Philosophy & Public Affairs

ORIGINAL ARTICLE OPEN ACCESS

The Perverse Incentives of Climate Integration: Why Researchers Can't Deliver What Funding Institutions Demand

Eric B. Winsberg[1,2]

[1]Department of History and Philosophy of Science, University of Cambridge, UK | [2]Department of Philosophy, University of South Florida, Tampa, Florida, USA

Correspondence: Eric B. Winsberg (winsberg@usf.edu; ew653@cam.ac.uk)

Received: 14 August 2025 | **Revised:** 20 September 2025 | **Accepted:** 22 September 2025

Funding: This work was supported by the British Academy.

Keywords: climate change | modeling | perverse incentives | scientific funding

ABSTRACT

Research funders increasingly require integration of future climate projections across health, agriculture, fisheries, and development economics, creating perverse incentives: institutions demand what current climate science cannot reliably deliver. I use "perverse incentive" here in its standard economic sense: an incentive that unintentionally produces counterproductive behavior, rather than implying ill will on the part of funders. Climate models designed for global, long-term analysis are being misapplied for short-term, regional uses beyond their validated scope. This paper identifies three problems arising from this mismatch: maladaptation in scientific labor allocation, erosion of trustworthiness through representational overextension, and representational risk from harmful signaling and normalization of inappropriate methodological norms. Researchers include climate projections not because they are justified, but because they are required, transforming models from tools of inquiry into performances of compliance. This threatens both scientific integrity and the legitimacy of science underwritten by democratic norms. Three institutional reforms are proposed to realign incentives with epistemic responsibility and ensure climate science serves as a reliable policy foundation rather than mere signaling.

1 | Introduction

In recent years, research funders worldwide have increasingly required investigators in diverse scientific fields—health sciences, agriculture, fisheries, infrastructure, development economics, and more—to integrate climate change projections into their research. Proposals studying wildfire-affected health outcomes in Canada, crop disease in East Africa, or fish stock fluctuations in New Zealand are now expected to account for how climate change will influence these phenomena over the coming decades. This requirement is often explicit: funding calls from agencies like the Canadian Institutes of Health Research (CIHR), the U.S. National Institutes of Health (NIH), and New Zealand's National Institute of Water and Atmospheric Research (NIWA) stipulate that researchers must "quantify climate-related risks" or "demonstrate how climate change will affect" their study systems. In some cases, proposals that omit climate projections are deemed ineligible for review. (See Table 1 for many examples.)

At first glance, this trend appears commendable. Climate change is a defining challenge of the twenty-first century. Its effects are likely to be far-reaching, and responsible research should aim to anticipate and mitigate its harms. The shift toward climate-integrated science seems to reflect a long-overdue recognition that climate change is not merely an environmental issue but a structural condition shaping health, agriculture, food security,

https://doi.org/10.1111/papa.70006

TABLE 1 | Climate projection requirements in public research funding initiatives.

Discipline	Funding agency	Quotation or summary	Source
Health	NIH (U.S.)	"As a part of this CCHI, this NOSI encourages applications that address the impact of climate change on health and wellbeing over the life course, including the health implications of climate change in the United States and globally."	https://grants.nih.gov/grants/guide/notice-files/NOT-ES-22-006.html
Agriculture/ Forestry	USDA–NIFA (U.S.)	"USDA's research and statistical agencies are adapting their programs to support the science and innovation needed to address the challenges climate change poses, while adjusting their operations to ensure reliability of the critical information they supply. The Economic Research Service has been expanding the resources it allocates to developing data products, enhancing modeling capabilities, and producing new research products that inform discussions of how to facilitate farm- and sector-level adaptation to changing climate conditions and risks."	www.sustainability.gov/pdfs/usda-2024-cap.pdf
Ocean Fisheries	NOAA (U.S.)	"The Climate, Ecosystems, and Fisheries Initiative (CEFI) is a cross-NOAA effort to build the nation-wide, operational ocean modeling and decision support system (System) needed to reduce impacts, increase resilience and help marine resources and resource users adapt to changing ocean conditions."	https://www.fisheries.noaa.gov/resource/document/noaa-climate-ecosystems-and-fisheries-initiative-fact-sheet
Freshwater/ Marine Science	NIWA (New Zealand)	NIWA guidance highlights that researchers "must account for potential climate-change impacts" when working in freshwater environments.	https://niwa.co.nz/freshwater/climate-change-freshwater-impacts-assessments
Public Health/ Health	ZonMw (Netherlands)	"Consideration of the long-term sustainability of the proposed solutions [regarding the delivery of health care systems in the E.U.] and their environmental impact, including the promotion of greener health practices and adaptation to climate change."	https://www.zonmw.nl/sites/zonmw/files/2024-11/thcs-jtc-2025-call-text-v1.0.pdf
Public Health/ Health	UBC Heath—Ministry of Health Research Seed Grant Program	Using insights from behavioral science, how can public health messaging be enhanced to improve public knowledge and implementation of actions that protect against the risks of climate hazards?	https://health.ubc.ca/interdisciplinary-health-research/ministry-health-research-seed-grant-program
Public Health/ Health	Canadian Institutes of Health Research (CIHR)	To convene interest and/or rights-holders to co-identify priorities and/or investigate climate change mitigation and adaption efforts to reduce the health impacts of populations in Canada and/or globally, including the inequitable distribution of the burdens of climate change on health between/across population groups.	https://cihr-irsc.gc.ca/e/52605.html
Public Health/ Health	Canadian Institutes of Health Research (CIHR)	The CIHR Centre for Research on Pandemic Preparedness and Health Emergencies, will be launching a funding opportunity to catalyze research aligned with the climate change priorities identified in the 2022 CPHO Annual Report and companion document, "*Generating Knowledge to Guide Public Health Action on Climate Change in Canada*".	https://cihr-irsc.gc.ca/e/53233.html

Note: This table summarizes various public research funding bodies and their explicit or implicit requirements that applicants incorporate projections of climate change into their research. Quoted language is drawn directly from the respective program calls or strategic documents. It is worth noting that in the time between when work on this paper began and the time it will appear in print, the science funding situation in the United States has dramatically changed, and the American examples, above, predate the change. Here I'll simply note that this is a recent, and certainly not global, change, And also that many of the lamentable changes in the US science funding situation are in part (though of course only in part) are plausibly an over-reaction to an inchoate recognition of some of the problems I am documenting in this paper and the erosion of public trust in science that it causes.

economic development, and more. To the extent this integration enhances foresight, responsiveness, and policy relevance, it is a welcome development.

Yet, there is a growing dissonance between the expectations embedded in these funding requirements and the actual epistemic resources that climate science can reliably provide, particularly regarding the temporal and spatial scales at which most applied research is conducted.[1] Despite their sophistication, current state-of-the-art climate models—such as those in the Coupled Model Intercomparison Project (CMIP)—are not designed to yield reliable projections over short time horizons (e.g., 10–20 years) or at fine-grained regional or sub-regional levels. Nor are they tailored to the domain-specific causal pathways that matter for health outcomes, agricultural productivity, fishery dynamics, or infrastructure degradation. This is especially true when precipitation plays a key causal role, as it often does, because global climate models have not adequately solved the problem of accurately estimating the effect of increased carbon forcing on local precipitation distributions. The probabilistic and scenario-based nature of climate model output, coupled with deep uncertainties in model structure and parameterization, makes them ill-suited for many mandated applications. As Nissan et al. [1] note in the context of development planning, climate science "is not currently well placed to deliver information of the kind or quality often assumed in sectoral decision-making."

This epistemic mismatch creates three interlocking problems. First, it distorts the allocation of scientific labor and funding. Non-climate researchers are tasked with incorporating climate projections that climate scientists themselves have not been funded or incentivized to produce at the required scales or with domain-specific fidelity. Often, researchers might resort to citing a poor-quality paper with limited uncertainty quantification in a pinch. When fisheries scientists, health economists, or agricultural planners are required to use climate information that is either unavailable or unreliable, they are placed in an impossible position—one that undermines their ability to pursue rational public research policy. Imagine being asked to incorporate projections of wildfire smoke over the next 12 years into a model forecasting future health outcomes, with no source providing responsible forecasting or associated uncertainty estimates for these inputs. It is a difficult position while trying to satisfy the dual demands of funding agencies and epistemically responsible science. While it may seem that "satisfying funding agencies" is something scientists should avoid when it conflicts with being epistemically responsible, this perspective is naïve. Ultimately, satisfying funding agencies is what scientists do, and when funding agencies broadly agree on their requests, these "asks" effectively shape what the public is *paying scientists to do.*

Second, the mismatch threatens to erode the trustworthiness of science. In the absence of curated, domain-relevant, and adequately qualified climate information, researchers may resort to using model outputs that are poorly suited to their purposes but carry the appearance of rigor and credibility.[2] This results in what we might call *representational overextension*: the use of climate model outputs in contexts where their *evidential value is weak*, but their *rhetorical force is strong*. This dynamic risks distorting research findings and undermines the public's ability to distinguish between well-supported scientific claims and those that merely borrow the language of climate science. This risks turning science, which should inform adaptation to climate change, into rhetorical exercises about the dangers of climate science. Even those who believe such rhetoric is valuable should recognize it should not be confused with or collapsed into the former.

Third, this dynamic introduces a form of representational risk [2–4] that is both epistemic and ethical. When a climate projection is included in a policy-relevant model—such as a health economic analysis of a public health intervention or a cost–benefit analysis of prescribed burns—it signals two things: first, that climate change is likely to exacerbate the problem under study in specific, modelable ways; and second, that, without radical global climate mitigation, this exacerbation is *inevitable*. However, in many cases, neither claim is justified by the available evidence. Instead, it arises from a structural incentive: researchers include climate projections not out of trust in their accuracy, but due to institutional compulsion. This form of signaling masquerades as representation, distorting policy discussions, overshadowing alternative mitigation strategies, and lending unearned epistemic authority to speculative claims, ultimately granting credibility to the poor-quality methods scientists resort to in order to meet funding agency demands. When climate models are used in this way—as obligatory but poorly understood inputs—they risk sending unintended and undesirable signals, for example, that future increases in wildfire smoke are unavoidable, when they might be mitigated through prescribed burns (more on this below). It suggests that climate change will exacerbate the problem under study in specific, modelable ways, and that, without radical global mitigation, this worsening is unavoidable. Even if this signal is weak compared to other policy inputs, its rhetorical effect can still crowd out attention to tractable interventions—such as land-management practices that reduce wildfire risk—by creating the impression that local strategies are futile. They may also imply that health researchers—people viewed by the public as experts—believe that prescribed burns are a poor smoke/health trade-off and should be avoided. Why else would they build models that assume wildfire smoke will definitely worsen in the near future? Yet this claim (that prescribed burns are a bad trade-off) is at best controversial and likely not something health researchers would endorse under normal circumstances if their expertise were solicited in good faith.

This paper seeks to diagnose and analyze this climate epistemology mismatch. It offers a philosophical analysis of an institutional and epistemic problem—one at the intersection of science policy, public reasoning, and the ethics of representation. Using examples from public health, fisheries science, and development economics, I argue that the current landscape of climate-integrated research is marked by a structural issue: a growing burden on applied researchers to use climate information that is not epistemically secured, responsibly curated, or adequately funded. This mismatch is not merely a technical issue; it poses a threat to the trustworthiness, utility, and legitimacy of science in public life.

The analysis unfolds in four parts. In Section 1, I examine the problem of *maladaptation and misallocated epistemic labor*: how expectations placed on non-climate researchers exceed the epistemic capacities of current climate modeling and why this necessitates a rethinking of funding priorities. In Section 2, I address *representational overextension and the erosion of trust in science*, showing how the obligatory use of climate projections in weak-evidence contexts can undermine public trust in scientific institutions. Section 3 develops the concept of *representational risk via harmful signaling*, focusing on how the inclusion of climate projections can shape policy narratives in misleading or ethically problematic ways. Finally, in Section 4, I provide *institutional and normative recommendations*: how funding bodies, modeling centers, and researchers can realign their practices to support responsible, trustworthy, and decision-relevant climate-informed science.

The stakes are high. Climate change is not going away, nor is the imperative to incorporate it into applied research. However, doing so responsibly requires more than good intentions or the desire to signal that climate change is a real danger. It necessitates an honest reckoning with the limitations of current models, a commitment to curating their outputs with epistemic humility, and an institutional framework that does not impose demands that climate science cannot meet. Without such reforms, we risk transforming the downstream use of climate modeling from a tool of knowledge into a performance of compliance—one that may do more harm than good.

2 | Maladaptation and Misallocated Scientific Labor

In today's funding landscape, a curious inversion has emerged. Climate scientists themselves emphasize the coarse spatial resolution, structural uncertainties, and inter-model spread of Earth System Models—especially for socially relevant outputs such as precipitation, net primary productivity, or regional ocean biogeochemistry. Yet researchers in public health, agriculture, fisheries, and infrastructure are increasingly required to apply these same models to short-term, localized, and decision-sensitive questions. Ironically, these are precisely the scales where model reliability is weakest.

Health researchers are asked to project how wildfire smoke and related morbidities will change over the next 12 to 20 years. Fisheries biologists must estimate how climate change will affect streamflow, variability, and extreme events. Agricultural economists are expected to evaluate crop viability under "likely near-term climate conditions," including rainfall timing, drought and flood risk, and soil loss. These tasks all require projections at very fine spatial scales—precisely where current models are least reliable. Nevertheless, such requirements are now embedded in the funding calls of major agencies in Canada, the United States, the United Kingdom, Australia, New Zealand, Sweden, and elsewhere (see Table 1).

The concern is not that researchers are being told to ignore climate change, but that they are required to integrate forms of knowledge that are often unavailable or available only in ways that are poorly curated, vetted, or funded for confident use in decision-making. This amounts to a misallocation of epistemic labor: applied researchers are expected to use climate projections, while climate modelers are neither funded nor incentivized to produce outputs tailored for those applications. Worse, climate models are typically more precise than accurate. They generate detailed datasets that can appear convincing even when their epistemic credentials are weak—tempting downstream researchers to treat them as decision-ready.

This section explores the implications of this mismatch. Climate-integrated research places heavy epistemic demands downstream—in fields such as health, fisheries, and agriculture—without parallel investment upstream in producing, interpreting, and validating model outputs for those needs. Applied researchers are left either to stretch global-scale projections beyond their evidentiary warrant or to simulate compliance by gesturing toward climate scenarios whose relevance is mostly cosmetic.

2.1 | The Epistemic Focus of Climate Science

Contemporary climate science—particularly the construction and interpretation of general circulation models (GCMs) and Earth system models (ESMs)—has made extraordinary progress in characterizing large-scale, long-term climate trends. Models participating in initiatives such as the Coupled Model Intercomparison Project (CMIP) are finely tuned to simulate the effects of greenhouse gas forcing on global temperatures, ocean heat uptake, sea-level rise, and other aggregated metrics [5]. The models' primary outputs are assessed over multi-decadal timescales (e.g., 2030–2100) and global or continental spatial scales. Their core strength lies in identifying large-scale climate trajectories under different emissions scenarios, not in producing detailed, locally specific, short-term predictions, especially not predictions of the rich set of causal variables typically required to fuel the models of downstream researchers in the fields mentioned above.

Climate scientists themselves acknowledge these limits. Even at large spatial and temporal scales, variables such as temperature, precipitation, and circulation patterns show inter-model variation and systematic biases. The problems worsen as resolution increases and timescales shorten. As Nissan et al. [1] note in the context of development planning, "the nature of the information that climate science is currently able to provide is often poorly matched to what users in the development sector assume or require." They highlight three main issues: internal variability dominates trends on decadal scales; downscaling introduces new uncertainties; and many socially relevant impacts depend on compound or threshold events poorly captured by current models. Despite this, funders and development actors often assume the availability of precise, reliable projections at 5-, 10-, or 20-year horizons for specific locales.

These assumptions extend into funding calls across other domains. For example, the Canadian Institutes of Health Research instruct applicants to "integrate climate-related risks or adaptation opportunities" into proposals on illnesses affected by smoke—even when the project's primary aim is biomedical or epidemiological. Likewise, the U.S. Department of Agriculture's National Institute for Food and Agriculture (NIFA) returns proposals without review if they fail to address "climate-related risks or mitigation strategies." In both cases, researchers must incorporate climate information for timeframes such as 10–20 years—precisely the window where forced trends are modest and natural variability dominates.

The problem is not only that CMIP-class models produce coarse or probabilistic outputs, but that they are not validated for these decision contexts. Climate scientists generally resist making deterministic near-term predictions about wildfire frequency, fish stocks, or the burden of climate-sensitive diseases. Yet downstream researchers are expected to use such projections—not to advance understanding, but to satisfy institutional mandates and funding criteria. The result is an epistemic incoherence with real-world consequences. The climate science literature is full of caveats about the weak decision relevance of fine-scale projections, yet these warnings coexist with datasets, "downscaled" models, and individual studies that appear to invite precisely such uses.[3]

My argument is general, but I will illustrate it with two examples: smoke and fish. Climate change will likely affect both wildfire smoke and fish populations in freshwater and oceans. Yet these outcomes also depend on many small-scale variables, often acting in opposite directions. Fire is shaped by land use, management, and human ignition as much as by climate. Likewise, New Zealand's river fish will be influenced by agricultural intensification, zoning, housing policy, and riparian buffer enforcement at least as much as by climate-driven streamflow changes [6].

2.1.1 | Fire

Let us begin with fire. Sanderson and Fisher [7] highlight serious concerns about the reliability of projections of future wildfire smoke exposure, especially when used to guide applied research or public policy. Fire activity responds to climate change in nonlinear and contingent ways, shaped not only by temperature and precipitation but also by vegetation dynamics, land management, and human ignition and suppression. Because of substantial interannual variability, distinguishing long-term climate-driven trends from natural fluctuations is difficult. Emissions models add further uncertainty, relying on assumptions about fuel loads, ignition sources, suppression efforts, and climate–vegetation feedbacks. These uncertainties are especially acute on decadal timescales, where internal variability often outweighs the forced signal from greenhouse gas emissions. Sanderson and Fisher caution that incorporating such projections into health research or land-use planning without these caveats risks producing misleading or overly confident results.

Kloster and Lasslop [8] show that CMIP5 Earth System Models performed poorly in simulating historical fire occurrences, with little agreement across models. In general, higher temperatures create conditions favorable to wildfires, as reflected in risk metrics based on temperature, wind, moisture, and fuel availability [7, 9, 10]. Yet Sanderson and Fisher caution that these metrics, derived from historical data, may not hold in the future. Fires may ultimately behave unpredictably as climate conditions and human activities evolve.

It is crucial to recognize that future wildfires are sensitive to human activities beyond carbon emissions, including forest management policies and practices. A growing body of research suggests that forest management, including traditional Indigenous burning practices and "prescribed burns," plays a key role in wildfire dynamics [11].

2.2 | Modeling Fire

Despite the clear limitations of climate modeling in forecasting future wildfires, health researchers across a wide array of funding opportunities are strongly encouraged to incorporate future wildfire changes into their research. I will discuss one example I am familiar with: the LEAP project [12]. The Lifetime Exposures and Asthma Outcomes Projection (LEAP) model aims to develop a comprehensive 'Whole Disease' model of asthma to guide resource allocation decisions in Canada. In 2019, a steering committee of allergists, respirologists, and health economists created a 'concept map' of childhood asthma to direct model development and identify key risk factors. The group decided to study the cost-effectiveness of HEPA air filters in reducing childhood asthma incidence in Canada. This study requires understanding the concentrations of pollutants affecting asthma onset at detailed regional scales. For this modeling project, the scale is Canadian census divisions (of which there are 296), and the timeframe is the next 12 years.

Pollutant concentrations are affected by wildfires, which can originate far away; for example, wildfires in northern British Columbia can influence smoke concentrations in Vancouver, and Alberta wildfires can affect Montreal. Many scientists involved in the LEAP project share the intuition, shaped by popular media and their disciplines, that wildfires will increase with climate change. Additionally, funding agencies encourage this perspective, leading LEAP scientists to model PM2.5 concentrations in each Canadian census division and how they will evolve over the next 12 years as a result of climate change.

Importantly, the LEAP project does not uncritically accept future $PM_{2.5}$ projections. Instead, it frames modeling as a deliberative, value-laden process where public participants contribute ethical and experiential knowledge to technical decisions. However, the constraints of the modeling exercise—and the broader funding and policy environment—shape the types of futures modeled and emphasized. The project received support from a grant from the Canadian Institutes of Health Research (CIHR) under the Climate Change and Health Initiative (https://cihr-irsc.gc.ca/e/53233.html), a program that explicitly promotes research on the health impacts of climate change. In this context, the decision to project a wildfire-derived $PM_{2.5}$ increase of 5.5% (with a range of 0%–11%) over a 13-year horizon—based on a U.S. modeling study—reflects both the best available proxy and a policy environment where climate change is treated as the dominant explanatory frame.[4]

Yet the challenges of modeling future wildfire smoke exposure, particularly over a short-term horizon like 2023–2036, are substantial. The chosen period is arguably too brief to capture the robust long-term effects of climate change on wildfire regimes—which are likely dominated by internal climate variability and other short-term factors—yet too long for high-confidence short-term forecasting. Moreover, the LEAP model calculates outcomes at the census district level—granular enough to model the cost-effectiveness of health interventions, but arguably too fine-grained to integrate meaningfully with climate-change-driven changes in fire behavior projections. We have reasonably good historical data (and reanalyzed machine-learning-based reconstructed data) at the census district level for fire smoke, but no effective way to extrapolate the effects of climate change on data this fine-grained. The required fine spatial resolution of predictions for wildfire $PM_{2.5}$ led the LEAP team to rely on historical averages using two models—CanOSSEM and RAQDPS—with imperfect agreement, and to make simplistic assumptions about future wildfire increases based on a U.S.-centric study [13]. Both CanOSSEM and RAQDPS inform the model, even though CanOSSEM is considered the "gold-standard" but it lacks fire-specific attribution. Only RAQDPS distinguishes between wildfire-sourced smoke and other sources, allowing for a climate-change-based multiplicative factor to be applied. The ultimate decision to use a U.S.-based projection adjusted for Canadian conditions and applied to RAQDPS was arguably the most reasonable given the evidence constraints and the sources available to the group, but it relies on extrapolation from one region to a very different one, and scaling that introduces significant epistemic uncertainty.

At the same time, modeling an increase in wildfire-attributable $PM_{2.5}$ may inadvertently reinforce the impression that escalating smoke exposure is an unavoidable consequence of climate change, rather than a partly tractable outcome of policy decisions regarding land use, suppression, and prescribed fire. Health researchers with expertise in the morbidity effects of smoke exposure are well positioned to intervene upstream—shaping forest management and mitigation strategies—not just responding to downstream health burdens. Yet within a funding landscape focused on climate-health linkages, the LEAP model, like many others, directs analytic attention toward downstream adaptive responses and individual-level vulnerability, potentially at the expense of modeling upstream structural and ecological drivers that may be more causally central. The result is a policy-relevant model that is unusually transparent about its assumptions and limitations (in a positive sense!) but still constrained by the epistemic tools and incentive structures available to publicly funded health scientists working on climate change.

2.2.1 | Fish

What about fish? As the National Oceanic and Atmospheric Administration (NOAA) website emphasizes, climate change significantly impacts marine and Great Lakes ecosystems and fisheries, as well as the communities and economies that depend on them. The NOAA Climate, Ecosystems, and Fisheries Initiative (CEFI) has set an ambitious agenda: to "build the end-to-end operational modeling and decision support system needed to provide the information and capacity resource managers and stakeholders need to reduce impacts and increase resilience in a changing climate" (NOAA Fisheries, n.d.) [14]. This initiative reflects a broader trend—observed across many funding agencies—of embedding climate change projections into applied domains such as fisheries management.

Yet this raises a crucial question: are Earth System Models (ESMs) and their outputs adequate for these purposes? The answer, at present, appears to be no. As Kearney et al. [15] point out, "projections of LMR [living marine resource]–relevant metrics such as net primary production can vary widely between ESMs, even under identical climate scenarios." Recent versions of the CMIP protocol have introduced ESM variants with expanded biogeochemical modeling and representations of lower-trophic-level processes, but they continue to suffer from critical limitations. Many models fail to represent the taxonomic and functional diversity of phytoplankton and zooplankton, omitting traits such as energy density, lipid content, or acidification sensitivity that are essential to higher-order consumers. Their coarse spatial and temporal resolutions further limit their ability to resolve the fine-scale, life stage-specific biophysical interactions that shape fish recruitment, survival, and migration. As Kearney et al. note, this mismatch between the physical resolution of ESMs and the ecological resolution required for fisheries management significantly impairs their utility.

These limitations are corroborated by NOAA's internal assessment. A recent workshop report on modeling fish stocks in a changing climate—convened by NOAA scientists—underscores the challenges of translating climate projections into reliable forecasts of fisheries outcomes [16]. Participants noted that general circulation models (GCMs) often fail to resolve key oceanographic features that drive fish distributions, such as coastal upwelling systems or mesoscale eddies. Ecosystem models that couple climate inputs with biological processes must contend not only with uncertain biogeochemical drivers but also with nonlinear and often poorly understood ecological responses. The report also highlights the lack of integration between climate modelers and fisheries scientists, resulting in the misapplication of climate outputs or the omission of key uncertainties. Most tellingly, the report concludes that current modeling efforts are insufficiently equipped to support operational decision-making about fish stocks, especially at the scales and time horizons relevant to adaptation planning.

Fisheries managers are also studying the impact of climate change on fish populations in New Zealand rivers. Projecting the effects of climate change on freshwater fish in New Zealand rivers and streams is even more challenging due to the complex interplay between climatic, hydrological, ecological, and human systems. While temperature and streamflow are key drivers of fish physiology, habitat suitability, and life-cycle timing, these variables alone do not determine species persistence. Climate change interacts with many other factors—land use change, water abstraction, pollution, invasive species dynamics, and conservation policy—all of which shape the physical and biological conditions of freshwater ecosystems. Most modeling exercises, including those like Canning et al. [6], must make strong assumptions about these non-climatic variables: treating them as fixed, slowly changing, or too uncertain to meaningfully project. This creates a fundamental limitation on the scope and reliability of any forecast, as Canning et al. document in a lengthy discussion section.

Moreover, even the climate-driven variables central to these models—stream temperature and flow—are difficult to predict at the spatial and temporal resolution needed for ecological inference. Hydrological downscaling relies on climate model precipitation inputs, which are notoriously uncertain over New Zealand's complex terrain. Local flow regimes depend not just on climate but also on soils, vegetation, topography, and human water use—all of which introduce error and uncertainty. Finally, fish do not respond passively to environmental gradients; their distributions are shaped by migration barriers, competition, predation, disturbance regimes, and stochastic events like floods and droughts. These factors are rarely captured in statistical species distribution models and are likely to change in non-linear ways as the climate warms. While projections can highlight plausible directions of change and identify species at potential risk, they fall short of offering robust, policy-relevant predictions of future freshwater fish populations.

2.2.2 | Modeling Fish

The study by Canning et al. [6] received partial support from the Wellington Fish and Game Council, and the authors, along with many data sets and tools in the study, have close ties to NIWA's Climate Change Vulnerability Assessment (see Table 1). Like many conservation organizations, NIWA operates in a policy environment where climate change is a primary focus for funding and research prioritization. Consequently, the study emphasizes projecting the impacts of end-of-century climate scenarios—modeled using downscaled global circulation models and hydrological models—on the spatial distributions of native and nonnative freshwater fish. The authors conclude that under a high-emissions pathway (RCP8.5), up to nine native fish species could face extinction or near-extinction, while distributions of trout—New Zealand's valued sport fish—may decline by 30%–40%. They argue that such losses could be prevented by mitigating climate change and improving land use, and recommend proactive adaptive management strategies to address predicted range shifts and interspecific interactions.

However, the challenges in predicting freshwater fish responses to climate change—such as significant uncertainty in hydrological downscaling [17, 18], limited ecological realism in species distribution models [19, 20], and the exclusion of critical non-climatic drivers like land use change, pollution, water abstraction, and governance—suggest that the conclusions drawn are both contingent and highly uncertain. The structure of the funding and policy landscape incentivizes researchers to frame biodiversity threats and management solutions primarily in terms of climate scenarios, even when those scenarios represent only one aspect of a more complex causal system. Consequently, research like Canning et al.'s may overemphasize the role of CO_2 emissions and climate mitigation while neglecting proximate, actionable human drivers of freshwater degradation. This reflects a perverse incentive in the funding system: researchers are rewarded for aligning their work with climate priorities, even when the tools for assessing climate impacts are too coarse, and when more immediate ecological and policy variables—though harder to predict—may be more easily influenced by political power and could be more relevant to conservation outcomes.

The situation is complex, and I want to be clear to avoid misinterpretation. While Canning et al. may be insufficiently critical of the climate-hydrological modeling that underpins their work, one cannot fault them for downplaying the role of non-climatic variables in shaping fish populations by the century's end. Nor can one criticize them for being overly optimistic about the reliability of their findings. They clearly outline a long list of limitations in a discussion section that spans nearly four pages. However, this extensive discussion does not prevent them from including graphics that strongly suggest we possess detailed knowledge about how each representative carbon pathway (RCP 2.6–8.5) will impact individual fish species. The juxtaposition of this graphic (my Figure 1) with the discussion section, along with the other sources of uncertainty they do not address, illustrates a consequence of the perverse incentives I want to highlight in this paper.

2.3 | The Absence of Curated, Fit-For-Purpose Climate Information

One might hope that, given the growing demand for actionable climate knowledge by downstream scientists (in health, fisheries, etc.), climate science institutions would be well-funded to provide *interpretable, decision-relevant, and domain-specific model outputs* for health, agriculture, and related fields. However, this is not the case. While there are commendable initiatives (e.g., the Copernicus Climate Change Service in Europe, NOAA's Regional Integrated Sciences and Assessments program in the U.S.), they are small relative to the scale of the challenge and unevenly distributed across regions and sectors.

In most contexts, researchers must navigate model output repositories (e.g., CMIP archive data, downscaled climate scenarios) with minimal guidance regarding their adequacy for specific purposes. They must decide which models to use, which emissions scenarios to select, how to interpret ensemble spread, and how to communicate uncertainty—all without the necessary infrastructure or training for such interpretive work. The result is a proliferation of inconsistent practices: some researchers use ensemble means; others cherry-pick scenarios; still others present climate projections without incorporating their uncertainty at all. This is not willful misconduct. It is a rational response to institutional incentives coupled with epistemic scarcity.

Suppose you are a program officer or policy adviser tasked with designing a funding call for research on climate-sensitive challenges—say, the impact of changing precipitation extremes on agriculture and infrastructure in Bangladesh. You want to decide whether to require applicants to incorporate future climate projections in their work. In principle, the output of flagship Earth System Models (ESMs), now amounting to tens of petabytes of simulated climate data, should offer exactly what you need: daily precipitation values for Bangladesh, projected through the end of the century under various emissions scenarios. The CMIP6 archive, accessed through the Earth System Grid Federation (ESGF), contains just such data, neatly formatted, well documented in technical terms, and openly accessible.

But the apparent precision and completeness of this data is deceptive. From the ESGF metadata, you can find the model name, scenario, temporal resolution, units, and download links—but nothing tells you whether daily precipitation over Bangladesh is a field in which models have high skill. You do not learn that many models misplace the South Asian monsoon rainbands,

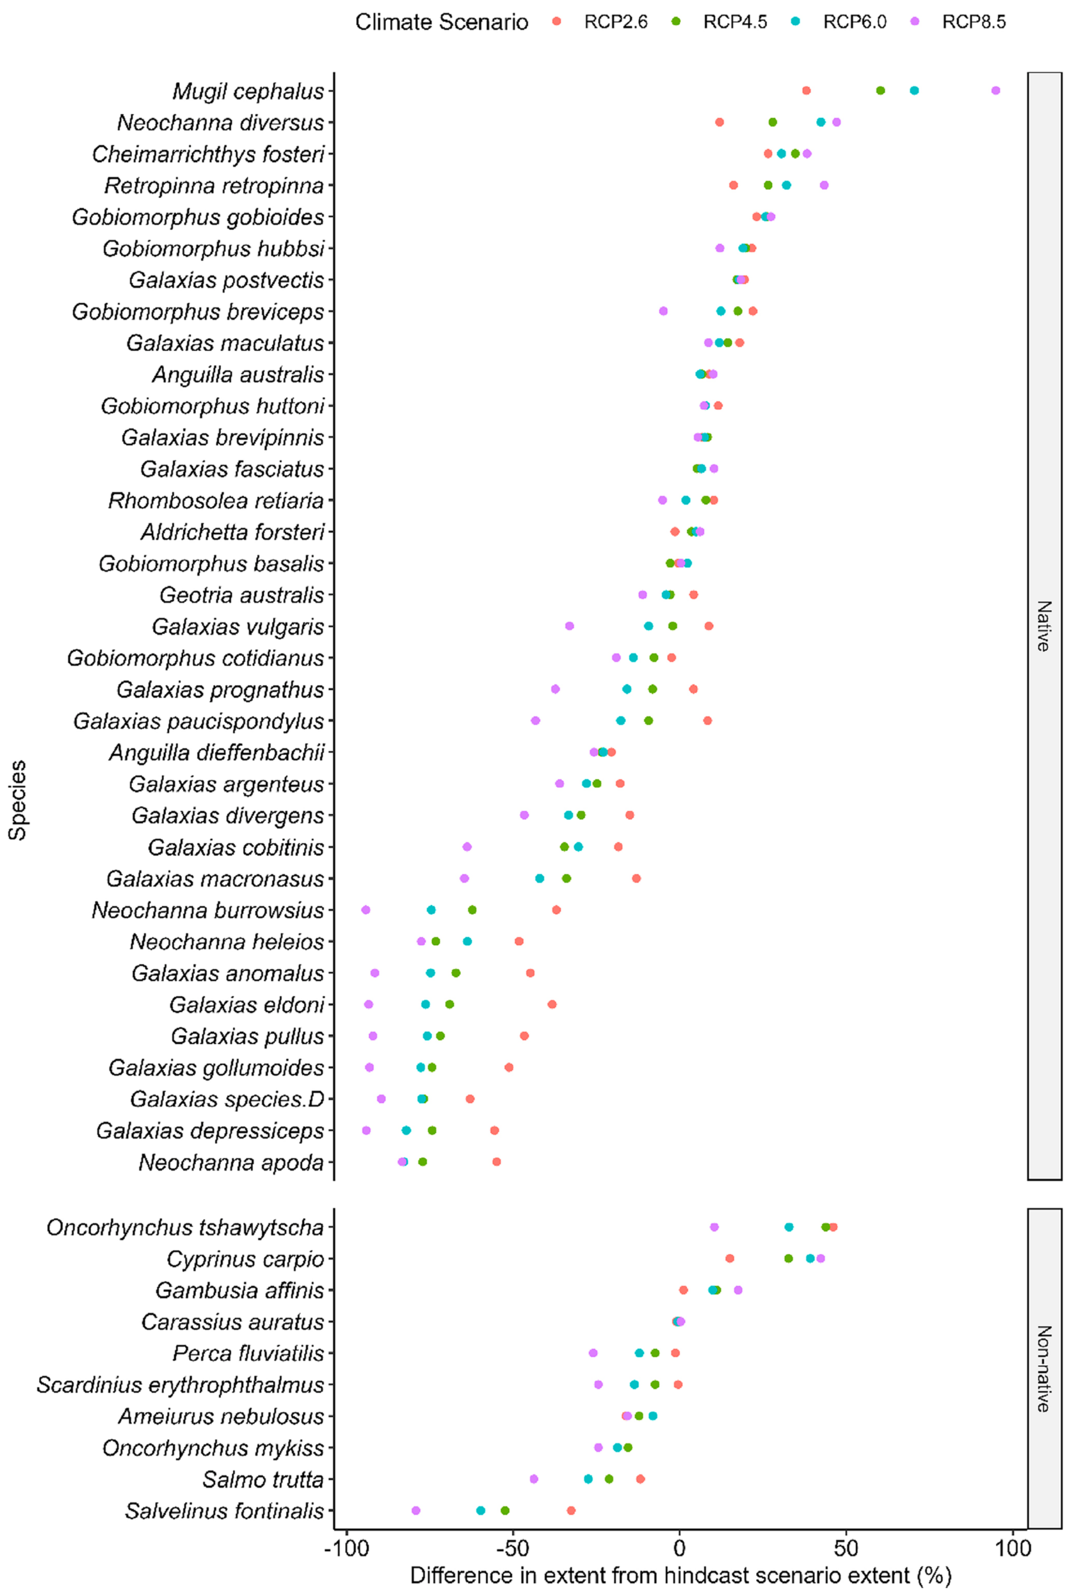

FIGURE 1 | The extent of the modeled hydrological network (%) across New Zealand predicted to be occupied by a given freshwater fish for four end-of-century climatic scenarios relative to the extent in the hindcast climate scenario (measured as percentage difference). This figure is adopted from Canning et al. [6].

that biases in seasonal totals can exceed 50%, or that changes in daily rainfall extremes vary in sign across models. You see no warnings that natural variability will swamp the forced signal for much of the century, or that the coarse resolution of ESMs means they miss the mesoscale convective systems that generate many of the region's most damaging rainfall events. In short, you are given the data without the epistemic context needed to judge whether it is fit for your purpose.

This opacity is not primarily the result of negligence or secrecy—climate modeling centers do invest heavily in data curation. But the focus of that curation is almost entirely technical: ensuring consistent file formats (netCDF), standardized variable names (CF conventions), and uniform metadata so that climate scientists can combine and compare outputs. The system is built to serve other modelers—after all, the original intent of CMIP was to improve future models, not to make model outputs useful to outsiders—not upstream users in different scientific domains or in policy. Model adequacy-for-purpose assessments are sporadic, buried in specialized literature, and almost never embedded into the archives themselves. Funding structures reward improving model physics, increasing resolution, and running more experiments; they rarely earmark significant resources for creating "epistemic metadata"—structured guidance on where models are strong, where they are weak, and how that maps onto the needs of other disciplines.

The result is a systematic risk of over-interpretation by outsiders. A non-specialist could easily treat the multi-model mean of CMIP6 daily precipitation in Bangladesh as a robust forecast, bake it into a health economics or fisheries model, and present the results with an aura of scientific authority—when in fact the projections may be no more than loosely constrained guesses for that particular variable, region, and timescale. This is not a rare edge case: the same structural problem applies to many regional, process-specific climate quantities that appear in the CMIP archive. Without sector-specific evaluation and clear adequacy signals, funding calls that require "accounting for climate change" risk compelling grantees to incorporate low-quality inputs simply because they are available and look precise.

This example is not unique. Most CMIP-class ESMs are too coarse to resolve the structure and peak winds of severe tropical cyclones. Raw track and frequency data are available in ESGF, but without warnings about resolution limits or the need for specialized downscaling, an outsider could treat these as reliable projections for coastal engineering design. ESMs simulate large-scale fire weather indices but not ignition, fuel loads, or smoke dispersion at decision-relevant scales. Coupling to air quality models is possible, but rarely evaluated in a coordinated way for future scenarios. The CMIP output looks detailed, yet without clear communication of these limits, it's easy to overstate confidence in future $PM_{2.5}$ exposure. ESM ocean components often have large biases in upwelling strength and nutrient flux. Fisheries models that take climate inputs exist, but few have been robustly driven by multi-model ensembles with quantified uncertainty. ESGF doesn't flag the low confidence in regional productivity changes, so a funding call could inadvertently require applicants to base their work on unreliable projections. Across these cases, the same pattern emerges: technical curation without embedded adequacy guidance leaves non-specialists unable to judge whether model outputs for their variable, region, and timescale are robust, marginal, or effectively guesswork.

In other words, the culture, structure, and funding arrangements of contemporary climate science are not optimized to produce what upstream users most need: not just access to climate model outputs, but usable knowledge about their strengths and limitations in specific contexts. Bridging that gap would require dedicated funding for sustained cross-disciplinary evaluation, embedding adequacy-for-purpose tags in the data infrastructure itself, and cultivating norms that treat such guidance as part of the modeling community's responsibility. Until then, the petabytes of climate model output will remain, for most outsiders, technically curated but epistemically opaque—and upstream actors will have to navigate this opacity at their own risk.

The epistemologist might argue that what is missing is not just information, but *second-order information*: knowledge about when and how first-order climate projections can be trusted for specific uses. While climate scientists may possess this second-order knowledge, it is often tacit, underfunded, or isolated within specialist communities. Thus, the social epistemologist would assert that what is lacking is the institutional scaffolding[5] necessary for effective knowledge production. Few funding mechanisms support the careful curation of model outputs for interdisciplinary use, reflecting a structural failure to recognize the representational fragility of climate projections across contexts. As Parker [21] and I [5] emphasize, climate models are not general-purpose tools; they are representational constructs with specific scope conditions and contextual limitations. When those conditions are ignored, the models' outputs become epistemically vulnerable—even when technically sound.

2.4 | Incentivized Compliance and Research Maladaptation

In this context, the requirement to integrate climate projections takes on a performative character. Researchers include climate scenarios not because they find them reliable, but because they are institutionally compelled to do so. Over time, this distorts the incentive landscape: it rewards researchers not for epistemic responsibility, but for formal compliance. Worse, it may displace genuine efforts to adapt to climate risks. In public health, excessive focus on future climate-driven disease burdens can distract from present-day vulnerabilities and low-cost interventions. In agriculture, reliance on poorly grounded yield projections can lead to premature technological lock-in or misaligned investment strategies. In fisheries, using incompatible climate-ecological models can justify policies that over- or under-protect stocks. These are not merely theoretical risks; they are increasingly the unintended consequences of well-meaning but epistemically incoherent institutional design.

3 | Trust and Overextension: When Climate Models Do Too Much

The demand for climate-integrated research, particularly in health, agriculture, and resource management, is often driven by an appeal to trust: decision-makers and the public are told that science is equipping society to prepare for climate change, that projections are responsibly incorporated into research and policy, and that future-oriented planning is grounded in rigorous evidence. But what happens when the epistemic machinery behind these projections is less robust than the institutional practices that rely on them? What happens when the appearance of rigorous projection outpaces the actual reliability of the underlying models?

This section argues that current practices in climate-integrated research are generating representational overextension: the use of climate models in contexts where their evidentiary authority is weak, but their rhetorical authority is strong. This overextension is not always intentional, but it is structurally incentivized. It poses a serious threat to one of the most critical features of science in public life: its trustworthiness.

3.1 | Climate Models as Policy Tools and Public Symbols

Climate models were not originally designed to produce the kinds of projections now required in many applied fields. The general circulation models (GCMs) and Earth system models (ESMs) used in international modeling exercises like CMIP are built to address large-scale scientific questions: How will global mean temperature respond to different emissions scenarios? What feedback loops exist between the carbon cycle, the ocean, and the atmosphere? How do aerosols affect planetary energy balance?

These models are invaluable for such questions. However, they were not constructed, validated, or calibrated to make regional predictions on 10–20-year horizons about variables such as wildfire ignition likelihood, zoonotic disease emergence, populations of specific species in fisheries stocks, or crop-specific agricultural yields. Nor were they designed to serve as direct inputs for cost-effectiveness analyses of public health interventions or infrastructure investment decisions. Yet, that is increasingly how they are being used.

One reason for this shift is institutional: funding agencies now mandate climate integration. Another reason is rhetorical: model outputs often resemble data—high-resolution maps, probabilistic time series, or heatmaps of projected health risks. These outputs can carry the authority of empirical observation, even when they are fundamentally scenario-based extrapolations built on coarse-grained assumptions and structurally uncertain models.

The danger is not only that decision-makers or the public may misinterpret these outputs; scientists themselves—particularly those outside climate science—may mistake precision for accuracy or representation for reliability. This is the epistemic heart of the problem. Model outputs can become what we might call *epistemic artifacts of trust*: objects treated as evidence not because of their underlying validity, but due to their formal resemblance to other trusted objects.

When these artifacts guide high-stakes decisions, their epistemic weaknesses become ethical vulnerabilities. They can justify policies that misalign with the actual state of knowledge, disregard stakeholders' values, or prioritize adaptation strategies unsupported by the best available evidence.

3.2 | Pseudo-Precision and the Corruption of Salience

A particularly dangerous form of representational overextension occurs when model outputs exhibit what we might call pseudo-precision: numerical detail or spatial resolution that exceeds the model's validated range. For example, downscaled CMIP projections may provide temperature or precipitation estimates at a 10 km grid resolution for the year 2036. This apparent precision masks profound uncertainties—structural uncertainty about model dynamics, internal variability that overshadows forced trends, and parameter uncertainty not constrained for local scales [22].

Nevertheless, pseudo-precision is often reified in downstream applications. A health economist might use such a projection to simulate the changing burden of asthma in a region over a 15-year span. A fisheries scientist might rely on model-projected sea surface temperature anomalies to estimate spawning shifts in a coastal species. These researchers are not acting irresponsibly per se—they are responding to institutional mandates. However, the outputs they rely on may have exceptionally low salience for the questions at hand.

This disconnect reflects a breakdown in *epistemic salience*: the alignment between what a model represents and what a user perceives it to represent. Ideally, model outputs are salient to the user's epistemic aims—they represent phenomena relevant to the user's decision problem, at the right level of abstraction, and with interpretable uncertainty. In many current cases, the salience is artificially induced: it arises not from the model's fidelity to the decision-relevant target, but from institutional pressure to use climate projections as a sign of sophistication, credibility, or relevance.

3.3 | Trustworthiness Versus Trust: A Philosophical Distinction

While Stephanie Harvard and I [2] emphasize that the duty to clarify model adequacy for purpose falls under the general responsibility to avoid foreseeable harms, there are good reasons to treat this as one of scientists' role responsibilities—core obligations tied to their social role as experts. A well-functioning society depends on some degree of trust in science (e.g., [23–25]), and being trustworthy is arguably essential for scientists' authority in public policy, education, technology, and culture ([23], p. 71). Trust can be misplaced; what matters is trustworthiness. Prominent accounts define trustworthiness as reliability plus a morally valenced factor, such as a genuine commitment to the trustor's interests (McLeod 2023; Hawley 2014) [26, 27]. This distinction clarifies why it is a category error to describe a thermometer as "trustworthy": it may be reliable, but it cannot possess morally valenced commitments. For scientists, maintaining trustworthiness requires avoiding epistemic practices known to be unreliable. Making a model—or its outputs—available without investigating its adequacy for purpose or informing relevant users of known limitations exemplifies such unreliability.

Hills (2023) [28], however, argues that even a sophisticated reliability-plus-commitment account is insufficient to fully capture trustworthiness. She characterizes trustworthiness as a form of *responsibility*, noting that many tasks we entrust to others are too complex and unpredictable to articulate in advance, requiring trustees to exercise judgment and negotiate value-laden trade-offs. Modeling work exhibits precisely these features: modelers face unforeseeable challenges and must choose representational strategies that best serve the model's intended purpose (Harvard et al. 2020; Winsberg 2012) [29, 30]. According to Hills, trustworthy scientists take responsibility for what is entrusted to them (e.g., an epistemic project) and display "the right kind of values, motivations, and judgments" (Hills 2023, 758) [28]. Her account supports our (Harvard and Winsberg 2022) [2] claim that ensuring model adequacy-for-purpose is a moral-epistemic duty.

Both Hills (2023) [28] and Hawley (2014) [26] provide useful insights. Hills emphasizes that trustworthiness in modeling involves responsibility for an open-ended task, requiring continuous critical evaluation of adequacy for purpose and promoting models only for genuinely suitable purposes. Hawley adds that the trust relationship between science and the public has a contractual element: the public funds and defers to scientists with the expectation that they will produce knowledge serving public interests. While this implicit contract cannot specify all details in advance, scientists have a duty to clarify those details in ways the public is unlikely to reject. In both cases, trustworthiness necessitates ongoing stewardship of the uses and limitations of scientific models.

4 | Representational Risk and Harmful Signaling

The previous sections illustrate how institutional expectations have outpaced the epistemic capacities of climate modeling, leading to maladaptive scientific labor and erosion of science's trustworthiness. Here, we focus on the representational risks arising from the signaling function of models in public reasoning and institutional discourse. The issues we highlight extend beyond wasted epistemic resources and trust erosion; they can lead to specific harms from sending false signals about policy-relevant information.

The problem is not only that projections might be inaccurate or that uncertainty is poorly communicated. The mere inclusion of climate projections—especially in contexts with weak evidential bases—can send misleading signals regarding the nature and severity of risks, the applicability of climate models, and the quality of scientific consensus. These signals can reshape public understanding, distort deliberation, and misallocate political and moral attention.

This section draws on the concept of *representational risk*, introduced in recent work on modeling and public communication, elaborating on its interaction with *institutional signaling incentives*. The argument is not merely epistemological; it is also political and ethical. In contexts where models serve dual roles—as tools for inquiry and as badges of relevance—their misuse can constitute epistemic misrepresentation, even without explicit falsehoods.

4.1 | Models as Signals: The Case of Prescribed Burns

Work on prescribed burns and wildfire management offers a vivid example. Representing uncertainty involves value-laden judgments with ethical implications. Despite the complex determinants of wildfire, many sources emphasize that future wildfires will inevitably increase in frequency and severity due to climate change. In public health modeling, where health outcomes are sensitive to air quality, a wide range of uncertainty might lead decision-makers to underestimate the seriousness of wildfire smoke issues, potentially discouraging investments in climate adaptation or mitigation measures. A related worry is that modeling the full range of possibilities may not be informative for decision-making regarding interventions. Conversely, depicting wildfires as necessarily intensifying could suggest they are beyond the influence of local policies, undermining support for local interventions like prescribed burns. Prescribed burns and traditional Indigenous burning practices are increasingly recognized for their role in wildfire mitigation [11]. However, these practices face regulatory hurdles due to health concerns from smoke exposure, with the health and economic impacts of prescribed burns remaining inconclusive.

D'Evelyn et al. recommend collaboration between environmental and health scientists to optimize prescribed burns, minimizing negative impacts on health and equity. However, could representing wildfire as necessarily intensifying undermine this goal? There is a need to consider the 'signaling power' of models here. By "here," I refer not to modeling work explicitly studying the desirability of prescribed burns. I am referring to modeling efforts, such as cost–benefit analyses of health interventions related to diseases exacerbated by poor air quality. The representational choices made by *these modelers* could inadvertently signal to decision-makers that interventions to control wildfires do not exist (when they do) or are on balance harmful (when this is uncertain). This could send misleading signals about prescribed burns and hinder essential research to establish their net benefits. The crucial issue often revolves around 'smoke trade-offs': the idea that there is always a trade-off between the harms caused by smoke during prescribed burns and the harms (hopefully) prevented by those burns concerning (hopefully) averted wildfires. Health economists could advance this initiative by collaborating with respirologists and other experts to develop models relevant to forest managers, regulators, and health decision-makers. Key questions include the relative costs and benefits of prescribed burns compared to uncontrolled wildfires across different populations and conditions. In the future, health economists could model the cost-effectiveness of prescribed burns and associated policies, including protections for vulnerable populations during burns.

Instead, by incorporating increasing wildfire smoke into health policy or economic models as an input rather than an output, health researchers could signal that, *in their expert opinion*, wildfire management is hopeless—perhaps due to unfavorable health trade-offs of prescribed burns. It is easy to imagine a council meeting where well-meaning individuals attempt to overcome regulatory hurdles to implement prescribed burns to improve health outcomes, only to face opponents armed with published papers by health policy experts and health economists using models that assume increasing wildfire smoke. These opponents will argue that relevant domain experts have a pessimistic view of controlling wildfires. In reality, the opposite may be true; health policy experts and health economists might know little about the impact of climate change on wildfire smoke but are responding to powerful structural incentives to include these assumptions in their models.

4.2 | Balancing Ethical Implications: Considerations for Health Modeling

In my experience speaking with health modelers, many health scientists perceive increasing wildfires as inevitable and beyond the influence of local adaptation policies. This perception conflicts with insights from climate science and forest management, which acknowledge significant uncertainties and the potential impact of human activities on wildfire activity. Ethical modeling, and importantly, ethical science funding infrastructure, requires transparency about the limits of scientific knowledge and honesty about uncertainties, even when they pose challenging policy questions. Overstating confidence in model projections can send inappropriate signals and hinder the development of effective interventions.

The challenges faced in developing the LEAP model [2] prompt broader considerations regarding health economic modeling in the context of climate change:

1. *Limitations of Current Models*: Forecasting regional events affecting health is subject to significant limitations. Climate models may overpromise their ability to predict local events, and health economists should critically evaluate the disclaimers and limitations noted by climate scientists.
2. *Impacts of Human Activities*: The impacts of climate change are uncertain and intertwined with local human activities. Understanding the role of human interventions, such as forest management practices, is crucial for modeling and policy development. Some may see value in emphasizing the link between wildfire and climate change, believing it will prompt North Americans to take climate change seriously if they perceive it as an immediate threat. However, this could downplay the effectiveness of local interventions like prescribed burns. Scientific integrity and a genuine interest in public health necessitate transparency about the limits of science and honesty about uncertainty, even when it raises difficult political and ethical questions. Climate scientists have valid reasons to fear that wildfires will intensify in the future, but there is no conclusive evidence that this is unavoidable or insensitively tied to policy decisions, nor a crystal ball for predicting the strength of intensification.
3. *Cascading Uncertainties*: As health researchers are increasingly required to integrate climate change considerations into their work, particularly in areas like respiratory health and asthma onset, a key challenge arises: What scientific resources are available to support claims about future wildfire smoke exposure? In jurisdictions like British Columbia or California, where wildfire smoke is a significant concern, researchers assessing future health burdens or the cost-effectiveness of mitigation measures (such as HEPA filters or relocation strategies) must rely on projections of smoke exposure over decadal timescales. Yet they often lack expertise in fire ecology or climate modeling and depend on a small and unevenly distributed set of studies and data products—produced largely outside the health sciences—by groups (e.g., "environmental epidemiologists") who do not possess genuine climate science expertise or access to detailed information about the quality of data available for download from CMIP.

Three studies dominate the current landscape of decadal wildfire smoke forecasting. Liu et al. [13], working with the U.S. Forest Service, use dynamic vegetation models coupled with empirical fire and fuel moisture models to estimate fire emissions across the western United States under mid-century climate scenarios. Their projections suggest that fire-related emissions could increase by approximately 50% by the 2050s compared to early 21st-century baselines. The Four Twenty-Seven (2020) report—commissioned by a climate risk analytics firm—provides a global gridded dataset of wildfire potential at ~25 km resolution, projecting significant increases in the length of high fire-risk seasons in key regions such as California, Australia, and southern Europe. A third study, a working paper from the National Bureau of Economic Research (NBER), attempts to quantify the future mortality burden of smoke exposure using statistical models linked to climate-emission trajectories.

These studies are largely the only tools available to interdisciplinary researchers tasked with incorporating climate-driven wildfire risk into downstream areas such as health research, epidemiology, and urban planning. This is especially true given the lack of centralized infrastructure for interpreting climate model outputs—let alone translating those outputs into exposure metrics relevant to public health. Consequently, researchers in these fields often adopt the outputs of fire modeling studies uncritically, treating them as established decision-support tools.

In summary, the limited number of publicly available wildfire smoke projection tools now serve as de facto infrastructure for applied health research in climate-sensitive domains. Their use has become "accepted best practice" in applied fields. However, these tools carry significant representational and epistemic risks. Addressing this situation requires not only technical improvements in fire modeling but also institutional reforms that clarify when, where, and how climate projection products should—and should not—be used across scientific domains.

4.3 | Up and Down the Food Chain

The risk of overstating confidence in model-based projections compounds as one moves up and down what might be called the *modeling food chain*. Each layer in this chain—from physical climate models to ecosystem simulations to public health assessments—can introduce new uncertainties. Yet these uncertainties are often not transparently propagated through the chain. While climate modelers may include careful caveats about the limits of global models or regional downscaling, these caveats often dissipate as model outputs are passed downstream to environmental modelers, who use them to estimate phenomena such as future fire behavior or ecosystem responses. In turn, environmental modeling products are taken up by researchers in applied health domains, who may treat these outputs as fixed inputs—narrowing the range of projected wildfire smoke exposure values and conducting sensitivity analyses without fully engaging with the deeper uncertainties upstream.

The dynamic also works in reverse. Consider the widely cited study by Maji et al. (2022) [31], which estimates the number of premature deaths associated with prescribed burns, aiming to quantify trade-offs between different fire management strategies. They estimate excess mortality from exposure to $PM_{2.5}$ and ozone generated by prescribed fires in the southeastern United States, explicitly acknowledging several sources of uncertainty in their methodology. For instance, they note that their mortality outcomes differ from prior studies "due to varying choices in the selection of concentration–response functions" and recognize that existing models, such as the Global Exposure Mortality Model (GEMM) and the Integrated Exposure–Response (IER) function, assume equal toxicity of $PM_{2.5}$ regardless of its source. They also state that their air pollution simulations rely on emissions inventories and chemical transport models that, while state-of-the-art, have known limitations. Importantly, Maji et al. are transparent about the fact that their estimates should not be interpreted as precise forecasts but rather as scenario-based approximations intended to inform public health planning.

My intention is not to criticize the Maji et al. paper, which responsibly characterizes its own limitations, but to highlight that norms can emerge in related sciences where it becomes standard to simply plug values from a paper like Maji et al. into a downstream model, glossing over the disclaimers in the original work. Thus, the study exemplifies a broader pattern in environmental health modeling: the tendency for complex, uncertain upstream inputs—such as emissions inventories, chemical transport simulations, and source-agnostic concentration–response functions (CRFs)—to be passed through successive modeling layers in ways that may obscure upstream uncertainty. While Maji et al. employ reputable CRFs, such functions assume linear or log-linear relationships [32] between pollutant exposure and mortality, despite evidence that the true exposure-response relationship may be nonlinear and population-specific. Their models also fail to account for the chemical complexity of wildfire smoke, which includes a wide array of reactive and potentially more toxic co-pollutants beyond $PM_{2.5}$ and ozone. A growing body of research indicates that wildfire smoke may be more harmful to human health than equivalent concentrations of $PM_{2.5}$ from urban or industrial sources due to its distinct chemical composition. For example, Reid et al. [33] find that wildfire-specific $PM_{2.5}$ is associated with higher rates of respiratory hospitalizations than non-wildfire $PM_{2.5}$, suggesting greater toxicity per unit mass. Additional studies identify elevated levels of volatile organic compounds (VOCs), polycyclic aromatic hydrocarbons (PAHs), and other harmful pollutants in wildfire smoke, complicating attempts to model its health effects using generalized concentration–response functions [34]. Recent studies indicate that wildfire smoke, as it ages, dramatically changes its composition concerning some of the most dangerous pollutants regarding long-term health risks.

Furthermore, population vulnerability to smoke is highly heterogeneous; using uniform effect estimates across counties may mask significant differences in baseline health, housing conditions, and capacity to mitigate exposure. Finally, the analysis abstracts from the policy context in which prescribed burns might occur. As others have noted, the health consequences of burns depend not just on emissions but on co-occurring interventions—such as public warnings, evacuations, and clean-air shelters—that are rarely modeled but may significantly alter outcomes.

More fundamentally, the use of such functions presupposes a framing of health harms that focuses exclusively on premature mortality. However, the public may value other outcomes—including quality of life, economic security, or disruptions to daily activity—and these values are rarely captured in modeling exercises. Moreover, even if we accepted the accuracy of every scientific input, the health effects of prescribed burns would remain deeply contingent on implementation. A prescribed burn accompanied by timely evacuation and public health messaging may pose minimal risk, whereas one executed without such supports may yield far greater harms. These social and institutional contingencies are rarely modeled and even more rarely made visible in decision-support outputs.

These limitations do not discredit the Maji et al. study, but they highlight the challenges of building trustworthy, policy-relevant knowledge when cascading uncertainties are insufficiently tracked across the modeling food chain. Health researchers rely on overly simplistic climate forecasts of wildfire smoke, which climate researchers recognize as limited, while forest management policy researchers use all-cause mortality estimates from prescribed burning programs based on CRFs that health researchers know to be overly simplified. Urgent integration of diverse background knowledge from different disciplinary groups is needed to manage representational risk effectively.

What we sometimes (but of course not always, I am not making a universal claim here) observe is a cumulative dynamic in which uncertainty is compressed at each step in the modeling process, even as the outputs are applied in increasingly policy-relevant and normatively charged contexts. This can create a misleading sense of convergence—not because uncertainties have been resolved or value-laden decisions have been settled, but due to the replication of simplifying assumptions across model layers. The result is an end product with a manufactured and inflated social significance: projections that seem rigorous and quantified yet lack transparency about their representational limits and normative blind spots. When these projections are used in fields like health science under cross-cutting climate integration mandates, the risk is not just scientific error—it is the erosion of science's trustworthiness as a guide for public action.

So why are they included? In part, because they serve as signals. They convey that researchers are "taking climate seriously," that the model is forward-looking and responsible, and that it aligns with funder mandates or policy priorities. The inclusion of a climate scenario becomes more than a representational choice; it becomes a rhetorical one. More importantly, this occurs due to structural and institutional mandates in the science funding apparatus.

However, such signaling can overshadow alternative policy framings. If the model indicates that "climate change will increase wildfire risk in this region over the next 12 years," even when that claim lacks robust evidential support, it may divert attention from other pressing drivers—such as legacy fuel build-up, poor fire management, or zoning policy. Once included, a model frames the problem space, informing stakeholders about plausible interventions, salient risks, and credible futures. If people believe that prescribed burns "result in such and such all-cause mortality," they may assume decisions about their implementation have an all-things-considered answer agreed upon by experts.

4.4 | The Ethics of Modeling Under Epistemic Uncertainty

A deeper ethical question arises: What are scientists obligated to signal? More importantly, what should the structure of science funding incentivize in terms of scientific signaling? When a model is included in a public-facing analysis, it is not just a technical object—it is a claim about the state of the world and the structure of uncertainty. Including a projection implies that it is decision-relevant, carries evidential weight, and represents a justified expectation.

When researchers include projections they know to be weakly supported—or for which they lack domain-specific interpretive capacity—they risk engaging in what may be termed *epistemic bad faith*. This is not deception in the usual sense but a subtler failure to honor the norms of scientific representation: including only what is justified, explaining what is uncertain, and clearly distinguishing between credible inference and symbolic compliance.

These failures are typically not the fault of individual researchers. They arise from institutional design: from funder mandates, evaluation criteria, and the political economy of scientific prestige. Yet, the ethical burden remains. If science is to maintain its role as a trustworthy guide for public reasoning, it must resist the temptation to use models as signals of responsibility and correct political alignment rather than as tools of inquiry. Scientists, scientific institutions, and funding agencies invite public trust. Simply attaching disclaimers to model results or acknowledging uncertainties does not always suffice, especially when institutional incentives or communicative framing amplify the perceived authority of model outputs. Trustworthiness in science requires more than cautious language; it demands attention to the representational risks inherent in transmitting model-based knowledge, including the potential for outputs to be misunderstood, overinterpreted, or embedded in policy decisions in ways that exceed their evidentiary warrant. In short, responsibly inviting trust entails anticipating and managing the moral and epistemic consequences of how scientific claims circulate and are acted upon. In a science-funding context, inviting public trust obligates institutions to refrain from incentivizing the erasure of these safeguards.

5 | Institutional and Philosophical Implications

The preceding sections have revealed a structural tension in the expectation that climate science should support applied research and public policy. On one hand, there is a growing consensus—among scientists, policymakers, and funders—that the effects of climate change must be integrated into all domains of inquiry: health, agriculture, fisheries, infrastructure, and beyond. On the other hand, there is insufficient systematic investment in ensuring that the climate information used for this integration is reliable, interpretable, and decision relevant at the necessary spatiotemporal scales.

This tension gives rise to maladaptive institutional behavior: researchers are required to use model outputs whose validity for their domain is weak; funders reward symbolic climate integration over epistemic rigor; and models are employed not for what they can legitimately represent but for what they can signal. The result, as argued, is a loss of epistemic integrity, a risk of public mistrust, and a distortion of the norms that make science a reliable guide to policy.

In this closing section, I shift from diagnosis to prescription. What institutional reforms and philosophical principles are needed to mitigate these risks? How can we create a research ecosystem in which climate integration is epistemically responsible, publicly trustworthy, and democratically legitimate?

I begin with three institutional recommendations, each grounded in the philosophical analysis above.

5.1 | Fund the Curation of Model Outputs, Not Just Model Construction

The first and most urgent recommendation is straightforward: *climate science funders must support the curation, interpretation, and evaluation of model outputs for decision-relevant domains*. Currently, climate modeling programs invest heavily in model development (e.g., new parameterizations, increased spatial resolution, multi-model ensembles) but comparatively little in what might be called epistemic translation—the process by which model outputs are made usable, interpretable, and appropriately constrained for use in non-climate domains.

Without this investment, downstream researchers—those in health, agriculture, fisheries, and infrastructure—are left to interpret complex model outputs without guidance regarding fitness-for-purpose. As Nissan et al. [1] emphasize, the result is widespread misuse: "Users are required to make difficult decisions about emissions scenarios, models, and downscaling methods, often with little or no formal training in climate science." This situation is not just suboptimal; it is structurally unethical. Responsibility for interpreting complex epistemic artifacts is offloaded onto researchers who lack the resources or training to do so, while the institutions that produce those artifacts face no incentive to ensure their responsible uptake.

In practical terms, this reform would involve:

- *Dedicated funding* for interdisciplinary boundary organizations that interpret and translate model outputs for specific applied domains.
- *Formal epistemic audits* of model applications—that is, expert evaluations of whether a given use of model output is well supported by the underlying data and model structure.

- *Metadata standardization* for climate models, including documentation about validated use cases, known limitations, and appropriate decision horizons.

This infrastructure already exists in skeletal form—for example, in the UK Climate Projections (UKCP), the Copernicus Climate Change Service, or NOAA's RISA programs—but these efforts are underfunded and unevenly distributed. What is needed is not only technical capacity but social-epistemological clarity: a recognition that the meaning of a model output depends not just on its numerical content but on the institutional context in which it is interpreted and used.

5.2 | Shift From Mandated Inclusion to Justified Relevance

The second recommendation concerns research policy: *funding agencies should shift from requiring climate change integration as a formal criterion to demanding a justification for its relevance.* Rather than asking every researcher to include a climate projection regardless of its appropriateness, agencies should require researchers to explain whether and how climate change is relevant to their question—and, more importantly, if climate change is potentially relevant, whether climate science has adequate tools for delivering the inputs needed to incorporate climate change into a scientific investigation.

This shift would incentivize epistemic responsibility over rhetorical compliance. It would allow researchers to exclude climate projections when they are genuinely irrelevant or uncertain while rewarding those who invest in understanding and contextualizing the climate information they include.

Importantly, this is not a call to retreat from climate-integrated research. Rather, it is a call to recalibrate the institutional incentive structure through which climate is integrated. Instead of treating climate projection inclusion as a checkbox—something required to signal relevance—funders should evaluate whether the inclusion reflects epistemic justification and appropriate characterization of uncertainty. Rather than requiring that a funded project in health science integrate the effects of climate change, one could require that a grant application demonstrate due diligence in determining whether the project's goals can responsibly incorporate climate change given the quality of available evidence. This requirement would be much easier to meet if the requirements of Section 4.1 were being fulfilled by climate science groups.

5.3 | Recognize and Manage Representational Risk as a Core Norm of Science Governance

The final recommendation is conceptual and institutional: *representational risk must be recognized as a primary concern in science governance.* As Harvard and I have shown, inadequate models can do more than produce wrong results. The consequences of using such models can extend beyond epistemic error to include moral, social, and institutional damage. Institutions must treat the deployment of model projections as a practice subject to ethical and epistemic norms—not just technical ones. Just as there are standards for data privacy, human subject research, and conflicts of interest, there should be institutional standards for representational propriety: norms governing when and how projections are used, how their uncertainty is communicated, and what their inclusion implies about the state of knowledge.

This is especially important for domains where model outputs directly influence public decisions—for example, cost–benefit analysis, regulatory impact assessments, or health policy modeling. In such contexts, representational risk can have real-world consequences: inappropriate model use can justify policies that harm vulnerable populations, entrench poor infrastructure choices, or produce distracting information that diverts attention from more pressing risks.

5.4 | Science, Legitimacy, and Democratic Accountability

This analysis leads us to a broader philosophical point. The authority of science in public life depends not just on its empirical success but on *its legitimacy as a form of reasoning in a democratic society. That legitimacy, in turn, depends on the transparency, proportionality, and accountability of its representations.*

Climate change poses unique challenges for democratic governance. Its risks are long-term, global, and deeply uncertain. To respond effectively, societies must rely on science. But if the models that inform that response are used in ways that exceed their evidentiary basis—especially in contexts affecting public goods, resource allocation, or regulatory decision-making—then the authority of science is jeopardized. Not because its models are "wrong," but because the institutions that deploy them fail to manage representational risk carefully.

This paper has argued that a central institutional failure lies at the intersection of climate science and its allied domains: a mismatch between what models can reliably provide and what applied researchers are expected to deliver. That failure leads to three cascading harms:

1. *Maladaptation* in the allocation of scientific labor and funding.
2. *Erosion of epistemic trustworthiness* through overextension and pseudo-precision.
3. *Distortion of policy discourse* via signaling incentives that reward symbolic compliance over epistemic responsibility.

Each of these failures is remediable—but only if we take seriously the philosophical and institutional implications of using models not just as representations but as signals. Climate change is too important to be left to performance. If we want science to inform our future, we must build institutions that support its integrity—not just its image.

6 | Conclusion: Reclaiming Responsible Representation

This paper has examined a growing yet underappreciated tension in the current landscape of publicly funded science: the expectation that researchers in health, agriculture, fisheries, infrastructure, and other domains will incorporate climate change projections into their work, even when those projections are not epistemically secure or decision-relevant at the necessary scale. I have shown that this expectation generates a mismatch—a form of institutional incoherence—between the epistemic capacities of climate science and the representational demands placed on applied researchers.

Three interrelated problems follow from this mismatch. First, maladaptation: when non-climate researchers are expected to integrate projections that are unavailable, unreliable, or epistemically irrelevant, scientific labor is misdirected, and funding is allocated in ways that do not reflect where the need for foundational knowledge is most acute. The applied sciences are burdened with representational responsibilities that climate sciences have not been funded—or structurally equipped—to fulfill.

Second, the erosion of trustworthiness: climate models, once used cautiously to inform long-term global trends, are now included in cost-effectiveness studies, health policy models, and regional planning analyses for purposes they were not designed to serve. These inclusions are often driven not by genuine epistemic fit but by institutional mandates or rhetorical incentives. The result is a subtle but pervasive corruption of public trust: not through outright misinformation, but through epistemic overreach masquerading as rigor.

Third, representational risk through harmful signaling: when climate projections are included in policy-relevant research, they function not only as representations but also as signals—to funders, policymakers, and the public—that climate change is relevant, that its effects are predictable, and that the research is forward-looking and responsible. But when these signals are not backed by reliable evidence, they mislead. They distort public reasoning, redirect attention from more plausible causal pathways, and crowd out mitigation strategies better supported by available knowledge.

These failures are not primarily the fault of individual researchers. They are predictable consequences of institutional design. They call for institutional reforms that foreground the concept of representational risk not as an abstract exercise, but as a practical framework for understanding and correcting the epistemic and ethical dynamics of science in public life.

I have proposed three such reforms. First, climate science must invest in the curation and domain-specific translation of model outputs, not just in their technical development. Second, funders must shift from mandates that require climate integration regardless of context to justification-based evaluation that prioritizes epistemic relevance and uncertainty awareness. Third, the scientific community and its institutions must recognize representational risk as a first-order norm, governing how models are selected, interpreted, and communicated in public-facing research.

More broadly, I argue for a reorientation of climate-integrated science around a set of epistemic virtues that are increasingly at risk: humility, transparency, proportionality, and context sensitivity. These are not merely scientific ideals; they are also democratic ones. In a society where policy decisions are mediated through expert knowledge and model-driven reasoning, the public must be able to rely on scientific institutions to represent what they know—and do not know—with care.

This reorientation is especially urgent given the political and rhetorical power of climate change. Few scientific concepts are as symbolically potent or normatively loaded. It is precisely because climate models carry such weight that their use must be disciplined by a philosophy that takes representational risk seriously. Otherwise, models become tools of compliance rather than inquiry—tokens of institutional relevance rather than instruments of epistemic support.

I do not argue against the integration of climate science into applied research; quite the opposite. My concern is that, under current institutional arrangements, this integration often occurs in hollow ways. If climate change is to be taken seriously—as it must—then we must also consider the epistemic conditions under which its projected impacts can responsibly inform other domains.

Ultimately, the authority of science in a democratic society depends not only on its empirical adequacy but also on its representational integrity, which primarily stems from scientific institutions, particularly funding agencies. It relies on institutions that align incentives with epistemic justification and on practices that communicate uncertainty without undermining the urgency for action. Above all, this requires building a research ecosystem where models are used not to *perform* our concern about climate change but to enact it.

Acknowledgments

I would like to thank Kate Johnson, Stephanie Harvard, and other members of the LEAP model team for the opportunity to engage with their project and for the many discussions that inspired this paper, and also to thank Stephanie for helpful comments on earlier drafts. I would also like to thank Rachel McCrary and Monica Morrison at NSF NCAR for extremely helpful conversations about climate model adequacy regarding wildfire projection. And I would like to thank all the members of the audience at a presentation of related work at Stanford University in 2024. And I would like the thank the British Academy for the financial support that made this project possible.

Conflicts of Interest

The author declares no conflicts of interest.

Data Availability Statement

Data sharing not applicable to this article as no datasets were generated or analysed during the current study.

Endnotes

[1] The concept of evaluating climate models as "adequate for purpose" originates from [21, 35]. An early philosophical paper highlighting limitations of climate models for local and short time scales is [22]. Extensive discussion is found in [5]. A recent paper underscoring these issues in development economics is [1]. A comprehensive overview of how climate models exceed their adequacy in related sciences is provided in [36], along with many other papers addressing these deficiencies.

[2]This sometimes relies on a single paper that domain experts would likely consider insufficiently supported or that downplays or overlooks uncertainties.

[3]See (Winsberg, Morrison, and Harvard) for more details on why these downscaled models are often inadequate for such purposes.

[4]Ultimately, after consulting various sources, the LEAP team decided to rely primarily on [13].

[5]See, for example, [37–39] for discussions of institutional scaffolding in social epistemology.

References

1. H. Nissan, L. Goddard, E. C. De Perez, et al., "On the Use and Misuse of Climate Change Projections in International Development," *WIREs Climate Change* 10, no. 3 (2019): e579, https://doi.org/10.1002/wcc.579.

2. S. Harvard and E. Winsberg, "The Epistemic Risk in Representation," *Kennedy Institute of Ethics Journal* 32, no. 1 (2022): 1–31.

3. E. Winsberg and S. Harvard, *Scientific Models and Decision Making* (Cambridge University Press, 2024), https://www.cambridge.org/core/elements/scientific-models-and-decision-making/B7AC2159C941E7D0A08D9981FC8822F1.

4. E. Winsberg, S. Harvard, and J. Symons, "Value Judgments in a Covid-19 Vaccine Model," *Social Science and Medicine* 286 (2021): 1–6.

5. E. Winsberg, *Philosophy and Climate Science* (Cambridge University Press, 2018).

6. A. D. Canning, C. Zammit, and R. G. Death, "The Implications of Climate Change for New Zealand's Freshwater Fish," *Canadian Journal of Fisheries and Aquatic Sciences* 82 (2025): 1–15, https://doi.org/10.1139/cjfas-2024-0127.

7. B. M. Sanderson and R. A. Fisher, "A Fiery Wake-Up Call for Climate Science," *Nature Climate Change* 10, no. 3 (2020): 175–177.

8. S. Kloster and G. Lasslop, "Historical and Future Fire Occurrence (1850 to 2100) Simulated in CMIP5 Earth System Models," *Global and Planetary Change* 150 (2017): 58–69.

9. E. J. Hanan, M. C. Kennedy, J. Ren, M. C. Johnson, and A. M. S. Smith, "Missing Climate Feedbacks in Fire Models: Limitations and Uncertainties in Fuel Loadings and the Role of Decomposition in Fine Fuel Accumulation," *Journal of Advances in Modeling Earth Systems* 14, no. 3 (2022): e2021MS002818.

10. M. C. Kirchmeier-Young, F. W. Zwiers, N. P. Gillett, and A. J. Cannon, "Attributing Extreme Fire Risk in Western Canada to Human Emissions," *Climatic Change* 144 (2017): 365–379.

11. S. M. D'Evelyn, J. Jung, E. Alvarado, et al., "Wildfire, Smoke Exposure, Human Health, and Environmental Justice Need to be Integrated Into Forest Restoration and Management," *Current Environmental Health Reports* 9, no. 3 (2022): 366–385, https://doi.org/10.1007/s40572-022-00355-7.

12. S. Harvard, R. Carter, S. H. Chong, et al., "The 'Values in Modelling' Framework for Patient and Public Involvement in Health Economics Modelling: Development and Application in the LEAP Model Project," (forthcoming).

13. Y. Liu, Y. Liu, F. Joshua, et al., "Projection of Future Wildfire Emissions in Western USA Under Climate Change: Contributions From Changes in Wildfire, Fuel Loading and Fuel Moisture," *International Journal of Wildland Fire* 2022 (2022): 1–13, https://doi.org/10.1071/WF20190.

14. NOAA Climate, Ecosystems, and Fisheries Initiative | NOAA Fisheries. (n.d.), https://www.fisheries.noaa.gov/topic/climate-change/climate,-ecosystems,-and-fisheries.

15. K. A. Kearney, S. J. Bograd, E. Drenkard, et al., "Using Global-Scale Earth System Models for Regional Fisheries Applications," *Frontiers in Marine Science* 8, no. August (2021): 1–27, https://doi.org/10.3389/fmars.2021.622206.

16. North Pacific Fishery Management Council, "Climate Scenarios Workshop Report," September (2024).

17. G. Blöschl, "Changing Climate Both Increases and Decreases European River Flood Discharges: Regional Patterns Over Five Decades," *Scientific Reports*, ahead of print (2019), https://doi.org/10.1038/s41598-019-44073-z.

18. M. P. Clark, "Characterizing Uncertainty of the Hydrologic Impacts of Climate Change," *Current Climate Change Reports* 2 (2016): 55–64, https://doi.org/10.1007/s40641-016-0034-x.

19. J. Elith and J. R. Leathwick, "Species Distribution Models: Ecological Explanation and Prediction Across Space and Time," *Annual Review of Ecology, Evolution, and Systematics* 40 (2009): 677–697, https://doi.org/10.1146/annurev.ecolsys.110308.120159.

20. K. L. Yates, P. J. Bouchet, M. J. Caley, et al., "Outstanding Challenges in the Transferability of Ecological Models," *Ecological Modelling* 33, no. 10 (2018): 790–802.

21. W. S. Parker, "II—Confirmation and Adequacy-For-Purpose in Climate Modelling," *Aristotelian Society Supplementary Volume* 83, no. 1 (2009): 233–249.

22. R. Frigg, L. A. Smith, and D. A. Stainforth, "The Myopia of Imperfect Climate Models: The Case of UKCP09," *Philosophy of Science* 80, no. 5 (2013): 886–897.

23. M. J. Brown, *Science and Moral Imagination: A New Ideal for Values in Science* (University of Pittsburgh Press, 2020).

24. S. John, "Epistemic Trust and the Ethics of Science Communication: Against Transparency, Openness, Sincerity and Honesty," *Social Epistemology* 32, no. 2 (2018): 75–87, https://doi.org/10.1080/02691728.2017.1383533.

25. S. A. Schroeder, "Using Democratic Values in Science: An Objection and (Partial) Response," *Philosophy of Science* 84, no. 5 (2017): 1044–1054, https://doi.org/10.1086/694497.

26. K. Hawley, "Trust, Distrust and Commitment," *Noûs* 48, no. 1 (2014): 1–20, https://doi.org/10.1111/j.1468-0068.2012.00889.x.

27. C. McLeod, "'Trust'. In the Stanford Encyclopedia of Philosophy, Fall 2023," in *Zalta and Uri Nodelman*, ed. N. Edward (Metaphysics Research Lab, Stanford University, 2023), https://plato.stanford.edu/archives/fall2023/entries/trust/.

28. A. Hills, "Responsibility and Trustworthiness," *Ethics* 133, no. 4 (2023): 746–774, https://doi.org/10.1086/723733.

29. S. Harvard, G. Werker, and D. Silva, "Social, Ethical, and Other Value Judgments in Health Economics Modelling," *Social Science and Medicine* 253 (2020): 1–9.

30. E. Winsberg, "Values and Uncertainties in the Predictions of Global Climate Models," *Kennedy Institute of Ethics Journal* 22, no. 2 (2012): 111–137, https://doi.org/10.1353/ken.2012.0008.

31. S. Maji, A. M. Vicedo-Cabrera, Y. Liu, Y. Liu, and M. L. Bell, "Health Impacts of Exposure to Emissions From Prescribed Vegetation Fires Inthe Southeastern United States," *Environmental Research Letters* 17, no. 6 (2022): 064008, https://doi.org/10.1088/1748-9326/ac6c5e.

32. K. Johnson, C. Carlsten, S. Harvard, E. B. Winsberg, and A. Adibi, "The Impact of Time Scale on the Cost-Effectiveness of Air Cleaners for Reducing Exposure to Air Pollution From Wildfires," *American Journal of Respiratory and Critical Care Medicine* 211, no. Abstracts (2025): A5413, https://doi.org/10.1164/ajrccm.2025.211.Abstracts.A5413.

33. C. E. Reid, M. Brauer, F. H. Johnston, M. Jerrett, J. R. Balmes, and C. T. Elliott, "Critical Review of Health Impacts of Wildfire Smoke Exposure," *Environmental Health Perspectives* 129, no. 5 (2021): 055001, https://doi.org/10.1289/EHP8900.

34. M. L. Wilgus and M. Merchant, "Clearing the Air: Understanding the Impact of Wildfire Smoke on Asthma and COPD," *Healthcare (Basel)* 3, no. 12 (2024): 307, https://doi.org/10.3390/healthcare12030307.

35. W. S. Parker, "Model Evaluation: An Adequacy-For-Purpose View," *Philosophy of Science* 87, no. 3 (2020): 457–477.

36. E. Winsberg, M. Morrison, and S. Harvard, "Climate Models for Guiding Action and Informing Cost-Benefit: Adequacy, Inadequacy, and the Ethics of Downstream Use of CMIP Data," (2025), accessed October 7, 2025, https://philsci-archive.pitt.edu/26855/.

37. H. Andersen and N. J. Nersessian, "Conceptual Innovation and the Organization of Interdisciplinary Research," *Philosophy of Science* 86, no. 5 (2019): 1256–1269, https://doi.org/10.1086/705547.

38. K. Rolin, "The Social Epistemology of Research Groups," *Philosophy of Science* 87, no. 5 (2020): 1049–1060, https://doi.org/10.1086/709828.

39. S. Wagenknecht, N. J. Nersessian, and H. Andersen, "Philosophical Perspectives on Social Constraints and Norms in Science," *Synthese* 192, no. 10 (2015): 3229–3248, https://doi.org/10.1007/s11229-015-0730-6.

Biography

Eric B. Winsberg is British Academy Global Professor of History and Philosophy of Science at the University of Cambridge and Professor of Philosophy at the University of South Florida. His Principal research interests are in the philosophy of science, the philosophy of climate science, and the role of models and simulations to guide policy, especially in climate and public health.