Try Planning! But Whose Plan?
The hidden choices in planning for climate.

Alice Lépissier, PhD
Green Finance AI | Plan4Planet Organizing Committee

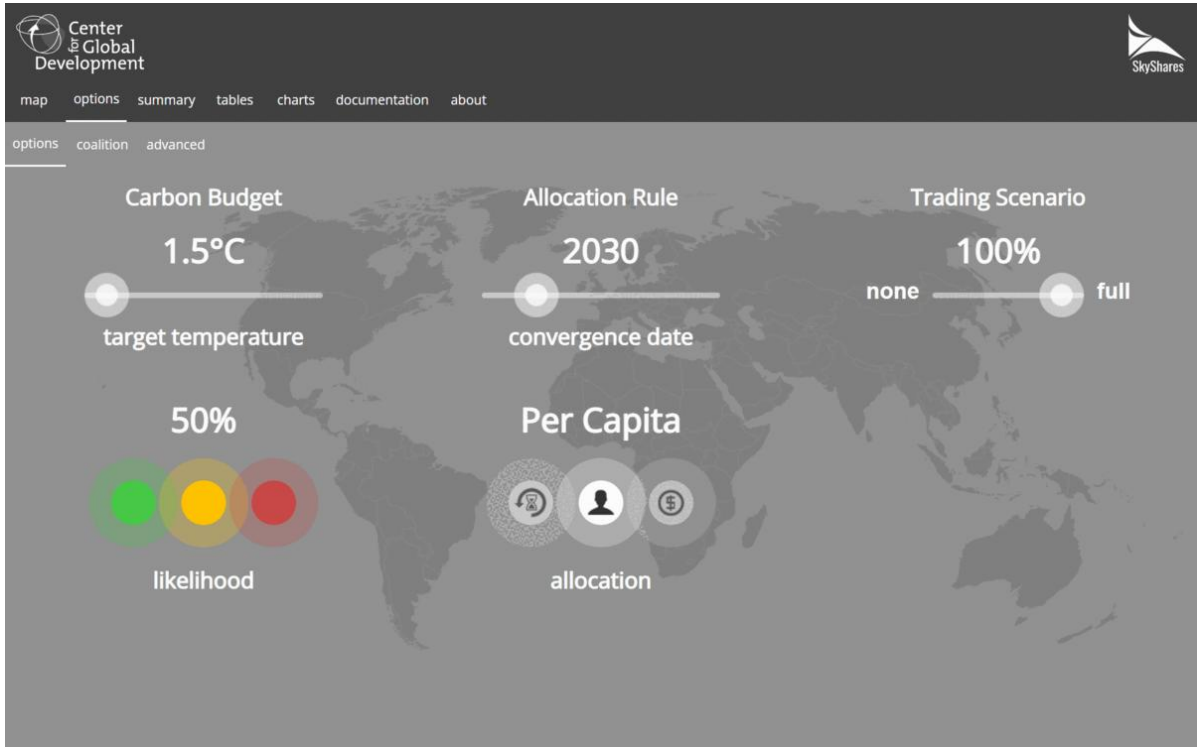
Planning for a Better Planet, AAAI Fall Symposium
5 November 2026



Plan4Planet
AI planning for climate decisions.



Every modeling choice turned out to be a fairness choice



SkyShares, Center for Global Development, 2013 (https://skyshares.io)

SkyShares

CGD CENTER FOR GLOBAL DEVELOPMENT

Pick a rule for sharing the global carbon budget: equal per person, by today's emissions, or converging over time.

See what every country's bill comes to.

Change one parameter and Nigeria's cost triples.

Planning sits downstream of estimation and prediction

Estimation
econometrics
What is the effect?
Output: a causal parameter

Prediction
machine learning
What will happen?
Output: a forecast

Control
planning
What should we do?
Output: a policy

The first two build the model. The third decides.

Everything they get wrong flows into the plan, and the third adds a choice the other two never have to make: **whose welfare counts.**

Climate change is a "wicked" problem

- No definitive formulation: how you frame it already decides what counts as a solution
- No stopping rule: only better or worse, never right or wrong
- Every attempt is a one-shot operation: it changes the world and cannot be undone

“The planner has no right to be wrong.”

Rittel and Webber, 1973

Climate is super wicked

- Time is running out
- The people who caused it are in charge of fixing it
- Our institutions discount the future

Levin et al., 2012

Climate change is three problems in one



A physical problem

Warming is set by cumulative emissions. There is a finite carbon budget left, and every ton anyone emits is a ton someone else cannot.



A collective-action problem

The benefits of cutting emissions are global and spread over centuries. The costs are local and land this decade.



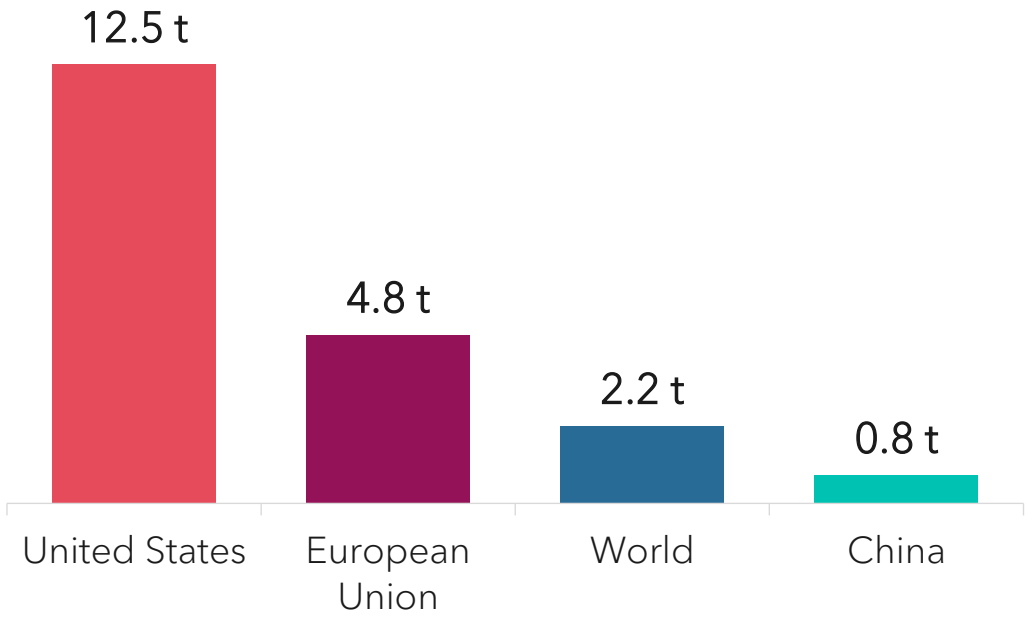
A justice problem

Those who caused it, those who suffer most, and those who can pay are three different groups.

The justice problem has three faces: across people, across time, and across uncertainty.

The people least responsible are the most exposed

CO2 per person per year, average since 1751



The whole of Africa: about 3 percent of cumulative emissions.

3.3 to 3.6 billion people live in places highly vulnerable to climate change, overwhelmingly in poorer countries (IPCC, 2022)

The 1992 UN climate treaty's answer:

Common but differentiated responsibilities

- Everyone must act, but not equally.
- Rich countries cut first
- Rich countries pay for adaptation in poor countries
- Since 2022, a fund for loss and damage: harms that can no longer be adapted to

How much should the future count?

Stern Stern Review, UK Treasury, 2006

A near-zero rate of pure time preference: treat future people almost as equal as us

→ Cut hard, now

Nordhaus Nobel Prize in Economics, 2018

Rates closer to what markets pay

→ Ramp up slowly

The discount rate converts harm in 2100 into a value today.

The entire disagreement between Stern and Nordhaus was ethical, and it lived in one parameter.

In the standard planning model (the Markov decision process, where rewards are added up over time), that parameter is gamma, and it is set so the sum converges.

Climate is Knightian uncertainty, not risk

Risk vs. uncertainty Frank Knight, 1921

Risk: the probabilities are known. A dice roll, an insurance table.

Uncertainty: the probabilities cannot be assigned. The system is changing underneath you.

Climate is deep uncertainty: no agreement on the model, the probabilities, or how to value the outcomes.

Three ways to decide

Expected value: best on average, given the odds

Maximin: make the worst case as good as possible

Minimax regret: least chance of kicking yourself afterward

The planner's default is expected value: the one criterion that needs the odds. Under uncertainty, it is a bet on probabilities nobody has.

Every planner aggregates along the same three axes

	In the objective	The choices	What climate calls it
Across people	Social welfare function	Utilitarian, prioritarian, or protect the worst off	*Justice across space*
Across time	Discount rate	Stern vs. Nordhaus; γ in your MDP	*Justice across generations*
Across states of the world	Decision criterion	Expected value, maximin, minimax regret	*Deep uncertainty*

Underneath all three sits the unit: dollars, structures, hectares, lives. You cannot write an objective without taking a position on every row.

The planner's defaults are normative positions

Σ	One scalar reward, added up with equal weights	→ **That is utilitarianism**
$\mathbf{E}[\,]$	An expectation over uncertainty	→ **That assumes the probabilities can be trusted**
γ	A discount factor chosen for convergence	→ **That is a verdict on how much the future matters**

They look like technical settings. They are ethical choices made by default.

A famous thought experiment. Have your plans already taken a side?



The veil of ignorance Rawls, 1971

Design the rules of society before you know who you will be in it: rich or poor, born now or in 2100. What rules would you choose?

Rawls

You would protect the worst off, because you might be them

→ **Maximin**

Harsanyi

You would maximize your expected outcome, because you are equally likely to be anyone

→ **Expected value → utilitarianism**

Every expected-reward objective is a vote for Harsanyi. Did you mean to cast it?

The unit you optimize decides whose losses matter

- Minimize structures lost: a mansion and a trailer count the same
- Minimize dollars lost: the crews move to the wealthy side of town

Hallegatte and colleagues, World Bank, Unbreakable

Count disasters in well-being rather than in assets and the picture inverts. A poor household losing its roof suffers more than a rich household losing a car, even though the car is worth more.

Optimizing an average hides the tail.

Distributional justice, in statistical terms, is attention to the lower quantiles, and maximin is optimization of the bottom one.

Two kinds of humility, and planners already have one of them

Humility about the model

You already do this

- Admit the probabilities are a range, not a point (Manski, partial identification)
- Robust planning, risk-sensitive objectives
- Plan over a set of models, not one
- Adaptive pathways with monitored triggers

Humility about the objective

What climate and social science people bring

- Say whose welfare you weighted, and by how much
- Show what changes under a different ethics
- Prefer plans that hold up across several
- Keep a record someone affected can contest

Goal 4 of this symposium: surface the tensions that formal reasoning cannot resolve on its own.

Five questions to carry through the next three days

1. Whose losses does your objective count, and with what weights?
2. How much does the future count, and who chose γ?
3. If you swapped expected value for minimax regret, would your plan change?
4. Where do your probabilities come from, and would you bet on them under a changed climate?
5. Which decisions should never be fully automated?

Methodological humility about the model. Normative humility about the objective.

Final thought:
The objective function is where your ethics lives. Write it on purpose.