Final Response

AD-MERGE 2.0: An Integrated Assessment Approach to Assess World Energy and Climate Policy Options

General response and overview of the revision

We thank the Reviewers for their constructive evaluations of the manuscript and for the detailed suggestions that helped us improve the clarity, documentation, and interpretation of AD-MERGE 2.0. We provide a revised manuscript and a marked-up manuscript and an additional supplementary material. The revised manuscript addresses three cross-cutting issues raised across the reviews: (i) the interpretation of optimal adaptation as an idealized benchmark rather than a direct policy prescription, (ii) the need for clearer documentation of the damage, adaptation, climate, and welfare modules, and (iii) the interpretation of energy-system and carbon-removal results given the current model structure. We have updated the model code and reporting materials in Zotero and GitHub accordingly.

To avoid repetition in this consolidated response, issues that recur across Reviewer comments are addressed once in full and then referenced where relevant. In particular, the revised framing of optimal adaptation is detailed in Response 2.1, and subsequent responses refer back to that clarification rather than repeating the same caveat. Similarly, the expanded documentation of core model equations and parameters is described in Response 2.5 and is referenced when related questions arise elsewhere.

The main revisions are summarized as follows:

- We corrected the energy accounting used in Figure 8 and clarified the primary/final energy definitions.
- We clarified that CH4 and N2O emissions in Figure 7 refer to anthropogenic emissions.
- We added caveats on the near-term divergence between no-adaptation and optimal-adaptation scenarios.
- We moderated statements implying direct region-specific policy design and reframed the adaptation results as stylized macro-regional benchmark insights.
- We clarified the assumption of optimal adaptation regarding idealized implementation and perfect foresight.
- We removed the conceptual adaptation-cost figure from Appendix D and replaced it with text that distinguishes empirically derived calibration.
- We added a additional context to report the main adaptation and damage parameters, the all main module equations, and the Negishi-weighting structure.
- We clarified that BECCS deployment is conditional on broad regional biomass/resource potentials and techno-economic assumptions, not on an endogenous land-use allocation module and it is under consideration for future studies.

Response to Reviewer #1

We thank Reviewer #1 for the positive assessment of the manuscript and for the line-by-line comments that improved the transparency of the energy-system results, emissions accounting, and scenario interpretation.

Comment 1.1: Pumped-hydro storage and hydro reservoirs

Reviewer comment: 243 - "PHS [pumped hydro storage] dominates global electricity storage when hydro reservoirs are excluded." I don't know what this means.

Response: We agree that the original wording was ambiguous. The intended distinction was between conventional hydropower reservoirs, which provide dispatchable generation and operational flexibility but are typically reported as generation capacity, and explicit electricity-storage technologies such as pumped-hydro storage (PHS) and batteries. We revised the sentence to clarify the accounting boundary.

> **Revised methodology text:** Excluding conventional hydropower reservoirs from the storage accounting, pumped-hydro storage (PHS) provides the largest share of global electricity storage capacity. While conventional hydropower reservoirs offer significant dispatchable energy, they are typically classified as a primary generation source.

Comment 1.2: Anthropogenic CH4 and N2O emissions in Figure 7

Reviewer comment: Figure 7 - please clarify that this refers to anthropogenic emissions only of methane and nitrous oxide.

Response: We agree with the need for this clarification. The earlier caption used the terms total CH_4 and total N_2O in a way that could be read as including natural emissions. We revised the caption to state explicitly that the panels report anthropogenic emissions. We also clarified the split between endogenous energy-related CO_2 and CH_4 and exogenous non-energy CO_2, CH_4, N_2O, and other non-CO_2 gases in the methodology section.

> **Revised caption text:** Figure 7. Projection of anthropogenic emissions of climate change key drivers; (a) Total GHG emissions (b) energy-related CO_2 emissions (c) total N_2O emissions, and (d) total CH_4 emissions comparison across various scenario.

Comment 1.3: Primary and final energy accounting in Figure 8

Reviewer comment: "In the Baseline scenario, primary energy supply reaches 1221 EJ and final energy use reaches 1050 EJ by 2100...By 2100, fossil fuels will account for 75% of the total primary energy supply, with coal leading at 28%, closely followed by oil at 27%"

Something seems wrong with the calculation of final or primary energy shown in Figure 8, which explains this very high ratio of final to primary energy in 2100. Globally, final energy is about 2/3 of primary energy, with most of the difference lost in electricity generation, with contributions from transmission and distribution of electricity and gas, and oil refinery energy use. This ratio can increase over time with renewable electricity deployment, but that isn't what is driving this ratio so high here. In the 2015 time period it looks like the final energy is equal to the primary energy, and natural gas final energy exceeds primary energy, which can't be the case, especially as this is a major fuel input to the power sector.

Response: Thanks to the reviewer for point this. The previous version of Figure 8 reflected an inconsistency in how final energy was assembled from model outputs, particularly with respect to transformation inputs and energy-sector own use. This issue could yield implausible comparisons, including final energy appearing too close to primary energy and natural gas final energy exceeding natural gas primary energy. We corrected the accounting used to construct Figure 8 and revised the corresponding numerical statements. We also added a short note on the primary/final energy definitions and on the direct-equivalent method for non-combustible electricity sources. Additional validation is also provided in Appendix E.

Revised Figure 8b:

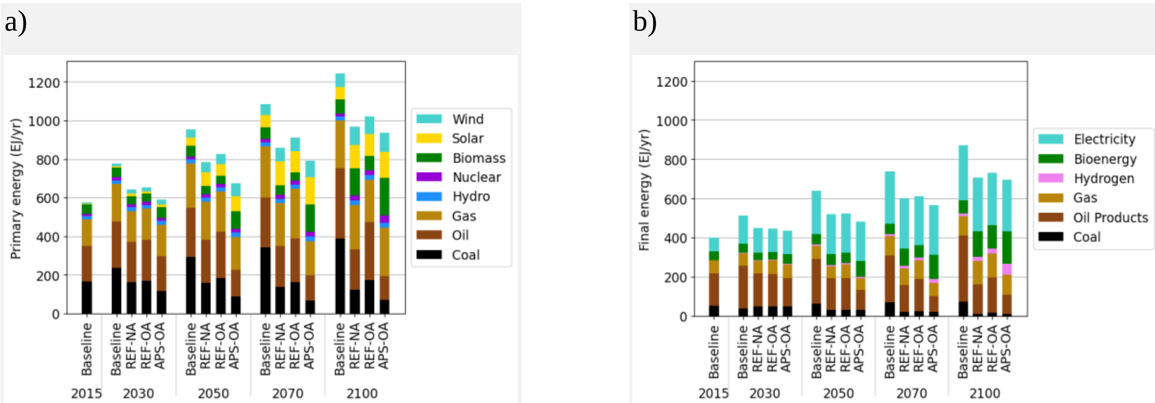

Updated Figure 8. (a) Primary energy supply mixes by sources (b) final energy use by sources for all scenarios

Revised results text: Primary energy is defined as energy inputs to conversion and transformation processes, including fuels used for electricity generation, whereas final energy represents energy delivered to end-use sectors after accounting for conversion losses and energy-sector own use.

Revised results text: Consistent with IPCC AR6 WGIII reporting conventions (IPCC, 2022), we use the direct-equivalent method for non-combustible electricity sources, counting one unit of electricity from non-combustible sources as one unit of primary energy, rather than converting output to a fossil-fuel-equivalent input through an assumed efficiency.

Revised numerical statement: In the Baseline scenario, primary energy supply reaches 1221 EJ and final energy use reaches 875 EJ by 2100.

Revised numerical statement: Electrification accelerates as the share of electricity in final energy rises from 29% in the Baseline to 38% in REF-NA, 36% in REF-OA, and 40% in APS-OA, reflecting lower decarbonization costs in the power sector and rapid advances in renewable and storage technologies. Hydrogen also moves from a marginal role in the Baseline scenario, where it contributes 2% of final energy, mainly gray hydrogen, to a more central position in the policy scenarios, doubling to 5% in REF-NA and REF-OA and reaching 9% in APS-OA by 2100.

Revised numerical statement: Throughout the century, the share of biomass in global final energy use rises from 5% in the Baseline to 14% in REF-NA, 11% in REF-OA, and 19% in APS-OA by 2100.

We will also update the archived code and output files on GitHub and Zenodo so that the corrected figure, text, and reproducibility materials are internally consistent.

Comment 1.4: Near-term divergence between no-adaptation and optimal-adaptation scenarios

Reviewer comment: Figure 10 - The spread between the scenarios' estimates of climate damages in 2030 seems a bit wide. The baseline is 2% of GDP, and the lowest two are under 1%. This just doesn't seem like it's enough time for the adaptation scenarios to have built the infrastructures necessary to address climate impacts. Can the authors comment on what the timelines are here? Another consideration is that it is 2026 right now, so even though the model base year is 2015, it doesn't make sense to have that level of divergence taking place in years that in the real world are historical. I'm not suggesting re-running or revising the scenarios, but the appropriate caveat should be provided if the scenarios are going to have this level of divergence for a variable that in the real world would take decades of action to achieve. Or perhaps starting the chart in 2040 would address this.

Response: We agree that the near-term divergence requires clearer interpretation. The NA and OA cases are constructed as illustrative bounding scenarios. OA assumes deployment of the model-implied cost-effective adaptation portfolio from 2020 onward, rather than a continuation of observed adaptation practice since the 2015 base year. Near-term reductions mainly reflect the model's damage–adaptation mechanics under these bounding assumptions. Reactive adaptation can be scaled within a model period, whereas proactive adaptation accumulates gradually through long-term investments; although these proactive investments become increasingly important over the long term, they can still generate benefits within a 10-year model period. We revised the text and caption to state this explicitly and adjusted the figure framing to start in 2020.

> **Extended results text:** The divergence between the NA and OA scenarios by 2030 reflects their roles as bounding cases rather than near-term forecasts. NA represents a counterfactual no-adaptation pathway, while OA illustrates an idealized optimal-adaptation pathway. In the OA scenario, adaptation is deployed from 2020 onward at a model-implied cost-effective level for each impact category, whereas NA assumes zero protective adaptation. Early reductions in climate change impacts are driven by both reactive and proactive adaptation. Reactive measures can be scaled within a single model period. Proactive measures, such as investments in seawalls, irrigation infrastructure, and resilient energy and health systems, accumulate more gradually through depreciable investments, although they still yield benefits within a 10-year model period. Because these scenarios abstract from real-world policy and implementation frictions, their differences around 2030 should be interpreted as illustrative of the model's internal damage-adaptation dynamics.

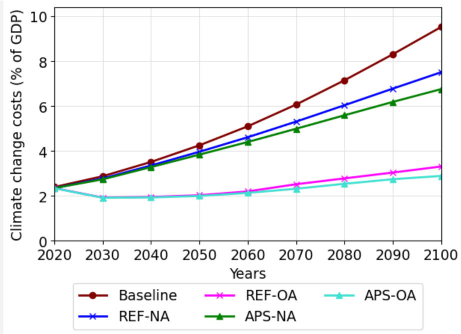

Updated Figure 10. (a) Global total climate change impacts.

Comment 1.5: Potential use of the model for optimal policy analysis

Reviewer comment: The final question that I have pertains to the use of this model for charting out optimal policies that balance the costs of mitigation against the benefits of avoided damages, taking into account the cost of that adaptation. It seems like this model has all of the necessary pieces to do that sort of analysis, but I don't see any commentary about this, and the scenarios weren't constructed so as to estimate it.

Response: We agree that AD-MERGE 2.0 framework supports welfare-consistent analysis of mitigation, adaptation expenditures, and residual damages. The present manuscript is primarily a model-development and validation paper. Its scenario set is designed to illustrate model behavior under policy pathways to provide a basis for full exploration of globally optimal joint mitigation-adaptation policies.

That said, the Reviewer's point is well taken. We strengthened the manuscript to explicitly state that the framework can be used for more comprehensive joint mitigation–adaptation policy assessment. At the same time, a systematic exploration of "optimal policy" design under alternative objectives, constraints, and instrument sets is a substantial standalone exercise; we, therefore, identify it as a planned follow-on analysis.

> **Extended discussion/conclusion text:** Beyond the illustrative scenario analysis presented here, the integrated welfare framework enables analysis of jointly optimal mitigation-adaptation pathways that balance mitigation expenditures, adaptation investments and spending, and residual damages under alternative objectives and constraints. The REF scenario applies an economy-wide carbon-price path that weighs marginal mitigation costs against the marginal benefits of avoided climate damages over time. The comparison between REF-NA and REF-OA allows us to study the interplay between adaptation and mitigation, demonstrating how the inclusion of adaptation options shifts the optimal carbon price and overall economic burden. A standalone analysis integrating joint mitigation and adaptation policies is a potential extension of this research and is intended for future work.

Response to Reviewer #2

We thank Reviewer #2 for the detailed and constructive critique. The main revisions motivated by this review concern the interpretation of aggregate adaptation in an IAM, the distinction between theoretical benchmark results and policy prescriptions, and the documentation of model equations and parameters.

Comment 2.1: Scientific and policy interpretation of aggregate regional adaptation

Reviewer comment: Adaptation is presented as a key contribution of this paper, and the authors draw substantive policy conclusions from adaptation results, including claims about 'region-specific policy design' (abstract, l. 11) and 'optimal policy approaches tailored to the unique needs and capabilities of each region' (l. 566-567). However, the paper does not address a prior question: why does modelling adaptation as an aggregate welfare-optimizing choice variable over regions like 'Africa' or 'Other Asia' produce scientifically credible or policy-relevant outputs?

Response: We agree that adaptation is not implemented in practice at the scale of aggregate regions such as Africa or Other Asia, and the original wording could be read as overstating the spatial precision and direct policy prescriptiveness of the model results. We revised the manuscript to make this scope explicit. The scientific value of the adaptation results lies not in reproducing local adaptation processes, but in representing broad interregional differences in exposure, adaptive response, and mitigation-adaptation trade-offs within a unified welfare-consistent IAM framework.

The policy relevance of the results is therefore limited to broad strategic differentiation: identifying where adaptation is likely to play a larger or smaller role relative to mitigation, where residual damages remain structurally high even under idealized adaptation, and where international finance and burden-sharing mechanisms may be most relevant. The results should not be interpreted as prescribing local or country-specific adaptation portfolios within aggregated regions.

> **Revised abstract text:** The results underscore critical trade-offs and synergies between adaptation and mitigation strategies, highlighting regional heterogeneity in adaptation needs, residual damages, and mitigation-adaptation interactions, alongside the integration of clean energy technologies.
>
> **Revised discussion/conclusion text:** These marked regional differences in climate impacts, adaptation effectiveness, and mitigation burdens should be interpreted as indicating broad differences in strategic emphasis across macro-regions rather than prescribing policy portfolios for individual countries within them. In this sense, the framework is most informative for comparative assessment of regional mitigation-adaptation trade-offs. At the same time, because adaptation is represented as an idealized regional aggregate under optimal implementation, these findings should be understood as stylized macro-regional insights, not as direct blueprints for local adaptation planning.

Comment 2.2: Adaptation is local, institutionally embedded, and constrained by adaptive capacity

Reviewer comment: A central and well-established insight in adaptation science is that adaptation is inherently local, context-specific, and institutionally embedded (Smit & Wandel, 2006; Adger et *al., 2007; IPCC AR6 WG2). There is no universal optimal adaptation level derivable from aggregate economic parameters. The UNEP Adaptation Gap Report (2025) documents a 'widening gap' between adaptation needs and realized action in many developing countries, which are precisely the regions where the model predicts the largest benefits (Africa, Other Asia). The calibration procedure (Appendix D) implicitly assumes that adaptation costs and effectiveness can be meaningfully expressed as a fraction of GDP at the regional scale, and that 'optimal adaptation' is achievable given sufficient investment. This follows the approach of earlier IAM-based adaptation models in the literature (de Bruin et al., 2009; Agrawala et al., 2011), and as such represents a reasonable starting point. However, to produce genuinely meaningful policy implications, the modeling community needs to start moving beyond this type of representation toward frameworks that are connected with empirical findings on how adaptation actually happens — taking into account characteristics such as adaptive capacity, institutional context, path dependency, and local governance conditions. I am aware this is a difficult challenge, but there are promising avenues in this direction, including adaptation pathways approaches new collection of papers coming in (https://www.nature.com/collections/gfebjgbjid) and efforts to couple IAMs with more bottom-up behaviourally rich simulation models (Filatova et al., 2025).*

Response: We thank the reviewer for this constructive comment. We agree that adaptation is inherently local, context-specific, and institutionally embedded, and that these features cannot be fully represented by a reduced-form regional variable in an IAM. We also agree that the results should not be interpreted as directly achievable policy targets.

At the same time, the revised manuscript clarifies the intended contribution of the present study. The contribution is not to claim that aggregate regional optimization can reproduce how adaptation occurs in practice, but rather to improve the empirical foundation of adaptation representation within the IAM framework. While adaptation was already explicit in earlier AD-MERGE formulations and in some other IAMs, regional damage and adaptation relationships in the IAM literature have remained only limitedly revised. One contribution of the present study is therefore to recalibrate these relationships using a broader and more recent empirical basis across regions and impact categories. Importantly, the adaptation inputs are derived separately for the corresponding impact categories, so that the calibration is based on internally coherent damage, adaptation cost, and adaptation effectiveness information within each category. Broader determinants of real-world adaptation, such as adaptive capacity, institutional quality, financing constraints, governance conditions, path dependency, and behavioural uptake, are only reflected indirectly insofar as they are embedded in the empirical relationships used for calibration. They are not represented as explicit dynamic drivers in the current framework.

We also clarified that the explicit adaptation structure remains valuable within this scope. Representing reactive and proactive adaptation as explicit variables allows adaptation to function as a dynamic decision variable in the model, so that its timing and composition can interact with mitigation, evolving damages, and regional welfare outcomes over time.

In this sense, the present study is positioned as an intermediate but necessary step to strengthen the empirical and structural foundation of the adaptation module before richer institutional and behavioural dimensions are incorporated more explicitly. This integration also set the stage for future coupling with more granular, bottom-up behavioral simulations.

Extended Section 2.5: AD-MERGE 2.0 builds on the stock-and-flow framework of de Bruin et al., 2009; Agrawala et al., 2011). The "optimal adaptation" serves as a theoretical benchmark rather than a policy forecast. As noted by the IPCC (2022) AR6 and UNEP (2025), real-world adaptation is constrained by local institutional barriers and varying adaptive capacities. By modeling adaptation as an explicit capital stock—rather than a reduced-form percentage of GDP—our approach captures the opportunity costs and investment inertia inherent in climate resilience. This structural detail is essential for analyzing the energy-climate-adaptation nexus, as it reveals how adaptation investments may compete with the capital flows required for mitigation (Bastien-Olvera & Moore, 2021).

Extended Section 2.5 caveat: Accordingly, the resulting optimal adaptation levels should be interpreted as benchmark outcomes within the model structure under idealized implementation assumptions, rather than as directly achievable policy targets.

Extended limitations text: While the current study explored optimal adaptation, future research is needed to quantify adaptation gaps and limits and to represent them more explicitly across regions. In particular, richer formulations should account for suboptimal implementation, financing and institutional constraints, path dependency, and broader behavioral responses (van Maanen, 2023). Some of these dimensions are partly embedded in the empirical relationships used here, but they are not yet modeled as explicit dynamic drivers. Extending the macroeconomic structure to better capture sectoral and behavioral adjustment and linking IAMs with more behaviorally rich or bottom-up simulation frameworks, would also help represent broader forms of adaptation beyond the current reduced-form specification (Filatova et al., 2025).

Comment 2.3: Reframing the contribution as theoretical benchmark analysis

Reviewer comment: To strengthen the contribution of this paper, the authors should either: (a) demonstrate, with explicit reference to the empirical literature, that the structural assumptions are reasonable approximations and constrain policy conclusions accordingly; or (b) reframe contribution as modelling exercise looking at theoretical upper bounds under idealized conditions, not as achievable policy targets. The macro-level outputs the model can legitimately support — aggregate damage cost comparisons across scenarios and the identification of highly vulnerable regions — are genuinely valuable. But the claims to region-specific policy design cannot be sustained. This concern is acknowledged at L. 563-564 ('these results are obtained on the assumption of optimal adaptation, an idealized scenario in practice'), but a single end-of-paper caveat is insufficient given the policy framing used throughout Sections 4.3 and 5. This caveat needs to be made more explicit in the framing of the paper and its contribution.

Response: We thank the reviewer for this important clarification. We agree that the policy interpretation of the adaptation results needed to be narrowed and made explicit much earlier in the manuscript. We, therefore, followed option (b) and reframed the contribution more clearly as a macro-regional modeling exercise that provides benchmark results under idealized adaptation assumptions, rather than directly achievable policy targets. In this revised framing, the principal value of the analysis lies in comparing aggregate damage burdens, adaptation expenditures, residual damages, and mitigation costs across scenarios, and in identifying highly vulnerable regional burden profiles. We agree that stronger claims regarding direct region-specific policy design were not sufficiently supported, and these have been removed or moderated throughout the text.

We also agree that the caveat regarding the assumption of optimal adaptation should not appear only in the concluding discussion. We revise the abstract, Section 2.5 caveat (provide in response to comment 2.2), Section 4.3, and the discussion/conclusion (in response to comment 2.1) to reflect this interpretation consistently throughout the manuscript.

> **Extended abstract sentence:** Each of the mitigation scenarios is studied with and without optimal adaptation, an idealized assumption of perfect foresight used to establish theoretical benchmarks.
>
> **Revised Section 4.3:** Furthermore, this analysis is conducted at the macro-regional scale. The resulting outputs are therefore most informative for the comparative assessment of aggregate damage burdens, mitigation-adaptation interactions, and residual damages across regions, rather than for prescribing locally specific adaptation strategies. The next section examines these regional dynamics in greater detail.

Comment 2.4: Conceptual adaptation-cost figure and empirical calibration

Reviewer comment: A related issue: Figure D1b — the adaptation cost curve — is explicitly described in its own caption as 'conceptual' with 'shapes of the curves theoretical and not reflecting specific empirical estimates.' If the figure used to illustrate the adaptation calibration logic has no empirical grounding, the paper's claim to have empirically recalibrated adaptation is weakened. The paper needs to be transparent about what is empirically derived versus theoretically assumed throughout the adaptation framework. Where empirical adaptation investment data exist — for instance for European countries (Cortes-Arbuez et al., 2025) — these should be used to benchmark whether modeled adaptation expenditure levels are plausible.

Response: Thank you for this comment; you raise an important point. We agree that confusion may arise from including this illustrative figure. As such, we have removed this figure in the revised draft as it detracted from the empirically calibrated Figure D1a. We agree that if we had not removed Figure D1b, the reader might have been confused about what was empirically derived and what was theoretically assumed.

Regarding the empirical calibration of adaptation expenditure, the paper relies on detailed protection cost data from the CIAM model (Diaz, 2016) and from the Aqueduct Flooding tool, which uses the GLOFRIS model (Ward et al., 2020; Winsemius et al., 2013; World Resources Institute, 2020). Estimates from other literature and expert judgment are also used. This is described in detail in the forthcoming paper on the impacts and adaptation calibration. A full description of this process would not be possible in this paper; it requires a standalone paper.

Though a comparison of our adaptation cost estimates, and empirical estimates would be useful, such a comparison is complex. No empirical data exists for all regions and impacts covered in the AD-MERGE model. The closest possible comparison is with the United Nations Environment Programme (2023,2025a) Adaptation Gap Reports. In these reports, adaptation costs are estimated for each impact and then summed for the developing nations in each region.

They find similar costs for the same impacts, but because different impacts are considered in their paper, the total adaptation costs differ (UNEP, 2023). For example, the Adaptation Gap Reports (UNEP, 2023, 2025b) include agriculture impacts, which lead to higher impacts in Other Asia countries than in AD-MERGE, whereas AD-MERGE includes a more comprehensive assessment of health impacts, leading to higher impacts in Africa. Additionally, our adaptation costs are for all nations in a region, not only developing countries, making their estimates not directly comparable to ours in some regions.

> **Revised Appendix D:** Drawing on the data sources summarized in Table D1, complemented by insights from the broader adaptation literature, we construct temperature-dependent profiles for (i) the fraction of gross climate damages that can be feasibly avoided through adaptation and (ii) the associated adaptation expenditures, by impact category and region. For each region and impact category, we further identify the optimal adaptation level at which the sum of residual damages and adaptation expenditures is minimized.
>
> We have compared the adaptation cost estimates to those reported in the Adaptation Gap Reports (UNEP 2023, 2025a,b). These reports estimate adaptation costs by first assessing costs for each impact category and then summing them to obtain total needs for developing countries. The estimates are broadly comparable for overlapping impacts, but the total adaptation costs differ because the reports include different impact categories and cover developing countries only, whereas AD-MERGE 2.0 covers all countries within each region.
>
> The optimal protection and cost profiles are then aggregated to obtain regional benchmarks for (i) the share of gross damages avoided by adaptation, (ii) the decomposition of total climate-related costs into residual damages and adaptation expenditures, and (iii) the split between flow and stock adaptation at several temperature levels. These calibration targets are summarized in Fig. D2. The parameters governing the CES protection function are calibrated such that the simulated protection levels, residual damages, and adaptation expenditures reproduce the regional benchmarks over the calibration temperature range.

Comment 2.5: Documentation of model components and parameters

Reviewer comment: A broader concern across the paper is that key model components are insufficiently documented for a model description paper. The main text frequently defers to earlier publications (Bahn & Kypreos, 2003; Bahn et al., 2019) for core structural elements, and the appendices, while useful, do not fully compensate. The paper is clearly written with appropriate conciseness, and I am not asking for an exhaustive technical manual in the main text. However, at minimum, the appendices — or a linked online documentation resource — should provide enough information for a reader to understand the main modeling without having to track down earlier papers. The adaptation and damage framework is a case in point. Equation D2 governs the entire adaptation structure via parameters σ, β1,r, β2,r, and β3,r (l. 679–680), yet none of these values are reported anywhere. The 'maximum adaptation level' mentioned at l. 301 is stated but neither quantified nor explained. A reader cannot assess the plausibility of the adaptation results, connect them to current literature, or understand their sensitivity without access to these parameters. The same issue applies more broadly: the welfare function parameterization and the Negishi weights are not reported, and the climate module equations are not provided. The authors should either *expand the appendices or provide a clearly signposted link to online technical documentation that fills these gaps.*

Response: We thank the reviewer for this important comment. We agree that, for a model description paper, readers should be able to understand the core structure and parameterization of the model.

To address this, we add supplementary material that documents the model more fully. In particular, the supplementary material report:

1. the climate module equations.
2. the Negishi weighting structure and its calculation process.
3. the main parameters governing the adaptation and damage modules.

We also add a clearly note in the manuscript directing readers to the relevant supplementary sections where these model elements are documented.

The added notes:

Section 2.1: The formulation of the damage module is provided separately in Appendix D, while the formulations of the other modules are presented in the Supplementary Material S1-S3.

Appendix D: The regional calibration parameters $\alpha 1,r$,...,$\alpha 4,r$ are provided in Table D2.

The parameters governing the CES protection function, $\beta 1,r$, $\beta 2,r$, $\beta 3,r$, and σ are calibrated such that the simulated protection levels, residual damages, and adaptation expenditures reproduce the regional benchmarks over the calibration temperature range (see Table D3).

Minor Comments:

- ### Selection criteria for IAM comparison in Table 1

Reviewer comment: Table 1: The selection of IAMs for comparison is not justified. Why these specific models? Some (e.g. GCAM, REMIND) have been substantially updated and the version used for comparison is not stated. Please add a brief justification of the selection criteria.

Response: We agree. We revised the text accompanying Table 1 to clarify that the models were selected because they are widely used IAM frameworks relevant to the dimensions emphasized in this study: energy-system representation, climate damages, and adaptation. We also specify model versions or source references where applicable, so that the comparison can be interpreted in light of continuing model development.

Extended Section 2.1: Table 1 compares AD-MERGE 2.0 offers an overview of how AD-MERGE 2.0's mathematical formulation, climate and damage specifications, carbon dioxide removal (CDR), adaptation treatment, and regional coverage compare with those of other widely used IAMs. The table is not intended to provide an exhaustive taxonomy of IAMs, but rather to situate AD-MERGE 2.0 relative to established frameworks used in climate-policy assessment. An important comparison emphasizes whether existing models provide impact and adaptation representations that are suitable for incorporating adaptation explicitly, including corresponding estimates of gross damages, avoided damages, residual damages, and adaptation costs. Where relevant, the model versions or source references used for the comparison are specified. The information is synthesized from review studies on IAM capabilities, including Füssel (2010), Nikas et al. (2018), and Keppo et al. (2021), and from the Integrated Assessment Modeling Consortium database and model documentations (IAMC, 2024). It is further complemented by model-specific publications and documentation, including de Bruin (2014), Bahn et al. (2019), Calvin et al. (2019), Fujimori et al. (2017), Emmerling et al. (2016), Aaheim et al. (2018), Hilaire and Bertram (2020), Krey et al. (2020), Zhao et al. (2020), Baumstark et al. (2021), Morrow et al. (2022), Wei (2025), and the main methodological papers associated with each model.

- **Use of existing damage assessments and benchmarking against COACCH**

Reviewer comment: Appendix D / Table D1: The paper estimates its own regional damage functions rather than using existing sector-specific, country-level approaches such as COACCH (Van Der Wijst et al,. 2023) . Given that Table D1 draws on many of the same underlying data sources that such frameworks use, why was an existing damage assessment not used or more systematically benchmarked against? This choice should be explicitly justified.

Response: We agree that this modeling choice should be justified more explicitly. In the revised manuscript, we clarify that our objective was not simply to adopt an existing damage assessment, but to derive a set of regionally aggregated damage and adaptation relationships that are consistent with the AD-MERGE 2.0 structure and can be used to calibrate the model's explicit adaptation module. While frameworks such as COACCH provide valuable estimates of climate damages and are highly relevant as underlying evidence, they are not directly available in a form that provides the corresponding explicit calibration of adaptation measures (except for sea-level rise).

Extended Section 2.5: Existing damage assessments provide useful benchmarks for different parts of the calibration. Empirical macroeconomic studies such as Burke et al. (2015, 2018) and meta-analyses such as Howard and Sterner (2017) and Tol (2024) inform the broader comparison of aggregate damage-function severity, while sectorally disaggregated assessments such as Bosello et al. (2021) provide more direct evidence for impact-specific damage calibration. However, these studies generally do not provide explicit adaptation-cost and effectiveness estimates; they are primarily used to benchmark climate damages, while adaptation calibration relies on more limited impact-specific evidence and model-specific assumptions. The objective here is to derive regionally aggregated damage and adaptation relationships that are consistent with the AD-MERGE 2.0 structure and suitable for calibrating both reactive and proactive adaptation, thereby allowing these adaptation responses to be represented explicitly as decision variables and studied jointly with mitigation and climate damages.

- **Presentation of APS-NA in Figure 10 and adaptation decomposition**

Reviewer comment: Scenarios / Figure 10: APS-NA does not appear in the energy figures (Figs. 7-9), which is justified at l. 367-369 because APS-NA and APS-OA share identical energy trajectories. However, showing APS-NA explicitly alongside APS-OA in Figure 10 and the adaptation decomposition figures (10b-10c) would allow clearer isolation of the marginal effect of optimal adaptation under the more ambitious mitigation pathway. It is not evident why this comparison is currently absent.

Response: We agree that the presentation should make the relationship between NA and OA cases more transparent. The figure decomposition is designed so that the no-adaptation case can be inferred from residual damages plus avoided damages, while the optimal-adaptation case decomposes gross damages into residual damages, adaptation expenditures, and avoided damages. We revised the caption and discussion to clarify this interpretation for both APS and REF.

> **Revised figure caption:** (a) Global total climate-change impacts; (b) impact components in the REF scenario; (c) impact components in the APS scenario; and (d) regional climate impacts in 2100 as a percentage of GDP.
>
> **Extended results text:** Figure 10b-10c also allow interpretation of the corresponding APS-NA and REF-NA scenarios. In the absence of adaptation, total climate-change costs would consist of the sum of residual damages and avoided damages shown here, since the avoided-damage component represents the damages prevented by optimal adaptation. Thus, APS-NA can be inferred from Fig. 10c as the level of climate-change costs that would prevail under the APS mitigation pathway without adaptation, while APS-OA shows how those gross damages are redistributed into residual damages, adaptation expenditures, and avoided damages once optimal adaptation is introduced. The same interpretation applies to REF-NA relative to REF-OA in Fig.10b. Therefore, comparing REF and APS scenarios further shows that stronger mitigation in the APS scenario lowers both the gross damages (in NA case) and the residual damages (in OA case).

- ## Broader validation beyond emissions

Reviewer comment: Section 3.1: The model is calibrated to 2015 and validated against observed 2015 and 2020 Baseline emissions (l. 326-330). It would strengthen the paper to also assess whether other outputs, such as GDP, align with historical data.

Response: We agree that the validation discussion can be strengthened by providing additional information beyond emissions. In the revised manuscript, we to extended content to Appendix A section describing the GDP and population assumptions used in the model, and the calibration of the main drivers of the model (Appendix E). This provides a broader view of the model's empirical grounding and better document how the baseline aligns with historical development patterns.

> **Extended Appendix A:** Population growth and economic trajectories are central determinants of future climate challenges, shaping both mitigation needs and adaptation capacity. Demographic and income pathways influence vulnerability: larger and less affluent populations are generally more exposed to climate risks and have fewer resources to cope with impacts (O'Neill et al., 2014). In AD-MERGE 2.0, socioeconomic inputs are based on SSP2 projections of population and GDP from 2015 to 2100 (Cuaresma, 2017; KC et al., 2024; IIASA, 2024). Figure A1 presents the regional population and GDP trajectories used in the calibration.
>
> **Added Appendix E:** In addition to the long-term consistency checks discussed in Section 3.1, we also evaluate whether the model reproduces observed or benchmark values in the base year and near-term period. This validation is particularly important because AD-MERGE 2.0's base year is 2015 and then projected forward through an intertemporal

optimization framework. Therefore, the calibration must reproduce the main socioeconomic, energy-system, and emission conditions at the beginning of the simulation horizon before being used for long-term scenario analysis…

Response to Reviewer #3

We thank Reviewer #3 for the positive evaluation of the manuscript and for the important comment on the interpretation of land use and BECCS deployment. This comment helped us clarify the current model boundary and the implications for interpreting carbon-removal scenarios.

Comment 3.1: Representation of land use and BECCS

Reviewer comment: The manuscript highlights the role of negative emission options such as Bioenergy with Carbon Capture and Storage (BECCS), particularly under mitigation scenarios where BECCS appears to play an important role. Given the strong link between BECCS deployment and land use dynamics, it would be helpful if the authors could clarify how land use change is represented in the AD-MERGE 2.0 framework.

It would be helpful to clarify whether the model represents land availability and competition (e.g., food, bioenergy, forests), how BECCS feedstock supply is determined, whether afforestation and bioenergy expansion share a land-use framework, and whether land constraints limit BECCS deployment. Clarification would help interpret the feasibility of the reported BECCS scenarios.

Response: We thank the reviewer for this important comment. We agree that the interpretation of BECCS deployment in AD-MERGE 2.0 requires a clearer explanation of how biomass supply and land-use interactions are represented in the current model version.

In AD-MERGE 2.0, BECCS is represented as one of the carbon removal technologies within the enhanced energy system module. However, the model does not include an explicit endogenous land-use framework that allocates land across competing uses such as food production, bioenergy crops, and forests. Consequently, land availability, land-use change, and land competition are not modeled directly, and afforestation and bioenergy expansion are not linked through a shared land-allocation structure. Instead, biomass availability for BECCS is represented in an aggregated way through broad regional resource potential and techno-economic assumptions, consistent with the model's level of spatial aggregation across 15 world regions. Therefore, BECCS deployment is not determined through an explicit endogenous simulation of feedstock supply under competing land demands, and detailed land constraints do not bind BECCS through a dedicated land-use module.

We agree that this point should be stated more explicitly in the manuscript, particularly because the current text notes the presence of carbon removal technologies in the energy system module but does not sufficiently clarify the interpretation of BECCS feasibility in the absence of explicit land-use competition. At the same time, we note that the manuscript already identifies the integration of additional modules and constraints as a direction for future model development, including limitations related to the current representation of climate responses and the scope of modeled processes.

We will therefore revise the manuscript in two ways. First, we will add a clarification in the methodology section stating that BECCS is represented through broad regional biomass/resource potential rather than an endogenous land-use allocation framework. Second, we will strengthen the limitation/future research discussion to state more directly that integrating an explicit land-use module would improve the assessment of BECCS feasibility, land competition, and interactions with other land-based mitigation options.

Extended Section 2.4 text: BECCS is represented within the energy system module using broad regional resource potential and techno-economic assumptions. However, AD-MERGE 2.0 does not include an explicit endogenous land-use framework that allocates land across competing uses such as food production, bioenergy crops, and forests. Accordingly, land-use change, land competition, and land availability are not modeled directly, and afforestation and bioenergy expansion are not linked through a shared land-allocation structure. CDR deployment should therefore be interpreted as conditional on aggregated regional potential assumptions rather than as a full assessment of land-use feasibility under explicit competing land demands.

Revised conclusion/future research text: Furthermore, AD-MERGE 2.0 can benefit from a more explicit representation of land-use dynamics and their interactions with energy systems and climate policies. In the current model, biomass availability for options such as BECCS is represented through broad regional resource potentials and techno-economic assumptions, rather than through an endogenous land-allocation framework. Consequently, land-use change, land availability, and competition across food, bioenergy, and forests are not modeled directly. Incorporating these dynamics would improve the assessment of biodiversity, food security, carbon sequestration, and the feasibility of biomass-based mitigation pathways. In this respect, IAMs such as GCAM and IMAGE explicitly represent land-use dynamics within the modeling framework (Kyle et al., 2011; Doelman et al., 2018), while MESSAGE and WITCH achieve this integration when linked to the GLOBIOM land-use model (Emmerling et al., 2016; Krey et al., 2020) or REMIND when linked to MAgPIE (Hilaire and Bertram, 2020). Together, these approaches enable a more comprehensive representation of cross-sectoral interactions between energy systems, land use, and climate.

Extended Appendix F text: The BECCS and DACCS deployment levels reported in this study should be interpreted in light of the model structure described in the methodology. In particular, BECCS deployment reflects the broad regional biomass/resource potential and techno-economic assumptions adopted in AD-MERGE 2.0. Accordingly, the reported BECCS trajectories should be interpreted as model-based mitigation outcomes within the current aggregated framework. DACCS is also not subject to direct land-use constraints, and its deployment remains conditional on assumed technology costs, energy-system integration, and overall mitigation requirements. The reported CDR levels should therefore be read as scenario outcomes conditional on the current model structure.

Closing statement

We believe that these revisions substantially improve the manuscript by clarifying the scope of the adaptation results, strengthening the documentation expected for a model-description paper, correcting the energy-accounting presentation, and making the interpretation of BECCS and other CDR trajectories more transparent. We thank the Reviewers again for their constructive comments.