RETHINK PRIORITIES

Earth Systems Risks and the Case for Funding Stabilization Research

Global Health and Development Department

Last updated: 2026-June-16 • Report

AUTHORS

Abbie Clare, Ruby Emerson, Britta Jewell, Thomas R Vargas

MANAGER

Abbie Clare

DIRECTOR

John Firth

Clare, A., Emerson, R., Jewell, B., Vargas, T. R., & Firth, J. (2026). *Earth Systems Risks and the Case for Funding Stabilization Research*. Rethink Priorities.

Global Health and Development at Rethink Priorities

The Global Health and Development Department of Rethink Priorities is a multidisciplinary research and consulting group that seeks to disseminate actionable insights and drive positive change in the fields of global health, international development, and climate change.
https://rethinkpriorities.org/global-health-and-development

For all queries, please contact ghd@rethinkpriorities.org.

We thank Outlier Projects for commissioning and funding this research report. The views expressed herein are not necessarily endorsed by Outlier Projects.

© 2026 Rethink Priorities

RETHINK PRIORITIES
530 Divisadero St. PMB #796
San Francisco, California 94117
United States of America



About Rethink Priorities
Rethink Priorities is a research and implementation group that identifies pressing opportunities to make the world better. We act upon these opportunities by developing and implementing strategies, projects, and solutions to key issues. We do this work in close partnership with foundations and impact-focused nonprofits.
www.rethinkpriorities.org

Contents

Editorial Note 4
Executive Summary 5
The Evidence on Climate Is Shifting 5
The Human Harms are Large, Even on Conservative Estimates 5
Mitigation and Carbon Dioxide Removal Alone are Unlikely to Prevent These Harms 7
There are Candidate Approaches Worth Investigating 8
Bottom Line 8
Introduction and Framing 10
Climate Change in the EA Landscape 10
New Evidence on Earth Systems Destabilization 10
The Big Policy Question: Are Mitigation and CDR Enough to Avert Significant Harm? 12
Structure of the Rest of the Report 13
Estimating the Human Harms of Earth Systems Destabilization 13
Using BOTECs to estimate mortality arising from four destabilizing pathways 13
Interventions to Address Earth Systems Destabilization 24
Mitigation and CDR are Unlikely to Keep Warming Below 1.5°C 25
The World is Likely to Spend Extended Periods Above 1.5°C, Increasing the Risk of Earth Systems Destabilization 26
Extended Overshoot may Cause Some Earth Systems to Permanently Destabilize 26
Synthesis 27
Overview and Deep Dives on the Tractability and Neglectedness of Earth Systems Stabilization Approaches 28
Mainstream Arguments Against Earth Stabilization Interventions 30
Stratospheric Aerosol Injection (SAI) 32
Marine Cloud Brightening and Mixed-Phase Cloud Thinning 37
Ice Sheet Stabilization: Basal dewatering and thermosiphons 40
Contributions and Acknowledgments 44

Editorial Note

This report was commissioned by Outlier Projects and produced by the Global Health and Development Department of Rethink Priorities (RP). The findings are based on approximately six weeks of desk research, concluding in June 2026. We revised the report for publication. Outlier Projects does not necessarily endorse the views or conclusions expressed herein, nor do the experts we interviewed or the organizations with which they are affiliated.

In this report, we explore the potential human harms caused by Earth systems destabilization and assess whether there is a case for the effective altruism (EA) community to fund Earth systems stabilization research. We focus on four key destabilization pathways: albedo loss, ice-sheet instability, Atlantic Meridional Overturning Circulation (AMOC) disruption, and warming-induced emissions. Furthermore, we evaluate three candidate interventions—Stratospheric Aerosol Injection (SAI); Marine Cloud Brightening (MCB) and Mixed-Phase Cloud Thinning (MCT); and Ice-Sheet Stabilization—using an importance-tractability-neglectedness (ITN) framework. To inform our analysis, we reviewed the scientific literature and conducted roughly 20 expert interviews.

Throughout this report, we refer to concepts and frameworks specific to the effective altruism community. Notably, we utilize back-of-the-envelope calculations (BOTECs) to estimate the scale of potential human harm and determine philanthropic "willingness to pay" (WTP). These calculations rely on Coefficient Giving's (CG) framework, which we depend on for various relevant inputs and valuations. For example, we apply a 2000x cost-effectiveness bar to evaluate potential grants. As of publication, CG's bar has fallen to 1000x. Under this lower bar, our WTP estimates for averting deaths in this cause area would exactly double, making funding here even more attractive.

We tried to flag major sources of scientific debate and uncertainty throughout the report. The estimates and BOTECs provided are meant to be directionally correct for philanthropic consideration rather than to act as precise climate models, and we are open to revising our views based on new information as the evidence base evolves.

Executive Summary

This report was commissioned by Outlier Projects to assess whether there is a case for the effective altruism (EA) community to increase funding for Earth systems stabilization research. In this context, Earth systems stabilization is defined as direct interventions aimed at keeping critical climate subsystems intact during the many decades that the world is likely to spend above 1.5°C of warming. The EA community has historically deprioritized climate change on the dual reasoning that (a) it does not pose a direct existential risk and (b) it is not financially neglected.

Our research finds that evidence arising since the last Intergovernmental Panel on Climate Change (IPCC) Assessment Report (AR6) materially weakens both premises as applied to Earth systems destabilization, specifically, and that there is a strong case to increase EA funds flowing towards research and governance work on Earth systems stabilization approaches.

The Evidence on Climate Is Shifting

A growing body of evidence, much of it published since the IPCC's Sixth Assessment Report (AR6), suggests that several of Earth's critical systems may be destabilizing faster than mainstream models projected. We surveyed this literature and conducted roughly 20 expert interviews. The post-AR6 evidence across four key pathways (albedo loss, ice-sheet instability, Atlantic Meridional Overturning Circulation [AMOC] disruption, and warming-induced emissions) points in a broadly consistent direction: destabilization appears to be accelerating, and several of the systems involved show signs of self-reinforcing dynamics that could persist even under aggressive emissions cuts. The strength of evidence varies across pathways, and there is legitimate scientific debate on several fronts, but there seems to be a consistent direction of travel: the climate is changing faster, and the consequences are likely arriving sooner and more severely than projected.

Two structural features of this evidence base are particularly significant. First, current climate models do not include representations of the interconnected dynamics of ice-sheets, most known tipping-element interactions, or several major amplifying feedback loops. Second, where evidence is contested, the defining feature of the systems involved is potential irreversibility on human timescales. This combination is precisely the situation in which an expected-value framework, rather than a most-likely-case framework, is appropriate, provided one is transparent about the wide uncertainty ranges involved.

The Human Harms are Large, Even on Conservative Estimates

To estimate the scale of potential human harm, we constructed back-of-the-envelope calculations (BOTECs) for each of the four destabilizing pathways. **It is important to note that these are not climate models.** They are deliberately simple, EA-style calculations designed to ask: Is the plausible scale of harm from Earth systems destabilization large enough to justify a dedicated philanthropic fund in the low-to-mid hundreds of millions of dollars? The four pathways fall into two groups with different calculation methods. For the temperature-dependent pathways (albedo loss and warming-induced emissions), we calculate the philanthropic willingness to pay (WTP) for averting 0.2°C of additional warming: this is the gap between the lowest plausible stabilization temperature in mainstream scenarios (~1.7°C) and the 1.5°C threshold beyond which tipping points become substantially more likely. For the sea

level rise (SLR) and AMOC pathways, we calculate WTP for averting a central-scenario amount of harm.

Our estimates are conservative by design. They count only direct mortality attributable to excess heat, human displacement, and/or food scarcity. They exclude morbidity, deaths from other causes, and the economic and social damages driven by agricultural disruption, ecosystem loss, disasters, and human displacement-driven instability, among other potential impacts. Moreover, we do not account for the compounding effects of interactions between these four pathways: each Earth system is modeled independently, when in reality these systems are nonlinear and highly interactive. Compared to estimates from our interviewees and people working in the Earth systems stabilization space, our numbers are typically lower than what they would predict on the basis of extrapolating current trends. Finally, it is worth noting that the numbers presented below are calculated on the basis of what EA-aligned philanthropists might theoretically be justified in committing to grants that would help advance the science of Earth systems stabilization or the policy environment that makes it feasible. This calculation is different from calculating the cost-effectiveness of actually deploying these approaches at scale.

Overall, the numbers are striking. We present WTP numbers at three levels:

- **Level 1**: **The maximum theoretical spend in EA terms, if each $ was certain to provide highly cost-effective returns**. We compute the total philanthropic value of the lives lost using Coefficient Giving (CG)'s framework (each disability-adjusted life year [DALY] is valued at $100,000), then divide by their cost-effectiveness bar of 2000x (meaning each $1 of grant money must produce at least $2,000 of philanthropic value to be worth funding). The result is the theoretical ceiling: the most a funder with CG-style valuations could spend on a cause area if they were certain that every $ spent would have a $2000 return.
- **Level 2: An estimate of what the EA-aligned philanthropic community as a whole could justify committing, adjusting for chance of success.** Real grants do not guarantee outcomes, so we discount our Level 1 estimate twice. First, we discount by the probability that funded research actually leads to a deployed intervention delivering the target outcome; we set this at 10%, reflecting that the chain from research through governance to deployment is long, even when the underlying science is mature. Second, by the share of credit attributable to EA-aligned funders rather than to governments or mainstream climate philanthropy; we set this at 50%, anchored on the observation that the space is currently sparsely funded (field-wide spend is under $10M/year for the relevant interventions), so EA-aligned giving in the low hundreds of millions over the next five years would plausibly account for at least half of relevant philanthropic spending during this initial window. Combined, these two discounts create a Level 2 estimate that is 5% (50% × 10%) of Level 1.
- **Level 3: Calculates what one EA-aligned funder could justify committing as a fund or portfolio of grants.** It is quite rare that the estimated cost-effective donations within a given cause area will be picked up by a single donor. We therefore apply a further 50% discount to capture the share of EA-aligned activity attributable to a single large funder, given that we expect a relatively small number of EA-aligned funders to be active in this space. This ultimately leads to the Level 3 being 2.5% of Level 1 WTP figures (10% × 50% × 50%).

Table 1 EA summarizes the results of our WTP calculation at each level for the four pathways, expressed in millions of USD. Per-pathway, the EA community WTP ranges from $263 million to $445 million and single-funder WTPs from $131 million to $223 million. Despite only valuing mortality impacts for each pathway, the analysis supports a fund in the low-to-mid hundreds of millions of dollars.

Table 1: Pathway-specific WTP estimates at Levels 1-3, expressed in millions of USD

	Albedo Loss	Warming-Induced Emissions	Sea Level Rise	AMOC Collapse
Central Damage Pathway	0.2°C of averted warming	0.2°C of averted warming	1.5 m of sea level rise by 2100	15% probability of AMOC collapse by 2100
Level 1 Theoretical maximum WTP	6,160	6,160	8,900	5,250
Level 2 EA community WTP	308	308	445	263
Level 3 Single EA-donor WTP	154	154	223	131

Mitigation and Carbon Dioxide Removal Alone are Unlikely to Prevent These Harms

These figures become more urgent when set against the realistic trajectory of global emissions policy. Mitigation and carbon dioxide removal (CDR), while essential, are unlikely on their own to keep warming below the 1.5–2.0°C band in which several Earth systems become at risk of destabilization. Even the International Energy Agency (IEA)'s most optimistic Net-Zero-by-2050 scenario peaks around 1.65°C and only returns to below 1.5°C late in the century. Moreover, this scenario is dependent on widespread CDR deployment, which the IEA itself describes as unproven at scale. Even full implementation of all current nationally determined contributions (NDCs) would result in warming of 2.3–2.5°C by 2100 (UNEP Emissions Gap Report 2025), confirming that the gap between pledges and the 1.5°C target is structural, not merely a matter of implementation speed. The world will therefore plausibly spend many decades, possibly the rest of this century, within the temperature range where multiple destabilizing processes can occur.

It is this logic, rather than pessimism about decarbonization, that motivates the question at the heart of this report: should we invest in researching complementary approaches that could

help keep critical Earth systems intact while the longer-term impacts of mitigation and CDR take effect?

There are Candidate Approaches Worth Investigating

We assessed three candidate Earth systems stabilization interventions under an importance-tractability-neglectedness (ITN) framework. All three show some promise, but significant research gaps remain across the board. The case for funding does not rest on any single intervention proving viable. Instead, it rests on the value of resolving the underlying scientific and governance questions to understand if there are genuine options available.

- *Stratospheric Aerosol Injection (SAI).* This approach is the highest-leverage and most near-term feasible of the solar radiation modification (SRM) options: volcanic-eruption analog and modeling work converge on the conclusion that ~8–16 million metric tons of sulfur dioxide (SO_2) injected annually could deliver ~1°C of cooling. The physical science is reasonably mature; however, the binding constraints are governance, equity, and political legitimacy, all of which receive far less funding than the natural-science work. Termination shock and unintended regional effects (monsoon disruption, ozone depletion) remain real risks, but these must be weighed against the counterfactual of continued warming without intervention. We see philanthropic dollars on the social-science, political legitimacy, or governance sides as having particularly high marginal value.
- *Marine Cloud Brightening (MCB) and Mixed-Phase Cloud Thinning (MCT).* Lower-leverage globally than SAI, but potentially well-suited as bridging technologies for regional protection (Arctic sea ice, ice shelves, coral reefs). Real-world proof-of-concept exists for MCB via shipping-emission cloud tracks. Both are very small fields (~$5–10M/year). MCT's evidence base for climate-relevant effects rests almost entirely on a single modeling study, and the field currently lacks the research infrastructure to absorb large amounts of funding without parallel capacity-building. Modest, well-targeted philanthropic investments could nonetheless help build this capacity and accelerate research timelines.
- *Ice-sheet stabilization (basal dewatering, thermosiphons).* These approaches aim to slow West Antarctic Ice Sheet (WAIS) retreat by raising basal friction. The intervention space is highly neglected: cumulative philanthropic deployment is under $20 million. The intervention's effect at scale is unproven, and no published evidence exists on the amount of sea level rise it could avert. However, the cost of producing the evidence to determine feasibility (estimated at $100–120M over 10 years) is modest relative to the potential consequences. Such research could either validate the approach or redirect resources toward more promising alternatives.

Bottom Line

The human harms from Earth systems destabilization are plausibly extraordinarily large. Even our deliberately conservative analysis, which counts only direct mortality, supports a single-actor philanthropic program budget of $131 million to $223 million per pathway at CG's cost-effectiveness bar. The true figure is likely substantially higher once broader mortality, morbidity, economic damages, and second-, third-, and higher-order effects are included—for example, climate-driven migration and the conflict it can drive, agricultural and food-system disruption, infrastructure damages, etc. At the same time, the current portfolio of mitigation

and CDR responses may not be able to protect at-risk Earth systems from the temperature overshoot that is now virtually certain. Meanwhile, the candidate stabilization interventions, while showing some promise, face unanswered questions about feasibility, safety, and governance that relatively modest research investments may help resolve.

Our main takeaway is that an EA-aligned fund focused on Earth systems stabilization research and governance is justified based on existing evidence.

Introduction and Framing

Climate Change in the EA Landscape

The effective altruism (EA) community has historically deprioritized climate change as a top-tier cause area. There are broadly two reasons for this. The first is a judgment that climate change is unlikely to pose a direct existential risk to humanity, with prominent EA organizations classifying it as less pressing than other priorities such as AI safety or engineered pandemics (Ord, 2020; Hilton, 2022). The second is that climate change is not financially neglected: global climate funding from public and private sources reached approximately $1.3 trillion (USD) annually as of 2021–2022 (Buchner et al., 2023; UNFCCC, 2024). However, there is increasing cause to reconsider both these assumptions as the evidence base is changing.

The conventional framing of climate risk centers on temperature. Under this framing, the question is: at what level of average global warming do dangerous impacts occur? The IPCC's sixth assessment report (AR6) identified 1.5°C above pre-industrial levels as a critical threshold, beyond which the risks of triggering climate tipping points rise substantially (IPCC, 2023; McKay et al., 2022). This framing has been valuable and remains important.

However, a growing body of evidence, much of it published since AR6, suggests that several of the Earth's critical systems may be destabilizing faster than these assessments anticipated and, moreover, that these Earth systems may continue to destabilize even in the event of rapid decarbonization. This is what makes this potential cause area distinct from the typical climate challenges; we do not know enough about potential nonlinear impacts and interactions to know how and when to intervene beyond mitigation measures.

The purpose of this report is to assess the evidence and investment case for these claims from a philanthropic perspective. **We ask: Is there a case for an EA-aligned organization to set up a fund focused on Earth systems stabilization technologies? To preview our answer, we think that there is sufficient evidence to do so.**

New Evidence on Earth Systems Destabilization

The Earth's climate is not a single system that responds uniformly to changes in global average temperature. It is a collection of interconnected but distinct Earth systems: ice sheets, ocean circulation patterns, atmospheric reflectivity, and terrestrial carbon stores, each of which could be pushed into new, self-reinforcing states by the cascading effects of human emissions. Recent peer-reviewed research suggests that several of these systems are showing signs of destabilization at lower temperatures and on faster timescales than previously expected (McKay et al., 2022; McKay, 2024; Wunderling et al., 2024). For example, the planet is now accumulating heat at nearly twice the rate climate models had predicted (Mauritsen et al., 2025), and new observational studies published this month provide evidence that the Atlantic Meridional Overturning Circulation (AMOC) is weakening faster than the average of current model projections (Portmann et al., 2026).

There are also reasons to believe that current climate models may not fully capture the scale of these risks. Models largely do not represent the interactions between destabilizing Earth systems. For example, the current generation of climate models used in IPCC assessments

(known as CMIP6) does not include coupled dynamic ice sheets,[1] rendering them unable to represent the tipping dynamics of these elements or their links to other components of the climate system.

More broadly, Earth system models continue to disagree on how tipping elements behave and interact, creating significant uncertainty about cascading risks (Communications Earth & Environment, 2025). A comprehensive review of tipping point interactions found that the majority of known linkages between tipping elements are destabilizing in nature, meaning that the crossing of one threshold can lower the thresholds at which others are crossed (Wunderling et al., 2024; McKay et al., 2022). Beyond interactions, a major assessment identified 27 amplifying feedback loops relevant to climate change, several of which are not fully incorporated into current models (Ripple et al., 2023). The significance of these limitations is increasingly recognized: for the first time, the IPCC's Working Group I contribution to AR7 will include a dedicated chapter on tipping points and high-impact events (IPCC, 2025), and the IPCC and WCRP co-convened a workshop (IPCC, 2025) in November 2025 to identify knowledge gaps in this area.

While the studies and interviews cited above point toward accelerating risks, there is legitimate scientific debate on several fronts. Some researchers have found no significant decline in the AMOC over the past 60 years, based on air-sea heat flux data (Terhaar et al., 2025). Moreover, while record-low planetary albedo has been identified as a driver of the 2023 temperature surge, the causes of declining low-cloud cover also remain under active investigation (Goessling et al., 2024), with recent work suggesting that these changes may be consistent with established expectations for cloud feedback rather than an unexpected shift (Zelinka et al., 2026). Finally, the thresholds at which specific tipping points may be crossed carry wide uncertainty ranges (McKay et al., 2022), and scientists disagree about whether some of the more alarming projections are well-constrained enough to inform policy.

We engage with this debate throughout the report, and we do not treat contested findings as established fact. However, uncertainty about the precise probability of these outcomes does not resolve the question of how seriously they should be taken. The defining features of risks like ice-sheet collapse, AMOC disruption, and runaway albedo loss is that they are (1) potentially irreversible, and (2) capable of causing harms of high magnitude. In contexts where the severity of outcomes is very high but probabilities are genuinely uncertain, an expected value approach is appropriate: one that asks not only how likely this is, but what the product of likelihood and magnitude is, and whether that changes how we should act.

This is the analytical lens we adopt throughout this report. We are not climate scientists, and this report reflects approximately six weeks of desk research supplemented by expert interviews, rather than years of specialist expertise. Where evidence is contested, we are transparent about the range of estimates. Where probabilities are uncertain, we test our conclusions across a range of assumptions. Our aim is not to resolve any scientific debate in this area, but to assess whether the ratio of risk to potential magnitude of harm is sufficient to justify a reconsideration of how the philanthropic community views Earth systems stabilization in terms of cost-effective giving.

[1] In plain language, this means that ice sheets are treated as largely static, meaning they cannot simulate feedback loops between ice loss, ocean circulation, and regional climate changes.

The Big Policy Question: Are Mitigation and CDR Enough to Avert Significant Harm?

A growing number of climate scientists and policy practitioners have argued that emissions mitigation and carbon dioxide removal (CDR), while essential, may not on their own be sufficient to avert the most severe human harms from climate change, particularly those arising from Earth system feedbacks that may persist even under deep emissions cuts (Global Tipping Points, 2023; Steffen et al., 2018).[2]

This concern is reinforced by the growing body of evidence on the "CDR gap": national plans for carbon dioxide removal fall substantially short of the levels required in scenarios that limit warming to 1.5°C (Lamb et al., 2024), a target that recent assessments increasingly suggest may no longer be achievable without overshoot (Geden et al., 2024). Most novel CDR approaches remain at early stages of deployment, with current capacity orders of magnitude below what modeled pathways require (ibid; IPCC, 2022).

With this in mind, several major scientific institutions have called for serious research and investment into Earth systems stabilization interventions, meaning the direct management and re-stabilization of at-risk Earth systems through approaches such as stratospheric aerosol injection (SAI), marine cloud brightening (MCB), and glacier stabilization. The US National Academies of Sciences, Engineering, and Medicine (NASEM) concluded that "a strategic investment in research is needed to enhance policymakers' understanding of climate response options," recommending a transdisciplinary research program of $100–$200 million over five years (NASEM, 2021, p. 17). A Nature editorial the same year argued that the risks and benefits of solar geoengineering approaches needed to be evaluated, urging the scientific community to give such research "a chance" (Nature, 2021). On glacier stabilization specifically, Moore et al. (2018) argued that targeted engineering interventions to slow ice-sheet collapse should be seriously explored as a complement to emissions reduction, and a 2024 white paper from the University of Chicago's Climate Systems Engineering initiative set out a research vision for glacial climate intervention, arguing that determining the feasibility of engineered interventions to reduce sea level rise is now urgent (MacAyeal et al., 2024).

Other scientists and institutions strongly dispute this position, arguing that the risks of such interventions are poorly understood, that they may create moral hazard by reducing pressure for emissions cuts, and that governance challenges are prohibitive. Over 600 scholars from 65 countries have signed an open letter (Solar Geoengineering Non-Use Agreement, 2022) calling for an international non-use agreement on solar geoengineering (Biermann et al., 2022); the

[2] Various experts we talked to underscored the idea that several Earth systems are already exhibiting changes that are largely independent of atmospheric CO_2 concentrations, driven instead by internal feedbacks that will persist regardless of future emissions trajectories. Take the West Antarctic Ice Sheet (and Thwaites Glacier in particular): Brent Minchew (Arête Glacier Initiative) noted that the region's instability is now "largely decoupled . . . from atmospheric CO_2 concentrations and warming," adding that these "internally driven instability mechanisms . . . will continue no matter what we do with climate afterwards. Nothing short of another ice age can stop it." Looking at Arctic summer sea ice, Charlotte DeWald (Arctic Stabilization Initiative, 2025) warned that "we will likely see the first ice-free Arctic summer months within 10 years under any baseline warming scenario." Loss of summer sea ice would push regional temperatures up an additional 1.5°C, representing a risk that "persists even if we were to decarbonize instantly." In another example, Marc Macias Fauria (Scott Polar Research Institute) highlighted the irreversible nature of permafrost thaw. Even if humanity could theoretically reverse global warming and cool the permafrost, he explained, "all the matter that had been accessed by soil organisms would already have decomposed and entered the atmosphere, and it would not return to the soil carbon pool for millennia."

Convention on Biological Diversity maintains a de facto moratorium on geoengineering deployment (UNEP, 2010); and African nations have also called for non-use (Biermann & Gupta, 2024).

Structure of the Rest of the Report

To remind the reader, our central question is whether the emerging evidence on Earth systems destabilization, when assessed rigorously, justifies philanthropic investment in Earth systems stabilization research and governance. We do not begin with a predetermined answer. Instead, we work through the following logical steps, each within a corresponding report section:

1. Estimating the human harms of Earth systems destabilization: We assess the current state of evidence on Earth systems destabilization across four key pathways: albedo loss, ice-sheet instability, AMOC disruption, and emissions-induced warming from natural feedbacks. We produce back-of-the-envelope calculations (BOTECs) to estimate the potential scale of associated human harms for each. **These are not climate models,** merely relatively simple EA-style calculations that we think are directionally correct within an order of magnitude of interest for EA-aligned funders.
2. We examine whether the current portfolio of climate responses, centered on mitigation and CDR, is likely to be sufficient to address these harms.
3. We evaluate three case study Earth systems stabilization approaches against an importance–tractability-neglectedness (ITN) framework, considering potential to reduce harm, feasibility of conducting further research, and the risks they themselves may introduce.

Estimating the Human Harms of Earth Systems Destabilization

Using BOTECs to estimate mortality arising from four destabilizing pathways

This section examines four of Earth's natural systems that are at risk of destabilization as a result of increased annual warming: albedo loss, sea level rise (SLR), AMOC disruption, and emissions-induced warming from natural feedback loops. We provide an overview of the evidence and emerging debates drawn from published literature and our expert interviews, and a set of BOTECs to estimate the human-level impacts of triggering these thresholds.[3]

Our BOTECs fall into two groups, structured by the underlying calculations:

- *Group 1* (albedo loss and warming-induced emissions) uses a temperature-dependent procedure because both pathways drive harm primarily through additional global

[3] A brief discussion of these BOTECs is in order to clarify their intended use and limitations. As non-climate scientists modeling phenomena that experts themselves are currently debating, we know our models are likely flawed in important ways. However, BOTECs of this sort (flaws and all) are common in the EA space to highlight the underlying assumptions required for an outcome to be true. The primary question here is whether our rough evaluation of damages justifies EA-aligned philanthropic investment in the millions of dollars. Because our spreadsheets are public, others can freely critique how we calculate mortality. Regardless, our exercise does not need to be scientifically "true" to demonstrate that such an investment is warranted, so long as our inputs are reasonable and our mortality numbers are plausible. Against guesses from our client and people working in the Earth systems stabilization space—some of whom disagreed in principle with our methods out of concern they were simplistic or based on limited evidence—we found that our estimates produced typically lower damages than they would predict simply on the basis of extrapolating current trajectories. In short, we think people who have thought harder about this would find our numbers are well within an order of magnitude from their own views.

warming. We calculate the willingness to pay (WTP) for averting 0.2°C of additional warming, the gap between the lowest plausible stabilization temperature in mainstream scenarios (~1.7°C) and the 1.5°C threshold beyond which tipping points become substantially more likely.
- *Group 2* (sea level rise and AMOC collapse) uses a different, bottom-up procedure because these pathways generate harm through different channels, namely displacement and food insecurity, that are not well captured by a warming-averted framing. For these, we estimate impact under a central guess and multiply through Coefficient Giving's WTP framework (below).

Table 2 below previews each pathway's WTP at three levels. Level 1 is the WTP at certainty: the theoretical ceiling, calculated by taking the total philanthropic value of averting the harm (using Coefficient Giving's framework, which values a DALY at $100,000) and dividing by our 2000x cost-effectiveness bar. Level 2 is the EA-aligned community WTP: Level 1 further discounted by the probability that funded research produces a deployed intervention delivering the target outcome (10%) and by the share of credit attributable to EA-aligned funders rather than to governments or mainstream climate philanthropy (50%). Level 3 is the single-funder WTP: Level 2 is further discounted by the share of EA-aligned activity attributable to one funder (50%), reflecting the small number of likely EA-aligned funders in this space. The framing section below walks through these parameters and their values in detail; each pathway's full description and BOTEC follow.

Table 2: Pathway-specific WTP estimates at Levels 1–3, expressed in millions of USD

	Albedo loss	Warming-Induced Emissions	Sea Level Rise	AMOC Collapse
Central damage pathway	0.2°C of averted warming	0.2°C of averted warming	1.5 m of sea level rise by 2100	15% probability of AMOC collapse by 2100
Level 1 Theoretical maximum WTP	6,160	6,160	8,900	5,250
Level 2 EA community WTP	308	308	445	263
Level 3 Single EA-donor WTP	154	154	223	131

All numbers in Table 2 should be read as best guesses given parameter inputs, not point estimates. Under many scenarios, the true WTPs will be substantially higher. For example, our BOTECs show that a scenario in which sea level rises by 2 meters justifies up to almost $2 billion, and a 40% probability of AMOC collapse justifies up to a similar number.

Procedures

We structure these BOTECs to project cumulative outcomes through the year 2100. This timeline accounts for the reality that many tipping points and feedback loops will plausibly impact us in the near term, while acknowledging that mapping human adaptation becomes increasingly speculative beyond the end of the century.

Two limitations of our BOTECs merit attention:

- *System interactions:* Our BOTECs attempt to independently model Earth system elements that are inherently nonlinear and highly interactive (see, for example, Dietz et al., 2021).[4] We do not have a principled way of combining them, so we consider them to be simply additive.[5]
- *Excluded impacts:* By only including direct human deaths from excess heat, we exclude deaths from other causes (e.g., disasters, agricultural system disruption), morbidity impacts, and other severe socioeconomic impacts recognized by the World Health Organization (WHO, 2023). We believe, from previous work, that our approach is likely to understate the real effects by at least double.

The combined effect of these limitations is that our damage estimates are likely to be more conservative than reality. Ultimately, we view our estimates as a defensible but conservative floor. While we ground our parameters in published literature and inputs from domain experts to the greatest extent possible, we acknowledge that emerging views, which will likely be reflected in the upcoming AR7, may shift these assumptions.

When estimating the cost-effectiveness of a grant in this area from the perspective of EA-aligned funders, we are specifically talking about philanthropic funding for grants that would help advance the science around Earth systems stabilization. This is not the same as the cost-effectiveness of actually deploying these approaches at scale. In the case of exploring the cost-effectiveness of these research grants, we calculate the Social Return on Investment (SROI) as Impact/Cost. The maximum Cost that still clears our 2000x cost-effectiveness bar is the willingness to pay (WTP) for the outcome. Because Impact depends on how much philanthropic value we actually capture in practice, we present WTP at three levels, working down from the theoretical ceiling. Level 1 is the WTP at certainty: the total CG (Coefficient Giving) philanthropic value of averting the harm, divided by the 2000x cost-effectiveness bar—the most a philanthropic actor with CG valuations on life and income could spend at the bar on a guaranteed outcome.

To determine this total philanthropic value, we borrow CG's framework. Notes on the larger framework are here (Favaloro & Berger, 2021), but for our purposes, it suffices to know that CG values a disability-adjusted life year (DALY) averted at $100,000, and that it does not discount DALYs over time or space. Therefore, the total possible philanthropic value is straightforward

[4] See, for example, Dietz et al. (2021).
[5] It seems to us plausible that they would combine to multiply their effects, but we are not certain by how much.

to calculate: expected deaths x 41 DALYs per death[6] x $100,000. Dividing by 2000x gets you the Level 1 WTP.

Level 2 is the WTP for the EA-aligned philanthropic community as a whole. We apply two discounts to Level 1. The first is the probability that a funded grant (e.g., to do more relevant research) actually produces or speeds up a deployed intervention delivering the target outcome. We set this at 10% (low-to-medium confidence; defensibly anywhere in a 4–20% range), reflecting both that some interventions like SAI are technically feasible and that the causal chain from research through governance to deployment is long. The second is the share of credit attributable to EA-aligned funders rather than to governments or mainstream climate philanthropy at this initial stage of research and policymaking. We set this at 50% (medium confidence), anchored on the fact that field-wide philanthropic spend on the relevant work is currently very small (~$5-10M/yr for MCB/MCT; under $20M cumulative for ice-sheet stabilization), so EA-aligned giving of low hundreds of millions over the next five years would plausibly account for roughly half of the dollars (and at least half of the credit) for progress on the research and policy fronts during this window. Combined, Level 2 receives a discount of 5% (10% × 50%).

Level 3 is the WTP that a single EA-aligned funder could justify committing to this space, deployed as a fund or portfolio of grants (individual grants plausibly in the single-digit millions). We further apply the share of EA-aligned activity attributable to one funder, set at 50%, reflecting the small number of likely EA-aligned funders in this space. The combined Level 3 discount is therefore 2.5% (10% × 50% × 50%). The EA community figure (Level 2) is our recommended primary anchor; it is what the fund is being designed against. The single-funder figure speaks to the SROI calculation that any one funder would run before committing, and would entail their own guesses and preferences around the discounts we highlight here. We spent the majority of our time conducting BOTECs around mortality, as we considered this to be the most important and certain input, while acknowledging that there are surely damages beyond it (e.g., infrastructural and economic losses). Our mortality numbers are our anchors.

We have high confidence in our Group 1 BOTECs at the certainty (Level 1) layer, where the calculation reduces to deaths per °C × DALYs per death × CG value per DALY; the underlying warming-to-mortality estimates are also in line with what some published work suggests is likely under these scenarios (Rennert et al., 2022; Bressler et al., 2021). We have medium confidence in our Group 2 BOTECs at the certainty layer, because of the length of the causal chain from displacement or food insecurity to mortality. Confidence in the Level 2 and Level 3 WTPs is lower and depends primarily on the two named discounts: p(success) (low-to-medium confidence) and EA credit share (medium confidence). Changing the information or assumptions regarding any of these inputs could significantly revise our estimates in either direction.

[6] Because CG weights DALYs per death by age, our average DALY estimate could be considered too high if we assume heat-related deaths primarily burden older populations. However, we retained our original DALYs per death figure after doing a quick Google search and finding emerging evidence that these fatalities are more widely distributed across age groups.

Group 1: Temperature-Dependent Pathways

Albedo Feedback

Albedo is the proportion of sunlight (solar radiation) reflected by a surface or object, often expressed as a percentage (IPCC, 2020). Clouds, snow, and ice usually have high albedo, whereas the ocean has low albedo (ibid). For the purposes of this report, albedo loss is defined as the *decline* in Earth's overall planetary reflectivity. This loss can be driven by the disappearance of bright surfaces (such as sea ice, glaciers, and low marine clouds) and their replacement by darker ones (like open water or bare ground). As the Earth reflects less solar radiation, it absorbs more energy, creating a feedback loop that directly reinforces global warming.

Earth's albedo is declining, acting as an amplifier of global warming. Multiple independent observational methods, ranging from two decades of "earthshine"[7] lunar measurements (Goode et al., 2021) to CERES satellite records (Loeb et al., 2021), confirm this loss of reflectivity, though the magnitudes differ across studies. CERES data indicate an increase in absorbed solar energy of roughly 1.7 W/m² since 2010 (Hansen et al., 2025), with more recent updates indicating the a 2.2 W/m² increase in absorbed solar energy between 2000 and 2005.[8] Hansen et al. (2025) characterize this additional heating as a combination of climate forcing and climate feedback,[9] noting that it is the equivalent of adding ~138 ppm of CO_2 to the atmosphere. We note that Phil Duffy, Chief Scientist at Spark Climate Solutions, has warned that this number is likely a typo, and should actually be 158 ppm.

The Earth's dimming appears to be primarily driven by changes in cloud cover (Loeb et al., 2021; Loeb et al., 2024), with additional contributions from reduced surface albedo as sea ice and snow cover retreat. The scale of this effect became visible in 2023, when global temperatures exceeded predictions even after accounting for greenhouse gases and El Niño (Schmidt, 2024). Goessling et al. (2024) attributed approximately 0.23°C of the 2023 anomaly directly to record-low planetary albedo driven by missing low clouds, largely accounting for this gap. However, the relative contributions of cloud feedback and reduced aerosols (e.g., post-2020 shipping-fuel sulfur cuts) to this additional warming remain active research questions (Goessling et al., 2024; Yuan et al., 2024; Zhang et al., 2025; Hansen et al., 2025; Watson-Parris et al., 2025).

Some aspects of declining reflectivity are self-reinforcing both globally and at regional scales, including Arctic-specific feedbacks documented by Chapin et al. (2005) (Hansen & Kharecha 2025; Forster et al., 2021).[10] Dakota Gruener, the CEO of the nonprofit climate initiative Reflective, summarizes this compounding threat, noting that satellite evidence shows the Earth "is just becoming less reflective . . . which means that it is increasing the amount of absorbed radiation. It's compounding on both sides of the ledger in scary ways." As the planet warms,

[7] The faint glow visible on the dark side of the moon.
[8] Phil Duffy cites the observed trend of 0.87 W/m² per decade from slide 6 of a 2025 AGU presentation by Norman Loeb at NASA Langley Research Center. Extrapolating out to 25 years, he calculates that 0.87 × 2.5 = 2.2 W/m².
[9] Strictly, not all of the albedo decline is feedback. The cryosphere and cloud components are self-reinforcing, but a portion of recent dimming reflects a one-time reduction in reflective sulphate aerosols, which is a forcing change rather than a feedback. The relative contributions remain contested.
[10] From the IPCC report: "For example, Arctic sea ice has been melting as the Earth warms, reducing the amount of reflected sunlight and adding to the initial warming (an amplifying feedback)." (Chapter 7, FAQ 7.1).

highly reflective surfaces (e.g., ice and snow) disappear, exposing darker surfaces (such as ocean, soil, and vegetation). Marc Macias-Fauria, a professor of physical geography at Cambridge University, highlights how this manifests on land: "as tall vegetation encroaches previously open tundra, you're darkening the ground, you're sticking out of the snow and even with snow on the ground, you're darkening the overall surface of the Arctic, which means you're warming it by enabling more of the sun's energy to get into the system and further warm it." These darker surfaces absorb more solar energy, which further heats the atmosphere, causing even more snow and ice melting and soil thawing. Simultaneously, warming drives cloud changes that further accelerate warming, compounding the effect (Hansen & Kharecha, 2025). This feedback is one of the key drivers of Arctic amplification; the phenomenon whereby the Arctic has warmed nearly four times faster than the global average since 1979 (Rantanen et al., 2022).[11] One expert we interviewed describes the severity of this localized acceleration: "the Arctic is really the canary in the coal mine of climate change: it's warming four times faster [and] sea ice is rapidly declining, creating the albedo feedback loop of exposed ocean."

Warming-Induced Emissions

Warming-induced emissions refer to greenhouse gases (primarily methane and carbon dioxide) released from natural ecosystems as a direct response to human-caused climate change. These emissions act as self-reinforcing, positive climate feedback loops: as anthropogenic greenhouse gases warm the planet, natural reservoirs unlock and release additional carbon, which in turn accelerates further global warming. The primary sources (Arctic permafrost thaw, tropical wetlands, forest dieback, and retreating ice sheets) are not hypothetical future risks, but ongoing and measurable phenomena:

- **Permafrost thaw** is already well underway. The northern permafrost zone holds ~1.5 trillion metric tons of organic carbon, roughly double what is currently in the atmosphere (Hugelius et al., 2014; Schuur et al., 2015). Both gradual thaw and an abrupt "thermokarst" collapse, where ice-rich ground gives way within years, expose previously-frozen organic matter to decomposition, releasing carbon dioxide and methane to the atmosphere. Marc Macias-Fauria highlighted the particular severity of nonlinear thawing: "What used to be 30 cm of soil that would thaw in a particular year ends up being several meters, just because large bodies of ground ice have melted."
- **Wildfires in the Arctic and boreal regions** are also already occurring and accelerating. For example, the frequency of extreme wildfires across boreal forests increased sevenfold between 2003 and 2023 (Cunningham et al., 2024), largely driven by hotter, drier conditions and warmer temperatures overnight.
- **Tropical wetlands** are highly sensitive to increased temperature and shifts in rainfall, and are already activating as a major feedback loop, significantly boosting methane output (Zhang et al., 2023). Phil Duffy emphasized that emissions from wetlands are "probably one of the main drivers of the recent increases in atmospheric methane," and they are already the world's largest source of global natural methane emissions (Saunois et al., 2020).
- **Ice-sheet retreat** in Greenland and Antarctica releases methane that was previously trapped in subglacial sediments and causes sea level rise (Hatton et al., 2026). Additionally, the loss of reflective white ice exposes dark ocean water, which absorbs

[11] While the ice-albedo feedback is a major contributor, particularly in summer months, modeling work suggests that temperature feedbacks (specifically the lapse rate feedback) may play an even larger role in the overall amplification (Pithan & Mauritsen, 2014).

more heat (the ice-albedo feedback loop), further accelerating warming (Kashiwase et al., 2017).

The human and climate consequences across these feedback streams are interconnected and potentially severe. Thawing permafrost destabilizes infrastructure and threatens to release pathogens into local water and food supplies. Amazon forest dieback would disrupt regional rainfall patterns, devastating agriculture and water access for millions. Melting ice sheets are the primary driver of long-term SLR, threatening coastal communities and low-lying island nations with displacement. At the climate level, each of these feedback loops compounds the others: more warming unlocks more carbon, which drives more warming. Unmitigated, gradual permafrost thaw alone could release an additional ~150 billion tons of carbon by 2100 (Natali et al., 2021).

Though warming-induced emissions are projected to be on "roughly similar scale of our largest anthropogenic emissions sectors today, and could prove to be more sensitive and have higher emissions" by the end of the century, according to Erika Reinhardt at Spark Climate Solutions, these feedback mechanisms are often only partially included, if at all, in major climate models and remaining carbon budgets. This gap means current estimates of how close the world is to breaching 1.5°C or 2°C of warming may be overly optimistic, and that feedbacks under climate overshoot are increasingly likely. Nonetheless, limiting warming to 1.5°C rather than 2°C could, for instance, avoid the degradation of roughly 2 million square kilometers of permafrost (Chadburn et al., 2017), which could represent the difference between a manageable greenhouse gas feedback loop and a runaway one.

Averting 0.2°C of Warming Due to Albedo Loss or Warming-Induced Emissions Implies an EA Community WTP of $308M and a Single-Funder WTP of $154M

Both albedo loss and warming-induced emissions drive harm primarily through additional global warming and the heat-related mortality it causes, so we use a single shared calculation for both pathways. The procedure asks: what would the philanthropic community be willing to pay to avert 0.2°C of additional warming, regardless of source? We use 0.2°C because it is the gap between the lowest plausible stabilization temperature in mainstream scenarios (~1.7°C) and the 1.5°C threshold beyond which Earth system tipping points become substantially more likely. The two pathways differ in the plausibility of attributing that much warming to each individually and in the interventions available, which we treat separately in each pathway's importance and tractability sections.

We anchor on the total CG philanthropic value of averting 1°C of warming, which we calculate to be $61.62 trillion. This figure is the product of three terms: the marginal cumulative excess deaths to the year 2100 from each additional degree of warming layered on top of the SSP2-4.5 baseline trajectory (15,028,339 deaths/°C); the DALYs lost per death (41); and the CG value per DALY ($100,000). Scaling linearly, the CG philanthropic value of averting 0.2°C is $12.32 trillion.

Dividing the $12.32 trillion CG value of averting 0.2°C by the 2000x cost-effectiveness bar gives a WTP at certainty of ~$6.16 billion. Applying the discounts noted about (community discounts — p(success) = 10% × EA credit share = 50%) yields an EA community WTP of $308 million. Applying the further within-EA share of 50% yields a single-funder WTP of $154 million. Halving either discount factor halves the corresponding WTP.

These figures tell us the total philanthropic spend justified to clear the 2000x bar for an intervention that averts 0.2°C. The same parameters apply to either albedo loss or warming-induced emissions; the pathways differ in the plausibility of attributing 0.2°C of averted warming to each individually.

Group 2: Independent Damages Pathways

Ice Sheet Instability and Glacier Dynamics

The Greenland and West Antarctic ice sheets, together with the world's mountain glaciers, are losing mass at rates that track the upper end of model projections (Slater et al., 2020). Indeed, some experts we talked with suggested that they are far exceeding predictions.[12] Glaciers alone lost 267 ± 16 gigatons per year between 2000 and 2019, accounting for 21% of observed sea level rise, with the rate of loss accelerating by 48 gigatons per year per decade (Hugonnet et al., 2021). Greenland's contribution to 2100 SLR is large but relatively well-constrained: even if temperatures hold at their 2000 to 2019 average, the ice sheet will likely still lose enough mass to raise global sea levels by at least 274 ± 68 mm before reaching a new equilibrium (Box et al., 2022; see also Aschwanden & Brinkerhoff, 2022; Rahlves et al., 2025), though the timeline of this entire rise will likely span hundreds of years.[13] By contrast, the wide uncertainty around 2100 SLR comes overwhelmingly from West Antarctica, which we turn to next.

In West Antarctica, the picture may be even more concerning. The West Antarctic Ice Sheet (WAIS) holds enough land ice to raise global sea levels by roughly 3.2 meters (Bamber et al., 2009), but it relies on floating ice shelves to hold it back from slipping into the ocean. Ocean-driven melting of this crucial system is projected to accelerate at triple the historical rate over this century, which appears unavoidable even under the most ambitious Paris Agreement targets (Naughten et al., 2023). As these shelves weaken, the ice behind them flows into the ocean much more quickly.

The outstanding question is no longer whether WAIS will lose significant ice, but how fast. Marine ice sheet instability (MISI), where retreat into deeper bedrock becomes self-reinforcing, is a well-established mechanism for rapid loss, and present-day mass loss rates from the Thwaites and Pine Island glaciers are already consistent with the early stages of collapse. Modeling suggests this retreat would continue even without further warming (van den Akker et al., 2025). Meanwhile, a recently identified tipping point in grounding-zone melting, where small ocean temperature increases trigger unbounded water intrusion beneath the ice, represents an additional vulnerability not yet captured in standard projections (Bradley & Hewitt, 2024).

Standard projections put Antarctica's contribution to 2100 SLR at up to 30 cm under mid-to-high emissions (Fox-Kemper et al., 2021, Ch. 9), but recent work documents a substantial low-end bias in this ensemble.[14] Calculations shared with us by Brent Minchew

[12] Emphasizing this alarming discrepancy, researcher Brent Minchew notes, "[E]ven if you just take current observed trends, things that we can measure for sure right now, and you run them forward in time, and you compare the model output for close to five degrees of warming, the models still underpredict the current observed trends with five times more warming in the system."
[13] Nias et al. (2023) suggest only about 34mm (interquartile range: 18–52mm) would arrive by 2100, with the remainder unfolding over subsequent centuries.
[14] It appears that most contributing models cannot reproduce observed mass loss when hindcast (Aschwanden et al., 2021).

(Arête Glacier Initiative) suggest exceeding 1 m of total SLR by 2100 is more likely than not under policy-relevant emissions scenarios, with West Antarctica driving the dominant share of the uncertainty, though this is unpublished.

The human impacts of this rise cannot be discounted. Kulp and Strauss (2019) estimate that 230 million people currently live on land below one meter above high tide. Even under conservative models, the potential scale of displacement is large; under the higher-end scenarios our experts suggest, it is catastrophic. Highlighting the extreme socio-economic vulnerability of these coastal communities, Melissa Barnes of Arête Glacier Initiative adds, "[A] third of the world's population is within a day's walk of the shore and half of that population makes less than $5 a day . . . you just start to do the math on that, it's astronomical."

Our BOTECs Estimate an EA Community WTP of $445M to Address 1.5 m SLR by 2100 (Single-Funder WTP $223M)

Our calculations can be found in this spreadsheet. We estimate excess deaths from displacement caused by sea level rise (SLR) of between 1 and 2 meters by 2100. We set this amount to match plausible estimates Brent Minchew shared with us, see here, but we flag that this is a single estimate.[15] We therefore think of this as a BOTEC reflecting this scenario rather than reflecting current agreement in climate science. To estimate human impacts, we model how displacement affects mortality. We start with the raw excess mortality ratio for internally displaced persons (IDPs) from Heudtlass et al. (2016), which is 2.5x. Because this is derived mostly from conflict-ridden areas with a unique combination of aggravating elements, we apply a discount[16] of 70% to account for what we would expect is the slower-onset, less violent nature of climate displacement. This yields an effective multiplier of ~1.45x,[17] meaning that for every 10,000 displaced people who would normally be expected to die in a given year, roughly 4,500 additional deaths occur due to displacement-related vulnerabilities; for example, overcrowded conditions, disrupted healthcare access, loss of livelihoods, and psychological stress. We also estimate the share of displacement occurring in low- versus high-income countries (LMICs vs. HICs), assuming that the share of exposed people in LMICs will be much higher than in HICs and that displacement rates will differ accordingly (roughly 30% in LMICs versus 5% in HICs). Though we apply the excess mortality multiplier equally across income groups, the burden falls overwhelmingly on lower-income countries, since these are the populations most likely to lack the infrastructure and institutional capacity to adapt.[18]

Using a central scenario of 1.5 meters of SLR by 2100, our BOTEC produces a total naive damage from displacement-driven excess mortality of $17.8 trillion in CG philanthropic value. Dividing by the 2000× cost-effectiveness bar gives a WTP at certainty of $8.9 billion. Applying our two community-level discounts (p(success) = 10%, EA credit share = 50%) gives an EA

[15] This is above what IPCC AR6 predicts even under the high emissions scenarios, though they acknowledge high uncertainty here.
[16] I.e., a reduction in the assumed effect size, based on our own subjective assessment.
[17] To calculate the adjusted mortality rate, we start with the initial estimate of 2.5 and subtract the baseline normal death rate (1.0) to isolate the excess portion ($2.5 - 1.0 = 1.5$). Next, we discount this excess by 70% (meaning we retain 30%). Multiplying the excess by 30% gives us our new excess rate ($1.5 \times 0.30 = 0.45$). Finally, we add the baseline back in to arrive at the new estimated total death rate of 1.45.
[18] Essentially, the HIC deaths are not contributing much to our overall mortality estimate, so the excess mortality multiplier consequently does not make a big difference in our results. It is possible that this changes over the medium term, however, as poorer countries become wealthier. This is an assumption we do not make, however.

community WTP of $445 million; applying the further within-EA share of 50% gives a single-funder WTP of $223 million.

We have the most uncertainty over the excess mortality multiplier and the length of elevated excess mortality risk. If either were off, our conclusions would move significantly in either direction.

AMOC Collapse

The Atlantic Meridional Overturning Circulation (AMOC) is a large ocean circulation system that moves warm, salty water northward along the surface of the Atlantic and cold, dense water southward at depth, acting as a major conveyor belt of heat through the climate system. It plays a critical role in regulating temperatures, rainfall, and sea levels, making its potential disruption one of the most consequential risks in climate science (Jackson et al., 2015).

The AMOC may be vulnerable to a tipping point, beyond which a rapid and self-sustaining collapse could occur that cannot be easily reversed, even if warming stopped. The key instability arises from the salt-advection feedback: a weakening circulation imports less salt to the North Atlantic, reducing surface density and further suppressing deep water formation in a self-amplifying loop. Freshwater forcing (e.g., from Greenland ice-sheet melt or river runoff) acts as the main destabilizing input. Some climate models show that once a tipping point is crossed, the AMOC does not slow down gradually but can collapse within decades (van Westen et al., 2024); mainstream models tend to project a gradual weakening rather than abrupt collapse (IPCC, 2023), but this is possibly explained by model biases (Liu et al., 2017).

Estimates of collapse probability have also evolved considerably over time. Earlier IPCC assessments (through AR5) regarded an abrupt AMOC collapse this century as unlikely, largely because climate models rarely produced full collapses under standard emissions scenarios. However, more recent work has pushed the probability estimates upward. A study by Ditlevsen and Ditlevsen (2023), analyzing sea surface temperature proxies as indirect AMOC indicators, used statistical early-warning signals to estimate a potential tipping point between 2025 and 2095 under continued forcing, although this methodology attracted significant methodological criticism (Ben-Yami et al., 2024; Morr et al., 2026[19]). Recent work by Portmann et al. (2026) emphasized that climate models have been consistently underestimating AMOC weakening, projecting a slowdown of approximately ~50% by 2100, even under a moderate emissions scenario.[20] A specific timeline for collapse lacks a clear scientific consensus, but it is increasingly treated less as a negligible tail risk and more as a plausible scenario that demands serious attention.

Our BOTECs Estimate an EA Community WTP of $263M to Address AMOC Collapse Risk (Single-Funder WTP $131M)

Our BOTEC estimates the impact of AMOC collapse on mortality. The estimate is structured around the probability of collapse by 2100, two impact pathways, and a probability-weighted expected value calculation. We view food insecurity, occurring as a result of temperature and

[19] Note that this is a non-peer-reviewed preprint.
[20] Portmann et al. specifically estimate, with 90% probability, a 51 ±8% weakening by 2100, in contrast to a 90% probability of a 32 ± 37% weakening over the same time horizon in the CMIP6's 2-4.5 baseline scenario.

precipitation changes induced by AMOC collapse, as the primary cause of mortality in this scenario. To do so, we estimate the number of people pushed into food insecurity by:

1. **Northern European cooling.** AMOC collapse is predicted to result in substantial cooling across Northern Europe, leading to shocks to the food supply that could lead to price increases.[21]
    a. To estimate human impacts, we take the predicted agricultural loss in Northern Europe and translate it into an assumed increase in food prices, which pushes an additional proportion of those populations into food insecurity.[22] One key assumption is the percentage of the population pushed into food insecurity as a direct result of agricultural productivity loss; we estimate this as ~1–10% of the overall population in HICs, but this could be conservative if the agricultural loss is sudden.[23] We then apply an increased risk of mortality (~15–30%) for food-insecure populations in high-income countries (Jalili et al., 2025).
2. **Monsoon system disruption across the Sahel region and South Asia.** AMOC collapse would severely weaken West African and South Asian monsoon systems, causing catastrophic rainfall deficits that would threaten food security.[24]
    a. To estimate human impacts, we estimate the population affected by monsoon changes in 2100, accounting for population growth. We then draw from the literature to estimate the additional population pushed into food insecurity as a result of reduced rainfall (~5–20%, Randell et al., 2022; IMF, 2020).

We excluded the following potential effects:

- **Sea level rise**. One expert also mentioned additional SLR effects,[25] but we chose not to include these because the effects seemed likely to be minor in comparison to food insecurity,[26] and we already model SLR separately.
- **Displacement and economic effects**. The same expert also mentioned that "such a profound environmental shift would trigger large-scale human migrations and economic instability, as nations struggle to adapt to the new climatic reality." We did not

[21] van Westen and Baatsen (2025) estimate drops in the average monthly winter temperatures of 10–30°C in northern Europe as a result of AMOC collapse, while Ritchie et al. (2020) estimate that this would also result in a ~80% reduction in arable land in Britain by 2080. We modeled these climate disruptions as a driver of food insecurity in northern and Western Europe as a function of price increases pushing additional people into food insecurity, with an associated increased risk of mortality.
[22] Note that we exclusively model the effect on populations in Northern and Western Europe and exclude populations that may depend on European food exports. This is a significant simplification of the economic dynamics of food pricing, but we expect that export regions may be able to change trade patterns to adapt to changes in European production.
[23] Our assumption is based on an expected loss of ~30% of European agricultural yields (Benton et al., 2017), which we believe would affect HICs as a function of price shocks. The relationship between such a substantial agricultural productivity loss and transitions into food insecurity is not well-established; we also assume that this stream of impacts affects high-income countries only, who would be better placed to manage price shocks (e.g., by increasing trade).
[24] Ben-Yami et al. (2024) found using a cross-model comparison that an AMOC collapse would cause substantial disruptions to the West African Monsoon (WAM), Indian Summer Monsoon (ISM), and East Asian Summer Monsoon (EASM), with shorter wet seasons and longer dry seasons. We focused on regions affected by the WAM and ISM zones and excluded the EASM, due to a much lower predicted change in rainfall (-4% vs. -20–30%).
[25] Expert: "A shutdown of the AMOC would cause a significant rise in sea levels along the North American Atlantic coast, as the current's weakening would reduce the dynamic slope of the ocean surface."
[26] Specifically affecting only a proportion of residents in high-income countries, who would be expected to absorb negative effects more easily than populations in LMICs.

explicitly model displacement or the economic consequences of a collapse because we chose to prioritize direct impacts on human mortality across all models.
- **Risk of additional warming.** Recent modeling work from Nian et al. (2026) indicates that AMOC collapse could trigger deep convection in the Southern Ocean, ventilating carbon-rich waters and releasing 47–83 ppm of CO_2 into the atmosphere.

The impacts are calculated conditional on the probability of a full collapse, with the expected value of averting a collapse calculated based on the expected value (probability of collapse x human impact given collapse). We do not explicitly take into account how these effects might be offset, e.g., by production increases in other locations or adaptations in international trade. Instead, we apply a highly uncertain assumption of the duration of elevated mortality risk for both effect streams (5–15 years for Northern European cooling and 10–40 years for Monsoon system disruption). The effects are heavily concentrated in the consequences of monsoon system disruption, largely because they (i) affect a greater number of people; (ii) populations affected are in LMICs; and (iii) we see less potential for adaptation to a long-term weather-induced shock than a price shock.

Using a 15% probability of collapse by 2100, we calculate the total probability-weighted naive damage as a result of AMOC collapse to be $10.5 trillion in CG philanthropic value. Dividing by the 2000× cost-effectiveness bar gives a WTP at certainty of $5.25 billion. Applying our two community-level discounts (p(success) = 10%, EA credit share = 50%) gives an EA community WTP of $263 million; applying the further within-EA share of 50% gives a single-funder WTP of $131 million.

Overall

Taken together, we think there is a plausible case for an EA-aligned philanthropic investment in the low-to-mid hundreds of millions of dollars to support an Earth system stabilization fund. If we take each pathway as a distinct opportunity, the EA community WTP ranges from $263 million to $445 million per pathway, with single-funder WTPs of $131 million to $223 million per pathway. Independent of the specific parameter choices, the central conclusion is that the philanthropic value at stake comfortably justifies a fund in this size range. Even if we are off by an order of magnitude, the case in favor of supporting Earth systems stabilization is strong.

A few takeaways from our preceding exercise:
- Our estimates are intentionally conservative, applying significant “haircuts” to displacement mortality and focusing strictly on direct deaths.
- By using mortality as our primary metric, we are surely understating the total human harm. Our decision to exclude morbidity, displacement costs, and broader socioeconomic destabilization means these figures should be viewed as a lower bound.
- An EA community WTP in the low-to-mid hundreds of millions and a single-funder WTP in the low hundreds of millions both align with the capacity of EA-aligned climate philanthropy. We would have to be off by three to four orders of magnitude (indicating human impacts are much smaller than we assume) to significantly reduce our estimates of cost-effective funding in this space.

Interventions to Address Earth Systems Destabilization

The preceding sections establish that significant human harms may arise from the destabilization of Earth systems across four distinct pathways. The question that follows is whether the current portfolio of climate responses, centered on emissions mitigation and

carbon dioxide removal (CDR), is adequate to avert these harms and, if not, what alternatives exist. Our overall assessment is that current strategies are likely insufficient and alternative approaches should be researched, such that they are deployment-ready in the event that they are needed. The following three sub-sections detail our reasoning for arriving at this conclusion.

Mitigation and CDR are Unlikely to Keep Warming Below 1.5°C

Rapid and deep reductions in greenhouse gas (GHG) emissions remain the single most important lever for limiting climate risk, and CDR is an essential complement—the IPCC has stated that its deployment is "unavoidable if net-zero CO_2 or GHG emissions are to be achieved" (IPCC, 2022, p. 47). However, neither is currently on track to prevent warming from exceeding 1.5°C, and there are serious questions about whether they can do so even with accelerated effort.

On mitigation, the gap between commitments and requirements is well-established. The IPCC AR6 found that limiting warming to 1.5°C requires global CO_2 emissions to decline by approximately 43% below 2019 levels by 2030 and reach net-zero by the early 2050s (IPCC, 2022, Ch. 1). Current Nationally Determined Contributions (NDCs) under the Paris Agreement fall far short: the UNEP Emissions Gap Report 2025 projects that full implementation of all NDCs would result in warming of 2.3–2.5°C by 2100 (UNEP, 2025). The IEA's World Energy Outlook 2025 projects approximately 2.5°C under current policy settings, and states that an overshoot of 1.5°C is now "inevitable" in every scenario it models (IEA, 2025).

On CDR, the challenge is both scale and speed. The State of Carbon Dioxide Removal Report (2024) found that after filtering scenarios for feasibility and sustainability, pathways limiting peak warming to approximately 1.7°C represent the lowest warming outcomes currently supported in the literature, and that limiting peak warming to 1.5°C with no or low overshoot is "no longer feasible" (Smith et al., 2024). Novel CDR methods currently remove approximately 1.3 million tons of CO_2 per year, less than 0.1% of what mid-century 1.5°C pathways require. The IPCC itself has described CDR at scale as "unproven" and stated that reversing an overshoot of 0.2°C or more would require deployment "at rates and volumes that might not be achievable" (IPCC, 2022, D. 1.2).

However, CDR is on a steep growth trajectory. Direct Air Capture (DAC) capacity increased approximately 873% from 2024 to 2025, with over 130 facilities in the pipeline (Budunis, 2024), and there is real policy momentum behind it. But even under optimistic cost projections ($150–$200/ton by 2030, down from $400–$1,000 today; Al Juaied & Whitmore, 2023), the gap between current removal and what Paris-aligned pathways require remains vast.

We take these developments seriously. But even with accelerated investment, CDR faces structural constraints that limit how fast it can scale: it is extraordinarily energy-intensive; permanent geological storage is unevenly distributed and slow to permit (Massei, 2025); its efficacy is partially self-defeating in a warming world, as the IPCC AR6 found that CDR is partially counteracted by CO_2 release from warming oceans and land reservoirs; and it operates on decadal timescales, requiring time to build facilities, develop supply chains, and progressively draw down atmospheric CO_2. Lamb et al. (2024) confirmed that no NDCs contain pledges to substantively scale novel CDR methods.

The World is Likely to Spend Extended Periods Above 1.5°C, Increasing the Risk of Earth Systems Destabilization

If the most optimistic feasible pathway still implies peak warming of approximately 1.7°C, the world will pass through the range (1.5–2.0°C) where several tipping elements have their lower thresholds, including the Greenland and West Antarctic ice sheets, abrupt permafrost thaw, and Barents Sea ice loss (McKay et al., 2022). The critical question is not just whether 1.5°C is breached, but for how long.

Even under the IEA's most optimistic Net Zero Emissions by 2050 scenario, peak warming exceeds 1.5°C for several decades, reaching approximately 1.65°C around 2050 and returning below 1.5°C only by 2100. Moreover, this return depends on widespread deployment of CDR technologies that the IEA itself describes as "currently unproven at large scale" (IEA, 2025). Under the more realistic Stated Policies Scenario, which reflects current policies, warming reaches approximately 2.5°C and *never* returns below 1.5°C. This means that under the best-case scenario, the world faces approximately 70 years above 1.5°C; under more likely trajectories, the overshoot is permanent.

This duration matters. Schleussner et al. (2024) showed that climate outcomes during an overshoot period are materially different from a world that avoids one, with lasting damage to ice sheets and sea levels that is not fully reversed even if temperatures eventually decline. This concern is consistent with the broader analytical framing of Steffen et al. (2018), who warned that Earth system feedbacks could lock the planet into a "hothouse" trajectory even if emissions cease, and with the Global Tipping Points Report (2023), which concluded that current governance and policy frameworks are inadequate to manage the risks of cascading tipping points.

Overall, the concern is that decades of exposure to temperatures above 1.5°C give destabilizing processes sustained time to accelerate, interact, and potentially cross thresholds from which recovery is slow or impossible.

Extended Overshoot may Cause Some Earth Systems to Permanently Destabilize

The most concerning dimension of this evidence is the possibility that some Earth systems may destabilize on self-reinforcing trajectories that persist regardless of subsequent human action to reduce global mean temperatures, effectively decoupling them from the emissions pathway.

On sea level rise, Naughten et al. (2023) found that ocean-driven melting of West Antarctic ice shelves is projected to accelerate at triple the historical rate and appears "unavoidable" even under the most ambitious Paris Agreement targets. van den Akker et al. (2025) found that present-day mass loss rates from the Thwaites and Pine Island glaciers are consistent with the early stages of marine ice-sheet instability and that this retreat would continue even without further warming. On warming-induced emissions, Park et al. (2025) showed that permafrost ecosystem carbon loss may persist even under net-zero and net-negative emissions scenarios, with methane emissions increasing due to the irreversibility of inundated areas. On AMOC, modeling evidence consistently shows that, once collapsed, AMOC does not recover until forcing conditions reverse well beyond the collapse threshold, with recovery timescales of

centuries to millennia even under favorable conditions (van Westen & Dijkstra, 2023; Jackson & Wood, 2018). Finally, regarding albedo loss, feedback loops from sea ice and snow loss are generally assessed as reversible on decadal timescales if warming is reduced (IPCC, 2021, Table 4.10), but the reversibility of the observed decline in low-cloud cover is unknown, as its causes are not yet well enough understood to determine whether the trend would reverse with reduced warming (Goessling et al., 2024; Zelinka et al., 2026).

In addition, while each feedback pathway described above may individually demonstrate some risk of irreversibility, there is growing evidence for cross-pathway interactions that could further exacerbate these trajectories. For example, Wunderling et al. (2024) report that the majority of known linkages between tipping elements are destabilizing, meaning that crossing one threshold can lower the thresholds at which others are crossed.[27]

An oft-cited counterargument is that additional warming after emissions reach net zero is likely negligible. MacDougall et al. (2020) found this in a multi-model intercomparison,[28] and the IPCC AR6 adopted it as its best estimate. If taken at face value, this implies that temperatures would stabilize once net zero is achieved, which would weaken the case for permanent destabilization. However, the uncertainty range is substantial (±0.3°C), and a subsequent review by many of the same researchers cautioned that several Earth system feedbacks, like permafrost carbon release, ocean circulation changes, and cloud responses, are not fully represented in the models used to estimate this commitment (Corner et al., 2023). In other words, the feedbacks most likely to push additional post-net-zero warming above zero are precisely the feedbacks that the models did not include. Moreover, even if temperatures do stabilize at net zero, the Zero Emissions Commitment finding does not address the *level* at which they stabilize: if net zero is achieved at 1.8°C, tipping points with thresholds below that temperature may already have been triggered.

Some tipping elements may also be more reversible than others. For example, Arctic summer sea ice loss is generally assessed as reversible on decadal timescales if warming is reduced, while ice-sheet collapse and permafrost carbon release are assessed as irreversible on centennial to millennial timescales (IPCC, 2021, Table 4.10). Current climate models do not fully capture tipping-point interactions, which means that risks could be either higher or lower than currently estimated. We note that model limitations should prompt humility rather than certainty in either direction. That said, the expected value framework adopted in this report is designed for precisely this kind of situation: where the central estimate is relatively reassuring, but the downside tail (in which stabilization occurs at a level that has already triggered irreversible changes, or in which Earth system feedbacks sustain warming beyond what emissions alone would produce) carries consequences severe enough to warrant serious attention even at moderate probabilities.

Synthesis

In summary: CDR is insufficient to prevent warming from exceeding 1.5°C on any plausible timeline (**high confidence**). The world is likely to spend several decades, and potentially the

[27] We note that these findings carry significant uncertainties, particularly regarding the pace of ice-sheet response, the magnitude of permafrost feedbacks, and the robustness of cascade interactions across different models. None should be read as deterministic predictions. However, the direction of evidence is consistent across multiple independent lines of research, and the potential irreversibility of the processes involved means that waiting for greater certainty carries its own risks.

[28] The model ensemble median of temperature change after net CO_2 emissions reach zero is a very slight cooling (-0.05°C).

rest of this century, above 1.5°C, a range in which multiple tipping elements have their lower thresholds (**high confidence**). Extended exposure to these temperatures increases the likelihood that some Earth systems will begin destabilizing on self-reinforcing trajectories (**medium confidence**), and there is a non-negligible probability that some of these changes will prove irreversible on human timescales even if temperatures are subsequently reduced (**medium-low confidence, but high consequence**).

None of this diminishes the case for mitigation and CDR, both of which are essential and must be pursued with maximum urgency. It does, however, suggest that an exclusive reliance on these approaches carries significant risk. CDR addresses the stock of CO_2 in the atmosphere over decades; it cannot protect at-risk Earth systems from the energy imbalance they are experiencing now.

It is for this reason that humans may need alternative approaches in addition to mitigation and CDR to avert significant climate-induced harms. Such approaches, here termed "Earth systems stabilization interventions," are arguably needed as a means of keeping critical Earth systems intact long enough for the longer-term atmospheric effects of CDR to take effect. This conclusion is consistent with the growing institutional consensus documented in the introduction to this report, including calls from the US National Academies of Sciences (2021), Moore et al. (2018), and the University of Chicago's Climate Systems Engineering initiative (MacAyeal et al., 2024) for serious research into Earth systems stabilization interventions.

Overview and Deep Dives on the Tractability and Neglectedness of Earth Systems Stabilization Approaches

Various research avenues and engineering approaches have been proposed to address each of the previous risk pathways. It is our view that it is reasonable to consider these options given the magnitude of the potential harms.

There are many candidate approaches. Rather than longlist them, at the direction of our client, we consider only three such strategies: Stratospheric Aerosol Injection (SAI), Marine Cloud Brightening (MCB) and Mixed-Phase Cloud Thinning (MCT), and Ice-Sheet Stabilization (specifically, basal dewatering and thermosiphons). We assess not only what is theoretically possible, but what is practical, underfunded, and capable of moving the needle on climate stability (according to an importance-tractability-neglectedness [ITN] framework) to determine their overall promisingness and capacity to absorb marginal funding effectively. A summary of our views, from our own desk research and interviews on these approaches, can be found in Table 3 below. The table below summarizes our current view based on the research and interviews conducted for this report:

Table 3: Earth systems stabilization interventions and our ITN assessment of overall promise[29]

Intervention Avenue	Importance	Neglectedness	Tractability	Overall Promisingness
Stratospheric Aerosol Injection (SAI)	Proven proof-of-concept for 0.5C cooling (volcanic eruptions). Models indicate SAI can reduce global temperature by 1–2C. No controlled outdoor experiments have been conducted.	Funding is growing, but SAI remains under-resourced compared to potential global leverage and necessary governance infrastructure.	The engineering is broadly feasible within 10–15 years, but the lack of field validation, governance frameworks, and political sensitivity constrains near-term tractability.	**High leverage, high risk.** SAI is the most capable method for rapid global cooling. Philanthropic dollars on the governance and equity side may have particularly high marginal value.
Marine Cloud Brightening (MCB) & Mixed-Phase Cloud Thinning (MCT)	Observational proof-of-concept for MCB and modeling support for MCT, but limited real-world evidence. Both show promise for regional stabilization.	Highly neglected. MCB receives an estimated ~$10M/year. MCT is even more neglected, with only one dedicated initiative (ASI) currently operating.	MCB achieved first outdoor field demonstration in 2025, though nozzle engineering challenges remain. MCT leverages existing cloud-seeding infrastructure, but atmospheric modeling of second-order effects remains immature. Both are rapidly reversible, reducing termination risk relative to SAI.	**High potential for targeted preservation.** Small philanthropic investments here could drastically accelerate research to protect critical, localized climate tipping points. MCT may have capacity constraints to address alongside increased funding.
Ice Sheet Stabilization (Basal Dewatering & Thermosiphons)	Directly targets one of the most severe climate risks (WAIS collapse and	Rather neglected. Only a handful of young organizations	Preliminary research to resolve key feasibility questions is tractable and fundable at	**High marginal impact.** A severely under-researched area where early philanthropic

[29] How do we think about neglectedness? Our approach is closely grounded in thinking developed by the effective altruism community (see these posts by Coefficient Giving, 2025 and Wiblin, 2019). In brief, we ask whether the attention a problem receives is commensurate with its potential scale, assuming, as we do throughout this report, that the first dollars into a particularly neglected cause are unusually high-impact. We assess this qualitatively, weighing a few key inputs: 1) the size of existing funding in the relevant space; 2) the mix of funders and actors involved (e.g., funders: government, philanthropic; actors: academic, private sector); and 3) the counterfactual probability that important work to address the problem would happen anyway in the absence of EA-aligned funding. Practically speaking, we treat an area as low neglectedness when existing funding is broadly proportionate to the expected scale of the issue (or greater). In this case, additional dollars would largely duplicate work or crowd out funding. We treat an area as high neglectedness when the opposite is true: the total amount of funding is small relative to the potential scale of the problem. As such, marginal dollars are unlikely to crowd out other actors, and additional funding would meaningfully advance the work in ways it could not otherwise. We save the label of medium for the cases in between, where investment is accelerating quickly or where the area is receiving comparatively more attention than otherwise similar problems but remains underfunded relative to its potential leverage.

	rapid SLR), though the intervention's ability to slow ice loss at scale is unproven.	are active, with cumulative philanthropic funding likely under $20 million.	~$100–120M over 10 years. A full pilot would cost hundreds of millions and likely require government support.	dollars are desperately needed to simply determine if stabilization is physically possible at scale.

Mainstream Arguments Against Earth Systems Stabilization Interventions

Throughout this report, we use the term "Earth systems stabilization approaches" to refer to strategies, research, and interventions that aim to directly advance the management or restabilization of at-risk Earth systems, rather than addressing the root cause of their destabilization (excess atmospheric greenhouse gases). This category encompasses solar radiation modification (SRM) techniques (such as SAI and MCB) and physical interventions targeting specific Earth systems, such as glacier and ice sheet stabilization. While CDR is sometimes grouped alongside these technologies under the broader umbrella of "geoengineering," we treat it separately here because CDR addresses atmospheric CO_2 concentrations directly, whereas the interventions assessed in this section act on the symptoms of warming (temperature, energy balance, ice dynamics) rather than the underlying cause. This distinction matters: the arguments for and against these interventions differ in important ways from those surrounding CDR.

Earth systems stabilization approaches remain controversial in contemporary climate policy, with some prominent scientists and governments calling for moratoriums. In December 2024, the EU's Group of Chief Scientific Advisors and European Group on Ethics recommended a "Europe-wide moratorium on using solar radiation modification technologies," concluding that such technologies" (Scientific Advice, 2024) represent an unacceptable level of risk for humans and the environment" in their current form (European Commission, 2024). This moratorium built on the Commission's earlier position, expressed by then-climate policy chief Frans Timmermans in June 2023, that "nobody should be conducting experiments alone with our shared planet." Multiple US states have passed bans on SRM, while hundreds of academics called for a Solar Geoengineering Non-Use Agreement in 2022 (Biermann et al., 2022), and the Global Tipping Points Report 2025 called for a moratorium on SRM deployment and large-scale experiments. A group of 42 scientists critiqued polar geoengineering approaches (including both SRM and physical interventions such as seabed curtains and ice sheet stabilization) as "environmentally dangerous" and argued that further research "would not be an effective use of limited time and resources" (Siegert et al., 2025), though this assessment has itself attracted formal rebuttals from researchers who argue that the risks of unchecked climate change make further research "a scientific and societal imperative" (Henry & Duffey, 2025; Moore et al., 2025). Both scientific and public opposition indicate that there could be strong headwinds for further research and deployment. We see the central critiques and counterarguments as follows:

Earth systems stabilization approaches fail to fix the underlying problem: Earth systems stabilization approaches treat the symptoms (temperature, ice loss) rather than the cause (CO_2 accumulation). For example, SRM does not reduce atmospheric CO_2, meaning ocean acidification continues unabated, potentially creating an illusion of safety. Similarly, glacier stabilization interventions would slow ice loss without addressing the warming that drives it.

Counterargument: No single intervention is expected to solve climate change alone. Earth system stabilization approaches are framed as complements to mitigation and carbon removal, not substitutes, intending to provide a bridging measure to protect critical systems while longer-term solutions take effect.

Moral hazard effects: If Earth systems stabilization interventions are perceived as a viable backstop, they may reduce political pressure for decarbonization.

Counterargument: Empirical research on this question is mixed. Merk and Wagner (2024) found no significant mitigation deterrence effect in a large-scale, revealed-preference Facebook experiment after being informed about SRM interventions. A broader literature review found that the moral hazard argument, while prominent in ethical debates, rests on largely theoretical rather than empirical foundations (Pamplany et al., 2020). Jebari et al. (2021) have argued that the "moral hazard" framing is itself unhelpful, proposing instead a "risk-response feedback" framework that better captures the complex interactions between different climate policy instruments.

Governance and ethical consent are insurmountable challenges: No international framework exists to decide who gets to deploy these interventions, at what scale, or on whose behalf. Unilateral deployment by one country or actor could impose climate effects (e.g., altered monsoon or precipitation patterns) on others who never consented. For physical interventions like glacier stabilization, governance is further complicated by the Antarctic Treaty System and questions of territorial jurisdiction.

Counterargument: This is likely one of the greatest barriers to deployment. However, governance frameworks are nascent rather than fully absent (e.g., the Alliance for Just Deliberation on Solar Geoengineering aims to shape governance around inclusive, cross-regional collaboration). The governance challenge is precisely the reason to invest in research now: building international norms and institutions before deployment pressure mounts is more tractable than doing so in a crisis.

Equity concerns around unintended regional impacts: The effects of these interventions are not uniformly distributed, with modeling suggesting that some regions could experience reduced precipitation even as global temperature falls. The populations least responsible for climate change may therefore bear disproportionate risks.

Counterargument: This risk is real and is one of the strongest arguments for caution. Irvine et al. (2019) found that a scenario with half-measure SAI substantially reduced climate risk for the vast majority of regions compared to unmitigated warming. However, "the vast majority" is not the same as "all", and the possibility remains that some vulnerable regions could face amplified risks from interventions (such as altered monsoon patterns) layered on top of the harms they are already experiencing from climate change itself. The relevant counterfactual is not these interventions versus a stable climate, but these interventions versus continued warming. The key question is therefore whether the net effect is beneficial for vulnerable populations, not just for the global average. This cannot be determined without substantially more research, nor without governance frameworks that give affected populations a meaningful voice in deployment decisions.

Epistemic humility: Climate systems are extraordinarily complex, and the history of large-scale environmental interventions (e.g., introducing species, damming rivers) is rife with unintended consequences that models failed to predict.

> *Counterargument:* Epistemic humility is an argument for more research, not less. Unmitigated climate change is itself arguably the largest uncontrolled experiment ever conducted on Earth systems, and its effects are already exceeding what many models predicted.

Distraction from proven solutions: Resources devoted to Earth stabilization research could be better spent on approaches with fewer downside risks, such as accelerating renewable energy, carbon removal, or adaptation.

> *Counterargument:* Current research budgets for Earth systems stabilization approaches are negligible relative to clean energy investment. Moreover, failing to understand the potential of these options leaves policymakers blind in the case of a climate emergency. The scale of funding recommended in this report (in the tens to hundreds of millions of dollars) is a fraction of what would be required to deploy even a single gigaton-scale direct air capture (DAC) facility, yet it could produce outsized returns by resolving critical questions about the safety, governance, and technical readiness of stabilization approaches. If these interventions prove viable, early research investment ensures they are available when needed. If these interventions prove unviable or too risky, relatively modest research expenditure now could clarify this before far larger sums are committed. In either case, the information value of the research is high relative to its cost.

Having made the case for investing in Earth systems stabilization research, it is worth noting that the specific pathways to moving the field forward remain uncertain. We now move to detailed assessments of the promise of three candidate interventions for philanthropic investment: stratospheric aerosol injection (SAI), marine cloud brightening (MCB) and mixed-phase cloud thinning (MCT), and ice-sheet stabilization. Within each, we structure our analysis according to an importance–tractability–neglectedness (ITN) framework, alongside potential risks and side effects.

Stratospheric Aerosol Injection (SAI)

SAI is an SRM intervention that involves dispersing reflective particles, most commonly sulfur dioxide (SO_2) or calcium carbonate ($CaCO_3$), into the lower stratosphere at altitudes of roughly 15–25 km. The particles effectively scatter incoming solar radiation back into space and reduce the amount of energy reaching Earth's surface. SAI would use small, high-altitude aircraft or alternative specialized delivery systems to maintain a persistent aerosol layer, effectively increasing albedo without altering atmospheric CO_2 concentrations.

SAI would directly address the albedo loss risk pathway by preventing temperature rise: as Arctic sea ice and land ice reflect less sunlight due to warming, SAI could partially counteract this feedback loop by reducing the radiative forcing that accelerates it. It could also reduce the risk of feedback loops on the warming-induced emissions pathway, slowing permafrost thaw and reducing greenhouse gas release from soils and wetlands. Overall, SAI could be thought of

as a band-aid, maintaining a certain level of planetary albedo while the world takes greater action to address the root causes of climate change.

What makes SAI distinct within SRM and climate intervention more broadly is its combination of speed, scale, and cost-effectiveness. Among the prominent SRM approaches, SAI is held as most proximately feasible (Parson & Keith, 2024). No other intervention can move global mean temperature as quickly or at comparable cost, but this also makes it risky if misapplied or abruptly terminated. Evidence exists both for and against deployment, and therefore the key debate is whether the risks outweigh the benefits. We have moderately high confidence that SAI would indeed bring down global temperatures, but are unsure that the governance and political challenges can be overcome for successful deployment in the near term.

Importance

Our take on importance: high importance of the problem and medium-high potential impact of the intervention, dependent on the balance between global and regional impacts. Models broadly agree that SAI could reduce global mean temperatures, with studies suggesting it could limit warming even under high-emissions trajectories or keep it below 1.5°C (Visioni et al., 2021; Richter et al., 2022). However, confidence in SAI's ability to reduce global average temperatures does not extend equally to its regional effects, which remain a major area of uncertainty and concern (see Risks and Side Effects, below). In terms of potential climate impact, SAI's ability to interrupt multiple tipping point feedback loops simultaneously places it in a category apart from most other interventions. Evidence in support of its potential impact includes the following:

- Volcanic eruptions demonstrate proof-of-concept that SAI works. The 1991 eruption of Mount Pinatubo in the Philippines injected approximately 20 million metric tons of SO_2 into the stratosphere, causing a 0.5°C reduction in global mean temperatures 18 months after the eruption (Soden et al., 2002).
- Models have high confidence that SAI could feasibly achieve a 1–2°C temperature reduction. Continuous annual injection of 8–16 million metric tons of SO_2 (roughly equivalent to a Pinatubo-scale eruption every year) would be required to cool the Earth by 1°C (Haywood et al., 2022). Diminishing marginal cooling is expected at higher injection rates, however, because particles grow larger and scatter solar radiation less efficiently (Niemeier & Timmreck, 2015).

Neglectedness

Our take on neglectedness: We have medium confidence that the space is moderately neglected, though much less neglected than alternative SRM approaches. Overall funding for SRM, of which SAI is the major component, increased almost threefold in 2025 compared to 2024 (Geoengineering Monitor, 2026), with public funding accounting for the majority of the increase. This trajectory means SAI is becoming progressively less neglected as a research area, though we think it remains significantly underfunded relative to its potential climate leverage and the governance infrastructure it would require, and suggests that this is a gap that philanthropy can help close.

Overall, the physical science is less neglected than the broader institutional groundwork that would be required for any legitimate deployment decision. This includes not only governance design and equity frameworks, but also efforts to support serious government engagement with SRM research, including the establishment of publicly funded, credible national and

multilateral research programs. At present, the field remains disproportionately shaped by private philanthropy and, increasingly, venture capital (see Tractability), rather than by the government-led research infrastructure that would be needed to build public trust and political legitimacy. Similarly, meaningful inclusion of climate-vulnerable countries in shaping any future deployment decisions may require not only academic equity frameworks but direct support for the political participation and institutional capacity of affected nations. It is therefore likely that philanthropic funding that targets the actions around catalyzing government leadership, building multilateral research infrastructure, and supporting political engagement from vulnerable countries could have high marginal impact.

Organizations currently working on SAI, with notable examples, can be broken down into the following categories:

- **Academic and research institutions:**
  - University of Chicago's Climate Systems Engineering Initiative, founded in 2023 and funded by $36 million in grants and private philanthropy, researches the benefits and risks of SRM, carbon removal, and glacial interventions, including research on policy and governance for SAI.
- **Government entities:**
  - The UK's Advanced Research and Invention Agency (ARIA) runs the £56.8 million Exploring Climate Cooling program, which funds 22 research projects spanning atmospheric physics, climate modeling, governance, ethics, and small-scale controlled outdoor experiments across a range of climate cooling technologies, including SAI, marine cloud brightening, and Arctic sea ice re-thickening (ARIA, 2025).
  - Separately, the UK's Natural Environment Research Council (NERC) funds the £10.5 million "Modelling environmental responses to SRM" program (2025–2030), which focuses on modeling the climate impacts of SAI and MCB using natural analogs such as volcanic eruptions and changes in shipping emissions. Together, these represent the largest public investment in climate-cooling research by any single country to date (UKRI, 2026).
- **Nonprofits:**
  - Reflective, incubated by Renaissance Philanthropy, aims to accelerate SAI research and development to support informed decision-making. Reflective's CEO, Dakota Gruener, estimated in an interview that it would take "roughly . . . $100 million to do the first [outdoor] experiment."
- **Private sector:**
  - Stardust Solutions, a US-Israeli for-profit startup, has raised a total of $75 million from venture capitalists to advance development of its SAI technology, with an ambitious goal of deployment by 2030.[30]

Tractability

Our sense is that the feasibility of further research is high. The physical mechanism is well understood, and the basic engineering is theoretically achievable within 10–15 years (Smith et al., 2024). Tractability is mostly constrained by the absence of an international governance

[30] Stardust's timeline and commercial model have attracted significant criticism from the research community: every researcher interviewed by MIT Technology Review was "adamant that such a deployment should not happen so quickly," (Temple, 2025) and the head of research at Reflective stated that Stardust's published materials "do not increase my confidence that Stardust will contribute meaningfully in the near future" (Talati et al., 2026).

framework, the political sensitivity of unilateral deployment risk, and the gap between models and field validation. Key points in support of deployment feasibility are:

- **Cooling effects would be seen more or less immediately**. SAI would begin suppressing global mean temperatures within months of sustained deployment, making it the "only option that could cool the planet within years" (Buck & Nicholson, 2023).
- **Technical challenges exist, but are broadly surmountable**:
    - Particle size control. Depending on aerosol injection strategies, they can warm the lower stratosphere and increase its water vapor content, which can trap heat and partially counteract the intended cooling. Effects would depend on the quantity, latitude, seasonality, and altitude of SO2 injection, and lower-altitude SAI would also require greater quantities (~60% more) to achieve the same cooling (Lee et al., 2023).
    - **Current-generation aircraft are not viable for high-altitude SAI**. However, developing a new airborne tanker for the job would likely not be technologically difficult or prohibitively expensive. Seasonal low-altitude, high-latitude injections would also be possible with existing aircraft, but such a strategy is less efficient than constant, subtropical applications (Duffey et al., 2025).
- **SAI is inexpensive relative to its climate leverage**. Smith and Wagner (2018) estimate that SAI would cost ~$2.3 billion per million metric tons of sulfur to deploy annually in the first 15 years of operation. While cheap relative to the costs of unmitigated climate change,[31] annual SAI costs at substantial climate impact would still run to approximately $18 billion[32] per degree Celsius of warming avoided, which would be "well beyond the financial reach of individuals, small states, or other non-state potential rogue actors and would instead be the exclusive domain of large national economies or coalitions" (Smith, 2020).

However, fundamental challenges with deployment also exist:

- **A substantial gap remains between model-based confidence and field validation**. No controlled outdoor experiment involving actual stratospheric injection has been conducted at research scale, and one prominent attempt ultimately failed to take place entirely.[33] This inherently limits the Technology Readiness Level (TRL) as high theoretical confidence in SAI's cooling effect has not been matched by progress in understanding how it would actually be implemented, controlled, or reversed.
- **Introduction of a new technology could lead to political conflict**. An anonymous expert emphasized in an interview that "adding a new powerful capacity that some countries can [deploy] in a world that is really very politically unstable is inherently a dangerous thing . . . there are risks of new technology inducing political instability, right up to war."
- **Privatization of investment in SAI potentially deepens mistrust, rather than alleviating it**. Keith and Visioni (2025) wrote a recent op-ed arguing that adding business

[31] The Organisation for Economic Co-operation and Development (OECD) estimates that global gross domestic product (GDP) could fall by 9% (OECD, 2025) as a result of climate damages in a "business-as-usual" scenario.
[32] In 2020 USD.
[33] SCoPEx (Stratospheric Controlled Perturbation Experiment) was a Harvard University project that aimed to release small amounts of particles into the stratosphere via a high-altitude balloon to study how they might reflect sunlight and cool the Earth. The experiment faced repeated delays and its planned 2021 launch in Kiruna, Sweden was ultimately blocked due to community opposition (Saami Council, 2021).

interests, profit motives, and rich investors to the funding landscape for SRM creates more cause for concern for those already mistrustful of SAI. The only way such startups make money is if someone pays for their services, which presents a reasonable fear that financial pressure could drive companies to lobby governments or other parties to use such tools based on financial interests, rather than objective analysis of risks and benefits.

Risks and Side Effects

Despite being the most well-studied SRM approaches, there are substantial uncertainties about SAI deployment, including:

- **Necessity for off-ramps:** If deployment were terminated suddenly, temperatures could rebound toward the levels they would have reached without the intervention (Jones et al., 2013), a process referred to as "termination shock." This could be very damaging for natural and human systems as the rate of warming would probably be much higher than that otherwise expected under anthropogenic climate change, meaning less time for adaptation (Parker & Irvine, 2018).[34,35]
- **Potential failure at the regional level:** Even if SAI successfully stabilizes global mean temperatures, internal climate variability can mask its effects on regional scales. Even with global temperatures stabilized by SAI, internal climate variability can cause continued regional warming: in one ensemble member of a climate model simulation, warming persisted across 49% of the global land surface in the decade following deployment, despite SAI working as intended at the global level (Keys et al., 2022). If local populations or policymakers interpret continued regional warming or extreme weather events as evidence that SAI has failed, this could generate political pressure for premature termination, precisely the scenario in which termination shock becomes most dangerous.
- **Ecosystem responses to altered light and precipitation dynamics:**
    - Depending on deployment strategy, SAI may dry tropical and monsoon regions like the Amazon and South Asia (Asutosh et al., 2025; Cheng et al., 2019).
    - Modeling finds that SAI would result in an abrupt deepening of the Antarctic ozone hole in the first 10 years, with recovery delayed by ~25–55 years (Tilmes et al., 2021).
    - SAI does not reduce the concentration of CO_2 in the atmosphere, which is the root cause of ocean acidification, leaving marine ecosystems vulnerable to harm.
- **Negative human health effects from injecting pollutants into the atmosphere:** Sulfate aerosols settling from the stratosphere could degrade surface air quality by increasing PM2.5 concentrations, while lower temperatures would suppress normal removal of particulates, and stratospheric ozone disruption would increase surface UV-B radiation (Tracy et al., 2022). The emerging evidence on net health effects is mixed. A 2026 modeling study found that maintaining temperatures at 1.5°C through SAI would yield a modest net 0.4% reduction in pollution-related mortality, driven by a 1.3% reduction in ozone-related deaths, but partially offset by a 0.9% increase in PM2.5-related deaths

[34] Liam St Louis from The Navigation Fund noted that: "The risk of termination shock is directly proportional to how much cooling was being done. A termination shock of 0.1°C after the first decade of slowly ramping up deployment would be within the range of natural variation, as a strong El Niño event can already raise global temperatures by more than that."

[35] It is worth noting that termination shock is a governance problem rather than an inherent physical objection to SAI: it can be avoided by ensuring that deployment, once started, is maintained and tapered gradually rather than abruptly ceased. However, designing the institutional frameworks to guarantee this over multi-decadal timescales is itself a formidable challenge.

(Wang et al., 2026). In an interview and related academic paper (Harding et al., 2024), David Keith argues that the health risks would be dwarfed by the climate benefits of reduced warming.[36] However, this could still mask regional disparities — for instance, the PM2.5 burden could fall disproportionately on populations already exposed to high levels of air pollution.

Marine Cloud Brightening and Mixed-Phase Cloud Thinning

Marine Cloud Brightening (MCB) is an SRM technique designed to reflect more sunlight back into space. It works by spraying fine sea salt particles to increase droplet numbers in marine clouds in the lower atmosphere, resulting in higher albedo and greater sunlight reflectivity. MCB could be deployed regionally (e.g., for targeted Arctic or reef protection) or, theoretically, globally. Mixed-Phase Cloud Thinning (MCT) works by seeding mid-to-high altitude clouds with particles that accelerate glaciation, causing the clouds to dissipate. Unlike most other SRM approaches, it targets outgoing infrared radiation rather than incoming sunlight, cooling the Earth by allowing more heat to escape.

Both interventions primarily address the albedo loss and sea level rise risk pathways. While they offer theoretically promising mechanisms to address catastrophic ice loss, both technologies remain in an underfunded, exploratory research phase. Real-world evidence of the benefits and risks of these technologies is scarce, leaving us with low confidence in their viability; yet the field experiments needed to generate such evidence face significant barriers to community approval precisely because that evidence does not yet exist.

Importance

Our take on importance: medium confidence that impact would be moderate. Arctic-specific feedbacks (e.g., sea ice albedo loss, permafrost thaw) represent some of the highest-leverage tipping points in the climate system. If effective, even modest changes over strategic ocean areas might meaningfully reduce Arctic sea ice loss and West Antarctic ice-sheet retreat. Key research includes:

- **Strong observational evidence for MCB proof-of-concept:** Substantial cloud brightening was observed from shipping emissions in the southeast Atlantic, with an effective radiative forcing of approximately −2 W/m² within the shipping corridor (Diamond et al., 2020). The 2020 IMO sulphur regulations, which reduced shipping emissions, provided a natural experiment in the reverse direction: the reduction in unintentional cloud brightening appears to have contributed to warming, though the precise quantitative magnitude of this effect remains under active investigation (Diamond, 2023; Watson-Parris et al., 2025).
- **Models agree that Arctic MCB would substantially cool the Arctic and maintain sea ice:** While the amount of aerosol required to achieve a given cooling varies by up to an order of magnitude across the models, estimated Arctic cooling is in the range of ~3–10°C (Henry et al., 2025). This wide range reflects substantial model disagreement on the magnitude of the effect, even though models agree on its direction. MCB also showed partial AMOC stabilization in two models where it was predicted to decline in strength.

[36] Keith: "I think we can be very confident that we quantitatively know a lot about the risks and that they seem to be small compared to the benefits."

- **Evidence is stronger for regional brightening, with substantial gaps to reach global scale**: Sea-salt aerosols injected near the ocean surface have a very short atmospheric lifetime, typically dissipating within a few days, such that the cooling effect would be highly localized to the injection site (Feingold et al., 2024). Furthermore, MCB requires specific cloud regimes to work, which are not uniformly present.
- **Modeling of seasonal polar MCT shows promise, but real-world evidence is thin.** One study predicted that MCT could offset about 25% of the expected increase in polar sea-surface temperature as a result of CO_2 doubling, accompanied by annual sea ice increases of 8% around the Arctic and 14% around Antarctica (Villanueva et al., 2022).[37] However, little practical work has been done beyond modeling, putting it well behind SAI and even MCB for field-readiness.

Neglectedness

Our take on neglectedness: high confidence that this space is moderately to highly underfunded. Philanthropic and government interest in MCB is growing quickly, meaning the field is no longer as neglected as it was several years ago, but funding is still very modest in comparison to the potential benefits. One expert we interviewed estimated estimated that the current budget for MCB is "probably $10 million a year [or] less" and "my sense would be you'd want to at least triple that . . . [to] accelerate the timeline for potential deployment." MCT receives even less attention than MCB, with only one recently announced initiative specifically focusing on it. Some notable research groups in the field include:

- The National Marine Science Centre of Southern Cross University, which conducts research on using MCB to cool and preserve the Great Barrier Reef in Australia, has an estimated $50 million in funding over the next five years.
- The University of Washington's Marine Cloud Brightening Research Program, funded by a mix of philanthropic foundations and venture capitalists, is a team of over 35 scientists conducting research on cloud brightening, including field experiments.
- The Centre for Climate Repair at Cambridge University and Delft University of Technology Climate Institute (TUDCI) in the Netherlands received joint funding (Centre for Climate Repair, 2023) from the Refreeze the Arctic Foundation in 2023 to investigate methods for conducting MCB.
- The Arctic Stabilization Initiative (ASI), incubated through Renaissance Philanthropy's Advanced Research for Climate Emergencies (ARC) program, recently raised $6.5 million in funding to research climate interventions targeting Arctic warming, with an initial focus on MCT. To our knowledge, ASI is the only dedicated MCT research initiative currently operating, underscoring the extreme neglectedness of this specific approach.

Capacity to absorb additional funding is also a substantial constraint for small fields. For MCT research in particular, on expert said that he believes the "bigger obstacle [is that] there's no capacity to absorb research funding in [the MCT] space at this point." The lack of research infrastructure for MCT is both a challenge and an opportunity; while current capacity limits immediate progress, the field's small scale means that relatively modest philanthropic investments could potentially lead to substantially sped-up development timelines.

[37] It is important to note that the MCT evidence base rests almost entirely on this single modeling study, and that the underlying physics of mixed-phase clouds, which are prevalent in the Arctic and contribute significant uncertainty to predicted radiative effects, remain poorly understood (Zamora et al., 2025). Confidence levels for MCT should therefore be understood as very low.

Tractability

Our sense of tractability in this space is medium to high. MCB's delivery mechanism is relatively simple and theoretically low-cost, but cloud system dynamics are poorly understood, and work to develop the required nozzles for deployment has so far been unsuccessful. For MCT, operational similarity to existing cloud-seeding programs is an advantage, and delivery infrastructure is well understood. However, efficacy at the scale needed to meaningfully cool the Arctic has not been demonstrated, and atmospheric modeling required to predict second-order effects remains immature. Key points in favor of feasibility include:

- **A recent outdoor experiment showed real-world capability of sea salt aerosol dispersion.** In the first outdoor field demonstration of MCB, researchers showed that sea spray aerosols generated from a moving vessel can travel upward and reach cloud base heights of 700–900 meters, providing proof-of-concept that real-world cloud perturbation experiments are technically feasible (Hernandez-Jaramillo et al., 2025).
- **Both interventions are rapidly reversible, reducing termination risk:** MCB salt particles wash out within days to weeks (Feingold et al., 2024); MCT aerosols even faster (Villanueva et al., 2022), making them lower risk to test than more persistent interventions.
- **Both may face fewer governance barriers than SAI**. Cloud seeding already operates in nine US states and in China, providing existing regulatory frameworks for MCT. For MCB, localized interventions are explicitly excluded from the proposed Solar Geoengineering Non-Use Agreement (Biermann et al., 2022), a potentially significant distinction for research and small-scale deployment.

On the other hand, challenges to tractability include:

- **Technical capability to make the required nozzles for MCB:** Ben Margetts (Climate Hub) emphasized the large gap in current technology and the lack of research capacity, saying that MCB "needs a nozzle that's an order of magnitude better at spraying sea salt into the air. MXB has only a handful of PhDs in the world who work on it. There's no team of engineers that do accelerated technological development." Another anonymous expert underscored that despite substantial existing work on the problem, the gap in nozzle technology remains a fundamental barrier.
- **Decades of research could be needed to assess the technical feasibility of MCB**: Feingold et al. (2024) outline necessary research that spans particle generation and delivery mechanisms, field experiments, laboratory facilities, modeling, regional climate impacts, detection and monitoring, and scale-up assessments, meaning that the necessary timeline for building empirical evidence could last decades.
- **Multilateral cooperation, liability, oversight, and community buy-in:** Most effective MCB targets would be over international waters, potentially requiring complex multilateral agreements. Moreover, deployment would require sustained fleets of vessels over vast ocean areas. As one expert said, "the bigger obstacle to MCB being useful is the politics . . . you really want to be able to [replace] the brightening of the shipping containers, which is a global deployment for which you need very different politics."
- **Indigenous rights and free, prior, and informed consent (FPIC).** Arctic-targeted interventions such as MCT and regional MCB would operate in or near territories where Indigenous peoples hold legal rights, traditional knowledge, and deep cultural ties to the ecosystems these technologies aim to modify. Any research or deployment in these regions would therefore need to secure the free, prior, and informed consent of affected Indigenous communities, consistent with the UN Declaration on the Rights of Indigenous Peoples (UNDRIP). This is not only an

ethical requirement but a practical one: as the Alameda MCB episode illustrates below, interventions that proceed without meaningful community engagement face significant risk of being halted.

An attempted 2024 field experiment in MCB led by the University of Washington in Northern California was quickly canceled (Bush, 2024) after city officials intervened, citing a lack of transparency and community engagement similar to the failed SCoPEx trial (see footnote 33). This episode illustrates concretely how the governance and public trust challenges described above can halt even small-scale, well-intentioned research, and suggests that building community buy-in and transparent engagement processes is a prerequisite for, not a parallel track to, field experimentation.

Risks and Side Effects

- **Risk of harmful weather pattern changes:** Because MCB cools the planet heterogeneously, it could alter large-scale atmospheric and oceanic circulation patterns, particularly if applied on a global scale (Feingold et al., 2024). Chen et al. (2025) found that cloud seeding over 5% of ocean area, while inducing 1°C of cooling, could also lead to an unintended "La Niña-like pattern of response" in both temperature and precipitation.
- **MCB could significantly impact ozone levels:** Some models also showed an unexpected phenomenon: changes in atmospheric circulation patterns and chemical processes, triggered by MCB deployment, could deplete both stratospheric and tropospheric ozone (Bednarz et al., 2025).
- **Potential for time-limited impact:** Wan et al. (2024) found that regional MCB that effectively reduced dangerous heat exposure under present-day conditions became largely ineffective or even counterproductive under mid-century warming due to AMOC weakening, implying that delayed research timelines could reduce the period during which MCB remains a viable intervention.

Ice Sheet Stabilization: Basal dewatering and thermosiphons

Above, we discussed that West Antarctic Ice Sheet (WAIS) retreat is plausibly underway, and that by 2100 SLR is likely to exceed baseline projections. By way of summary, we uncovered various plausible opportunities and strategies that seem relevant for investment and for organizations and individuals looking to arrest ice-sheet retreat.

The challenge is that glaciers like WAIS are essentially sliding into the ocean on a pressurized layer of meltwater that acts as a lubricant (Fricker et al., 2025). Removing some of that water should, in principle, raise basal friction and slow the slide of the glacier. Two methods are typically discussed: pumping water out through boreholes or freezing it in place (Climate Interventions, 2024) using passive thermosiphons (sealed pipes that leverage the surface/basal temperature gradient to extract heat without external power; Lockley et al., 2020).

To begin, we identified a series of potential candidate approaches. Lockley et al. (2020) appear to be the first to catalog geotechnical approaches to stabilize glaciers: seabed curtains or sills to block warm Circumpolar Deep Water from reaching grounding lines; basal dewatering or bed freezing to increase friction and slow ice flow; shear-margin stiffening with frozen ice; ice-shelf thickening via surface accumulation; and sub-shelf physical obstacles (anchors, artificial pinning points). Local SAI and MCB are also sometimes discussed, but we exclude them from

this section since they were already addressed above. We narrow our review here to **basal dewatering/thermosiphons**, because it is our impression that this approach (alongside seabed curtains) is one of the most promising approaches: it has some scientific attention, has an institutional home (Arête Glacier Initiative), and we qualitatively evaluate that these interventions have the capacity to absorb funding in the millions of dollars over the next five to 10 years; a key time period that could decide whether these technologies are promising enough to launch at scale.[38] Moreover, it seems less controversial than seabed curtains in terms of its potential ecological harm.

Importance

Our take on importance: high confidence that the underlying problem is important; medium confidence in the intervention's ability to address the problem. Our BOTEC estimates the SLR contribution we expect by 2100. Brent Minchew at Arête Glacier Initiative frames a more pessimistic lower bound: in materials shared with our team, he writes that "more than one meter of sea-level rise is likely, with more than two meters within the plausible range" based on observationally calibrated sensitivities, and that IPCC projections "consistently fall below observed trends" (Minchew, B., "Sea level forecasts, simplified" (draft).[39]

We are more unsure about the ability of the intervention to work at the magnitude required. We were unable to locate a published amount of SLR avoided for basal dewatering or thermosiphons. An assessment by UArtic Climate Interventions is pessimistic about the ability of its scalability, timeliness, and impact, though the data and calculations behind the scoring are not published.[40] The first peer-reviewed physics paper on water-pumping (Meyer et al., 2026) and forthcoming papers on thermosiphons should partially close this evidence gap over the coming months. Our sense is that there is simply a lack of information regarding the actual likely impacts of these technologies at scale, representing a significant research gap rather than a fundamental flaw of the approach.

Moreover, Fricker et al. (2025) provide some empirical evidence that subglacial lake drainage can measurably increase ice flow, so it makes sense that stopping drainage would reduce ice flow. However, this inference has not been directly tested. The relationship between subglacial water and ice dynamics is nonlinear and complex: a comprehensive review of subglacial lakes found that their behavior is conditioned by the hydrological, dynamic, and mass balance regime of the overlying ice mass, with drainage events producing variable and sometimes counterintuitive effects on flow (Livingstone et al., 2022).

More recently, the first continent-wide subglacial hydrology modeling found extensive channelization beneath both Pine Island and Thwaites Glaciers, with discharge fluxes of 45 and 69 m^3/s, respectively (Ehrenfeucht et al., 2025), underscoring how interconnected and dynamic these systems are. Removing water at one point in such a network could have currently unpredictable effects elsewhere, which is a reflection of the lack of data and experimentation rather than a fundamental physical limit. Developing a deeper understanding of subglacial hydrology is clearly a precondition for any intervention and would improve our ability to

[38] A key input to us forming this view is this white paper (MacAyeal et al., 2024) from the University of Chicago that highlights both seabed curtains and basal-hydrology interventions as worthy of attention.
[39] See also IPCC AR6 WG1 Chapter 9: Ocean, Cryosphere and Sea Level Change, whose SSP5-8.5 likely-range upper bound is 1.01. Minchew's framing puts our central BOTEC estimate at the conservative end of plausible (Fox-Kemper et al., 2021)
[40] The more detailed Lockley et al. (2024) treatment is a useful reference, but leaves much of the dewatering and thermosiphon analysis without supporting data.

forecast system behavior. In short, we think that producing further evidence on these dynamics could either speed up deployment or, if no evidence is found, help pivot resources to other more promising interventions. Indeed, a University of Chicago white paper acknowledges this uncertainty but makes the case for research into basal dewatering without explicitly guaranteeing its feasibility.

Neglectedness

Our take on neglectedness: high confidence that this space is relatively neglected. The Arête Glacier Initiative appears to be the main player. It is a young, small-budget nonprofit. Disclosed funders are Outlier Projects, Kissick Family Foundation, Grantham, and The Navigation Fund.[41] Based on disclosed funders and visible operating scale, our estimate is that cumulative philanthropic deployment to date is under $20 million.

After additional Google and LLM-supported searches, we found very little activity in stabilization engineering. The UChicago white paper cited above came out of the first-ever workshop on these ideas in October 2023, and the first peer-reviewed paper on dewatering physics has only just been accepted (Meyer et al., 2025). We identified a few major precedents for subglacial drilling on the Antarctic ice sheet, but they all seem to be related to observational science, namely trying to sample microbes. Our sense from talking to experts is that the intervention field is small and deliberately careful: a handful of scientists, working slowly because of the controversial nature of the topic. In short, it's likely that at the margin, additional grants can help produce high-impact findings both by funding new research and by enabling the constructive disagreement that the field's current size makes hard to sustain.

Tractability

Our sense of the tractability for this type of intervention is medium to high, since developing the technology and doing some pilots is a feasible engineering challenge, and in any case, the research to show whether it is possible is even more tractable.

Ice thickness in the fast-flowing trunks of the WAIS outlet glaciers we care about is ~800–1,200 m at Thwaites. The interior West Antarctic basins where the Lake Ellsworth borehole failed were deeper (~3 km), but that is more than the required target for outlets.

The Subglacial Lake Ellsworth program cost approximately US$13 million in 2012 (~US$18M in 2024 dollars) for a single ~3 km borehole that failed (Siegert et al., 2025). One researcher we spoke with suggested that an instrumented borehole drilled specifically for an intervention test could be roughly an order of magnitude cheaper. If so, the tens-to-hundreds of holes estimated by Siegert et al. (2025)[42] would cost in the tens of millions rather than the hundreds, bringing early-stage research within reach of a focused philanthropic program, though full deployment at scale would still require government partnership.

Before any large-scale pilot, smaller pilots and general ice-sheet/glacier research (which we think is more relevant to marginal philanthropic dollars) would be needed to understand

[41] We note that The Navigation Fund and Outlier Projects, who commissioned this report, are among Arête's disclosed funders; we do not believe this affects our assessment, but flag it for transparency.
[42] We note that Siegert et al. (2025) themselves are skeptical of the feasibility of glaciological intervention at this scale, concluding that the approaches they assessed are "environmentally dangerous" and "not feasible." This is a view we have engaged with in our earlier discussion of arguments against Earth stabilization approaches. Their cost data and engineering assessments are nonetheless valuable for scoping the scale of the challenge.

whether deploying these solutions is possible. Given that current cumulative philanthropic deployment in this space is under $20 million, and that the Arête Glacier Initiative and the University of Chicago's Climate Systems Engineering initiative represent the primary institutional homes for this work, we estimate that a 10-year research program to resolve key feasibility questions could plausibly be funded at the scale of $10 million per year for 10 years (~$100–$120M). We view this as a modest sum relative to the potential consequences of WAIS collapse, and think it is plausible that such a program could meaningfully resolve some of these questions.

Risks and Side Effects

Our understanding is that boreholes have very local effects and have limited biological effects, though there is little research into the microbes in basal areas. Siegert et al. (2025) and others flag, however, that subglacial drainage networks are complex and poorly characterized, and removing water in one location may redirect flow and accelerate sliding elsewhere. This makes sense and is likely to command significant attention from researchers before deployment.

Contributions and Acknowledgments

Abbie Clare, Ruby Emerson, Britta Jewell, and Thomas R Vargas jointly researched and wrote this report. Tom Vargas served as Project Lead. Abbie Clare also served as the Project Manager. John Firth supervised the project as Director.

Special thanks to the experts who took the time to speak with us, share data, and provide their insights: Melissa Barnes, Charlotte DeWald, Phil Duffy, Dakota Gruener, David Keith, Marc Macias-Fauria, Ben Margetts, Brent Minchew, Erika Reinhardt, Liam St Louis, and additional unnamed experts.

Special thanks to Shane Coburn for copyediting, to Thais Jacomassi for bibliography support, and to Elisa Autric for publishing the report online and assisting with dissemination.

Outlier Projects provided funding for and commissioned this report, but it does not necessarily endorse our views or conclusions. Furthermore, the experts we interviewed and the organizations with which they are affiliated do not necessarily endorse the views or conclusions expressed herein.