of what the world currently spends in education and defence (for the *Net Zero scenario*), or to one-and-half times the current defence expenditure (in the *Net Zero Transformation case*). Conversely, the three scenarios to which our model assigns the highest probabilities (*Climate Catastrophe*, *Climate Distress* and *Baseline*) are the ones associated with the lowest carbon taxes as a function of GDP.

Figure 16: The evolution of the carbon tax for the various scenarios in the Oxford Economics model, expressed as a fraction of the same-time GDP.

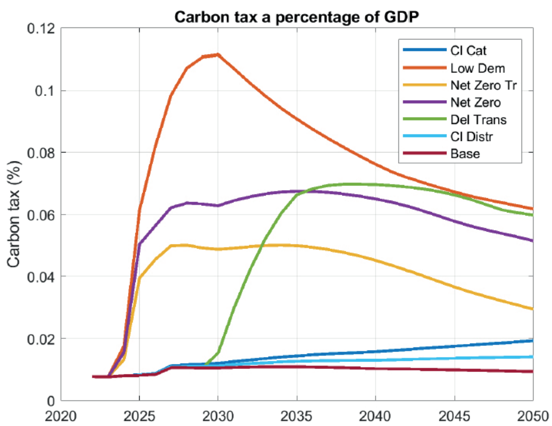

We can look at these probabilities from a different perspective. Given the no-controls emission intensity function that we have estimated, and given the data provided by Oxford Economics, we have shown how the associated abatement schedules, $\mu^i(t)$, can be extracted. These are shown in Fig 17. From Fig 17 one can see that all the 'optimistic' scenarios imply that in 25 years the economy will be 80% decarbonised. Not surprisingly, based on the current and persistent disconnect between recommendations and abatement actions, these decarbonization paths are given very low probabilities. Indeed, the best-fit abatement speeds, κ, associated with the various scenarios are displayed in Tab 18. As one can read off the graphs in Fig 6, the scenarios with high abatement speeds (high carbon taxes) have very low probabilities, and the maximum probabilities are found for the lowest values of the abatement speeds, κ.

Figure 17: The abatement functions, μ_t associated with the various scenarios.

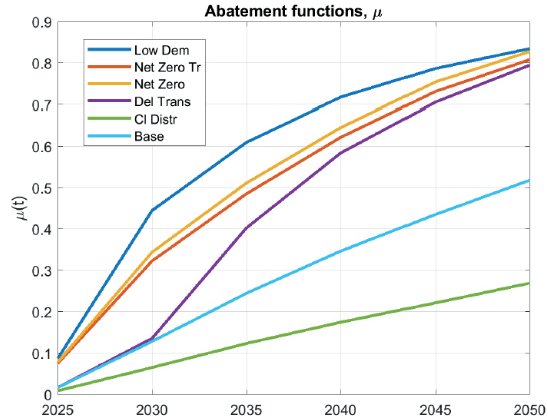

In sum: the high concentration of probability mass that we associate to the Oxford Economic scenarios stems from the 'polarised' choices of abatement policies associated with the various scenarios, with the very 'severe' scenarios (*Climate Catastrophe*, *Climate Distress* and *Baseline*) in one group (with an average κ of 0.013), and all the other, very 'optimistic', scenarios in the other group (with an average κ of 0.065). Since we have imposed that the probabilities should add up to one, a less polarised assignment of probabilities to scenarios would have arisen if some 'intermediate' abatement polices had been included in the Oxford Economic set. This would have 'shifted' some of the probability mass away from *Climate Catastrophic* and *Climate Distress*. The very optimistic scenarios (of the *Net Zero* family) would however have remained of low likelihood.

Table 18: The effective abatement speeds, κ, a sociated to the various Oxford Economic scenarios that best fit the abatement functions, $\mu^i(t)$, in Figure 17, with i denoting the scenario.

Cl Cat	Low Dem	NZ Tr	NZ	Del Tr	Cl Distr	Base
0.0001	0.080	0.061	0.065	0.055	0.012	0.028

6. Conclusions

We have presented several ways to associate unconditional probabilities with climate scenarios in general, and to the various possible abatement policies in particular. The results we have presented are very robust, in the sense that the different estimation methods gave very similar answers.

To arrive at these probabilities we have taken either an informative route, or the least committal (maximum-entropy) approach. With the former, we have assumed that there is imperfect but useful information in the economists' estimates of the optimal SCC. With the latter, we have simply estimated the most likely distribution compatible with the available information.

In both cases, we have anchored our probabilistic estimates to real-world data by requiring that the expectation of the SCC from our models should match an empirically observable proxy for this quantity, ie, the cost of traded carbon permits. The quantification of the policy uncertainty was achieved by finding the probability distribution for an effective abatement speed, that parsimoniously characterises a variety of abatement schedules.

Given the probability distributions we have obtained, we have derived the *unconditional* expected value of the temperature anomaly at the end of the century. If our strategy is valid, this would be a major contribution to the policy debate and to investment analysis, because up until now only *conditional* estimates for temperature anomalies have been provided (conditional, that is, on an exogenous abatement path).

We have shown that, whilst technologically still achievable, the lower Paris-Agreement target of 1.5°C has a negligible probability of occurrence; that also the 2.0°C is very unlikely; that there is a high probability (no less than 40%) that the planet will experience temperatures never seen by the human species (above 3.0°C); and that temperatures between 3.5°C and 4.0°C are unlikely, but far from impossible. For such high temperatures, the probability of tipping point thresholds being crossed significantly increases. This would bring us into an uncharted climate territory where current damage-function estimations would no longer be adequate.

Apart from the technical aspects, these high probabilities for high temperatures have been arrived at by observing the significant disconnect over the last 40 years between recommended and implemented abatement policies, and assuming that this disconnect will prevail in the future. Perhaps, by showing these unconditional estimates of high probabilities for dangerously high temperatures, our study can become self-refuting, in the sense that the gap between policy recommendations and action will be narrowed. We can only hope this to be the case, but recent international political developments do not offer much hope in this direction.

Finally, we have shown how our model can be used to associate probabilities with the climate scenarios offered by various commercial and public bodies. We have analysed the Oxford Economics scenarios in some detail, where a close match between the macrofinancial quantities projected by our engine and the Oxford Economics model was achieved. (Similar results, not reported, were obtained for the NGFS scenarios).

7. Appendices

A Derivation of the Maximum-Entropy Distribution

For a continuous random variable, X, with a strictly positive probability density function $\phi(x)$ with support Ω, the entropy of the distribution, $H(\phi)$, is defined as

$$H(\phi) = -\int_{\Omega} \phi(x) \cdot \log \phi(x) \cdot dx \tag{9}$$

Let

$$\int_{\Omega} f_i(x) \cdot \phi(x) \cdot dx = c_i \tag{10}$$

be a number of constraints (i = 1, 2, ..., n) on the moments or range of the distribution. Then there are strong information-theoretical reason to argue that the least-committal distribution that reflects these constraints is the distribution that maximises the entropy defined above, subject to the chosen constraints. The problem can therefore be solved by means of Lagrange multipliers, by creating a Lagrangean function, $L(\phi, \lambda_i)$, given by

$$L(\phi, \lambda_i) = H(\phi) + \sum_i \lambda_i \left(c_i - \int_{\Omega} f_i(x) \cdot \phi(x) \cdot dx \right) \tag{11}$$

To obtain an extremum of the Lagrangean, we differentiate with respect to the Lagrange multipliers, λ_i, perform a variational differentiation with respect to ϕ (·) and set all the derivatives to zero.

For the problem at hand, we want to ensure i) that the density of SCC integrates to 1 between 0 and TUL; and that the expectation of this distribution is equal to μ. The constraints therefore are

$$\int_0^{TUL} \phi(x) \cdot dx = 1 \tag{12}$$

$$\int_0^{TUL} x \cdot \phi(x) \cdot dx = \mu \tag{13}$$

The Lagrangean therefore is

$$L(\phi, \lambda_1, \lambda_2) = -\int_0^{TUL} \phi(x) \cdot \log \phi(x) \cdot dx + \lambda_1 \left(1 - \int_0^{TUL} \phi(x) \cdot dx \right) + \lambda_2 \left(\mu - \int_0^{TUL} x \cdot \phi(x) \cdot dx \right) \tag{14}$$

Setting to zero the functional derivative with respect to ϕ (·) gives

$$\log \phi(x) = -1 + \lambda_1 + \lambda_2 \cdot x \rightarrow \phi(x) = \exp(\lambda_1 - 1) \exp(\lambda_2 \cdot x) = K(\lambda_2) \exp(\lambda_2 \cdot x) \tag{15}$$

with $K(\lambda_2)$ a normalization constant (a function of λ_2). Imposing

$$\int_0^{TUL} K(\lambda_2) \exp(\lambda_2 \cdot x) = 1 \tag{16}$$

gives

$$K_{\lambda_2} = \frac{\lambda_2}{\exp(\lambda_2 \cdot TUL)} \tag{17}$$

The value for λ_2 is then obtained by imposing that

$$\mu = \int_0^{TUL} x \cdot K_{\lambda_2} \cdot \exp(\lambda_2 x) \cdot dx = \frac{TUL \cdot \exp(TUL \cdot \lambda_2)}{\exp(TUL \cdot \lambda_2) - 1} - \frac{1}{\lambda_2} \tag{18}$$

and solving for λ_2.

B The Weighted Average Abatement Function

We start from

$$\Delta C(T) = \int_0^T e(s)ds \tag{19}$$

and

$$e_{ind}(t) = \sigma(t) \cdot (1 - \mu(t)) \cdot y(t) \tag{20}$$

where we have dropped the subscript 'gross' for gross economic output, $y(t)$. These two equations can be combined to give

$$\Delta C(T) = \int_0^T w_s \cdot ds - \int_0^T w_s \cdot \mu_s \cdot ds \tag{21}$$

Define

$$w_t \equiv \sigma_t \cdot y_t \tag{22}$$

and rewrite Eq 21 as

$$\Delta C(T) = \int_0^T \sigma_s \cdot y_s \cdot ds - \int_0^T \sigma_s \cdot y_s \cdot \mu_s \cdot ds \tag{23}$$

Dividing both sides by $\int_0^T w_s \cdot ds$ one gets

$$\frac{\int_0^T w_s \cdot \mu_s \cdot ds}{\int_0^T w_s \cdot ds} = 1 - \frac{\Delta C}{\int_0^T w_s \cdot ds} \tag{24}$$

We recognise the LHS as the weighted average, $< \mu >$, of the abatement function μ_t, with weights given by the function w_t. We therefore write

$$< \mu >= 1 - \frac{\Delta C}{\int_0^T w_s \cdot ds} \tag{25}$$

Since the RHS does not depend on the abatement function, we conclude that, for given paths of economic output and no-control emission intensity, all the abatement functions with the same weighted average give rise to the same change in CO_2 concentration.

We have chosen the symbol w_t for the product $\sigma_t \cdot y_t$ to emphasise the role of weighting function played by this quantity. We can gain further insight by defining

$$e_s^* \equiv \sigma_t \cdot y_t \tag{26}$$

$$\Delta e_s \equiv \mu_t \cdot \sigma_t \cdot y_t \tag{27}$$

We recognise e_t^* the emissions that would occur for a given level of economic output if no active decarbonization steps were taken, ie, if all the energy needed for production were obtained from burning fossil fuels. With Nordhaus and Sztorc (2013), we call this a 'nocontrols' situation. The quantity

$$\Delta C^* = \int_0^T e_s^* ds \tag{28}$$

is therefore the increase in concentration from time 0 to time T that would occur under no controls. Equation 29 can therefore be rewritten as

$$< \mu >= 1 - \frac{\Delta C}{\Delta C^*} \tag{29}$$

showing that the weighted average abatement function is a function of the ratio of concentrations with the abatement schedule μ_t to the concentration that would obtain under no controls.

Define next

$$H(T) = \int_0^T \Delta e_s \cdot ds \tag{30}$$

Then we can write

$$\Delta C(T) = \int_0^T \sigma_s \cdot y_s \cdot ds - \int_0^T \sigma_s \cdot y_s \cdot \mu_s \cdot ds = \Delta C^*(T) - H(T) \tag{31}$$

and therefore $H(T)$ can be interpreted as the change in time T concentration because of abatement efforts, over and above the decarbonization brought about by the declining function σ_t: $H(T) = \Delta C^*(T) - \Delta C(T)$. Therefore the weighted average abatement speed can be written as

$$< \mu >= \frac{H}{\Delta C^*} \tag{32}$$

ie, as the fraction of the abatement-induced concentration reduction (H) to the no-controls concentration (ΔC^*).

Let $M(\Delta C)$ denote the set of all the abatement functions that, for given paths of economic output and no-control emission intensity, have the same weighted average, and let $m_s^1, m_s^2 \in M$. Then we have

$$\int_0^T w_s \cdot m_s^1 \cdot ds = \int_0^T w_s \cdot m_s^2 \cdot ds \tag{33}$$

This equation holds in particular for the constant function, $\overline{m}$, and, denoting by m_s the generic element of the set M, we have

$$\overline{m} = \frac{\int_0^T w_s \cdot m_s \cdot ds}{\int_0^T w_s \cdot ds} \tag{34}$$

ie, there exists a constant abatement function, $\overline{m}$ that produces the desired change in concentration.

A final comment about negative emissions. Negative emissions at one point in time can be modelled by a function $\mu_t > 1$. However, as long as the weighted average of the abatement function *over the whole period* is smaller than 1, even a function, μ_S, that is smaller than 1 everywhere can produce the correct time-T concentration. It is only if the weighted average concentration were greater than 1 that an instantaneous function, μ_t greater than 1 at least in some parts of the interval $[0\ ,\ T]$ would be needed.

C Shape-Preserving Shift of the Distribution of the SCC

In this Appendix we show how the distribution for the SCC can be shifted as to recover an exogenously assigned first moment in such a way as to retain the shape of the original (unshifted) distribution as much as possible.

We start from a distribution, $\phi'(x)$, about which we only know the normalization constraint, $\int_{\Omega'}\phi'(x) \cdot dx = 1$ (where Ω' denotes the finite support of $\phi'(x)$). We want to create a distribution, $\phi(x)$, such that, in addition to the normalization constraint, also the condition $\int_{\Omega} \phi(x) \cdot x \cdot dx = \mu$ should be satisfied. (Here μ denotes the observed cost of traded carbon permits, and Ω is a different, but still finite, support.) Of course, this additional constraint can be satisfied in an infinity of ways. Since we believe that there is information on the economists' distribution, we want to preserve its shape as much as possible. To do so, we discretise the integral above (and, for economy of notation, we denote the independent variable, the SCC, by x). We have

$$\int_{\Omega} \phi'(x)ds \approxeq \sum_i \phi'(x_i) \cdot \Delta x_i \tag{35}$$

The probability, $P(x_i)$, of the SCC having values between x_i and x_{i+1} is given by $P(x_i) = \phi'(x_i) \cdot \Delta x_i$. We impose that the same probability will apply for values of x in a different range, given by $R \cdot [x_i, x_{i+1}]$, with $R > 0$. To preserve the probability of the interval, the density function, $\phi'(x)$, will have to be changed to

$$\phi'(x_i) \to \phi(x) = \frac{\phi'(x_i)}{R} \tag{36}$$

When we do this, the normalization will be preserved by construction, and, if we define $\int_{\Omega'}\phi'(x) \cdot x \cdot dx = \mu'$, the ratio R that ensures that $\int_{\Omega'}\phi(x) \cdot x \cdot dx = \mu$ is simply given by $R = \frac{\mu}{\mu'}$. The construction ensures that, apart from a 'stretching' of the x axis, the shape of the original distribution – which we believe to convey useful information – is preserved.

D Modelling Uncertainty in Economic Growth

We model uncertainty in economic growth using the Jensen and Traeger (2014) modification of the influential 'long-term growth' Bansal and Yaron (2004), Bansal and Shaliastovich (2012) model. In this approach, the variability in economic outcomes arises from uncertainty in the growth process for the total factor productivity (TFP), $A(t)$, denoted by $g_A(t)$. This is given by:

$$A(t + \Delta t) = A(t) \exp\left(g_A(t)\Delta t\right) \tag{37}$$

where $g_A^{\det}(t)$ is the deterministic growth trend and $z(t)$ is a random growth shock. The deterministic component of the technology process is assumed to decay with time following the Nordhaus specification:

$$g_A^{\det}(t) = g_A^{\det}(0)\exp(\delta_a t), \qquad \delta_a = -0.005\text{yr}^{-1}, g_A^{\det}(0) = 0.076 \tag{38}$$

To capture the empirically observed strong time persistence of TFP, growth shocks are assumed to consist of two uncorrelated components

$$z(t) = x(t) + w(t) \tag{39}$$

where $x(t)$ is assumed to follow a Wiener process and $w(t)$ follows an Ornstein-Uhlenbeck process with reversion level of 0:

$$\begin{aligned} x(t+\Delta t) &= x(t) + \mu_x \Delta t + \sigma_x dZ_t^x \\ w(t+\Delta t) &= \mu_w(1-e^{-\theta\Delta t}) + w(t)e^{-\theta\Delta t} + dZ_t^w \end{aligned} \tag{40}$$

Here $dZ_t^x \sim N(0,1)$ and $dZ_t^w \sim N\left(0, \frac{\sigma_w^2}{2\theta}(1-e^{-2\theta\Delta t})\right)$ and $\mathbb{E}(dZ_t^x, dZ_t^w) = 0$.

When simulating the model, we discretise assuming a finite time step, Δ_t = 5 years.

When discretised, the Ornstein-Uhlenbeck process becomes an AR(1) process:

$$w_{t+1} = \mu_\epsilon + \gamma w_t + \epsilon_t$$

where the properties of γ, μ_ϵ and ϵ_t are readily developed from the corresponding values of θ and Δ_t:

$$\begin{aligned} \gamma &= e^{-\theta\Delta t} \\ \mu_\epsilon &= \mu_w(1-e^{-\theta\Delta t}) \\ \epsilon_t &\sim N(0,\sigma_\epsilon^2),\ \sigma_\epsilon^2 = \frac{\sigma_w^2}{2\theta}(1-e^{-2\theta\Delta t}) \end{aligned} \tag{41}$$

The volatilities, σ_x and σ_ϵ and the autocorrelation parameter γ are estimated by Jensen and Traeger (2014) so that $A(t)$ is consistent with empirical long run US data on the total factor productivity and consistent with Bansal and Yaron (2004). Both volatilities are set at 1.9%, while γ is determined for a five-year interval and is set at 0.5. The drift terms, μ_x and μ_ϵ, are developed by requiring that the overall mean of the growth rate of the TFP, $g_{A,t}$, should match the deterministic growth rate component, ie, we require that:

$$\mathbb{E}_t\left[A(t+\Delta t)\right] = A(t)\exp\left(g_A^{\det}(t)\Delta t\right)$$

We can use this to show that $\mu_x = -\frac{\sigma_x^2}{2}$ and:

$$\mu_\epsilon = -\frac{\sigma_\epsilon^2}{2}\frac{\sum_{p=0}^{T-1}\left(\frac{1-\gamma^{T-p}}{1-\gamma}\right)^2}{\sum_{p=0}^{T-1}\left(\frac{1-\gamma^{T-p}}{1-\gamma}\right)} \tag{42}$$

Here T is the (finite) long horizon over which we simulate the discrete model (typically T = 100)[20].

E Characterization of the Oxford Economics Scenarios

Each of the Oxford Economics scenarios is characterised by a narrative description. Ultimately, all these scenarios boil down to different paths for the state variables. We have shown that economic output, emissions and temperature explain a large proportion of the variability of all the Oxford Economics scenarios. For each scenario, we therefore report below both the narrative description (as provided by Oxford Economics) and the paths for the reduced set of descriptive state variables (economic output, emissions and temperature).

Climate Catastrophe (Cl Cat): This scenario explores the severe economic and environmental consequences of unmitigated climate change, where rising temperatures and extreme weather events significantly disrupt economies and societies. Annualised GDP growth declines from 2.59% for the period 2025-2030 to -0.14% for the period 2045-2050. Emissions keep on growing, albeit at a reduced pace (2.02% for the period 2025-2030 to 0.50% for the period 2045-2050). Temperatures increase at 1.6% at the beginning, and then at around 2.0% per year.

Baseline: Serving as a reference point, the baseline scenario assumes a continuation of current policies and trends without significant changes, providing a benchmark against which other scenarios are compared.

20 - This is a correction to the result presented in Jensen and Traeger (2014)

Annualised GDP growth declines from 2.40% for the period 2025-2030 to 1.79% for the period 2045-2050. Emissions decline (at a rate of 0.3% for the period 2025-2030 to 0.74% for the period 2045-2050). Temperatures increase at 1.4% at the beginning, and then at around 1.2% per year.

Climate Distress (Cl Distr): This scenario examines the economic impacts of delayed or insufficient climate action, leading to heightened physical risks and associated economic stresses. Annualised GDP growth declines from 2.50% for the period 2025-2030 to 0.80% for the period 2045-2050. Emissions increase but at a declining pace (at a rate 0f 0.94% for the period 2025-2030 to 0.18% for the period 2045-2050). Temperatures increase at 1.5% at the beginning, and then at around 1.7% per year.

Delayed Transition (Del Trans): This scenario considers the effects of postponing the implementation of climate policies, resulting in a more abrupt and potentially disruptive transition in the future. Annualised GDP growth starts at 2.53% for the period 2025-2030, declines to 1.40% for the period 2030-2035, and then pick up to approximately 2.30% for the remaining period. Emissions decrease slowly (-0.43%) for the period 2025-2030, but then the pace picks up very quickly, by 6-4% over the remaining period. Temperatures increase at 1.4% at the beginning, but at an ever-decreasing rate (0.23% for the 2045-2050 period).

Low Demand (Low Dem): This scenario analyses the economic implications of a global shift towards lower energy consumption and reduced demand for carbon-intensive goods and services. Annualised GDP growth starts from low value of 1.23% for the period 2025-2030, reaches a maximum of 2.57% between 2035 and 2040, and then declines to 2.00% for 2045-2050. Emissions decline but at a declining pace (at a rate 0f -9.10% for the period 2025-2030 to -2.48% for the period 2045-2050). Temperatures increase at 0.85% at the beginning, and then at around 0.6% per year.

Net Zero (NZ): This scenario explores the pathway and economic outcomes of achieving net-zero greenhouse gas emissions by a specified target date, involving significant transformations in energy production, consumption, and technology. Annualised GDP growth accelerates from 1.90% to 2.55%. Emissions decline steadily at a rate between 5-4%. Temperatures increase at 0.98% at the beginning, and then less and less (increase of 0.37% for 2045-2050).

Net Zero Transformation (NZ Tr): Building upon the Net Zero scenario, this narrative delves into the comprehensive societal and economic changes required to attain net-zero emissions, emphasising innovation, policy shifts, and behavioural changes. Annualised GDP growth increases steadily from 2.17% (2025-2030) to 2.87% (2045-2050). Emissions decline at a rate between 4.5% and 3%. Temperatures increase at 0.98% at the beginning, and then less and less (increase of 0.37% for 2045-2050).

References

- Acemoglou, D., P. Aghion, L. Bursztyn, and D. Hemeous (2012): "The Environment and Directed Technical Change," *American Economic Review*, 102(1), 131–166.
- Bansal, R., D. Kiku, and J. Ochoa (2016): "Price of Long-Run Temperature Shifts in Capital Markets," NBER Working Paper No. 22529.
- Bansal, R., and I. Shaliastovich (2012): "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," *Review of Financial Studies*, 26(1), 1–33.
- Bansal, R., and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59(4), 1481–1509.
- Black, S., A. A. Liu, I. Parry, and N. Vernon (2023): IMF Fossil Fuel Subsidies Data: 2023 Update, wp/23/169. International Monetary Fund.
- Burke, M., S. Hsiang, and E. Miguel (2015): "Global Non-linear Effect of Temperature on Economic Production," Nature, 527(7577), 235–239.
- Chen, C., W. J. Riley, I. C. Prentice, and T. F. Keenan (2022): "CO2 Fertilization of Terrestrial Photosynthesis Inferred from Site to Global Scales"," *PNAS – Environmental Sciences*.
- Cleveland, W. S. (1979): "Robust Locally Weighted Regression and Smoothing Scatterplots," *Journal of the American Statistical Association*, 74(368), 829–836.
- Dietz, S., and N. Stern (2015): "Endogenous Growth, Convexity of Damage and Climate Risk: How Nordhaus' Framework Supports Deep Cuts in Carbon Emissions," *The Economic Journal*, 125(583), 574–620.
- Epstein, L. G., and S. E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Journal of Finance*, 57(4), 937–969.
- Giglio, S., B. Kelly, and J. Stroebel (2021): "Climate Finance," *Annual Review of Financial Economics*, 13, 15–36.
- Greville, T. N. E. (1969): Theory and Applications of Spline Functions. Academic Press.
- Howard, H. P., and T. Sterner (2017): "Few and Not So Far Between: A Meta-Analysis of Climate Damage Estimates,," *Environmental and Resource Economics*, 68(1), 197–225.
- IPCC (2019): "IPCC Special Report on the Ocean and Cryosphere in a Changing Climate,"
- Jensen, S., and P. Traeger (2014): "Optimal Climate Change Mitigation under Long-Term Growth Uncertainty: Stochastic Integrated Assessment and Analytic Findings," *European Economic Review*, 69, 104–125.
- Kotz, M., A. Leverman, and L. Wenz (2024): "The Economic Commitment of Climate Change," *Nature*, 628, 551–557.
- Nordhaus, W. D., and P. Sztorc (2013): DICE 2013R: Introduction and User's Manual. Yale University, New Haven, Connecticut USA.
- Rebonato, R., D. Kainth, and L. Melin (2024): "The Impact of Physical Climate Risk on Global Equity Valuations," *Environmental and Resource Economics*, forthcoming, 1–34.
- Rebonato, R., D. Kainth, L. Melin, and D. O'Kane (2024): "Optimal Climate Policy with Negative Emissions," *International Journal of Theoretical and Applied Finance*, pp. 1–32.
- Rebonato, R., L. Melin, and F. Zhang (2025): "Climate Scenarios with Probabilities via Maximum Entropy and Indirect Elicitation," *EDHEC Climate Institution working paper and ssrn working paper*, available at SSRN: https://ssrn.com/abstract=5128228, pp. 1–52.
- Riahi, K., D. P. van Vuuren, E. Kriegler, J. Edmonds, B. C. O'Neill, S. Fujimori, and et al (2017): "The Shared Socioeconomic Pathways and Their Energy, Land use, and Greenhouse Gas emissions Implications: An Overview," *Global Environmental Change*, 42, 153–168.
- Stern, N. (2007): The Economics of Climate Change: The Stern Review. Cambridge University Press, Cambridge, UK.
- Tol, R. S. J. (2023): "Social Cost of Carbon Estimates Have Increased Over Time," *Nature Climate Change, Nature*, 13(6), 532–536.

- van Vuuren, D. P., E. Kriegler, B. C. O. Neil, K. L. Rihai, T. Carter, J. Edmonds, S. Allegatte, T. Kram, R. Mathur, and H. Winkler (2014): "A New Scenario Framework for Climate Change Research: Scenario Matrix Architecture," *Climatic Change*, 122, 373–386.
- WorldBank (2024): "State and Trends of Carbon Pricing," *World Bank and International Bank for Reconstruction and Development working paper*, http://hdl.handle.net/10986/41544, 1–76.

About EDHEC Climate Institute

About EDHEC Climate Institute

Institutional Context

Operating from campuses in Lille, Nice, Paris, London and Singapore, EDHEC Business School is ranked in the top ten European business schools. With more than 110 nationalities represented in its student body, some 50,000 alumni in 130 countries, and learning partnerships with 290 institutions worldwide, it is truly international.

EDHEC Business School has been recognised for over 20 years for its expertise in finance. Its approach to climate finance is founded on a commitment to equipping finance professionals and decision-makers with the insights, tools, and solutions necessary to navigate the challenges and opportunities presented by climate change. EDHEC has developed a significant research capacity on the financial measurement of climate risk, which relies on the best researchers in climate finance, and brings together experts in climate risks as well as in quantitative analysis.

The DNA of EDHEC's work has also resided, since its origin, in the ability to generate business ventures, by encouraging spin-offs based on the research work of its teams. EDHEC is currently involved in three ventures: Scientific Portfolio, Scientific Infra and Private Assets, and the soon-to-launch Scientific Climate Ratings.

Mission and Ambitions

The EDHEC Climate Institute (ECI) focuses on helping private and public decision-makers manage climate-related financial risks and make the most of financial tools to support the transition to a low-emission economy that is more resilient to climate change.

It has a long track record as an independent and critical reference centre in helping long-term investors to understand and manage the financial implications of climate change on asset prices and the management of investments and climate action policies.

The institute has also developed an expertise in physical risks, developing proprietary research frameworks and innovative approaches. ECI is also conducting advanced research on climate transition risks, with a focus on supply chain emissions (Scope 3), consumer choices, and emerging technologies.

As part of its mission, ECI collaborates with academic partners, businesses, and financial players to establish targeted research partnerships. This includes making research outputs, publications, and data available in open source to maximise impact and accessibility.

The EDHEC Climate Institute gratefully acknowledges the support that the Monetary Authority of Singapore (MAS) has provided to its green infrastructure research.

climateimpact.edhec.edu

For more information, please contact:

EDHEC Climate Institute

Maud Gauchon
maud.gauchon@climateimpactedhec.com

London
10 Fleet Place
London EC4M 7RB
United Kingdom
Tel +44 (0)20 7332 5600

Nice
393 Promenade des Anglais
06200 Nice
France
Tel +33 (0)4 93 18 78 87

Paris
16 - 18 rue du 4 Septembre
75002 Paris
France
Tel +33 (0)1 53 32 76 30

Singapore
One George Street
#15-02
Singapore 049145
Tel +65 (0)6438 0030

climateimpact.edhec.edu