International Institute for Applied Systems Analysis
IIASA www.iiasa.ac.at

A Bayesian-inferred physical module to estimate robust mitigation pathways with cost-benefit IAMs

Thomas Gasser[1], Artem Baklanov[1,2], Armon Rezai[1,3], Côme Chéritel[1,4], Michael Obersteiner[1,5]

[1] International Institute for Applied Systems Analysis, Laxenburg, Austria (gasser@iiasa.ac.at)
[2] Higher School of Economics University, Saint-Petersburg, Russia
[3] Department of Socio-Economics, University of Economics and Business, Vienna, Austria
[4] École Polytechnique, Palaiseau, France
[5] Environmental Change Institute, University of Oxford, Oxford, UK

Work supported by EU H2020 (grants #773421 and #820829) and FWF (grant P31796-N29)

A Bayesian-inferred physical module to estimate robust mitigation pathways with cost-benefit IAMs

Main results:

	DICE2016 GAMS	DICE2016 python	Continuous formulation	New climate module (prior mean)	+ New carbon module (prior mean)	+ Permafrost module (prior mean)	Prior median [90% range]	Posterior median [90% range]	Robust optimization
Peak ΔT (°C)	4.1	4.1	3.9	3.4	3.2	3.2	3.2 [2.6–4.0]	3.0 [2.4–3.5]	2.9 [2.0–4.0]
SCC in 2020 (USD / tCO2)	37	37	36	40	35	36	35 [18–68] mean=38	26 [13–43] mean=27	27
SCC under 2°C constraint	impossible				140	166		67 [18–287] mean = 95	97 (@50%) 134 (@66%)

Context: The DICE model

DICE is a cost-benefit IAM. It is a simple climate-economy model.

Causality: *Production* is split between *Consumption* (to create *Welfare)*, *Investment* (to sustain future *Production*), and *Abatement* (to reduce *Emissions* and prevent future *Damage* that would reduce future *Production*).

Concept: The model solves for two control variables (mitigation rate, and savings rate) so that cumulative discounted *Utility* is maximized, finding the optimum between cost of future *Damage* and cost of present *Abatement*.

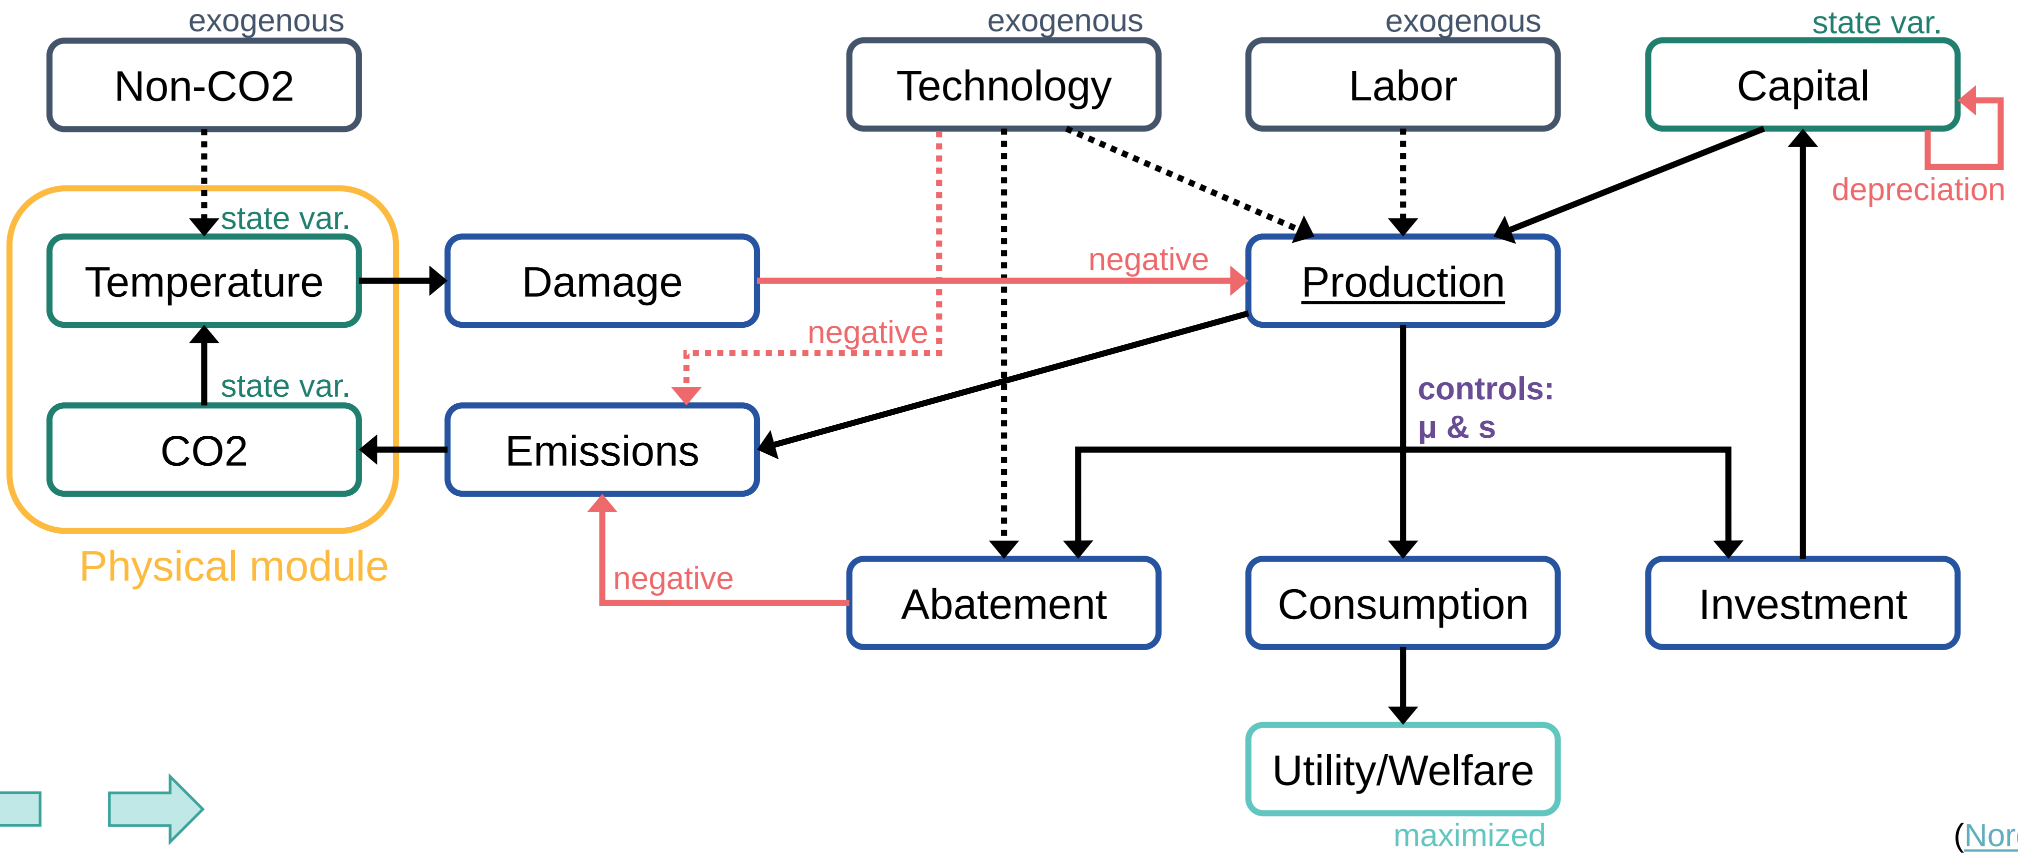

(Nordhaus, 2017)

Context: Social Cost of Carbon

Definition: Net present value of aggregate costs from one more tonne of carbon in the form of carbon dioxide (CO_2), conditional on a global emissions trajectory over time.

This includes absolutely all costs (as represented in a model), and notably:

- Costs from climate damages
- Gains from not mitigating the one tonne now
- Costs from mitigating it later on (if e.g. a temperature target must be reached)

Significance: "The most important single economic concept in the economics of climate change is the social cost of carbon (SCC). At present, regulations with more than $1 trillion of benefits have been written for the United States that use the SCC in their economic analysis." (Nordhaus, 2017)

Use: In a policy-making context, the SCC is used to monetize the hidden costs (negative externality) of emitting CO_2. It is a marginal metric (somewhat similar to GWP, but for the economic system).

For instance, if installing wind farms costs 25 USD per t of CO_2 avoided, and the SCC is estimated to be 40 USD $tCO2^{-1}$, then the net cost of the wind farms is -15 USD $tCO2^{-1}$, i.e. a net global economic gain.

Calculation: In DICE, it is technically calculated as the ratio of two Lagrange multipliers obtained after optimization. The numerator is the marginal impact on welfare of a unit of CO_2 emitted at time *t*. The denominator is the marginal impact on welfare of a unit of total consumption at *t*.

In other words, if along the optimal path an emission of one tonne of CO_2 occurs at time *t*, and one removes *x* USD to the consumption path also at *t*, such that the welfare value (i.e. the cumulative discounted utility) remains the same as in the optimal path, then *x* is exactly the SCC at time *t*.

Physical modelling : Climate module

Original DICE formulation:

$$T_{AT}(t) = T_{AT}(t-1) + \xi_1\left(F(t) - \xi_2 T_{AT}(t-1) \ - \xi_3\left(T_{AT}(t-1) \ - T_{LO}(t-1)\right)\right)$$

$$T_{LO}(t) = T_{LO}(t-1) + \xi_4\left(T_{AT}(t-1) \ - T_{LO}(t-1)\right)$$

More correct (continuous) formulation:

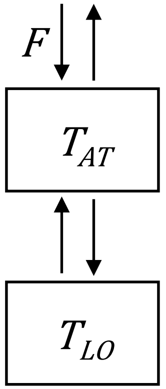

$$\frac{\delta T_{AT}}{\delta t} = \frac{\xi_1}{\delta t}\left(F(t) - \xi_2 T_{AT}(t) \ - \xi_3\left(T_{AT}(t) \ - T_{LO}(t)\right)\right)$$

$$\frac{\delta T_{LO}}{\delta t} = \frac{\xi_4}{\delta t}\left(T_{AT}(t) \ - T_{LO}(t)\right)$$

Typical formulation in climate science:

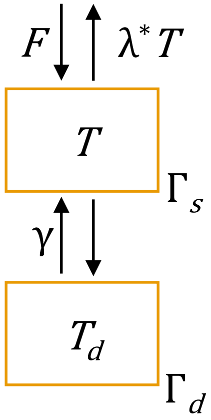

$$\Gamma_s \frac{\mathrm{d}T}{\mathrm{d}t} = F(t) - \lambda T(t) + \epsilon\gamma\left(T_d(t) \ - T(t)\right)$$

$$\Gamma_d \frac{\mathrm{d}T_d}{\mathrm{d}t} = \gamma\left(T(t) \ - T_d(t)\right)$$

The two systems are <u>equivalent</u>, though the (very wrong) <u>parameters need an update</u>:

	DICE2016	CMIP5	
$\Gamma_s = \delta t/\xi_1$	50	8.2 ± 0.9	W yr m^{-2} K^{-1}
$\lambda = \xi_2$	1.19	1.18 ± 0.37	W m^{-2} K^{-1}
$\epsilon\gamma = \xi_3$	0.09	0.87 ± 0.28	W m^{-2} K^{-1}
$\epsilon\Gamma_d = \delta t\, \xi_3/\xi_4$	18	134 ± 46	W yr m^{-2} K^{-1}

(Geoffroy et al., 2013)

Physical modelling: Ocean carbon module

Original DICE formulation (ocean + land):

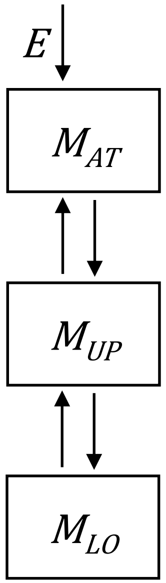

$$M_{AT}(t) = E(t) + \phi_{11} M_{AT}(t-1) + \phi_{21} M_{UP}(t-1)$$
$$M_{UP}(t) = \phi_{11} M_{AT}(t-1) + \phi_{22} M_{UP}(t-1) + \phi_{32} M_{LO}(t-1)$$
$$M_{LO}(t) = \phi_{23} M_{UP}(t-1) + \phi_{33} M_{LO}(t-1)$$

This is a linear system, which is inadequate to describe saturating effects and climate feedbacks.

Solution: explicitly represent key non-linear processes (and thus separate ocean and land)

Step 1: Ocean C

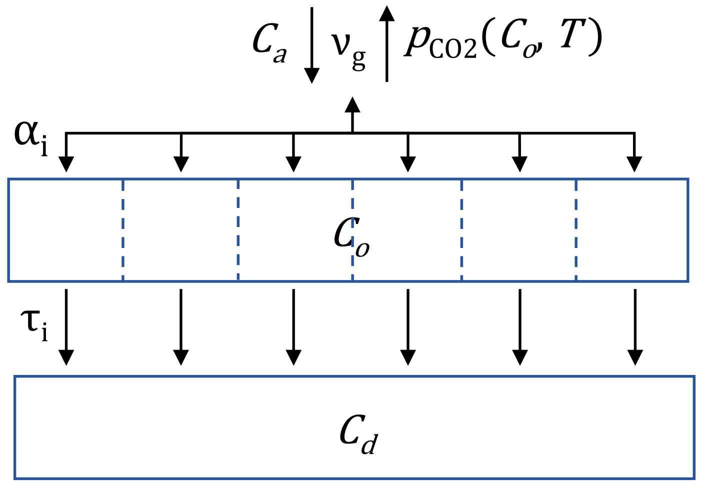

(Joos et al., 1996)

Key processes:

- Non-linear carbonate chemistry emulated with p_{CO2}, incl. climate feedback
- Complex (linear) dynamic obtained with 6 timescales (τ_i)

Physical modelling: Land carbon module

Step 2: Land C

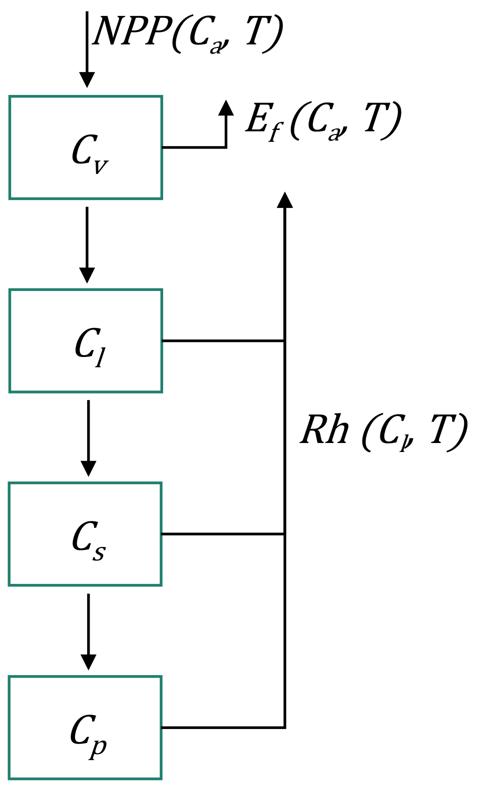

(Gasser et al., 2017)

Key processes:

- CO_2-fertilization of photosynthesis, incl. climate feedback
- Wildfire emissions
- Climate-dependent respiration, with "priming" effect
- Passive (= very slow) soil carbon pool
- Permafrost thaw (➔)

But land-use ignored!

Step 3: Permafrost C

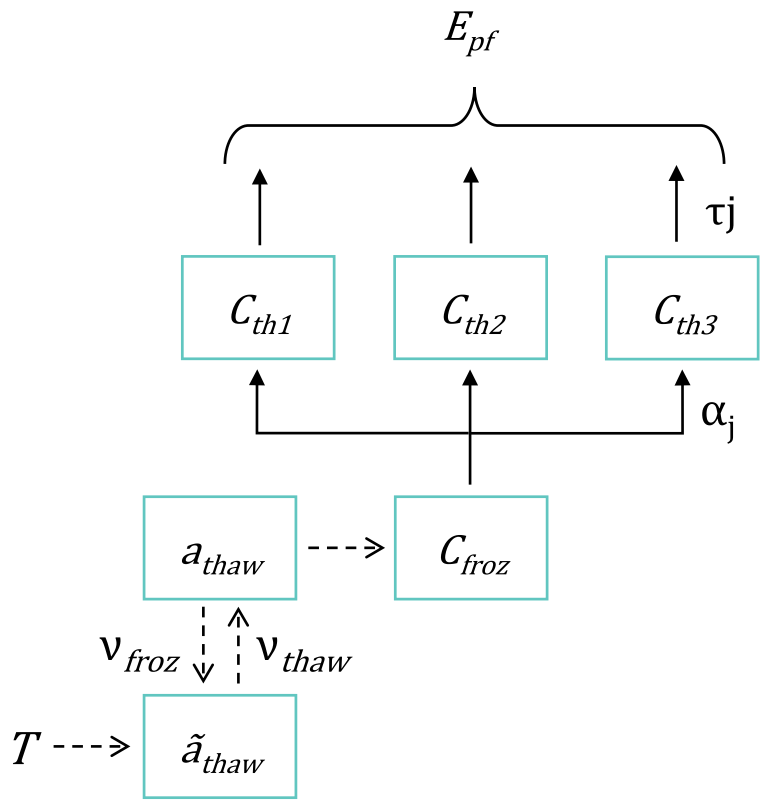

(Gasser et al., 2018)

Bayesian calibration : Overview

Principle: assimilate information from prior knowledge on parameters and observations, to derive posterior knowledge on both (➔).

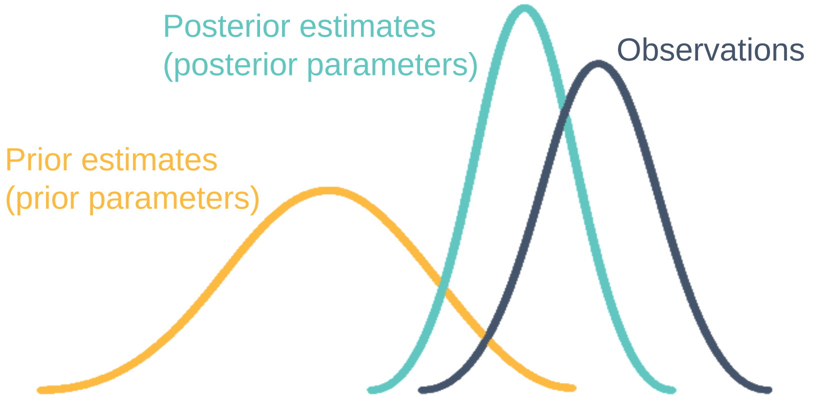

Prior distributions of 35 parameters (out of 63):

- 27 estimated with OLS regression on TRENDY or CMIP5 models, then taking multi-model average and standard deviation (16 models for climate, 11 for land preindustrial carbon, 7 for carbon transient responses)
- 8 taken from the literature
- assumed prior distribution depending on support: N if (-∞, +∞), Log-N if [0, +∞), Logit-N if [0, 1]

Prior drivers of the model (GSAT & CO_2):

- assumed auto-correlated timeseries (auto-correlation also estimated during calibration)

Prior distributions of 14 observations (= constraints):

- taken from the literature to cover all components (climate, land C, ocean C)

Bayesian inference:

- run with a full-rank ADVI (Automatic Differentiation Variational Inference) algorithm
- implemented in python with PyMC3 package

Bayesian calibration : Results (observations)

GSAT & CO_2 timeseries (inputs to the model):

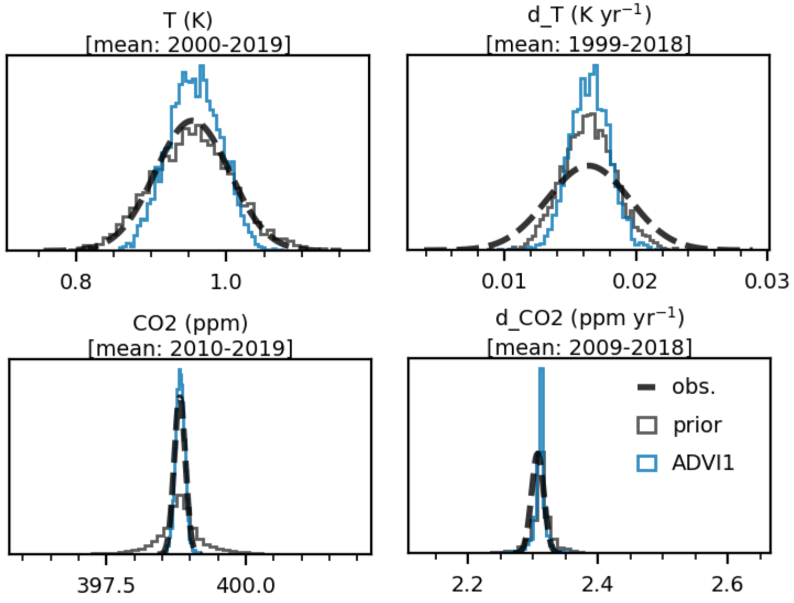

<u>Climate</u>: non-CO_2 ERF notably lower than IPCC:

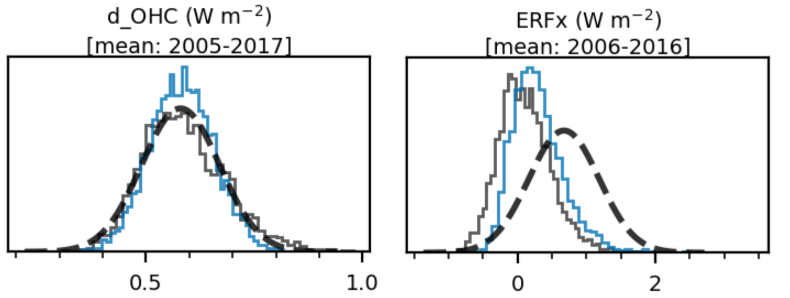

<u>Land C</u>: sink matches Global Carbon Budget, but TRENDY and IPCC estimates of preind. veg. carbon incompatible:

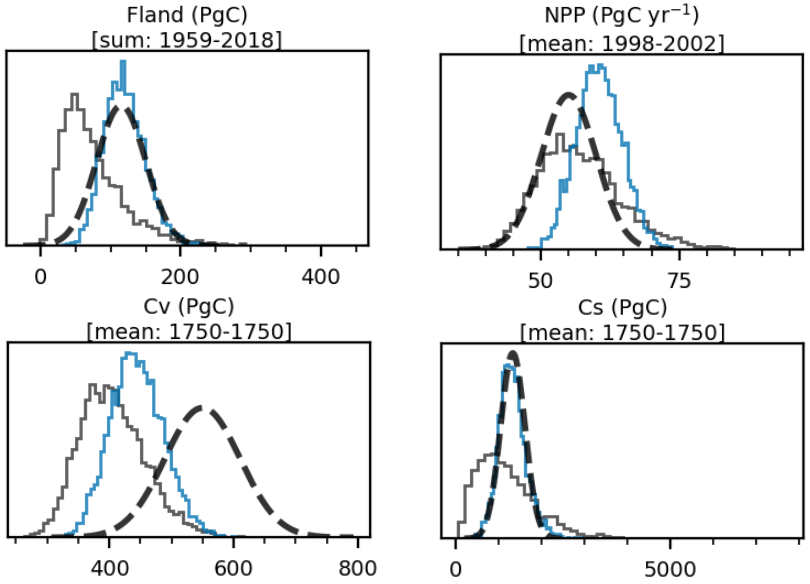

<u>Ocean C</u>: structure cannot match Global Carbon Budget, which is compensated by lower compatible CO_2 emissions:

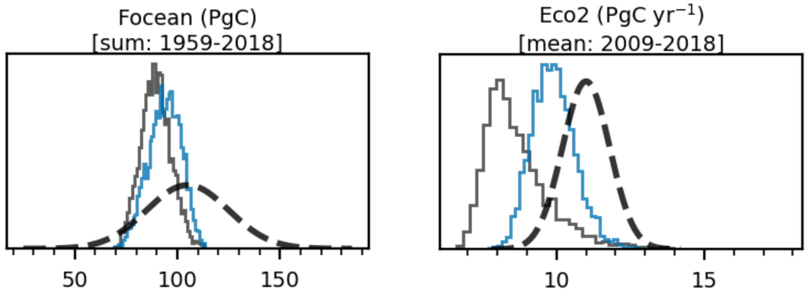

Bayesian calibration : Results (parameters)

Though many parameters only slightly affected by the calibration. Some significantly are, however.

Climate: extreme ECSs ($\mathrm{T2x}$) and slow timescales (THd) excluded:

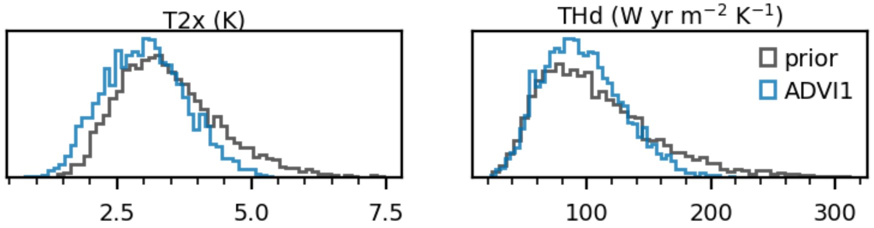

Land C: CO_2-fertilization (bnpp) increased to match GCB sink, preind. NPP ($\mathrm{npp0}$) adjusted, veg. turnover (vmort) reduced to increase preind. pool:

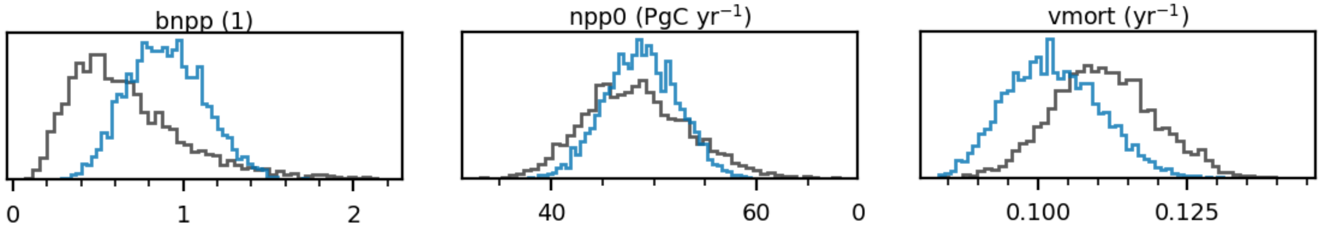

Ocean C: mixing layer depth multiplier (bdic) increased:

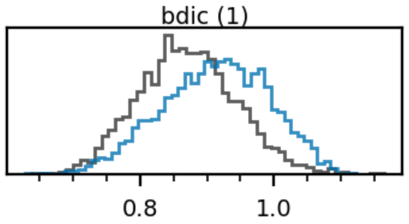

Additionally, the full-rank ADVI algorithm finds correlations (🡻) among parameters:

- anti-correlation of CO_2 radiative efficiency (phi) and ECS ($\mathrm{T2x}$)
- correlation of ECS and deep ocean heat uptake efficacy (eheat)
- anti-correlation of preind. NPP ($\mathrm{npp0}$) and fertilization factor (bnnp)
- correlation between soil C turnover times (vmet, $\mathrm{vrh1}$, $\mathrm{vcs2}$)

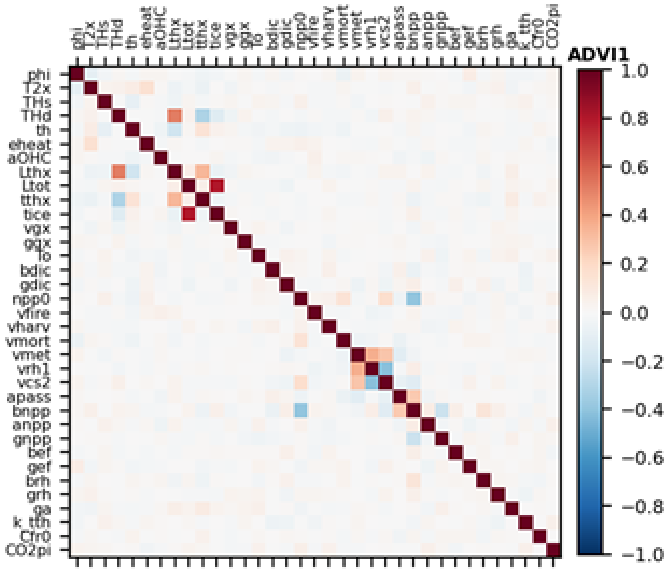

Bayesian calibration: Comparison to ESMs

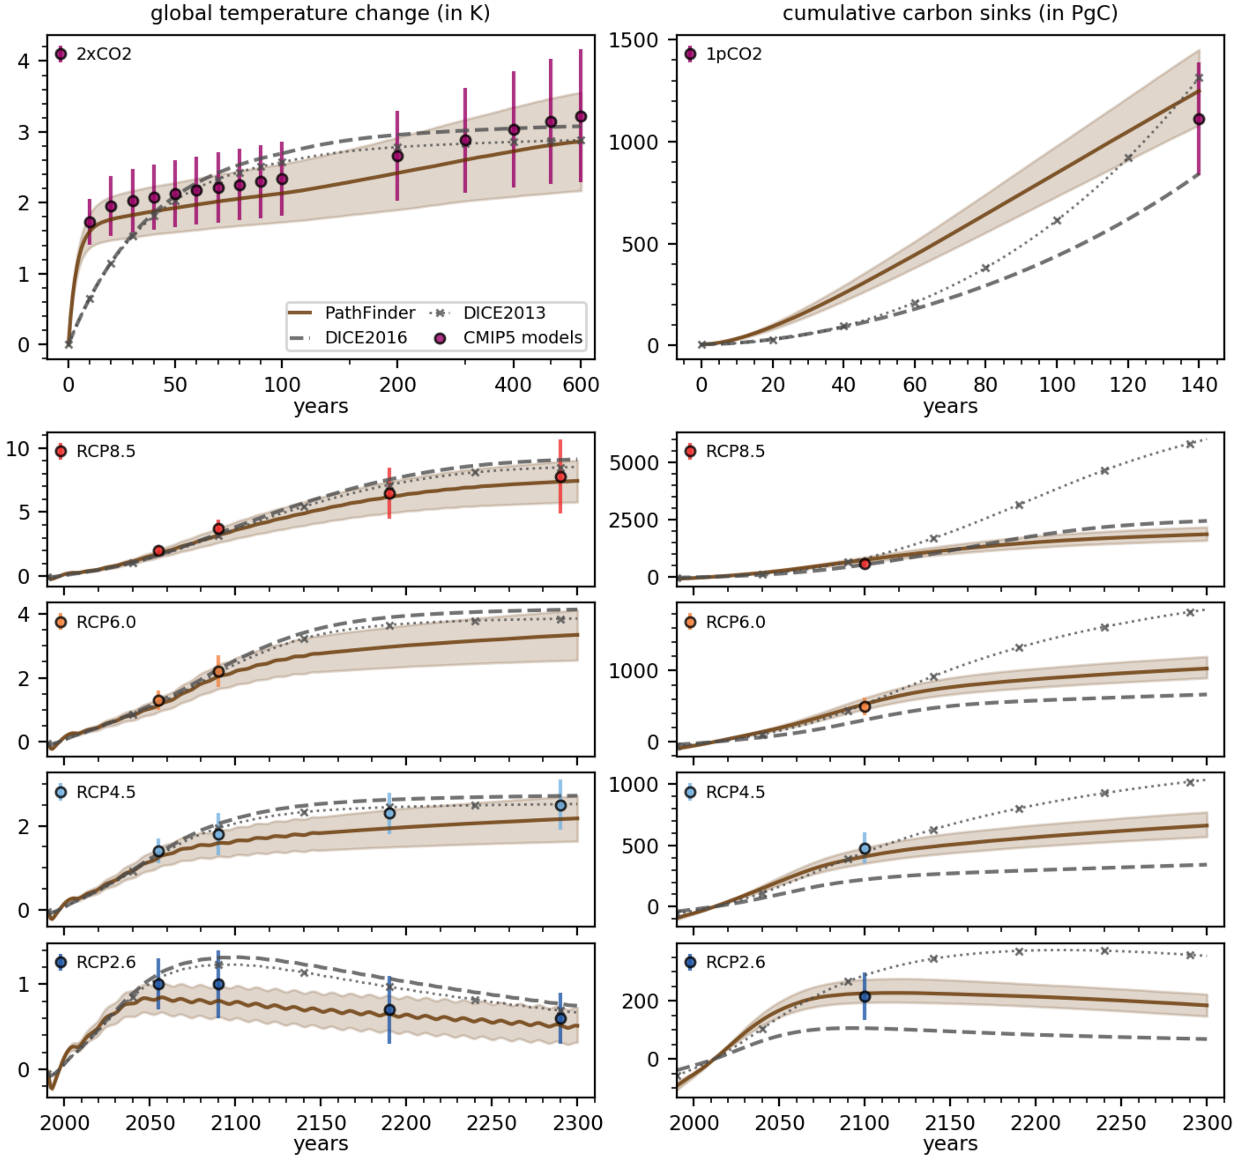

Concentration-driven runs to diagnose the climate response (left panels) and the carbon cycle (right panels).

Original module (gray lines):

- particularly wrong for low- or medium-warming scenarios.
- Same short and long climate timescales.
- Carbon-cycle slowed down to compensate lack of saturation and climate feedback.

New module (brown, named PathFinder):

- Slightly more optimistic climate
- More optimistic carbon-cycle.

Calibration on CMIP6 and analysis in the works.

Robust optimization: Caveat of Monte Carlo

Typical uncertainty analysis: Monte Carlo.

- Prior vs. posterior

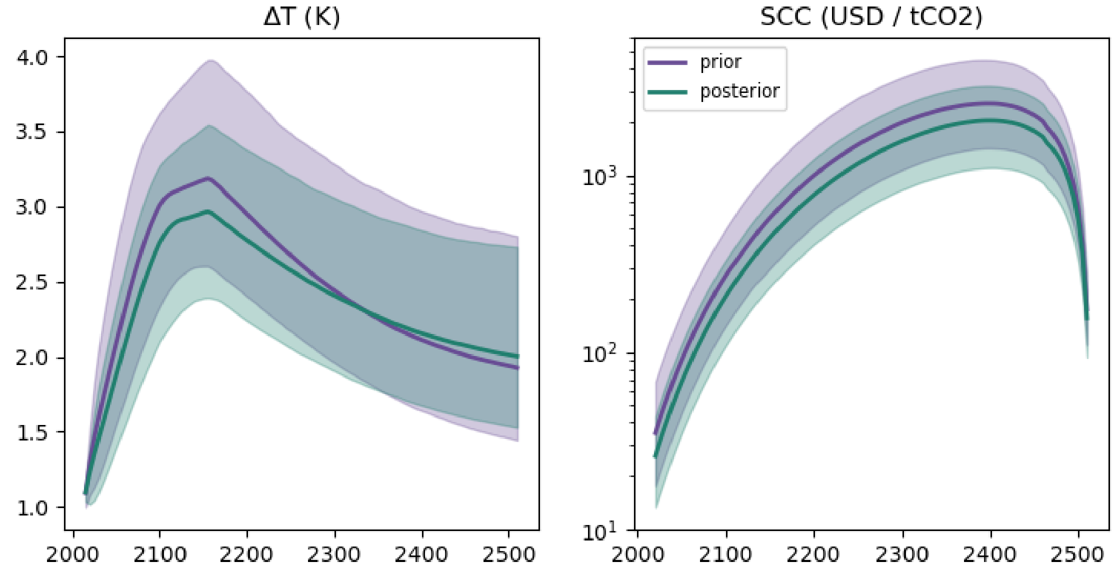

- Unconstrained vs. constrained

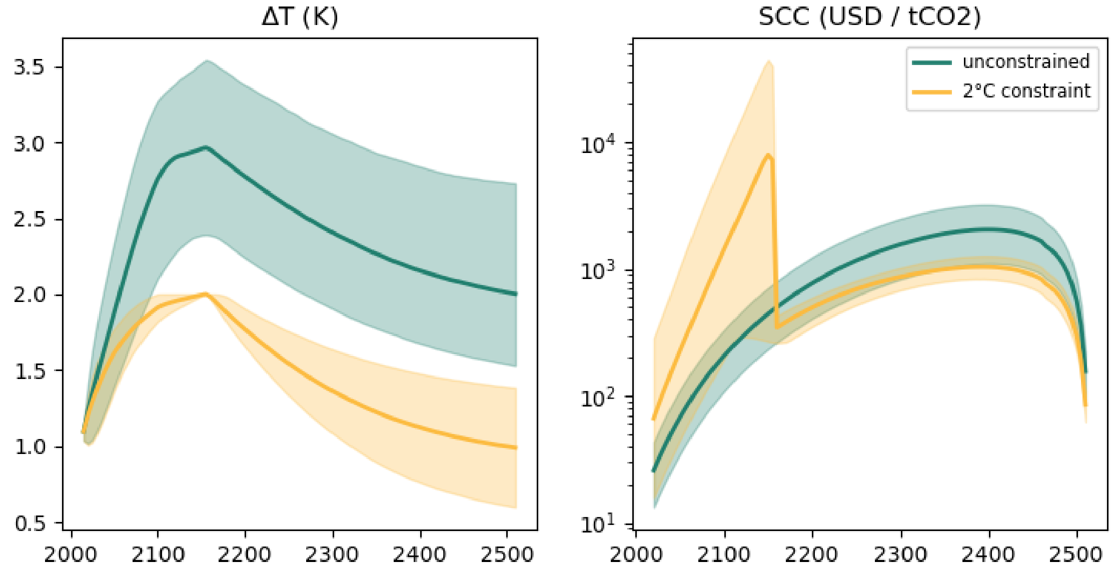

For an IAM, running a Monte Carlo ensemble merely requires solving several deterministic problems in parallel (here, 4000 different states of the world, with equiprobability).

Each solution is optimal for its own world, but assumes no uncertainty within that world, and ignores other worlds.

For the policy-maker, however, this represents 4000 solutions. Which one should be chosen? One could take the median (or average) solution, but is the median of optimal policies an optimal policy for the median world?

Does that properly account for physical uncertainty? (see e.g. the constrained case that peaks at 2°C with 100% certainty.)

Robust optimization: Robust SCC

Solution: finding a unique policy applied across all states of the worlds, by optimizing the cross-world average welfare

Preliminary results for robust optimization:

- generates a unique value of SCC, as the physical uncertainty moves back to the physical variables
- treats uncertainty consistently as a risk if one follows the unique policy
- appears to be very close to the average optimal policy when unconstrained
- differs from it when applying a constraint/target (which must be expressed in probabilistic terms)

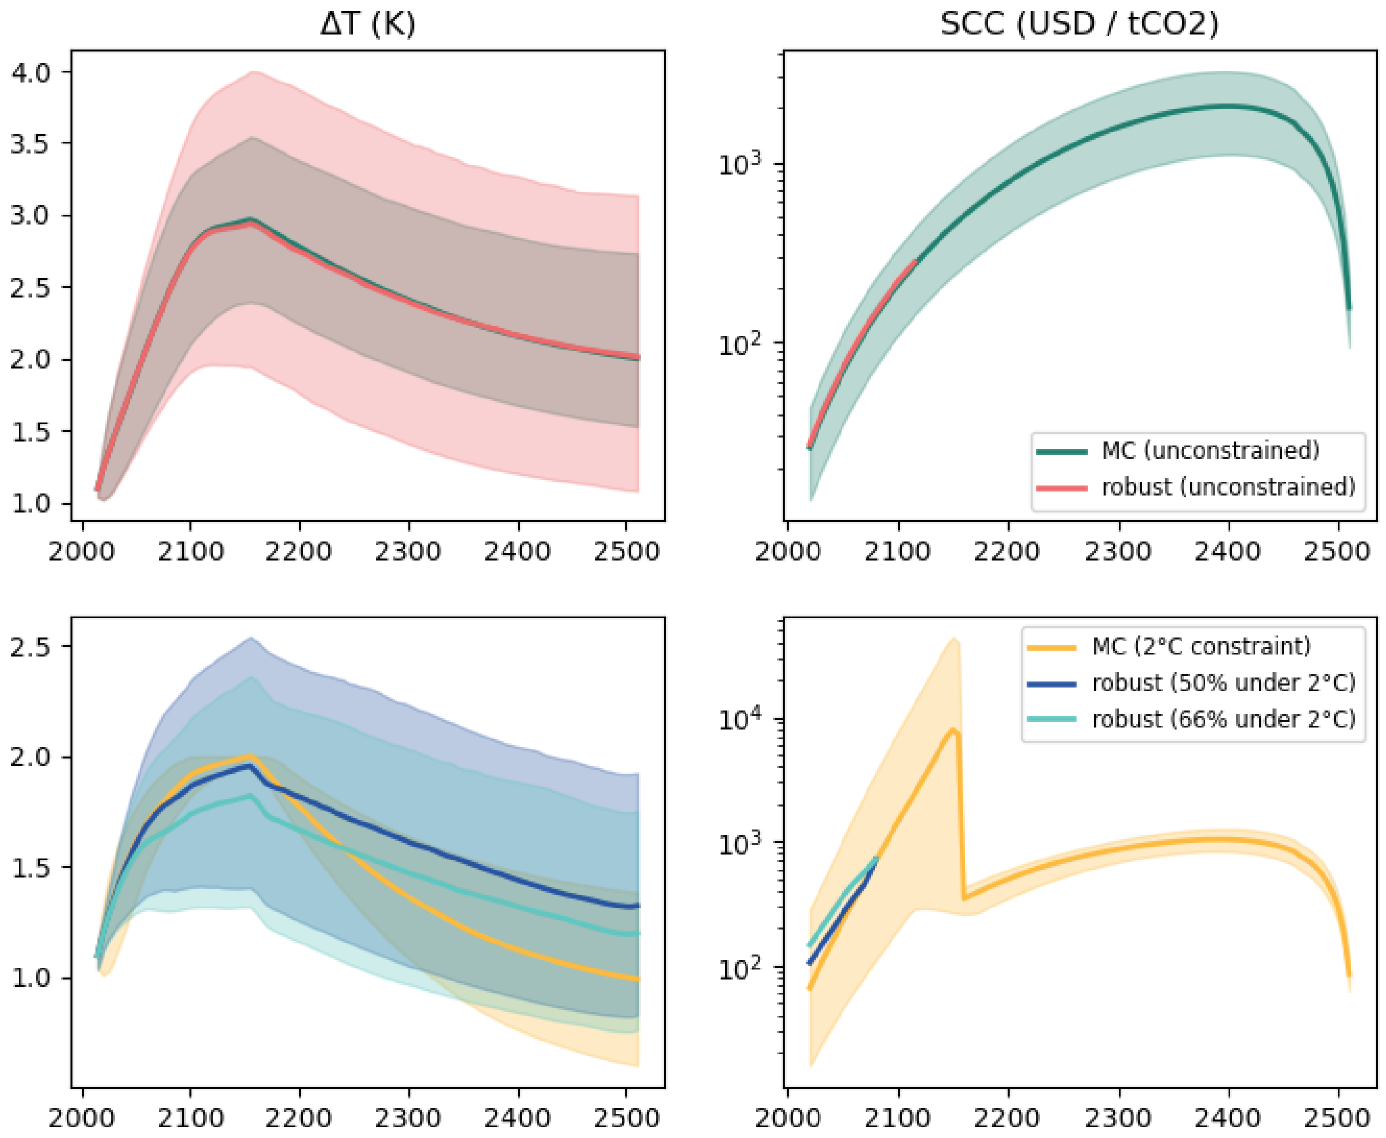

Take-home messages

- The original physical module of DICE is just wrong. It should never be used.
- Data and structural options for a simple (and yet accurate) physical module are available (but such module is not unique).
- Bayesian calibration is powerful magic.
- With a revamped and recalibrated physical module, in a Monte Carlo setup informed by observations, the SCC is lower than the original value.
- However, a Monte Carlo setup is inadequate to properly integrate uncertainty into decision-making. Robust policies can be designed by optimizing across future states of the world. Doing so leads to an increase in SCC for the 2°C target.