nature portfolio

Peer Review File

Optimal Life-Cycle Adaptation of Coastal Infrastructure under Climate Change

Corresponding Author: Dr Gordon Warn

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
The authors present a mathematical approach to consider adaptive (sequential) decision making when optimizing flood risk mitigation, which is based on Markov decision processes. This enables the identification of optimal risk mitigation strategies that account for the fact that there is significant uncertainty on climate change effects (here: sea level rise) and that future decisions can be made under reduced uncertainty.

The topic under investigation is highly relevant and the paper is very well written. It presents well-developed case studies, which clearly show the applicability of the methods. However, the paper falls a bit short in discussing assumptions and their implications on the results.
In the following I list points that I think should be addressed/discussed in order to provide more value to the reader. (Note that I do not expect the authors to include all this points into their analysis, but I would expect that these issues are discussed and that the obtained results are interpreted in view of these points.)

We have found that a major challenge with this type of analysis is the model of the time-variant uncertainty in the quantity of interest (here: sea-level rise) and the way information will become available in the future (e.g., new climate models, etc.). The authors choose to model the sea level by a Markov process, whose transition probabilities are learned from trajectories generated by the probabilistic climate models. The authors show that this Markov model leads to marginal distributions that are consistent with those of the input trajectories. This is surprising as I would expect that epistemic (knowledge) uncertainty should be part of the probabilistic climate modeling; and such epistemic uncertainty breaks the Markovian assumption. I would ask the authors to give a bit more information on the probabilistic climate model that is used and comment on the epistemic uncertainties in that model.

Related to the above is the fact that future improved climate models will – with high probability – lead to reduced uncertainty. This is not reflected in the modeling done by the authors. The model as it is assumes that the only additional information in the future is the observation of the sea level. Clearly, this is an underestimation of what additional information will be available in the future. So the question of how to quantify this increased knowledge seems essential here. Note that this is investigated in the scientific literature, see e.g. Völz and Hinkel 2023 for a review. I would be interested in reading from the authors how they think these factors affect the results that they obtain. (Given that more information is available in the future, the benefit of sequential decision making should increase.)

Importantly, whether it is considered in the analysis or not, new information will become available in the future, and future decisions (e.g., beyond 10 years from now) will utilize that information. Hence, for practical purposes the difference between a classical static optimization and the sequential decision optimization lies mainly in the decisions that are taken in the short- to mid-term (i.e. in the 10 years until the decisions are revisited). For this reason, it seems not entirely correct to compare to static strategies that only implement decisions, which might be pointed out in the discussion.

Even when assuming the stationary policy to not be updated over the entire life-span, the authors find for their example that the difference between the best static and the sequential policies are only in the order of a few percent in terms of the expected life-cycle cost. I missed a discussion by the authors as to why this is the case. As an input to this discussion: What we have found in studies on adaptation of flood mitigation is that for flood risk, the uncertainty is already high even in the static case and as a result, the objective function is pretty flat around the optimum. The additional uncertainty of the climate uncertainty does not change this much. This seems to be the situation also here, as the (aleatory) uncertainty in the storm surge is much larger than the uncertainty in the sea level rise.

Note that concepts of stochastic control have been previously suggested for optimal adaptive planning of flood risk mitigation considering climate uncertainty, e.g., Spackova and Straub 2017. Furthermore, the use of real-options analysis, which solves a related problem, also has been suggested for this application in the literature, e.g., Woodward et al 2014.

The authors emphasize multiple times the “global optimality guarantee”. This is true within the confines of the model they use. However, given the substantial simplifications of the model, as discussed above, this terminology can be misleading. They should restrict this description to the solution of the mathematical problem they pose, but not to the flood risk mitigation strategy.

In the abstract: “less informed management policies” is unclear. Less informed than what?

References:
Völz, V., & Hinkel, J. (2023). Climate learning scenarios for adaptation decision analyses: Review and classification. Climate Risk Management, 100512.
Špačková, O., & Straub, D. (2017). Long-term adaption decisions via fully and partially observable Markov decision processes. Sustainable and Resilient Infrastructure, 2(1), 37-58.
Woodward, M., Kapelan, Z., Gouldby, B. (2014). Adaptive Flood Risk Management Under Climate Change Uncertainty Using Real Options and Optimization. Risk Analysis 34, 75–92. doi:10.1111/risa.12088

Reviewer #2

(Remarks to the Author)
The goal of this paper is to demonstrate an approach using Markov decision processes and approximate dynamic programming to identify sequential decisions for coastal flood risk management. However, as currently framed, I do not think the paper is publishable in Nature Communications. The paper is currently primarily framed in terms of the method used, rather than the conceptual insights or results, but most of the technical advances of the paper have been done already (and are not captured by the literature review) or there is not much of a case for the benefit of the demonstrated approach over existing approaches in the literature. hat being said, if the paper were heavily re-framed (see below), it might be a different story.

General Comments:

The application of standard control methods, including dynamic programming, to coastal decision making, is not new, though much of this literature is not referenced in the paper. For example, Lickley et al (2014) used DP, and Shuvo et al (2022) use reinforcement learning to simulate socioeconomic development and adaptation to sea-level rise. In fact, even the classic Eijgenrraam et al (2014) paper used DP to identify sequential decisions.

Therefore, the main advantage of the method in the paper is the use of an approximate DP algorithm with learning about probabilities about likely states. But (in addition to this approximate approach not quite having as strong of a claim to global optimality as the paper presents), there is also the possibility of other closed-loop control approaches, such as direct policy search (see Quinn et al (2016) more generally and Garner and Keller (2018) for an application to levee heightening). DPS addresses the curse of dimensionality that the ADP approach in this manuscript suffers from (hence the stylization of the problem), and while DPS, as a heuristic optimization method, does not include the same optimality guarantees as DP, it is also not entirely clear why that is more important given the dimensionality of the relevant uncertainties, especially across the wide range of policy options, as well as the limited cost improvement in the identified solutions throughout the paper. There is also the dynamic adaptive policy pathways approach, which may not yield optimal results but may result in identifications of potential lock-in and be more flexible overall. If this version of the paper is to be publishable, quite a lot more work has to go into the literature review and better contextualization for why this is a useful method for the problem at hand (versus the way the paper is framed, which reads as though the method is the main focal point, as this takes up the bulk of the introduction and the discussion).

However, what is somewhat new and interesting is the role of the social cost of carbon in changing optimal solutions (I have yet to see this for this kind of decision problem), and I would encourage the authors to focus on these insights in any revision of the manuscript, particularly if they looked at SCC uncertainty. Another option would be to take the application more seriously and build out the range of policy options, rather than stylizing them to demonstrate the method; while this would run into the curse of dimensionality, this is a core issue with any DP method (and one that ADP is intended to at least partially address).

Specific Comments:

Line 50: The word "optimal" is used a lot without specifying what the objective is.

Line 54: There is a discussion of a variety of stakeholder objectives, but as far as I can tell none of these are taken into account beyond loosely interpreting environmental objectives as the SCC (though the SCC itself can be a fraught metric). What is the justification for focusing entirely on a single economic objective, even one which has the SCC included as a proxy for environmental considerations? In fact, a posteriori decision analyses, which are often multi-objective, can facilitate robustness analyses and stakeholder engagement in a way which a priori analyses cannot.

Line 56: A claim that criteria-based protection is currently the standard needs citation, and I'm not sure that it's entirely true, even if the USACE approach largely boils down to this assuming a cost-benefit test is passed. For example, the Netherlands has moved away from this type of design standard to an approach which regulates the expected loss of life due to flooding (see https://www.enwinfo.nl/publish/pages/183541/grondslagenen-lowresspread3-v_3.pdf).

Line 61: Does risk-based estimation need to focus on economic losses, or just the sampled papers?

Line 77: Optimization against an ensemble of stochastic realizations doesn't need to optimize the expected outcome, a robust optimization approach could be used which would optimize against a quantile (or any other summary statistic of interest).

Lines 85-88: I don't follow this discussion of adaptive approaches. The whole point of adaptativity is to make a flexible initial investment and then pivot based on learning about the more probable states of the world. The identification of adaptation tipping points serves effectively the same role as the learning in the manuscript's approach.

Line 124: The text makes it seem as though the presented approach is independent of the problems associated with sampling stochastic realizations of the relevant uncertainty space. But it's still an approximate dynamic programming method, and therefore requires samples, which may be more or less representative of the overall distribution, especially if the probability model is poorly specified. This is an issue which is under treated in this manuscript: in lines 228-229, there is a mention that 1,000 Monte Carlo samples were used without this being justified further (for example, what was the Monte Carlo standard error?), and very limited SSP scenarios were used. Also, despite the discussion of non-stationarity and the emphasis on fitting a model for higher-frequency mean SLR variations (without much justification), a standard variance is used for the future projections without justification. Why is this probability model appropriate?

Line 199: What kind of weather patterns should be captured in the mean SLR projections? This gets back to the justification for incorporating the higher-frequency variability at the expense of uncertainty across SSPs.

Line 214-216: For annual extremes (I assume that was the temporal resolution, though this was not made particularly clear), you shouldn't need to take away the tides, and in fact might end up under-estimating the extremes by only accounting for the storm surge and not the storm tides. You would need to for shorter horizons where the data are correlated due to tidal patterns, such as monthly.

Line 278-9: How do you know this is enough Monte Carlo samples? Was this number chosen because of a low Monte Carlo standard error, or because it was computationally tractable?

Line 481: I don't understand the discretization of SLR values. Is this per time step? What was the point of the Gaussian fit if you were just going to discretize over a grid?

Code Availability: The code should be made available to reviewers in preparation for public release upon potential publication, to be consistent with Nature policy.

References:

Eijgenraam, C., Kind, J., Bak, C., Brekelmans, R., den Hertog, D., Duits, M., et al. (02/2014). Economically efficient standards to protect the Netherlands against flooding. Interfaces, 44(1), 7–21. https://doi.org/10.1287/inte.2013.0721

Garner, G. G., & Keller, K. (2018). Using direct policy search to identify robust strategies in adapting to uncertain sea-level rise and storm surge. Environmental Modelling & Software, 107, 96–104. https://doi.org/10.1016/j.envsoft.2018.05.006
Lickley, M. J., Lin, N., & Jacoby, H. D. (2014). Analysis of coastal protection under rising flood risk. Climate Risk Management, 6, 18–26. https://doi.org/10.1016/j.crm.2015.01.001

Quinn, J. D., Reed, P. M., & Keller, K. (2017). Direct policy search for robust multi-objective management of deeply uncertain socio-ecological tipping points. Environmental Modelling & Software, 92, 125–141. https://doi.org/10.1016/j.envsoft.2017.02.017

Shuvo, S. S., Yilmaz, Y., Bush, A., & Hafen, M. (2022). Modeling and simulating adaptation strategies against sea-level rise using multiagent deep reinforcement learning. IEEE Transactions on Computational Social Systems, 9(4), 1185–1196. https://doi.org/10.1109/tcss.2021.3122282

Reviewer #3

(Remarks to the Author)
The study introduces key findings in coastal infrastructure adaptation to climate change, highlighting the effectiveness of Markov Decision Process (MDP)-based policies in cost reduction, especially for coastal marshes and reefs, saving up to

$0.018 million per meter compared to static policies. It emphasizes the MDP's capacity for achieving globally optimal solutions in complex scenarios, underscoring the importance of cost-effective adaptation strategies. The inclusion of environmental impacts, like carbon emissions, in the adaptation framework underscores the commitment to sustainable strategies. The dynamic and flexible MDP approach surpasses traditional models by accommodating uncertainties and real-time data, marking a significant improvement in the field. The methodological rigor and reproducibility of the study suggest its potential for broad application and validation. Overall, the research underscores the substantial cost and environmental benefits of MDP-based adaptation, proposing a notable advancement in coastal infrastructure planning, albeit suggesting further validation and comparison with existing models for a stronger foundation.
A suggestion I have for the authors for their future work is to investigate deep reinforcement learning techniques to manage the increasing complexity and dimensionality of the problem. This approach could allow for handling a larger set of adaptive actions and longer decision horizons, potentially leading to more substantial cost savings and improved policy effectiveness.

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
The authors' considerable effort to address the comments is highly appreciated. I am also happy with most responses and modifications to the manuscript. I have only two remaining points, which the authors should consider for the final manuscript:

1) I am not entirely satisfied with the response to my comment 1.3. What the authors write is correct, but my comment was directed towards the fact that the current practice is indeed a static policy within an open-loop planning framework (as the authors nicely put it). The authors comparison however assumes that the static policy is never changed, which makes it unnecessarily pessimistic. To be more specific, I refer to Figure 5: If my understanding is correct, both the MDP and the static policy optimization recommend the construction of F1 initially. If one only looks at the decisions taken within the nearer future, they are identical for the two approaches. The fact that the optimal static policy recommends not to build F2 later is irrelevant, as this decision will be revised later. But since this revision is not accounted for in the current calculations done in the paper, the estimated life-cycle costs for this policy are pessimistic. (If one were to compute the expected cost of the open-loop planning with static policy, that would be closer to the one of the MDP model).
(The fact that also the MDP model can also be updated a later point is irrelevant here, as the comparisons assume that dynamics underlying the MDP model are the true one.)
I would urge the authors to state this clearly in the manuscript, as this is an important point that is often misunderstood.

2) This is a minor point, but in line 261ff the authors write that the study "presents a general life-cycle adaptation framework under climate change that extends the MDP-based problem formulation …". This makes it sound like the MDP formulation was the current state of the art or practice. However, the current state of practice is the use of static policies in an open-loop planning framework, as discussed above. The current state of the art in science are POMDPs, which were considered in some of the references now included in the paper.

Some of the new text that the authors included involves very long sentences that are a bit hard to read (E.g., lines 183 – 193). It is recommended to check the manuscript for readability.

(Remarks on code availability)

Reviewer #2

(Remarks to the Author)
I have spent some time (and a few rounds) reading the revision of "Optimal Life-Cycle Adaptation of Coastal Infrastructure under Climate Change," and I unfortunately cannot convince myself that the authors have quite responded to a core critique from the initial review (both from myself and from Reviewer 1), which I will discuss more below. I unfortunately cannot recommend publication in Nature Communications (or a similar broad-audience journal) for this reason, though a major revision might be in order if this was a miscommunication or misinterpretation, rather than an attempt to avoid the critique.

Let me be more direct to avoid any potential miscommunication. The fundamental problem with this manuscript (for this journal; this may not apply to more discipline-specific journals) is that, for an article of this sort (I might frame this more in terms of decision analyses, but I think this applies more generally) to be useful to a broad audience, it should contain at least one of the following elements:

1) Development of a new method which addresses limitations of existing methods;
2) Application of an existing method to a problem which highlights new insights about the problem or its setting;
3) Application of an existing method to a problem which is a demonstrated advance on the state of the art (or the state of the practice, if it can be articulated why this method may be more implementable by practitioners than other methods which advance on the state of the practice but are not commonly used).

That this article fails to do any of these three was at the root of my comments about the framing of the paper (including the large gaps in the literature review) and the comment Reviewer 1 and I made about the overly stylized example (in terms of the decision space) and the corresponding lack of model fidelity to a "real" decision problem. The paper is applying an existing method (FRTDP) in what is hopefully a new context, and so the contribution is not in method development. The paper's framing does not make any claims to new insights (with the exception of the inclusion of the social cost of carbon in changing the cost-benefit analysis), but rather that the FRTDP algorithm improves on the state of the art through its adaptativity and by having optimality guarantees.

The problem with these claims is that there are many other adaptive methods; while several (such as DPS) do not have optimality guarantees, these are of limited practical application since one's model is never "correct," and the "optimal" solution to the modeled problem may be suboptimal for the "true" problem. In their revision, the authors have tried to address the issue of model error through hand-waving away the issue by claiming that, in principle, as the FRTDP algorithm has optimality guarantees, it would find the optimal solution for any given problem framing, so anyone could just select the method they prefer. The issue with this claim is that the benefits of the method are presented as an advance on other methods, but only compared directly to a static method (which is a strawman compared to the literature, though perhaps not compared to the practice; more on this below). It might even be that for this simplified problem (limited decision levers; independent [versus autoregressive] noise around the future projections; see also the response to my comment about sampling states of the world, which might be required for problems which are more complex), the FRTDP algorithm does indeed outperform other algorithms which have been used in the literature, but the paper does not make this comparison, instead relying on its characterization of other methods as a reason to prefer FRTDP. But the characterization of DPS, for example, is a bit glib: it claims (on lines 99-103) that DPS relies on heuristics (true, but may not practically matter), that the actions are "arbitrarily defined" and "explore only a limited subset of the policy space." But this is both true of the example application and is true of pretty much every model-based decision problem, not unique to DPS (one could set up a DPS problem to capture the same set of policy options as the FRTDP example). The problem with the stylization of the problem in this manuscript is that it is unclear whether, as a practical matter, the theoretical guarantees of the FRTDP algorithm make a difference when the problem becomes sufficiently complex, and this is impossible to understand from the manuscript.

This would not be a problem if the FRTDP algorithm addressed practical implementation problems with other adaptation methods, which are why many of these methods have not caught on in application; these range from framing complexity to computational complexity to transparency and communication complexity. But none of this is addressed by the manuscript other than comparing the number of samples the authors needed (which is conditional on the problem framing) to the number of all possible combinations, which is not the fair comparison. The best way to address this issue would be to include other analogous methods in the comparison to show that FRTDP improves on them in this case, and argue why this ought to generalize; it would also be appropriate to make the argument on the basis of the properties of the other methods, though the authors do not have a comprehensive enough literature review to do this, and I am afraid based on their characterization of DPS discussed above that this may not be a useful exercise in this case.

As a result, I feel that the issue of model fidelity and model error is a central problem for publication in a broadly read journal like Nature Communications, and one which the authors failed to respond to in their revision. This may be less of a problem for other journals which are more focused on disciplinary audiences, but I would not speak for them.

Rather than emphasizing the methods, the authors do have a route towards new insights, namely the role of monetizing future carbon emissions in changing decisions; in this case, the MDP formualtion and the FRTDP algorithm do provide a route towards identifying these effects, and the issue of "optimality" is less important. But the authors did actively choose to not reframe their manuscript in this fashion in the previous revision, so I am unsure if they are interested in doing so. However, that would in my view obviate the entire argument against publication I raised above (though others might emerge).

Specific Comments (these are somewhat limited because I flagged some before coming to the conclusion that the paper does require a much more thorough reframing for publication in Nature Communications):

Lines 14-17: This initial line in the abstract is the core argument about the focus on a static counterfactual, with very little evidence in the introduction to support it.

Lines 49-56: This sentence makes it seem as though the analysis would address a more complex problem formulation involving multiple interacting decision levers, which would actually be a relevant level of complexity for comparison to other methods, and one which is rarely done in the literature, but the actual examples do not do this.

Fig 1 Caption: Aren't there 21 climate projections used? Which SSP5-8.5 realization did you use in panel A?

Lines 99-103: Again, this is not quite a fair characterization of DPS, and I worry a bit about whether it is hyperbolic to make the method in the manuscript look better. DPS does not have to explore any more of a "limited" subset of the policy space than any other approach (though it may get more complex to formulate, implement, and communicate as the policy space expands, particularly to multiple levers), and often the state space mapping is formulated using exploratory analyses of relevant sensitivities and options, so is not any more arbitrary than the limited policy options used here.

Lines 108-112: Are you claiming that adaptive pathways necessitate a "few" forward simulations? There are no citations to justify this, or specifics given on what "a few" means.

Line 113: From a real-world perspective, there is no guarantee that even a theoretically "optimal" solution is optimal, hence the importance of model fidelity and the salience of the problem framing, which is totally missing from this manuscript despite its grandiose claims.

Lines 129-130: There are other examples of approaches in the literature for sea-level rise adaptation which also use MDP and closed loop controls (I pointed out some in the previous review) but the paper still reads as though it is claiming to be the first to have done so. This shows the limits to the improvement of the literature review in this revision. I get the sense from reading the paper that the authors are trying to overstate the novelty of the paper and the strength of its contribution, which is not needed (or helpful!).

Lines 138-139: I don't understand what you mean by "quite specific to the characteristics of their particular application scenarios." Do you mean that the decision problems are formulated to be specific to e.g. levee heightening or sea-level rise generally? Depending on what you mean, I am not sure how this is avoidable in any situation. If you mean that the example is specific to a particular case study, this seems like a more tenuous claim: details will always be specific to the particular case study, but it's rare that methods would only work in one location, isn't it?

Lines 149-153: This again makes it seem as though the paper will incorporate multiple levers in the same example, which would (I believe) be a strong contribution, but this is not done.

Line 290-296: The assumption that historical variability would be stationary in the future (despite possible changes to land-water storage or rapid ice sheet changes, which may be captured in a very limited way in the sampled climate projections) is actually quite a strong claim and deserves some justification.

Line 377 (and others): Is the improvement over the static baselines actually meaningful? The authors report decreased costs of several thousands to tens of thousands of US dollars (depending on the specific example) per meter or coastline, and it is not entirely clear what this means in the context of how stylized the problem is and the barriers to implementing adaptive solutions. Would these savings justify this approach? It might be better to calculate the overall savings for the case study (which could then be made into a rate), because right now the savings don't seem like much.

Lines 386-389: As mentioned above, the issue about the role of cost of carbon in e.g. incentivizing green infrastructure vs. grey infrastructure is an interesting one, and one which your framework is suitable to address. But it reads as somewhat of a throwaway in the current framing.

(Remarks on code availability)
I don't have MATLAB (and so did not replicate), but the workflow for reproducibility looks reasonable. While the README provides the steps for the reproducability workflow, it is not clear from the README how to use the code for a different problem, which is ok, but limits the code as a resource for the community beyond the reproduction.

Version 2:

Reviewer comments:

Reviewer #2

(Remarks to the Author)
The authors have addressed a number of my concerns, but I do believe the revised manuscript is still potentially understating how increased model complexity could impact the practicality of their approach. While the reasoning in the response to reviewers makes sense, most of it is not included in the discussion (just the point about how, in principle, this approach could serve as a benchmark, which I believe might be misleading if the desired model is not computationally tractable).

(Remarks on code availability)
The codebase looks reasonable and the workflow reproducible. The README needs to be improved but the authors have seen this in my prior review.

Notes on revisions made to manuscript Paper NCOMMS-23-48445

We would like to thank the editor and the reviewers for their insightful and constructive comments that helped us improve the manuscript. In response, we have significantly revised and extended the paper, including a more extensive discussion of the suggested framework and limitations, also alongside existing approaches, the role of uncertainty in its various forms and its influence on policy decisions, and the development of a POMDP model, on top of the MDP one, which can also include hidden states, related here to the uncertainty associated with the underlying climate scenario models. A point-by-point response to the reviewers' comments is provided in this document, and all new additions to the manuscript are highlighted in the paper and this document in blue. Due to the extensive revisions with respect to the aforementioned topics, the related currently available approaches and the differentiation of our work from the existing literature, we have also included here at the beginning relevant parts of the updated manuscript, for ease of reading and reviewers' convenience. Subsequently, all reviewers' comments are discussed independently and in detail.

Lines 57-170, 177-194: The identification of the optimal flood management policy is in general non-trivial, primarily due to various levels of associated uncertainties [1, 2], which can include (*i*) aleatoric (irreducible) uncertainty associated with the natural variability of climate change processes and extreme events, (*ii*) epistemic (reducible) uncertainty associated with imperfect knowledge of the underlying climate scenario models, related also to (*iii*) uncertainty regarding future knowledge and updates of climate models, (*iv*) epistemic uncertainty associated with the evolution of future trajectories given assumed climate scenario models, and (*v*) uncertainty associated with available observations in time. Depending on whether and how observations and information in time are used to guide policy plans and lead to more-informed decisions, the decision-making frameworks available in the literature can be broadly categorized into static/non-adaptive (no utilization of additional information in time) and dynamic/adaptive (additional information is utilized in time).

A conventional approach for climate change planning has been primarily driven by cost-benefit analysis (CBA). Typical CBA solves a static optimization problem where the optimal time of investments and the optimal infrastructure related actions are determined by maximizing the benefits of the sought policy benefits while minimizing the related investment costs over the planning horizon. Most of the work in the existing literature implement such static optimization frameworks, often within settings focused solely on levee heightening. The expected flood damages within such frameworks are determined by assuming a fixed base-case scenario of the evolution of the probability of failure [3] or by considering an ensemble range of future SLR and storm surge events to be representative of the range of possible future scenarios [2, 4–9]. Different solution techniques have been utilized to solve these static optimization problems, including genetic algorithms [4], dynamic programming, and non-linear programming [3].

The result of the static optimization approach, as defined here, is a single policy which is deemed optimal over the average, or any chosen percentile, of potential future climate trajectories; and this policy is to be implemented without further guidance by the actually evolving climate conditions observable in time. However, the actual climate trajectory may deviate significantly from the ensemble mean (or the chosen target percentile) owing to the substantial amount of uncertainty present in climate model projections [10], as illustrated in Figure 1A, which inevitably means that the identified policy is highly likely to be sub-optimal with respect to the actual climate trajectory that evolves in time. Thus, static optimization approaches solely rely on simulation results and do not utilize actual observations in time to update the knowledge regarding the associated risks and any of the (ii)-(v) related aforementioned uncertainties.

There are a few works in the literature offering adaptive planning approaches, as defined previously, which determine policy actions depending on the evolving, observed climate conditions in time. This existing literature on adaptive planning of infrastructure systems against climate change typically apply methodologies including direct policy search (DPS) [11, 12], real options analysis (ROA) [13, 14], and stochastic control methods [15–19]. With DPS, the proposed actions are condition/state-dependent, as observed in time, however, the expressions relating the observable state variables to the triggering actions are arbitrarily defined and heuristically trained. As such, DPS lacks optimality guarantees and explores only a relatively limited subset of the policy space. ROA also offers an adaptive planning framework, typically based on binomial decision trees with two possible actions/options at each node [14]. The action at each node is determined by computing action values backwards in time. This analysis represents a form of backward induction which can quickly become computationally intractable as the number of branches increases exponentially with the number of time steps, actions, and observations. Alternative approaches have also been suggested through adaptive pathway solutions [20–22] by identification, through explorative modeling, of adaptation tipping points in time. However, these methods rely on only a few forward simulations/trajectories of possible futures in time, to identify the tipping points, sharing many of the limitations of the cost-benefit and DPS methods, and making the adaptive pathway solutions likely to be sub-optimal with respect to the actual evolving conditions in time. Overall, these adaptive methods have been typically applied in levee construction/heightening applications, where a single levee exists or is constructed at one point in time and is then subsequently heightened in future time-steps.

Examples of adaptive planning optimization against climate change effects using stochastic control frameworks can be found in the relatively recent literature, where the sequential decision-making process is formulated as a Markov Decision Process (MDP) or Partially Observable MDP (POMDP), with state variables typically representing the related conditions, and decisions being the possible infrastructure actions. The MDP/POMDP solution determines optimal actions to be taken at each decision time-step, conditioned on the observed, or inferred for POMDPs, system states, and thus uniquely selected for the actual realized trajectory path, successfully optimizing the total overall expected costs/rewards over the decision horizon. MDP/POMDP approaches can offer superior and more general solutions, using the same models, with respect to other static and/or adaptive planning methods, as described above, since they are based on dynamic programming and closed-loop stochastic control foundations [23], that offer robust guarantees of global optimality of the action paths. A few notable works based on the MDP framework, also for levee heightening problems, can be seen in [15–17, 24]. Some of these works importantly also consider climate model uncertainty utilizing POMDPs and simple hydrological scenarios [16, 24]. However, the main limitation of standard MDP/POMDP solution approaches is the curse of dimensionality [25, 26], which often dictates an overly simplified representation of the system, and/or a crude approximate solution technique that does not offer optimality guarantees. The existing related methodologies in the literature largely follow these patterns and are quite specific to the characteristics of their particular application scenarios, with relatively few possible states at each point in time, while also employing approximate solution techniques, all in order to make a solution to the problem feasible. Utilization of the solution framework to other applications, and/or extensions is thus not straightforward, particularly also when a variety of possible actions needs to be considered.

In this work, we are presenting considerable new extensions and capabilities for adaptive planning approaches and a general and sophisticated solution framework for the optimal life-cycle adaptation of infrastructure under climate change. We are also relying on the firm mathematical foundations of closed-loop stochastic control through MDPs/POMDPs, while offering general, easily adaptable, and extensible implementations of infrastructure adaptation problems under climate change, with higher-fidelity representations, in comparison to currently available existing approaches, utilizing the IPCC models [10], accommodating and presenting a wide variety of possible actions, such as floodwalls, nature-based infrastructure, and seawalls, considering model uncertainty aspects, and incorporating social cost of carbon considerations in the analyses. Simultaneously, and in contrast to currently available relevant approaches in the literature, we are formulating the entire framework in such a way so that in can be generally supported by state-of-the-art, efficient MDP and POMDP optimization solvers with global optimality guarantees, based of course on the modeling assumptions used in the optimization process, enabling both accurate solutions and a finer representation of the problem settings based on higher-dimensional state-spaces and time/history dependent risks and costs, among others. Our suggested framework can also in principle accommodate all types of (ii)-(v) uncertainties mentioned earlier. As illustrated in Figure 1B, the integration of up-to-date annual information in relation to the current environment conditions significantly reduces the type (iv) uncertainty concerning the the evolution of the actually evolving trajectory, until the next observation becomes available. Although type (iii) uncertainties, regarding future knowledge and updates of climate models, are the only type that have not been directly addressed in this work, specific ways to incorporate these aspects in the framework are explained in the discussion section, together with still existing limitations and further possible future directions.

In the context of optimal adaptation of coastal infrastructure systems subjected to risks posed by storm surges and sea level rise in this paper, the states of our formulated MDP include the discretized SLR and storm surge levels, with the SLR states evolving probabilistically through transition dynamics computed from climate projections taken from the recent IPCC report (AR6) [10,27]. The problem states additionally include the infrastructure state, since the damage induced at any point in time is dependent not only on the SLR and storm surge levels but the implemented system protections as well. Unlike the existing, more simple adaptive policy settings of levee heightening, where the level of protection at any time is fully determined by the height of the levee, defined as the problem state, in our presented general setting, the infrastructure state and protection are determined by the entire history of the previously executed adaptation actions up to the current time step, given that these actions can completely now alter the system configuration. Although this essential characteristic for this type of problems is non-Markovian, given that the history of previously executed actions is now relevant, this issue has also been addressed in this work by appropriately and originally augmenting the state space in such a manner so as to enforce a Markovian formulation, which is also importantly generally consistent with state-of-the-art dynamic programming solvers.

We have further extended our framework by developing a POMDP model on top of the MDP one which can include hidden states, which are related to the uncertainty associated with the underlying climate scenario models.

Lines 215-220: In our suggested framework, the SLR levels are fully observed in time, while in our POMDP formulation, the underlying climate models, represented by the Shared Socio-economic

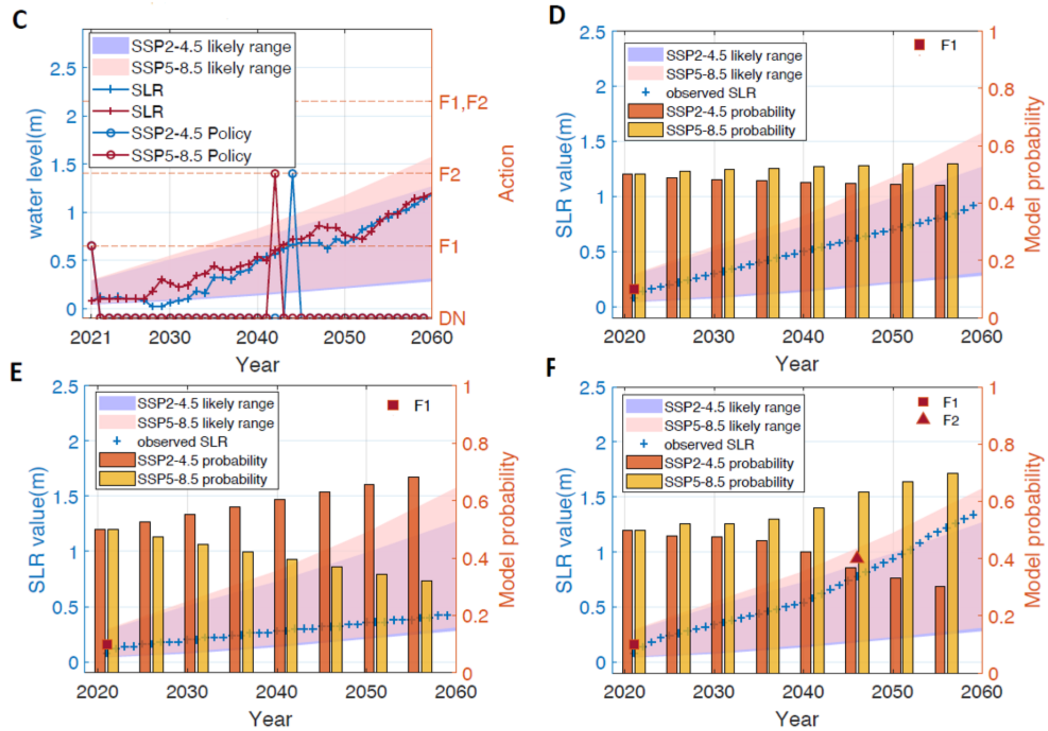

Figure 1: Effect of the model used for training policies (C), where two policies trained on different SSPs show very similar behavior under similar climate observations. The belief evolution over two SSP scenarios is shown at 5 years intervals based on three different observed SLR trajectories in time, in (D), (E) and (F). F1 and F2 in the legends of Figures (D), (E) and (F) correspond to the actions F1 and F2 being taken, respectively. (Figure 10 in the manuscript)

Pathways (SSPs), as provided in the most recent assessment report (AR6) of the Intergovernmental Panel on Climate Change (IPCC) [10], are hidden state variables which are only probabilistically inferred based on the SLR observations in time.

We show the detailed modeling of the POMDP framework in Section 4.2 in the paper and in the Supplementary materials. In Figure 1 D, E, F, we show the evolution of the beliefs over the considered climate scenario models in the formulated POMDP for three different possible trajectories, with the conclusions as reported in the paper:

Lines 504-513: Figure 10D shows a SLR trajectory where, at each time step, the mean of the possible SLR states is selected as the observed SLR. As seen, for such a SLR trajectory, the two models remain largely similarly probable. On the other hand, for a possible SLR trajectory where lower percentile SLR values are observed, as in Figure 10E, the belief over SSP2-4.5 gradually increases, as the likelihood of observing lower SLR states is higher under the SSP2-4.5 scenario than the SSP5-8.5 one. Accordingly, Figure 10F shows a SLR trajectory following a possible higher percentile, having higher likelihood of being observed under SSP5-8.5, and therefore the probability of SSP5-8.5 gradually increases as observations in time are collected.

A concise summary of the main contributions in this paper is provided as:

Lines 262-271: To summarize, this study makes the following significant contributions to the planning of climate-change-related risk mitigation:

1. Presents a general life-cycle adaptation framework under climate change that extends the MDP-based problem formulation to encompass a wide range of actions and to be suitable for state-of-the-art point-based MDP/POMDP solvers with global optimality guarantees.
2. Formulates a general POMDP problem, by considering climate model uncertainties updated in time through relevant observations.
3. Considers environmental impacts, the social cost of carbon, and related carbon emissions and uptake in the decision-making process.
4. Accounts for various nature-based infrastructure as potential flood risk mitigation measures, and demonstrates their effects in mitigating carbon emissions.

Response to the reviewers

Reviewer 1

The authors present a mathematical approach to consider adaptive (sequential) decision making when optimizing flood risk mitigation, which is based on Markov decision processes. This enables the identification of optimal risk mitigation strategies that account for the fact that there is significant uncertainty on climate change effects (here: sea level rise) and that future decisions can be made under reduced uncertainty.

The topic under investigation is highly relevant and the paper is very well written. It presents well-developed case studies, which clearly show the applicability of the methods. However, the paper falls a bit short in discussing assumptions and their implications on the results.

In the following I list points that I think should be addressed/discussed in order to provide more value to the reader. (Note that I do not expect the authors to include all this points into their analysis, but I would expect that these issues are discussed and that the obtained results are interpreted in view of these points.)

We thank the reviewer for the insightful and useful comments that helped us improve the manuscript, and the general appraisal of our work. We are responding to all the comments in detail below.

Reviewer Comment 1.1 — We have found that a major challenge with this type of analysis is the model of the time-variant uncertainty in the quantity of interest (here: sea-level rise) and the way information will become available in the future (e.g., new climate models, etc.). The authors choose to model the sea level by a Markov process, whose transition probabilities are learned from trajectories generated by the probabilistic climate models. The authors show that this Markov model leads to marginal distributions that are consistent with those of the input trajectories. This is surprising as I would expect that epistemic (knowledge) uncertainty should be part of the probabilistic climate modeling; and such epistemic uncertainty breaks the Markovian assumption.

I would ask the authors to give a bit more information on the probabilistic climate model that is used and comment on the epistemic uncertainties in that model.

Thank you for the comment. We have now provided more details about the used climate models in the paper, as shown below.

Lines 278-300: The SLR model is derived based on regional SLR projections using CMIP6 models from the AR6 of the IPCC [10] for the Battery tidal gauge [28] representing the general NYC area, without loss of generality. The SLR projections extending to 2150, relative to the 1995-2014 reference period, are obtained from [29]. These projections are based on various climate change and emission scenarios, represented by distinct Shared Socio-economic Pathways (SSPs), and SLR projections under three such SSPs are shown in Figure 3A. For each SSP model, an ensemble of 21 Global Climate Models (GCM) is considered in this work, in an effort to consider epistemic model uncertainties within a SSP scenario. A future SLR trajectory follows here a certain projection percentile within the likely range, which captures the long-term non-stationary SLR trend resulting from multiple factors, such as thermal expansion, ice sheet contributions, and glacial melt. However, these projections do not include the short-term variability inherent in the historical observations, which are contributions of various natural fluctuations. In order to also consider these variations in the future sea level rise projections, a regression model is fit to the observed historical data provided by the National Oceanic and Atmospheric Administration (NOAA) [30], as shown in Figure 3B, and the random noise present in the observations is utilized, assuming to follow a Gaussian distribution with a constant standard deviation of 3.18 cm, as also confirmed in the normal probability plot in Figure 3B. Indicative resulting sea level rise process simulations derived from the IPCC climate model projections, along with the historically observed variability, are illustrated in Figure 3C. The time-variant and uncertain sea level rise evolution is modeled by a Markov process with transition dynamics estimated based on these simulated trajectories.

Thus, the models are based on IPCC projections [10, 29] specifically downscaled for the Battery tidal gauge [28] representing the general NYC area. For each SSP model we used, we took the average over an ensemble of 21 Global Climate Models (GCM) for that SSP, in an effort to consider epistemic model uncertainties as well. The time-variant and uncertain sea level evolution is modeled by a Markov process with transition dynamics generated from trajectories simulated based on these SSP models, by following a certain percentile value within the likely range of the chosen SSP model and adding sea level variations based on historical data [30], as seen here in Figure 2.

The transition probabilities model estimated based on these trajectories is dependent on time (non-stationary). The probabilities are calculated for a specific SSP as shown here in Figure 3. The developed discrete-state non-stationary Markovian model (for a given SSP) is shown to lead to marginal distributions which are consistent with those generated from the probabilistic climate model (for the same SSP), as shown here in Figure 4, provided below for convenience.

Due to the non-stationarity of the transitions the model has become Markovian by augmenting the state-space with the year information as well. The state-space needs further augmentation for the problem to become Markovian in relation to the adaptation actions used, since they alter the system configuration. This is explained in Section 4.1.2 in the paper.

In this work, and inspired by the reviewers' comments, additional epistemic uncertainties, apart from taking the ensemble of GCMs for each SSP scenario, have now been considered directly in the POMDP case, which also quantifies the beliefs/ knowledge in time in relation to different SSP

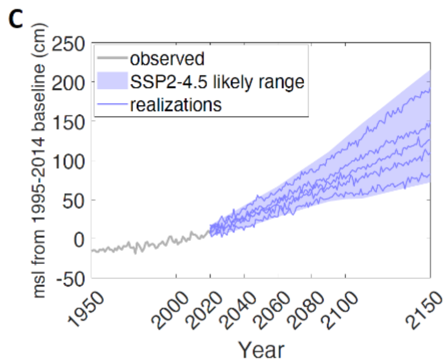

Figure 2: SLR realizations based on IPCC SSP2-4.5 projections along with the variability present in past observations. (Figure 3C in the manuscript)

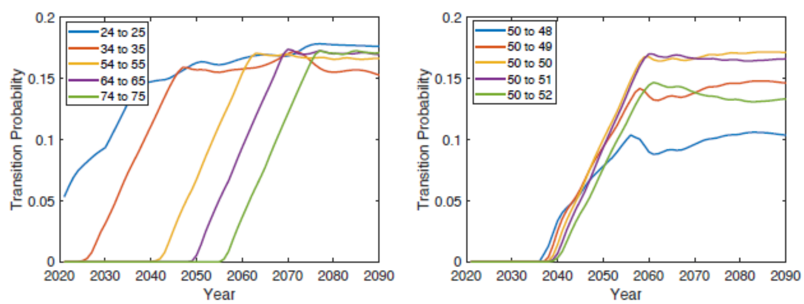

Fig. 3 Non-stationary transition probabilities for sea level rise states. Transition probabilities are smoothened with a moving average over a window of 5 years.

Figure 3: Non-stationary transition probabilities for sea level rise states. (Shown as Figure 4 in the supplementary materials)

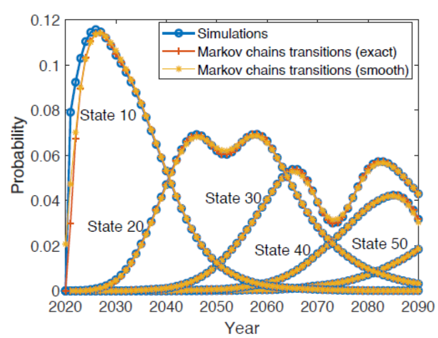

Fig. 4 Probability of being in a state starting from the initial state. Comparison of discrete Markov chain simulations based results with that obtained from continuous Monte-Carlo simulations. Smoothened transitions are obtained by moving average over a window of 5 years.

Figure 4: Probability of being in a state from the initial state (Figure 5 from supplementary materials)

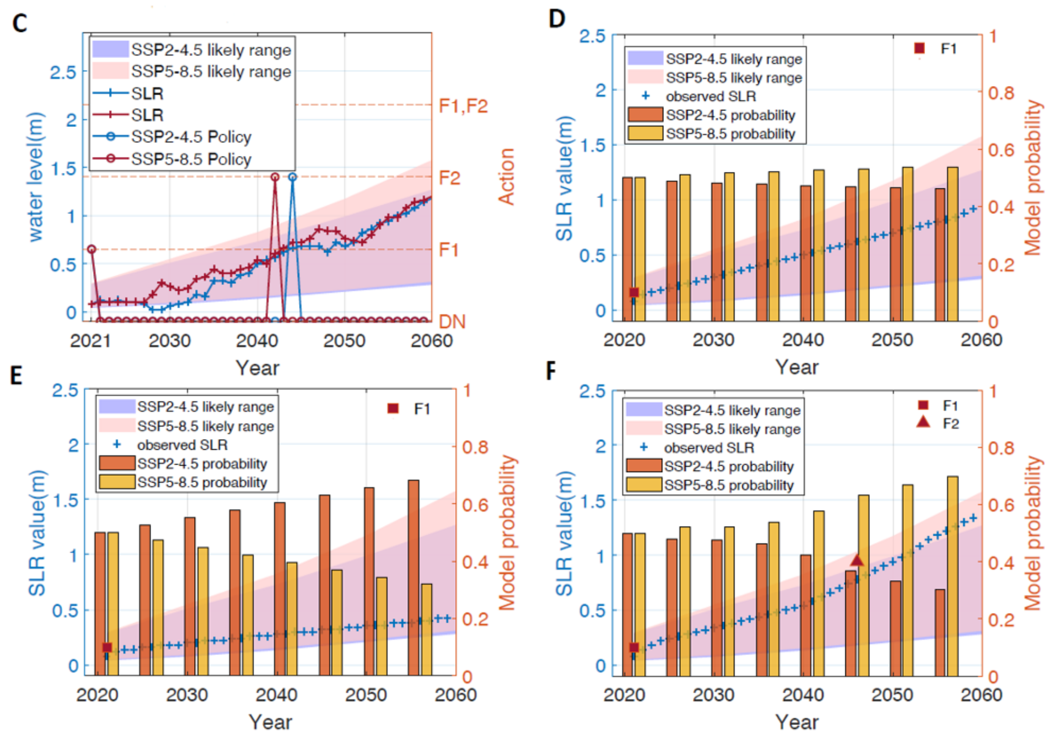

Figure 5: Effect of the model used for training policies (C), where two policies trained on different SSPs show very similar behavior under similar climate observations. The belief evolution over two SSP scenarios is shown at 5 years intervals based on three different observed SLR trajectories in time, in (D), (E) and (F). F1 and F2 in the legends of Figures (D), (E) and (F) correspond to the actions F1 and F2 being taken, respectively. (Figure 10 in the manuscript)

climate scenario models. To test the effect of SSP scenarios uncertainty, we have checked the performance of two MDP policies trained based on climate models associated with two different SSPs, and as seen in Figure 5C, they are performing quite similarly given similar future trajectories, implying that for the used horizon of 40 years the policy is not significantly influenced by the SSP model scenario, given also the generally similar/relevant predictions of the two SSP models in the studied timeframe.

We also show the evolution of the beliefs over the considered climate scenario models in the formulated POMDP for three different possible trajectories, as shown here in Figure 5 D, E and F, with the conclusions as reported in the paper:

Lines 504-513: Figure 10D shows a SLR trajectory where, at each time step, the mean of the possible SLR states is selected as the observed SLR. As seen, for such a SLR trajectory, the two models remain largely similarly probable. On the other hand, for a possible SLR trajectory where lower percentile SLR values are observed, as in Figure 10E, the belief over SSP2-4.5 gradually increases, as the likelihood of observing lower SLR states is higher under the SSP2-4.5 scenario than the SSP5-8.5 one. Accordingly, Figure 10F shows a SLR trajectory following a possible higher percentile, having higher likelihood of being observed under SSP5-8.5, and therefore the probability of SSP5-8.5 gradually increases as observations in time are collected.

The comment from the reviewer about different model uncertainties is very interesting, and more general, and we have now provided a taxonomy in the introduction as:

Lines 57-69: The identification of the optimal flood management policy is in general non-trivial, primarily due to various levels of associated uncertainties [1, 2], which can include (i) aleatoric (irreducible) uncertainty associated with the natural variability of climate change processes and extreme events, (ii) epistemic (reducible) uncertainty associated with imperfect knowledge of the underlying climate scenario models, related also to (iii) uncertainty regarding future knowledge and updates of climate models, (iv) epistemic uncertainty associated with the evolution of future trajectories given assumed climate scenario models, and (v) uncertainty associated with available observations in time. Depending on whether and how observations and information in time are used to guide policy plans and lead to more-informed decisions, the decision-making frameworks available in the literature can be broadly categorized into static/non-adaptive (no utilization of additional information in time) and dynamic/adaptive (additional information is utilized in time).

Although we have not directly addressed the type (iii) uncertainties in this paper, the presented framework can in principle deal with them as well, in addition to all others, and we are providing a relevant explanation and comments in the Discussion section as:

Lines 553-561: While this paper does not directly address the uncertainty associated with future knowledge and climate model improvements, with this additional possible future information further increasing the benefits of the sequential decision making framework, our presented approach is capable of also accommodating this feature by incorporating potential future model information in a similar manner as the different climate model scenarios have been considered in the POMDP framework. Of course, if the actual future model information deviates substantially from the assumed characteristics, the policies can be retrained, if needed/desired, as this is something that all optimization methods, including static policy ones, can do in an open-loop format.

Finally, epistemic model uncertainties can be also included in Markovian models through parametrization and state-space augmentation, as for example in [31]. The reason we are not following this updating approach in this paper, and we do not mention it in the Discussion section, is because it will probably be impossible to practically and usefully update parameters of GCMs from just a few regional observations in time. Along these lines, the comment by the reviewer about future updates of climate models is even more relevant and meaningful, and mentioned accordingly in the Discussion section.

Reviewer Comment 1.2 — Related to the above is the fact that future improved climate models will – with high probability – lead to reduced uncertainty. This is not reflected in the modeling done by the authors. The model as it is assumes that the only additional information in the future is the observation of the sea level. Clearly, this is an underestimation of what additional information will be available in the future. So the question of how to quantify this increased knowledge seems essential here. Note that this is investigated in the scientific literature, see e.g. Völz and Hinkel 2023 for a review. I would be interested in reading from the authors how they think these factors affect the results that they obtain. (Given that more information is available in the future, the benefit of sequential decision making should increase.)

Thank you for this insightful comment. Based on this we have added some remarks in the Discussion section as:

Lines 553-561: While this paper does not directly address the uncertainty associated with future knowledge and climate model improvements, with this additional possible future information further increasing the benefits of the sequential decision making framework, our presented approach is capable of also accommodating this feature by incorporating potential future model information in a similar manner as the different climate model scenarios have been considered in the POMDP framework. Of course, if the actual future model information deviates substantially from the assumed characteristics, the policies can be retrained, if needed/desired, as this is something that all optimization methods, including static policy ones, can do in an open-loop format.

A few additional comments here. Although, additional information in the future does indeed increase the benefits of sequential decision making, and the presented framework can, in principle, support this attribute, in various different ways, from policy retraining to closed-loop representations, which are worth exploring in the future, as also mentioned in the Discussion section, in this coastal application regarding sea level rise the observations of sea levels in time are quite informative and can meaningfully guide decisions. The reason for this is that sea level rise has a largely monotonic pattern which observations in time help to discover, out of important initial uncertainty. Because of this there is also a strong time dependence on the possible trajectory paths, which again can be discovered by observations. This fact is not however true for physical phenomena that exhibit important randomness and very loose time-dependence and/or patterns in time, requiring an extremely large number of sequential observations to infer any evolution in their characteristics, such as storm surges and riverine floods. In these latter cases, additional information in the future is probably the only or at least the most important feature regarding sequential observations and decisions in time, since data on storm surges or flooding events and similar phenomena are not that informative, for the reasons just explained. However, this is not the case with this sea level rise application.

Reviewer Comment 1.3 — Importantly, whether it is considered in the analysis or not, new information will become available in the future, and future decisions (e.g., beyond 10 years from now) will utilize that information. Hence, for practical purposes the difference between a classical static optimization and the sequential decision optimization lies mainly in the decisions that are taken in the short- to mid-term (i.e. in the 10 years until the decisions are revisited). For this reason, it seems not entirely correct to compare to static strategies that only implement decisions, which might be pointed out in the discussion.

New information in the future, which is important, as also mentioned before, can be utilized both in the static and sequential decision optimization approaches. In the crudest form, the MDP/POMDP model can be resolved every time substantially new/different model information is available in the future. However, there are possibilities to include this information in a pure closed-loop control, as briefly mentioned in the Discussion section, which is of interest to the authors for future investigation. Yet, the static optimization frameworks will always be equivalent open-loop-like control implementations when reassessing them every time new information is available. We have made a comment about this in the Discussion section, as:

Lines 553-561: While this paper does not directly address the uncertainty associated with future knowledge and climate model improvements, with this additional possible future information further increasing the benefits of the sequential decision making framework, our presented approach is capable of also accommodating this feature by incorporating potential future model information in a similar manner as the different climate model scenarios have been considered in the POMDP framework. Of course, if the actual future model information deviates substantially from the assumed characteristics, the policies can be retrained, if needed/desired, as this is something that all optimization methods, including static policy ones, can do in an open-loop format.

Reviewer Comment 1.4 — Even when assuming the stationary policy to not be updated over the entire life-span, the authors find for their example that the difference between the best static and the sequential policies are only in the order of a few percent in terms of the expected life-cycle cost. I missed a discussion by the authors as to why this is the case. As an input to this discussion: What we have found in studies on adaptation of flood mitigation is that for flood risk, the uncertainty is already high even in the static case and as a result, the objective function is pretty flat around the optimum. The additional uncertainty of the climate uncertainty does not change this much. This seems to be the situation also here, as the (aleatory) uncertainty in the storm surge is much larger than the uncertainty in the sea level rise.

This comment is related to the one about future climate models, in that for phenomena that exhibit important randomness and very loose time-dependence and/or patterns in time, such as storm surges and riverine floods, observations in time are not that informative, with respect to climate effects discovery and related risk updating, and as such sequential policies may not be as valuable without considering future model information. However, this is not the case in this coastal application where the observations of sea levels in time are quite informative, due to their time-dependent pattern, and can meaningfully guide decisions. Hence, the reason for the 2%-5% improvement of sequential policies over static ones in the presented applications in this work is mainly due to the fact that the optimal policy involves actions at the beginning of the considered horizon, or else an existing seawall already in place, which is also what the static policies identify and implement. The important storm surge risk of course plays a role for these optimal plans, but we believe that with higher fidelity and more available actions, the sequential policies will show more significant differences with respect to static policies, for coastal infrastructure adaptation. We will examine this hypothesis carefully in future works. Along this perspective, the fact that with only one action difference, since only one action was deemed optimal for the static policies in all cases, between static and dynamic policies, the cost difference is already 2%-5% shows the importance of sequential decision making and timely actions based on observed conditions. We have added the following to the Discussion section as shown below.

Lines 530-547: The efficacy of the MDP-based adaptive framework is guaranteed over other methods, for the same optimization problems posed, due to better-informed policies which are chosen according to the actual evolving climate conditions, and global optimality guarantees, mathematically defined. As expected, the MDP based policies thus lead to improvements, over the considered static policies, albeit marginal to moderate ones for the analyzed settings. This is due to the fact that the adaptive policies involve actions at the beginning of the considered horizon, or else an existing seawall already in place, which are also identified by the static policies. However, the fact that the static and dynamic policies differ by a single action, since only one action was deemed optimal for the static policies in all cases, a non-trivial $2\% - 5\%$ life-cycle cost savings is already realized, that shows the potential of sequential decision making and timely actions based on observed conditions. It is expected that with even higher fidelity and consideration of more available actions, a greater reduction in life cycle cost, relative to the static policy, will be seen. The significant aleatoric uncertainty associated with storm surge risk also plays a role for these optimal policies. However, the computed policies and the sequential actions in time seem to be primarily driven by the identified patterns of the evolving SLR process.

Reviewer Comment 1.5 — Note that concepts of stochastic control have been previously suggested for optimal adaptive planning of flood risk mitigation considering climate uncertainty, e.g., Spackova and Straub 2017. Furthermore, the use of real-options analysis, which solves a related problem, also has been suggested for this application in the literature, e.g., Woodward et al 2014.

We thank the reviewer for bringing these notable works to our attention. We have substantially revised the manuscript and enhanced the literature review and the related position and contributions of this work in it, providing now a new taxonomy of related approaches, and discussion in relation to the various sources of uncertainty present, as provided at the beginning of this documents and in manuscript lines: 56-168, 175-192, 202-238.

Reviewer Comment 1.6 — The authors emphasize multiple times the "global optimality guarantee". This is true within the confines of the model they use. However, given the substantial simplifications of the model, as discussed above, this terminology can be misleading. They should restrict this description to the solution of the mathematical problem they pose, but not to the flood risk mitigation strategy.

Thank you for the comment. Indeed, the optimality guarantees are with respect to the solutions of the problems in the mathematical optimization settings posed in this paper. This is of course true for all optimization applications in all fields, since an objective function that is supported by mathematical assumptions needs to be optimized. This discussion is however quite interesting and relevant here for these applications since the optimality guarantees are also important for the choice of the model used. Higher-fidelity models can describe the entire system and environment settings in a more realistic and refined way but cannot usually provide optimality guarantees, and vice versa for simpler modeling representations. We believe to have contributed to this dilemma with a framework that can be generally applicable and more detailed than many of the currently existing approaches in the literature, and at the same time can provide globally optimal solutions, for the mathematical setting in which it is posed, something quite challenging for many applications and methods in the literature. At the same time, higher-fidelity representations are useful and have lots to offer for decision-making in many cases, despite losing the optimality guarantees. In our future work, we will thus also explore this direction based on higher-dimensional MDP/POMDP representations that can be solved by deep reinforcement learning-based approaches, hopefully to near global optimality.

We have added this clarification, about the optimality guarantees pertaining only to the optimization problem as posed in this paper, at several locations throughout the paper (in blue in the paper), and we have also provided the following remarks in the Results and the Discussion sections:

Lines 461-463: As shown, the MDP results have guarantees to converge to globally optimal solutions, pertaining specifically to the MDP-based optimization formulation, even in complex cases...

Lines 602-612: As with all MDP/POMDP based frameworks, the computational complexity increases with an increase in the number of states, possible dynamic actions, and time horizon length, that at some point cannot be practically handled with MDP/POMDP solvers that provide guarantees of global optimality. This issue can be overcome, at the expense though of strong global optimality guarantees, by pursuing deep reinforcement learning-based solution techniques, allowing consideration of a larger number of adaptive actions and longer time horizons. Given this higher flexibility and fidelity, the identified adaptive polices may be expected to result in significantly higher cost savings in comparison to static cost-benefit analysis solutions, something we are aiming to explore in the future.

Reviewer Comment 1.7 — In the abstract: "less informed management policies" is unclear. Less informed than what?

Thank you for the comment. We have now revised the abstract as below to hopefully avoid confusions. We have now added a detailed section in the Introduction explaining different approaches, and defining static policies as the ones which do not utilize real-time observations in their solution process. On the other hand, dynamic/ adaptive frameworks use real-time observations and can potentially lead to better-informed decisions in time.

Lines 14-29: Climate change-related risk mitigation is typically addressed using cost-benefit analysis that evaluates mitigation strategies against a wide range of simulated scenarios and identifies a static policy to be implemented, without considering future observations. Due to the substantial uncertainties inherent in climate projections, this identified policy will likely be sub-optimal with respect to the actual climate trajectory that evolves in time. In this work, we thus formulate climate risk management as a dynamic decision-making problem based on Markov Decision Processes (MDPs) and Partially Observable MDPs, taking real-time data into account for evaluating the evolving conditions and related model uncertainties, in order to select the best possible life-cycle actions in time, with global optimality guarantees, for the formulated optimization problem. The framework is developed for coastal adaptation applications, considering a wide variety of possible actions, including various forms of nature-based infrastructure. Related environmental impacts of carbon emissions and uptake are also incorporated and social cost of carbon implications are discussed, together with several future directions and supported features.

Reviewer 2

The goal of this paper is to demonstrate an approach using Markov decision processes and approximate dynamic programming to identify sequential decisions for coastal flood risk management. However, as currently framed, I do not think the paper is publishable in Nature Communications. The paper is currently primarily framed in terms of the method used, rather than the conceptual insights or results, but most of the technical advances of the paper have been done already (and are not captured by the literature review) or there is not much of a case for the benefit of the demonstrated approach over existing approaches in the literature. That being said, if the paper were heavily re-framed (see below), it might be a different story.

We thank the reviewer for their useful comments that helped us improve the manuscript. We did extensive modifications to the paper, including incorporation of Partially Observable Markov Decision Processes (POMDPs) approaches, relating to model uncertainty, and we elaborated more on the contributions of the paper and its place in the literature. We feel that many of the reviewer's comments are related to the inexact nature of approximate dynamic programming solution techniques, that are abundant in the literature and can indeed generally lead to very suboptimal solutions, depending on what the solution algorithm does and how it is implemented. However, in our work, a dynamic programming method is uniquely suggested that truly provides global optimal solutions for large scale MDPs/POMDPs, for the optimization problem that is set, based on optimality bounds. We explain in detail in the paper how to set up the optimal adaptation framework of coastal infrastructure under climate change in such a context, in a general manner, and easily now transferable and extensible to other regions and similar applications and cases. To ease understanding and for reading convenience, we include here some additions in the paper, regarding the contributions of our work, and we also explain in more detail in this response how the dynamic programming formulation and solution work. Subsequently, all reviewer's comments are discussed independently and in detail.

Lines 145-170, 177-194: In this work, we are presenting considerable new extensions and capabilities for adaptive planning approaches and a general and sophisticated solution framework for the optimal life-cycle adaptation of infrastructure under climate change. We are also relying on the firm mathematical foundations of closed-loop stochastic control through MDPs/POMDPs, while offering general, easily adaptable, and extensible implementations of infrastructure adaptation problems under climate change, with higher-fidelity representations, in comparison to currently available existing approaches, utilizing the IPCC models [10], accommodating and presenting a wide variety of possible actions, such as floodwalls, nature-based infrastructure, and seawalls, considering model uncertainty aspects, and incorporating social cost of carbon considerations in the analyses. Simultaneously, and in contrast to currently available relevant approaches in the literature, we are formulating the entire framework in such a way so that in can be generally supported by state-of-the-art, efficient MDP and POMDP optimization solvers with global optimality guarantees, based of course on the modeling assumptions used in the optimization process, enabling both accurate solutions and a finer representation of the problem settings based on higher-dimensional state-spaces and time/history dependent risks and costs, among others. Our suggested framework can also in principle accommodate all types of (ii)-(v) uncertainties mentioned earlier. As illustrated in Figure 1B, the integration of up-to-date annual information in relation to the current environment conditions significantly reduces the type (iv) uncertainty concerning the the evolution of the actually evolving trajectory, until the next observation becomes available. Although type (iii) uncertainties, regarding future knowledge and updates of climate models, are the only type that have not been directly addressed in this work, specific ways to incorporate these aspects in the framework are explained in the discussion section, together with still existing limitations and further possible future directions.

In the context of optimal adaptation of coastal infrastructure systems subjected to risks posed by storm surges and sea level rise in this paper, the states of our formulated MDP include the discretized SLR and storm surge levels, with the SLR states evolving probabilistically through transition dynamics computed from climate projections taken from the recent IPCC report (AR6) [10,27]. The problem states additionally include the infrastructure state, since the damage induced at any point in time is dependent not only on the SLR and storm surge levels but the implemented system protections as well. Unlike the existing, more simple adaptive policy settings of levee heightening, where the level of protection at any time is fully determined by the height of the levee, defined as the problem state, in our presented general setting, the infrastructure state and protection are determined by the entire history of the previously executed adaptation actions up to the current time step, given that these actions can completely now alter the system configuration. Although this essential characteristic for this type of problems is non-Markovian, given that the history of previously executed actions is now relevant, this issue has also been addressed in this work by appropriately and originally augmenting the state space in such a manner so as to enforce a Markovian formulation, which is also importantly generally consistent with state-of-the-art dynamic programming solvers.

We have further extended our framework by developing a POMDP model on top of the MDP one which can include hidden states, which are related to the uncertainty associated with the underlying climate scenario models.

Lines 202-238: In our suggested framework, the SLR levels are fully observed in time, while in our POMDP formulation, the underlying climate models, represented by the Shared Socio-economic Pathways (SSPs), as provided in the most recent assessment report (AR6) of the Intergovernmental Panel on Climate Change (IPCC) [10], are hidden state variables which are only probabilistically inferred based on the SLR observations in time. A POMDP is thus denoted by the tuple $\langle S, A, P, O, P_O, R, \gamma \rangle$, where O represents the set of possible observations, and P_O is the likelihood function related to the observations, and a belief $\mathbf{b}$ is formed and updated over the climate model scenarios in time, representing a probability distribution over the considered SSP models and a sufficient statistic of the history of past SLR observations regarding model uncertainty. Hence, the decision-making problem in a POMDP setting starts with an initial belief over the considered models, and at each time step, an action is taken and an observation is made, with the belief over the models updated by Bayes' rule, as:

$$b\left(s'\right) = \frac{p\left(o|s',a\right)}{p\left(o|\mathbf{b}_s,a\right)} \sum_{s\in S} p\left(s'|s,a\right) b\left(s\right) \tag{1}$$

with $p\left(o|\mathbf{b}_s,a\right)$ being the normalizing constant:

$$p\left(o|\mathbf{b}_s,a\right) = \sum_{s'\in S} p\left(o|s',a\right) \sum_{s\in S} p\left(s'|s,a\right) b\left(s\right) \tag{2}$$

Relavent to the MDP case, the optimal value function for infinite horizon POMDPs also follows a recursive Bellman equation [32], as:

$$V^*\left(\mathbf{b}_s\right) = \max_{a\in A}\left\{\sum_{s\in S} b\left(s\right) R\left(s,a\right) + \gamma \sum_{o\in O} p\left(o|\mathbf{b},a\right) V^*\left(\mathbf{b}_{s'}\right)\right\} \tag{3}$$

Due to the continuity of the belief space, POMDP problems are in general more difficult to be solved than MDPs [33]. However, the optimal value function over the belief space can be shown to be piecewise-linear and convex [34], comprising of a finite number of hyperplanes. This attribute reduces the problem to determining this finite set of hyperplanes, otherwise referred to as α-vectors. Each such vector is associated with a unique optimal action and the optimal value function for the POMDP can then be obtained using dynamic programming and α-vector backups [32, 33, 35, 36]. In this work, the problem is generally set up so that state-of-the-art point-based, asynchronous dynamic programming solvers can be utilized here, and the Focused Real-Time Dynamic Programming (FRTDP) algorithm [37] is utilized, able to provide solutions with global optimality guarantees, for both MDP and POMDP problems, with more details provided in Section 4.

Further details are now provided, in this response, about how the dynamic programming formulation and related solution in this paper work:

In order to utilize state-of-the-art offline MDP/POMDP solvers, the probabilistic modeling information, including of transition dynamics, needs to be defined in a specific manner over the entire possible state space, as explained in the paper in a general fashion, directly applicable to all other relevant applications related to adaptation under climate change. Exact dynamic programming solutions, that can also be used for MDP/POMDP models, suffer for the curse of dimensionality for high-dimensional state spaces, as the ones in the paper with 887,041 number of states. Numerous Approximate Dynamic Programming (ADP) techniques have been therefore devised over the years, to provide a solution to these problems, however these ADP techniques can also be very crude and can result in very sub-optimal solutions. In this work, we are thus using a state-of-the-art point based asynchronous dynamic programming solver named Focused Real-Time Dynamic Programming (FRTDP), having an anytime solution performance, able to handle high-dimensional problems, and most importantly providing global optimality guarantees [37, 38]. FRTDP does not depend on Monte Carlo simulations, as the reviewer mentioned, but instead uses sophisticated exploration methods to identify all critical parts of the state-space domain. In our case, for example, in the coastal city application with two floodwalls described in the paper, FRTDP converges to optimal value function by updating 340,478 states, which consists of $\sim 38\%$ of the total state space requiring 5,964,844 backup operations. The algorithm converges when the

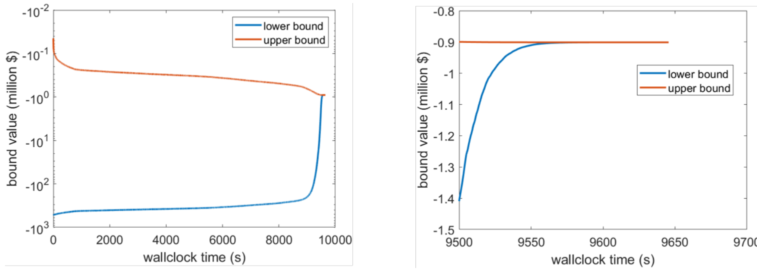

Figure 6: The upper and lower bounds of the optimal value function of the starting state (V*(s0))

Table 3 Simulation-based comparison between MDP and optimal static policies in coastal city and community settings. $E[C_T]$ is the expected total (monetary + carbon) costs in $ million per meter coastline. 95% confidence intervals are provided on the expected costs, assuming Gaussian estimators, based on 1,000 Monte Carlo simulations.

Setting	Flood Protections	$E[C_T]$ (95% C.I) MDP	$E[C_T]$ (95% C.I) Best static policy
Coastal city	Two floodwalls (no pre-existing seawall)	0.5632 $\pm$0.0105	0.5925 $\pm$0.0109($F1$)
	Green resistance zone and a floodwall (with pre-existing seawall)	0.0937 $\pm$0.0040	0.0992 $\pm$0.0038(GR)
Coastal community	Salt marsh and a floodwall	0.4473 $\pm$0.0042	0.4664 $\pm$0.0045(SM)
	Oyster reef and a floodwall	0.4215 $\pm$0.0047	0.4316 $\pm$0.0049(OR)

Note: F1, GR, SM, and OR represent the best static policy.

Figure 7: Table 3 in Supplementary Materials

upper and lower value function bounds reach an ϵ-tolerance, as shown in Figure 6, for the coastal city application with two floodwalls, with $\epsilon = 0.001$ units here. The figure shows the bounds approaching each other and finally converging after approximately 10,000 secs of solver run time. The sudden jumps of the lower bound is something common for such DP solvers, as also explained in [36, 39].

The converged globally optimal policy can then be used, providing a mapping from MDP/POMDP states to actions. In the optimal policy evaluation phase, i.e., to compute the expected total costs through simulations, 1,000 possible Monte Carlo trajectories have been utilized in our work. These simulations are deemed sufficient since they lead to low standard errors, as shown in the Table below (Table 3 in the supplementary materials file):

For additional explanation of the solution process, a schematic of the exploration process during the solution phase of FRTDP is shown in Figure 8 (Figure 5 in the supplementary materials file). Additional details can be seen in [26, 32].

Reviewer Comment 2.1 — The application of standard control methods, including dynamic programming, to coastal decision making, is not new, though much of this literature is not referenced in the paper. For example, Lickley et al (2014) used DP, and Shuvo et al (2022) use reinforcement learning to simulate socioeconomic development and adaptation to sea-level rise. In fact, even the classic Eijgenrraam et al (2014) paper used DP to identify sequential decisions. Therefore, the main

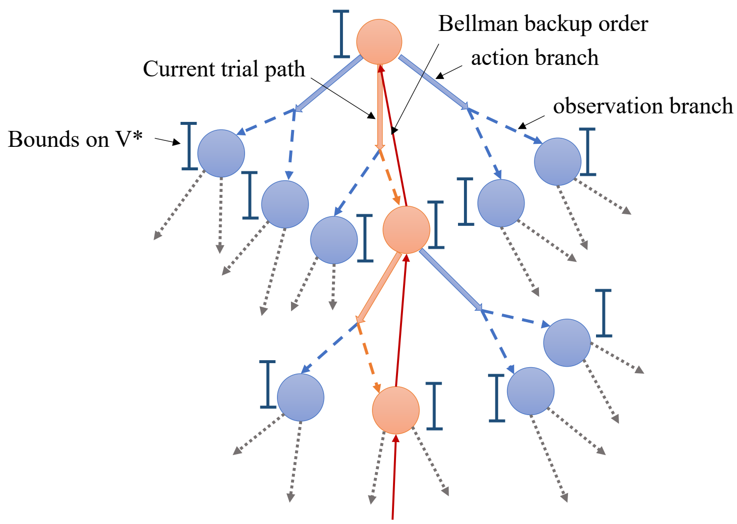

Figure 8: Search and Updating process in FRTDP (also provided as Figure 6 in the supplementary materials)

advantage of the method in the paper is the use of an approximate DP algorithm with learning about probabilities about likely states. But (in addition to this approximate approach not quite having as strong of a claim to global optimality as the paper presents), there is also the possibility of other closed-loop control approaches, such as direct policy search (see Quinn et al (2016) more generally and Garner and Keller (2018) for an application to levee heightening). DPS addresses the curse of dimensionality that the ADP approach in this manuscript suffers from (hence the stylization of the problem), and while DPS, as a heuristic optimization method, does not include the same optimality guarantees as DP, it is also not entirely clear why that is more important given the dimensionality of the relevant uncertainties, especially across the wide range of policy options, as well as the limited cost improvement in the identified solutions throughout the paper. There is also the dynamic adaptive policy pathways approach, which may not yield optimal results but may result in identifications of potential lock-in and be more flexible overall. If this version of the paper is to be publishable, quite a lot more work has to go into the literature review and better contextualization for why this is a useful method for the problem at hand (versus the way the paper is framed, which reads as though the method is the main focal point, as this takes up the bulk of the introduction and the discussion).

We thank the reviewer for this comment that assisted us in better explaining the contributions of this work and its overall position in the literature. We have significantly revised the manuscript accordingly, as below.

Lines 57-170, 177-194: The identification of the optimal flood management policy is in general non-trivial, primarily due to various levels of associated uncertainties [1, 2], which can include (i) aleatoric (irreducible) uncertainty associated with the natural variability of climate change processes and extreme events, (ii) epistemic (reducible) uncertainty associated with imperfect knowledge of the underlying climate scenario models, related also to (iii) uncertainty regarding future knowledge and updates of climate models, (iv) epistemic uncertainty associated with the evolution of future trajectories given assumed climate scenario models, and (v) uncertainty associated with available observations in time. Depending on whether and how observations and information in time are used to guide policy plans and lead to more-informed decisions, the decision-making frameworks available in the literature can be broadly categorized into static/non-adaptive (no utilization of additional information in time) and dynamic/adaptive (additional information is utilized in time).

A conventional approach for climate change planning has been primarily driven by cost-benefit analysis (CBA). Typical CBA solves a static optimization problem where the optimal time of investments and the optimal infrastructure related actions are determined by maximizing the benefits of the sought policy benefits while minimizing the related investment costs over the planning horizon. Most of the work in the existing literature implement such static optimization frameworks, often within settings focused solely on levee heightening. The expected flood damages within such frameworks are determined by assuming a fixed base-case scenario of the evolution of the probability of failure [3] or by considering an ensemble range of future SLR and storm surge events to be representative of the range of possible future scenarios [2, 4–9]. Different solution techniques have been utilized to solve these static optimization problems, including genetic algorithms [4], dynamic programming, and non-linear programming [3].

The result of the static optimization approach, as defined here, is a single policy which is deemed optimal over the average, or any chosen percentile, of potential future climate trajectories; and this policy is to be implemented without further guidance by the actually evolving climate conditions observable in time. However, the actual climate trajectory may deviate significantly from the ensemble mean (or the chosen target percentile) owing to the substantial amount of uncertainty present in climate model projections [10], as illustrated in Figure 1A, which inevitably means that the identified policy is highly likely to be sub-optimal with respect to the actual climate trajectory that evolves in time. Thus, static optimization approaches solely rely on simulation results and do not utilize actual observations in time to update the knowledge regarding the associated risks and any of the (ii)-(v) related aforementioned uncertainties.

There are a few works in the literature offering adaptive planning approaches, as defined previously, which determine policy actions depending on the evolving, observed climate conditions in time. This existing literature on adaptive planning of infrastructure systems against climate change typically apply methodologies including direct policy search (DPS) [11, 12], real options analysis (ROA) [13, 14], and stochastic control methods [15–19]. With DPS, the proposed actions are condition/state-dependent, as observed in time, however, the expressions relating the observable state variables to the triggering actions are arbitrarily defined and heuristically trained. As such, DPS lacks optimality guarantees and explores only a relatively limited subset of the policy space. ROA also offers an adaptive planning framework, typically based on binomial decision trees with two possible actions/options at each node [14]. The action at each node is determined by computing action values backwards in time. This analysis represents a form of backward induction which can quickly become computationally intractable as the number of branches increases exponentially with the number of time steps, actions, and observations. Alternative approaches have also been suggested through adaptive pathway solutions [20–22] by identification, through explorative modeling, of adaptation tipping points in time. However, these methods rely on only a few forward simulations/trajectories of possible futures in time, to identify the tipping points, sharing many of the limitations of the cost-benefit and DPS methods, and making the adaptive pathway solutions likely to be sub-optimal with respect to the actual evolving conditions in time. Overall, these adaptive methods have been typically applied in levee construction/heightening applications, where a single levee exists or is constructed at one point in time and is then subsequently heightened in future time-steps.

Examples of adaptive planning optimization against climate change effects using stochastic control frameworks can be found in the relatively recent literature, where the sequential decision-making process is formulated as a Markov Decision Process (MDP) or Partially Observable MDP (POMDP), with state variables typically representing the related conditions, and decisions being the possible infrastructure actions. The MDP/POMDP solution determines optimal actions to be taken at each decision time-step, conditioned on the observed, or inferred for POMDPs, system states, and thus uniquely selected for the actual realized trajectory path, successfully optimizing the total overall expected costs/rewards over the decision horizon. MDP/POMDP approaches can offer superior and more general solutions, using the same models, with respect to other static and/or adaptive planning methods, as described above, since they are based on dynamic programming and closed-loop stochastic control foundations [23], that offer robust guarantees of global optimality of the action paths. A few notable works based on the MDP framework, also for levee heightening problems, can be seen in [15–17, 24]. Some of these works importantly also consider climate model uncertainty utilizing POMDPs and simple hydrological scenarios [16, 24]. However, the main limitation of standard MDP/POMDP solution approaches is the curse of dimensionality [25, 26], which often dictates an overly simplified representation of the system, and/or a crude approximate solution technique that does not offer optimality guarantees. The existing related methodologies in the literature largely follow these patterns and are quite specific to the characteristics of their particular application scenarios, with relatively few possible states at each point in time, while also employing approximate solution techniques, all in order to make a solution to the problem feasible. Utilization of the solution framework to other applications, and/or extensions is thus not straightforward, particularly also when a variety of possible actions needs to be considered.

In this work, we are presenting considerable new extensions and capabilities for adaptive planning approaches and a general and sophisticated solution framework for the optimal life-cycle adaptation of infrastructure under climate change. We are also relying on the firm mathematical foundations of closed-loop stochastic control through MDPs/POMDPs, while offering general, easily adaptable, and extensible implementations of infrastructure adaptation problems under climate change, with higher-fidelity representations, in comparison to currently available existing approaches, utilizing the IPCC models [10], accommodating and presenting a wide variety of possible actions, such as floodwalls, nature-based infrastructure, and seawalls, considering model uncertainty aspects, and incorporating social cost of carbon considerations in the analyses. Simultaneously, and in contrast to currently available relevant approaches in the literature, we are formulating the entire framework in such a way so that in can be generally supported by state-of-the-art, efficient MDP and POMDP optimization solvers with global optimality guarantees, based of course on the modeling assumptions used in the optimization process, enabling both accurate solutions and a finer representation of the problem settings based on higher-dimensional state-spaces and time/history dependent risks and costs, among others. Our suggested framework can also in principle accommodate all types of (*ii*)-(*v*) uncertainties mentioned earlier. As illustrated in Figure 1B, the integration of up-to-date annual information in relation to the current environment conditions significantly reduces the type (*iv*) uncertainty concerning the the evolution of the actually evolving trajectory, until the next observation becomes available. Although type (*iii*) uncertainties, regarding future knowledge and updates of climate models, are the only type that have not been directly addressed in this work, specific ways to incorporate these aspects in the framework are explained in the discussion section, together with still existing limitations and further possible future directions.

In the context of optimal adaptation of coastal infrastructure systems subjected to risks posed by storm surges and sea level rise in this paper, the states of our formulated MDP include the discretized SLR and storm surge levels, with the SLR states evolving probabilistically through transition dynamics computed from climate projections taken from the recent IPCC report (AR6) [10, 27]. The problem states additionally include the infrastructure state, since the damage induced at any point in time is dependent not only on the SLR and storm surge levels but the implemented system protections as well. Unlike the existing, more simple adaptive policy settings of levee heightening, where the level of protection at any time is fully determined by the height of the levee, defined as the problem state, in our presented general setting, the infrastructure state and protection are determined by the entire history of the previously executed adaptation actions up to the current time step, given that these actions can completely now alter the system configuration. Although this essential characteristic for this type of problems is non-Markovian, given that the history of previously executed actions is now relevant, this issue has also been addressed in this work by appropriately and originally augmenting the state space in such a manner so as to enforce a Markovian formulation, which is also importantly generally consistent with state-of-the-art dynamic programming solvers.

Reviewer Comment 2.2 — However, what is somewhat new and interesting is the role of the social cost of carbon in changing optimal solutions (I have yet to see this for this kind of decision problem), and I would encourage the authors to focus on these insights in any revision of the manuscript, particularly if they looked at SCC uncertainty.

We appreciate the reviewer's comment recognizing the originality in considering the SCC in the climate adaptation decision-making problem. In this work, we assume that the current value of social cost of carbon is known, as it is typically set by relevant policy makers. However, the SCC value may change in the future, due to changing administrations, for example. This is related to additional information that may come in the future, and we have now made a note about this in the Discussion section, together with specifics of how this can be modeled in our framework.

Lines 570-577: In this work, we assume that the current value of the SCC is deterministically known, as it is typically set by relevant policy makers. However, the SCC value may change in the future, due to changing administrations, for example. This is also relevant to the previous discussion about additional information that may come in the future, and can be addressed within our framework, as explained earlier, either in closed-loop representations, by updating our knowledge of some representative potential SCC values in time, or by resolving the policy with the updated value of SCC, if substantial changes are suggested.

In this work, a sensitivity study has been performed to show the effect of different SCC values on the obtained optimal policies. As seen in Figure 9 (A), higher SCC values lead to higher frequency of construction of the higher floodwall. Increasing this cost also encourages earlier interventions, which is evident in Figure 9 (B), showing an increase in the probability of construction of the floodwall at the higher elevation earlier in the planning horizon, when this higher value of the social cost of carbon is used.

Reviewer Comment 2.3 — Another option would be to take the application more seriously and build out the range of policy options, rather than stylizing them to demonstrate the method; while

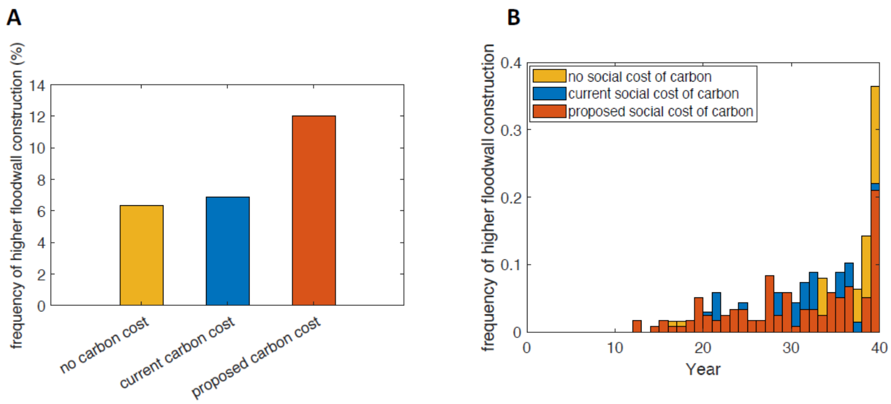

Figure 9: Effect of the social cost of carbon on the (A) frequency of construction, and (B) the time of construction of the higher floodwall. In (A), no carbon cost corresponds to the case without considering the social cost of carbon, current carbon cost corresponds to the case with the current value of $51 per ton, and proposed carbon cost corresponds to the case where the value is increased to the recently proposed value of $ 190 per ton. (Figure 10 in the manuscript)

this would run into the curse of dimensionality, this is a core issue with any DP method (and one that ADP is intended to at least partially address).

We thank the reviewer for the comment that assisted us in better presenting the contributions of our work, as below:

Lines 145-170, 177-194: In this work, we are presenting considerable new extensions and capabilities for adaptive planning approaches and a general and sophisticated solution framework for the optimal life-cycle adaptation of infrastructure under climate change. We are also relying on the firm mathematical foundations of closed-loop stochastic control through MDPs/POMDPs, while offering general, easily adaptable, and extensible implementations of infrastructure adaptation problems under climate change, with higher-fidelity representations, in comparison to currently available existing approaches, utilizing the IPCC models [10], accommodating and presenting a wide variety of possible actions, such as floodwalls, nature-based infrastructure, and seawalls, considering model uncertainty aspects, and incorporating social cost of carbon considerations in the analyses. Simultaneously, and in contrast to currently available relevant approaches in the literature, we are formulating the entire framework in such a way so that in can be generally supported by state-of-the-art, efficient MDP and POMDP optimization solvers with global optimality guarantees, based of course on the modeling assumptions used in the optimization process, enabling both accurate solutions and a finer representation of the problem settings based on higher-dimensional state-spaces and time/history dependent risks and costs, among others. Our suggested framework can also in principle accommodate all types of (ii)-(v) uncertainties mentioned earlier. As illustrated in Figure 1B, the integration of up-to-date annual information in relation to the current environment conditions significantly reduces the type (iv) uncertainty concerning the the evolution of the actually evolving trajectory, until the next observation becomes available. Although type (iii) uncertainties, regarding future knowledge and updates of climate models, are the only type that have not been directly addressed in this work, specific ways to incorporate these aspects in the framework are explained in the discussion section, together with still existing limitations and further possible future directions.

In the context of optimal adaptation of coastal infrastructure systems subjected to risks posed by storm surges and sea level rise in this paper, the states of our formulated MDP include the discretized SLR and storm surge levels, with the SLR states evolving probabilistically through transition dynamics computed from climate projections taken from the recent IPCC report (AR6) [10,27]. The problem states additionally include the infrastructure state, since the damage induced at any point in time is dependent not only on the SLR and storm surge levels but the implemented system protections as well. Unlike the existing, more simple adaptive policy settings of levee heightening, where the level of protection at any time is fully determined by the height of the levee, defined as the problem state, in our presented general setting, the infrastructure state and protection are determined by the entire history of the previously executed adaptation actions up to the current time step, given that these actions can completely now alter the system configuration. Although this essential characteristic for this type of problems is non-Markovian, given that the history of previously executed actions is now relevant, this issue has also been addressed in this work by appropriately and originally augmenting the state space in such a manner so as to enforce a Markovian formulation, which is also importantly generally consistent with state-of-the-art dynamic programming solvers.

In addition, this comment by the reviewer is also related to the very interesting discussion about optimality guarantees and model fidelity. Model fidelity can, in most cases, always be refined, and higher-fidelity models can describe the entire system and environment settings in a more realistic and refined way but cannot usually provide optimality guarantees, and vice versa for simpler modeling representations. We believe to have contributed to this dilemma with an original framework that can be generally applicable and more detailed than many of the currently existing approaches in the literature, and at the same time can provide globally optimal solutions, for the mathematical setting in which it is posed, something quite challenging for many applications and methods in the literature. At the same time, even higher-fidelity representations are useful and have lots to offer for decision-making in many cases, despite losing the optimality guarantees. In our future work, we will thus also explore this direction based on higher-dimensional, general MDP/POMDP representations that can be solved by deep reinforcement learning-based approaches, hopefully to near global optimality.

Reviewer Comment 2.4 — Line 50: The word "optimal" is used a lot without specifying what the objective is.

Thank you for the comment. We have revised the Introduction section, as below:

Lines 49-56: Owing to the significant uncertainties associated with the evolving climate demands and the limited resources for implementing various flood protection strategies at a given point in time, it is desirable to determine the optimal flood management policy, that can involve a synergistic implementation of different flood protections in time, that minimizes the present total life-cycle utility value, usually expressed in monetary units, including expected flood damage costs and risks, implementation and associated maintenance costs of flood protection measures, and any other desired quantities of interest.

We have also defined the objective function as:

Lines 204-214: In our optimal flood management setting, related costs at each time-step consist of flooding risks, maintenance of pre-existing flood protection assets, and implementation costs of new flood protection measures associated with the action being taken. Thus, the value function in Eq. **??** is given as:

$$V^*(s) = \max_{a \in A} \left[\left\{ R_i(s,a) + R_m(s,a) + \gamma \sum_{s' \in S} p(s'|s,a) R_f(s') \right\} + \gamma \sum_{s' \in S} p(s'|s,a) V^*(s') \right] \quad (3)$$

where R_i, R_m and R_f are costs associated with implementation of currently implemented flood measures, maintenance of existing measures and flood-induced damages, respectively. Carbon emissions associated with each one of these components and related implications are also considered in our framework. More details can be found in Section 4.

We have provided more details about the rewards in the Methods section 4.1.4.

Reviewer Comment 2.5 — Line 54: There is a discussion of a variety of stakeholder objectives, but as far as I can tell none of these are taken into account beyond loosely interpreting environmental objectives as the SCC (though the SCC itself can be a fraught metric). What is the justification for focusing entirely on a single economic objective, even one which has the SCC included as a proxy for environmental considerations? In fact, a posteriori decision analyses, which are often multi-objective, can facilitate robustness analyses and stakeholder engagement in a way which a priori analyses cannot.

In this work we have worked on an original MDP/POMDP optimization framework for optimal adaptation of infrastructure under climate change based on a single objective function. Please note that such a function can include various objectives that are appropriately combined into a scalar objective function output. As also mentioned before, any desired quantity of interest can be added to the problem as an objective. Most currently implemented decision schemes today work in a similar manner, as also explained in the next reply. However, multi-objective schemes also have lots of advantages and favorable properties. In their most straightforward implementation, someone can use our approach as presented in this paper in an outer-loop format, i.e., define quantities of interest, assign weights, devise a new objective function based on the weights, and solve the problem to optimality. Inner-loop approaches, where weights are also sought for Pareto front solutions, are, in general, an open problem for MDPs/POMDPs and we are actively working in this direction for our future work. We have added a relevant statement in the Discussion section as:

Lines 612-615: Moreover, the various considered objectives were aggregated in this work into a single scalarized objective function. In future work, we aim to extend our framework to consider relevant multi-objective formulations to the problem by balancing various tradeoffs among the constituent objectives.

As a final note, approaches where weights are also sought for Pareto front solutions are mostly formulated around static optimization problems, and through genetic algorithms solvers, that also have their own disadvantages and significant limitations.

Reviewer Comment 2.6 — Line 56: A claim that criteria-based protection is currently the standard needs citation, and I'm not sure that it's entirely true, even if the USACE approach largely boils down to this assuming a cost-benefit test is passed. For example, the Netherlands has moved away from this type of design standard to an approach which regulates the expected loss of life due to flooding (see https://www.enwinfo.nl/publish/pages/183541/grondslagenen-lowresspread3-v3.*pdf*).

Thank you for the reference and this comment that helped us to better express this point in the paper, as:

Lines 70-74: A conventional approach for climate change planning has been primarily driven by cost-benefit analysis (CBA). Typical CBA solves a static optimization problem where the optimal time of investments and the optimal infrastructure related actions are determined by maximizing the benefits of the sought policy benefits while minimizing the related investment costs over the planning horizon.

Cost-benefit analysis is what is mostly used for decision-making and decision-support today, particularly when federal agencies and authorities are concerned. Just for reference, this is just one indicative publication by the European Commission (Implementation of Cost Benefit Analaylis of Coastal Risk Management Prevention Measures against Natural Hazards (2016)) [40]. In the provided reference by the reviewer, a cost-benefit analysis is also followed, and as explained before, the loss of human life is one of the quantities of interest included in the objective function, expressing it, as customary, in monetary terms, in order to properly setup the problem. Please note that our framework can support all these approaches and can also offer a superset of relevant solutions.

Reviewer Comment 2.7 — Does risk-based estimation need to focus on economic losses, or just the sampled papers?

This question is related to the two previous ones and our replies there are also relevant to the discussion for this comment. Risk is generally defined as the probability of an event times the consequences of the event in some relevant unit. Any utility function can be used for this, and all utility functions are supported by our framework without limitations. In this work, we are using monetary units to set up the optimization problem, without any loss of generality.

Reviewer Comment 2.8 — Optimization against an ensemble of stochastic realizations doesn't need to optimize the expected outcome, a robust optimization approach could be used which would optimize against a quantile (or any other summary statistic of interest).

Thank you for the comment. We have now revised the text as:

Lines 83-94: The result of the static optimization approach, as defined here, is a single policy which is deemed optimal over the average, or any chosen percentile, of potential future climate trajectories; and this policy is to be implemented without further guidance by the actually evolving climate conditions observable in time. However, the actual climate trajectory may deviate significantly from the ensemble mean (or the chosen target percentile) owing to the substantial amount of uncertainty present in climate model projections [10], as illustrated in Figure 1A, which inevitably means that the identified policy is highly likely to be sub-optimal with respect to the actual climate trajectory that evolves in time. Thus, static optimization approaches solely rely on simulation results and do not utilize actual observations in time to update the knowledge regarding the associated risks and any of the (ii)-(v) related aforementioned uncertainties.

Just for completeness and better understanding of our approach, in MDP and POMDP formulations the expected value is usually considered in the objective function, for various practical and theoretical reasons, and mathematical formulations, proofs, convergence, etc., are based on this metric. If other quantiles need to be considered, this is possible, but it is easier to do this through the utility functions rather than the solution schemes which will require extensive new proofs and derivations. Please also note that robust optimization is usually related to the sensitivity of the solution, mostly through a variance metric, rather than the quantile used. Dynamic programming formulations for such robust approaches are also possible. Finally, in our MDP/POMDP solution scheme we do not use "an ensemble of stochastic realizations" to identify the optimal policy.

Reviewer Comment 2.9 — I don't follow this discussion of adaptive approaches. The whole point of adaptivity is to make a flexible initial investment and then pivot based on learning about the more probable states of the world. The identification of adaptation tipping points serves effectively the same role as the learning in the manuscript's approach.

Thank you for the comment. We have now extensively revised the manuscript text describing the contributions of this work and its position in the literature, as explained before. The particular discussion on adaptation tipping points is now given as:

Lines 109-118: Alternative approaches have also been suggested through adaptive pathway solutions [20–22] by identification, through explorative modeling, of adaptation tipping points in time. However, these methods rely on only a few forward simulations/trajectories of possible futures in time, to identify the tipping points, sharing many of the limitations of the cost-benefit and DPS methods, and making the adaptive pathway solutions likely to be sub-optimal with respect to the actual evolving conditions in time. Overall, these adaptive methods have been typically applied in levee construction/heightening applications, where a single levee exists or is constructed at one point in time and is then subsequently heightened in future time-steps.

Reviewer Comment 2.10 — Line 124: The text makes it seem as though the presented approach is independent of the problems associated with sampling stochastic realizations of the relevant uncertainty space. But it's still an approximate dynamic programming method, and therefore requires samples, which may be more or less representative of the overall distribution, especially if the probability model is poorly specified.

Indeed, the presented approach does not depend on sampling and Monte Carlo stochastic realizations. In this work, we are using a state-of-the-art point based asynchronous dynamic programming solver named Focused Real-Time Dynamic Programming (FRTDP) [37,38], having an anytime solution performance, able to handle high-dimensional problems, and most importantly providing global optimality guarantees. FRTDP uses sophisticated exploration methods to identify all critical parts of the state-space domain. In our case, for example, in the coastal city application with two floodwalls described in the paper, FRTDP can guarantee convergence to optimal policy by updating 340,478 states, which consists of $\sim 38\%$ of the total state space requiring 5,964,844 backup operations. The algorithm converges when the upper and lower value function bounds reach an ϵ-tolerance, as shown in Figure 10 below, for the coastal city application with two floodwalls, with $\epsilon = 0.001$ units here. The figure shows the bounds approaching each other and finally converging

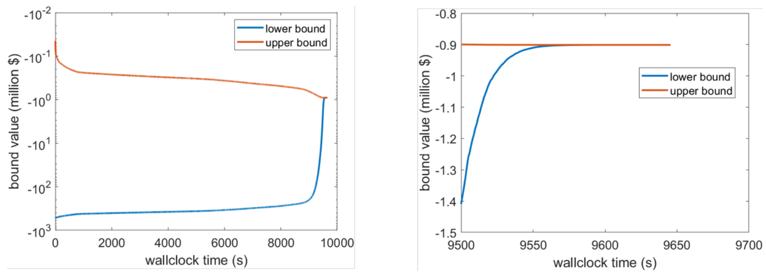

Figure 10: The upper and lower bounds of the optimal value function of the starting state (V*(s0))

Table 3 Simulation-based comparison between MDP and optimal static policies in coastal city and community settings. $\boldsymbol{E}[C_T]$ is the expected total (monetary + carbon) costs in \$ million per meter coastline. 95% confidence intervals are provided on the expected costs, assuming Gaussian estimators, based on 1,000 Monte Carlo simulations.

Setting	Flood Protections	$\boldsymbol{E}[C_T]$ (95% C.I) MDP	$\boldsymbol{E}[C_T]$ (95% C.I) Best static policy
Coastal city	Two floodwalls (no pre-existing seawall)	0.5632 ±0.0105	0.5925 ±0.0109($F1$)
	Green resistance zone and a floodwall (with pre-existing seawall)	0.0937 ±0.0040	0.0992 ±0.0038(GR)
Coastal community	Salt marsh and a floodwall	0.4473 ±0.0042	0.4664 ±0.0045(SM)
	Oyster reef and a floodwall	0.4215 ±0.0047	0.4316 ±0.0049(OR)

Note: F1, GR, SM, and OR represent the best static policy.

Figure 11: Table 3 in Supplementary Materials

after approximately 10,000 secs of solver run time. The sudden jumps of the lower bound is something common for such DP solvers, as also explained in [36, 39]). The converged globally optimal policy can then be used, providing a mapping from MDP/POMDP states to actions.

Reviewer Comment 2.11 — This is an issue which is under treated in this manuscript: in lines 228-229, there is a mention that 1,000 Monte Carlo samples were used without this being justified further (for example, what was the Monte Carlo standard error?), and very limited SSP scenarios were used.

The converged globally optimal policy can be used to provide a mapping from MDP/POMDP states to actions. In the optimal policy evaluation phase, i.e., to compute the expected total costs through simulations, 1,000 possible Monte Carlo trajectories have been utilized in our work. This has nothing to do with learning the optimal policy. These simulations are deemed sufficient since they lead to low standard errors, as shown in the Table below (Table 3 in the supplementary materials file):

Without loss of generality, two SSP scenarios have been used in this study, SSP2-4.5 and SSP5-8.5. The framework is, however, general and can of course be used with any SSP scenario. SSP2-4.5 has been used primarily for most reported results in the paper, and SSP5-8.5 has been chosen primarily

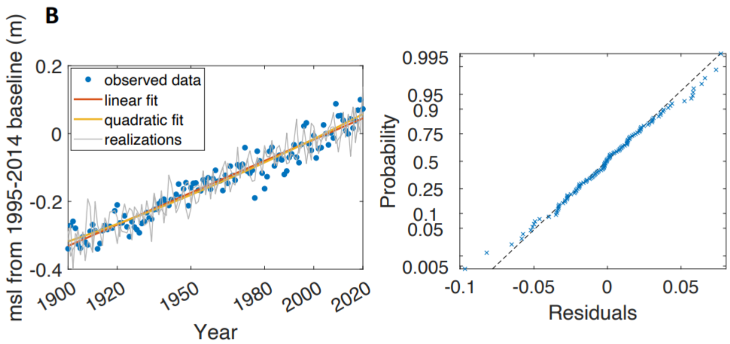

Figure 12: Observed variability in historical observations, the polynomials fitted to the observations, and sample hindcast realizations with observed variability distributed according to Gaussian distribution.(shown as Figure 3B in the manuscript)

to show how the policies differ with different SSP scenarios, as reported in Fig. 10 and in relation to all the now added POMDP results and relevant discussion about model uncertainty.

Reviewer Comment 2.12 — Also, despite the discussion of non-stationarity and the emphasis on fitting a model for higher-frequency mean SLR variations (without much justification), a standard variance is used for the future projections without justification. Why is this probability model appropriate?

We thank the reviewer for this comment that has helped us explain our modeling better. The nonstationary IPCC projections are used for our modeling and on top of that we are also modeling the natural variability of Sea Level, as recorded annually, based on historical data. This variability was found to be very well represented by a stationary zero mean Gaussian distribution with a standard deviation of 3.18 cm, as shown in Fig. 12. The Gaussianity assumption is also verified in Fig. 12. In Fig. 3C in the manuscript some possible future trajectories based on the IPCC model and the historical data randomness can be observed.

The paper has now been revised, as:

Lines 278-300: The SLR model is derived based on regional SLR projections using CMIP6 models from the AR6 of the IPCC [10] for the Battery tidal gauge [28] representing the general NYC area, without loss of generality. The SLR projections extending to 2150, relative to the 1995-2014 reference period, are obtained from [29]. These projections are based on various climate change and emission scenarios, represented by distinct Shared Socio-economic Pathways (SSPs), and SLR projections under three such SSPs are shown in Figure 3A. For each SSP model, an ensemble of 21 Global Climate Models (GCM) is considered in this work, in an effort to consider epistemic model uncertainties within a SSP scenario. A future SLR trajectory follows here a certain projection percentile within the likely range, which captures the long-term non-stationary SLR trend resulting from multiple factors, such as thermal expansion, ice sheet contributions, and glacial melt. However, these projections do not include the short-term variability inherent in the historical observations, which are contributions of various natural fluctuations. In order to also consider these variations in the future sea level rise projections, a regression model is fit to the observed historical data provided by the National Oceanic and Atmospheric Administration (NOAA) [30], as shown in Figure 3B, and the random noise present in the observations is utilized, assuming to follow a Gaussian distribution with a constant standard deviation of 3.18 cm, as also confirmed in the normal probability plot in Figure 3B. Indicative resulting sea level rise process simulations derived from the IPCC climate model projections, along with the historically observed variability, are illustrated in Figure 3C. The time-variant and uncertain sea level rise evolution is modeled by a Markov process with transition dynamics estimated based on these simulated trajectories.

The SLR is discretized in our framework with a step size of 2cm, and the higher fidelity of SLR is utilized efficiently by the optimal policy since this is what drives the decision-making process. The 2cm are deemed acceptable since past SLR observations are noisy with variations in the range of 3 centimeters. Of course, a cruder discretization would work as well, resulting also in cruder optimal policies in time. The discretization of the state-space, as explained in the manuscript, is needed for our framework, for the MDP/POMDP solution approaches to work to global optimality.

Reviewer Comment 2.13 — Line 199: What kind of weather patterns should be captured in the mean SLR projections? This gets back to the justification for incorporating the higher-frequency variability at the expense of uncertainty across SSPs.

Thank you for the comment. For the mean SLR projections, through the IPCC models, the term "weather patterns" is not relevant and we have now corrected this typo in the revised paper, which is now written as:

Lines 286-289: A future SLR trajectory follows here a certain projection percentile within the likely range, which captures the long-term non-stationary SLR trend resulting from multiple factors, such as thermal expansion, ice sheet contributions, and glacial melt.

The SLR is discretized in our framework with a step size of 2cm, and the higher fidelity of SLR is utilized efficiently by the optimal policy since this is what drives the decision-making process. In addition, in the revised version we have also directly modeled the uncertainty across SSP scenarios in a POMDP context, since the underlying model, or mixture of models, are hidden states that can never be fully observed.

Reviewer Comment 2.14 — Line 214-216: For annual extremes (I assume that was the temporal resolution, though this was not made particularly clear), you shouldn't need to take away the tides, and in fact might end up under-estimating the extremes by only accounting for the storm surge and not the storm tides. You would need to for shorter horizons where the data are correlated due to tidal patterns, such as monthly.

We apologize for any confusion, and we thank the reviewer for bringing this point to our attention. We have considered the tidal effects in our analysis, since a Generalized Extreme Value distribution has been fitted to the annual observed extremes with respect to the MHHW datum, after excluding the SLR trend in the observations.

Also, we define the temporal resolution of the MDP, in the manuscript as shown below, where states which correspond to the peak water levels are observed every year, and they suggest actions to be taken every year. It is also mentioned in Section 2.1 in the paper, where we define the simulated storm surge heights as annual maximum surge heights in **Line 293**.

Table 3 Simulation-based comparison between MDP and optimal static policies in coastal city and community settings. $E[C_T]$ is the expected total (monetary + carbon) costs in \$ million per meter coastline. 95% confidence intervals are provided on the expected costs, assuming Gaussian estimators, based on 1,000 Monte Carlo simulations.

Setting	Flood Protections	$E[C_T]$ (95% C.I) MDP	$E[C_T]$ (95% C.I) Best static policy
Coastal city	Two floodwalls (no pre-existing seawall)	0.5632 ± 0.0105	0.5925 $\pm 0.0109 (F1)$
	Green resistance zone and a floodwall (with pre-existing seawall)	0.0937 ± 0.0040	0.0992 $\pm 0.0038 (GR)$
Coastal community	Salt marsh and a floodwall	0.4473 ± 0.0042	0.4664 $\pm 0.0045 (SM)$
	Oyster reef and a floodwall	0.4215 ± 0.0047	0.4316 $\pm 0.0049 (OR)$

Note: F1, GR, SM, and OR represent the best static policy.

Figure 13: Table 3 in Supplementary Materials

Lines 326-328: The flood management problem is modeled here as an MDP where the states (S) correspond to the total peak water levels observed every year with certainty. These states guide the selection of the actions (A) over time.

Reviewer Comment 2.15 — Line 278-9: How do you know this is enough Monte Carlo samples? Was this number chosen because of a low Monte Carlo standard error, or because it was computationally tractable?

The converged globally optimal policy can be used to provide a mapping from MDP/POMDP states to actions. In the optimal policy evaluation phase, i.e., to compute the expected total costs through simulations, 1,000 possible Monte Carlo trajectories have been utilized in our work. This has nothing to do with learning the optimal policy. These simulations are deemed sufficient since they lead to low standard errors, as shown in the Table below (Table 3 in the supplementary materials file):

Reviewer Comment 2.16 — Line 481: I don't understand the discretization of SLR values. Is this per time step? What was the point of the Gaussian fit if you were just going to discretize over a grid?

The discretization of the state-space, as explained in the manuscript, is needed for our framework, for the MDP/POMDP solution approaches to work to global optimality, and is also customary for many other dynamic programming approaches. The entire original and general formulation in the paper is described based on this feature. The discretization is mentioned in detail in Section 4.1.1 and the related transition probabilities are explained in Section 4.1.3. The annual time steps are also part of the discretized state-space. In the Supplementary Materials document, we are showing some of our validation results, making sure that the discretized model accurately follows the underlying continuous modeling.

The underlying continuous model is based on the nonstationary IPCC projections and on top of that we are also modeling the natural variability of Sea Level, as recorded annually, based on historical data. This variability was found to be very well represented by a stationary zero mean Gaussian distribution with a standard deviation of 3.18 cm, as shown in Fig. 3B. The Gaussianity assumption is also verified in Fig. 3B in the paper. In Fig. 3C in the paper some possible future trajectories based on the IPCC model and the historical data randomness can be observed.

Reviewer Comment 2.17 — Code Availability: The code should be made available to reviewers in preparation for public release upon potential publication, to be consistent with Nature policy.

The MDP/POMDP solver used in this work is publicly available at [38] and described in detail here [39]. All data used in this work are also publicly available and appropriately cited in the paper. We are now also publicly sharing with reviewers all our relevant MDP and POMDP model codes in a supplementary software file.

Reviewer 3

Reviewer Comment 3.1 — The study introduces key findings in coastal infrastructure adaptation to climate change, highlighting the effectiveness of Markov Decision Process (MDP)-based policies in cost reduction, especially for coastal marshes and reefs, saving up to \$0.018 million per meter compared to static policies. It emphasizes the MDP's capacity for achieving globally optimal solutions in complex scenarios, underscoring the importance of cost-effective adaptation strategies. The inclusion of environmental impacts, like carbon emissions, in the adaptation framework underscores the commitment to sustainable strategies. The dynamic and flexible MDP approach surpasses traditional models by accommodating uncertainties and real-time data, marking a significant improvement in the field.

The methodological rigor and reproducibility of the study suggest its potential for broad application and validation.

We thank the reviewer for the appreciation of our work, and the recognition of the originality and flexibility of the presented framework and its potential for broad application.

Reviewer Comment 3.2 — Overall, the research underscores the substantial cost and environmental benefits of MDP-based adaptation, proposing a notable advancement in coastal infrastructure planning, albeit suggesting further validation and comparison with existing models for a stronger foundation.

We thank the reviewer for this comment that has assisted us in improving the manuscript. We have substantially revised the manuscript and enhanced the literature review and the related position and contributions of this work in it, providing now a new taxonomy of related approaches, and discussion in relation to the various sources of uncertainty present, as provided at the beginning of this documents and in manuscript lines: **56-168, 175-192**.

We have also added a note in the Discussion section about future potential directions, including a numerical comparison with existing adaptive methods, as:

Lines 615-620: Finally, the presented framework here can be further compared in the future with other formulations and adaptive solutions in the literature. Since it can provide globally optimal solutions for related optimization problems, as developed in this work, it can also serve as an additional point of reference for assessing the accuracy and performance of newly suggested methodologies.

Reviewer Comment 3.3 — A suggestion I have for the authors for their future work is to investigate deep reinforcement learning techniques to manage the increasing complexity and dimensionality of the problem. This approach could allow for handling a larger set of adaptive actions and longer decision horizons, potentially leading to more substantial cost savings and improved policy effectiveness.

We thank the reviewer for this suggestion. Higher-fidelity models can describe the entire system and environment settings in a more realistic and refined way but cannot usually provide optimality guarantees, and vice versa for simpler modeling representations. We believe to have contributed to this dilemma with a framework that can be generally applicable and more detailed than many of the currently existing approaches in the literature, and at the same time can provide globally optimal solutions, for the mathematical setting in which it is posed, something quite challenging for many applications and methods in the literature. At the same time, higher-fidelity representations are useful and have lots to offer for decision-making in many cases, despite losing the optimality guarantees. In our future work, we will thus also explore this direction based on higher-dimensional MDP/POMDP representations that can be solved by deep reinforcement learning-based approaches, hopefully to near global optimality. We also share the reviewer's view that a higher-fidelity model has a high probability to lead to more substantial cost savings and improved policy effectiveness, particularly if solved appropriately.

References

[1] R. J. Lempert, S. W. Popper, and S. C. Bankes, *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis.* Santa Monica, CA: RAND Corporation, 2003.

[2] R. L. Sriver, R. J. Lempert, P. Wikman-Svahn, and K. Keller, "Characterizing uncertain sea-level rise projections to support investment decisions," *Plos One*, vol. 13, no. 2, p. e0190641, 2018.

[3] C. Eijgenraam, J. Kind, C. Bak, R. Brekelmans, D. den Hertog, M. Duits, K. Roos, P. Vermeer, and W. Kuijken, "Economically efficient standards to protect the netherlands against flooding," *Interfaces*, vol. 44, no. 1, pp. 7–21, 2014.

[4] R. L. Ceres, C. E. Forest, and K. Keller, "Optimization of multiple storm surge risk mitigation strategies for an island city on a wedge," *Environmental Modelling & Software*, vol. 119, pp. 341–353, 2019.

[5] Y. Miura, H. Qureshi, C. Ryoo, P. C. Dinenis, J. Li, K. T. Mandli, G. Deodatis, D. Bienstock, H. Lazrus, and R. Morss, "A methodological framework for determining an optimal coastal protection strategy against storm surges and sea level rise," *Natural Hazards*, vol. 107, pp. 1821–1843, 2021.

[6] Ward, Philip J and Jongman, Brenden and Aerts, Jeroen CJH and Bates, Paul D and Botzen, Wouter JW and Diaz Loaiza, Andres and Hallegatte, Stephane and Kind, Jarl M and Kwadijk, Jaap and Scussolini, Paolo and others, "A global framework for future costs and benefits of river-flood protection in urban areas," *Nature Climate Change*, vol. 7, no. 9, pp. 642–646, 2017.

[7] J. H. Kwakkel, "The exploratory modeling workbench: An open source toolkit for exploratory modeling, scenario discovery, and (multi-objective) robust decision making," *Environmental Modelling & Software*, vol. 96, pp. 239–250, 2017.

[8] M. Gong, R. Lempert, A. Parker, L. A. Mayer, J. Fischbach, M. Sisco, Z. Mao, D. H. Krantz, and H. Kunreuther, "Testing the scenario hypothesis: An experimental comparison of scenarios and forecasts for decision support in a complex decision environment," *Environmental Modelling & Software*, vol. 91, pp. 135–155, 2017.

[9] R. Lempert, "Scenarios that illuminate vulnerabilities and robust responses," *Climatic change*, vol. 117, pp. 627–646, 2013.

[10] Arias, Paola and Bellouin, Nicolas and Coppola, Erika and Jones, Richard and Krinner, Gerhard and Marotzke, Jochem and Naik, Vaishali and Palmer, Matthew and Plattner, G-K and Rogelj, Joeri and others, "Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change; Technical Summary," 2021.

[11] J. D. Quinn, P. M. Reed, and K. Keller, "Direct policy search for robust multi-objective management of deeply uncertain socio-ecological tipping points," *Environmental modelling & software*, vol. 92, pp. 125–141, 2017.

[12] G. G. Garner and K. Keller, "Using direct policy search to identify robust strategies in adapting to uncertain sea-level rise and storm surge," *Environmental Modelling & Software*, vol. 107, pp. 96–104, 2018.

[13] M. Woodward, Z. Kapelan, and B. Gouldby, "Adaptive flood risk management under climate change uncertainty using real options and optimization," *Risk Analysis*, vol. 34, no. 1, pp. 75–92, 2014.

[14] A. Wreford, R. Dittrich, and T. D. van Der Pol, "The added value of real options analysis for climate change adaptation," *Wiley Interdisciplinary Reviews: Climate Change*, vol. 11, no. 3, p. e642, 2020.

[15] M. J. Lickley, N. Lin, and H. D. Jacoby, "Analysis of coastal protection under rising flood risk," *Climate Risk Management*, vol. 6, pp. 18–26, 2014.

[16] O. Špačková and D. Straub, "Long-term adaption decisions via fully and partially observable markov decision processes," *Sustainable and Resilient Infrastructure*, vol. 2, no. 1, pp. 37–58, 2017.

[17] R. Hui, J. Herman, J. Lund, and K. Madani, "Adaptive water infrastructure planning for nonstationary hydrology," *Advances in Water Resources*, vol. 118, pp. 83–94, 2018.

[18] J. D. Herman, J. D. Quinn, S. Steinschneider, M. Giuliani, and S. Fletcher, "Climate adaptation as a control problem: Review and perspectives on dynamic water resources planning under uncertainty," *Water Resources Research*, vol. 56, no. 2, p. e24389, 2020.

[19] S. S. Shuvo, Y. Yilmaz, A. Bush, and M. Hafen, "Modeling and simulating adaptation strategies against sea-level rise using multiagent deep reinforcement learning," *IEEE Transactions on Computational Social Systems*, vol. 9, no. 4, pp. 1185–1196, 2021.

[20] T. D. Ramm, C. S. Watson, and C. J. White, "Strategic adaptation pathway planning to manage sea-level rise and changing coastal flood risk," *Environmental Science & Policy*, vol. 87, pp. 92–101, 2018.

[21] J. H. Kwakkel, M. Haasnoot, and W. E. Walker, "Comparing robust decision-making and dynamic adaptive policy pathways for model-based decision support under deep uncertainty," *Environmental Modelling & Software*, vol. 86, pp. 168–183, 2016.

[22] M. Haasnoot, J. H. Kwakkel, W. E. Walker, and J. Ter Maat, "Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world," *Global Environmental Change*, vol. 23, no. 2, pp. 485–498, 2013.

[23] D. Bertsekas, *Dynamic programming and optimal control: Volume I.* Athena scientific, 2012, vol. 1.

[24] M. Pozzi, M. Memarzadeh, and K. Klima, "Hidden-model processes for adaptive management under uncertain climate change," *Journal of Infrastructure Systems*, vol. 23, no. 4, p. 04017022, 2017.

[25] C. P. Andriotis and K. G. Papakonstantinou, "Managing engineering systems with large state and action spaces through deep reinforcement learning," *Reliability Engineering & System Safety*, vol. 191, p. 106483, 2019.

[26] T. Smith and R. Simmons, "Point-based pomdp algorithms: Improved analysis and implementation," *arXiv preprint arXiv:1207.1412*, 2012.

[27] G. Garner, T. H. Hermans, R. Kopp, A. Slangen, T. Edwards, A. Levermann, S. Nowicki, M. D. Palmer, C. Smith, B. Fox-Kemper *et al.*, "Ipcc ar6 wgi sea level projections," 2022.

[28] "Station Home Page - NOAA Tides amp; Currents — tidesandcurrents.noaa.gov," https://tidesandcurrents.noaa.gov/stationhome.html?id=8518750, accessed 09-19-2023.

[29] G. Garner, T. Hermans, R. Kopp, A. Slangen, T. Edwards, A. Levermann, S. Nowicki, M. Palmer, C. Smith, B. Fox-Kemper, H. Hewitt, C. Xiao, G. Aalgeirsdóttir, S. Drijfhout, N. Golledge, M. Hemer, G. Krinner, A. Mix, D. Notz, I. S. Nurhati, L. Ruiz, J.-B. Sallée, Y. Yu, L. Hua, T. Palmer, and B. Pearson, "IPCC AR6 Sea Level Projections," Aug. 2021. [Online]. Available: https://podaac.jpl.nasa.gov/announcements/2021-08-09-Sea-level-projections-from-the-IPCC-6th-Assessment-Report

[30] NOAA, "Sea level trends - battery, ny - station id: 8518750," 2024, accessed: 2024-05-16. [Online]. Available: https://tidesandcurrents.noaa.gov/sltrends/sltrends_station.shtml?id=8518750

[31] P. G. Morato, K. G. Papakonstantinou, C. P. Andriotis, J. S. Nielsen, and P. Rigo, "Optimal inspection and maintenance planning for deteriorating structural components through dynamic Bayesian networks and Markov decision processes," *Structural Safety*, vol. 94, p. 102140, 2022.

[32] K.G. Papakonstantinou, C.P. Andriotis, Charalampos P and M. Shinozuka, "POMDP and MOMDP solutions for structural life-cycle cost minimization under partial and mixed observability," *Structure and Infrastructure Engineering*, vol. 14, no. 7, pp. 869–882, 2018.

[33] K.G. Papakonstantinou and M. Shinozuka, "Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part I: Theory," *Reliability Engineering &amp System Safety*, vol. 130, pp. 202–213, Oct. 2014. [Online]. Available: https://doi.org/10.1016/j.ress.2014.04.005

[34] E. J. Sondik, "The optimal control of partially observable markov processes," Ph.D. dissertation, Stanford University, 1971.

[35] K.G Papakonstantinou, M. Shinozuka, "Optimum inspection and maintenance policies for corroded structures using partially observable markov decision processes and stochastic, physically based models," *Probabilistic Engineering Mechanics*, vol. 37, pp. 93–108, 2014.

[36] K. G. Papakonstantinou and M. Shinozuka, "Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part II: POMDP implementation," *Reliability Engineering & System Safety*, vol. 130, pp. 214–224, 2014.

[37] T. Smith and R. G. Simmons, in *Focused Real-Time Dynamic Programming for MDPs: Squeezing More Out of a Heuristic.* Proceedings in AAAI Conference on Artificial Intelligence, 2006. [Online]. Available: https://api.semanticscholar.org/CorpusID:906058

[38] "GitHub - trey0/zmdp: Planning algorithms for problems with uncertain world state and action outcomes (POMDP and MDP models) — github.com," https://github.com/trey0/zmdp, [Accessed 17-09-2023].

[39] T. Smith, *Probabilistic planning for robotic exploration.* Carnegie Mellon University, 2007.

[40] "Implementation of Cost Benefit Analysis of Coastal Risk Management Prevention Measures against Natural Hazards," 2016, accessed: 2024-05-16. [Online]. Available: https://civil-protection-humanitarian-aid.ec.europa.eu/system/files/2017-11/ecoshaz_guidelines_1.pdf

Notes on revisions made to manuscript Paper NCOMMS-23-48445

We would like to again thank the editor and the reviewers for their comments and suggestions to further enhance the manuscript. In response, we have included a discussion on the novelty, impact, and limitations of our methodological approach, also explaining related features and comparisons with alternative methods. A point-by-point response to the reviewers' comments is provided in this response, and all revisions and additions to the manuscript are highlighted in the manuscript and this response in blue. For convenience, we have included here at the beginning a brief discussion outlining the novelty and current limitations of the methodology. Subsequently, all reviewers' comments are discussed independently and in detail.

The existing MDP/POMDP approaches for life-cycle adaptation under climate change effects, such as those in [1–4], have mainly, if not exclusively, successfully addressed the levee heightening problems that typically involve only a single state metric. Building on this initial work, our work here proposes a more general life-cycle adaptation framework that extends the MDP-based problem formulation to now encompass a wide range of action types, including system altering ones, and several state metrics, and is also appropriately designed for state-of-the-art MDP/POMDP solvers with global optimality guarantees, requiring specific modeling formulations. In the context of optimal adaptation of coastal infrastructure systems subjected to risks posed by storm surge and sea level rise, the states of the formulated MDP in this work include the incoming flood levels, along with the current infrastructure state and setting, since the damage induced at any point in time is dependent not only on the flood levels but the implemented system protections as well. In the currently existing, more specific adaptive policy applications of levee heightening, the level of protection at any time is fully determined by the height of the levee, which is defined as the only problem state. However, in our more general setting, the infrastructure state and protection are determined by the entire history of the previously executed adaptation actions up to the current time step, given that these actions can completely alter the system configuration. This characteristic makes the problem essentially non-Markovian, given that the optimal action at any point in time now depends on the history of previously executed actions rather than on the current state alone, as is the case if Markovian.

In the current work, and to realize a more general framework, the non-Markovian issue has been appropriately and originally addressed by augmenting the state space in such a manner as to enforce Markovian properties. Thus, we conjecture that the presented problem formulation is *general* in that it can include any type of action that can be used to mitigate climate-related risks. Furthermore, we have tested different nature-based solutions, such as salt marshes and oyster reefs, along with grey infrastructure, such as floodwalls, in a novel and first-time implementation in a MDP-type framework. We would also like to highlight the insights on the policies with grey and green infrastructure that are obtained by originally including current and future carbon emissions and carbon costs as part of the objective function in basing the decisions in our framework. Additionally, we expand the developed framework to incorporate uncertainty associated with the underlying climate models and show the evolution of model beliefs over time and its effects on the obtained policies.

The formulation in this work allows the problem to be solved exactly with dynamic programming (DP) solvers, a first in this general modeling setting. For the current work, we employed a state-of-the-art MDP/POMDP solver, FRTDP, which provides global optimality guarantees and can thus also serve as a critical benchmark for comparison in all possible applications and solution approaches. However, solving the problems with these guarantees imposes certain constraints on the dimensionality of state and action spaces that can be presently considered. Yet, this work, we believe, sets the stage for new adaptive climate change-related solutions that can inspire the further development of DP solvers that are able to scale even more (for example, factorized DP solvers, dedicated finite horizon solvers, etc.). Deep reinforcement learning (DRL) also presents another promising approach suitable for this modeling framework, which we are actively exploring. However, for ourselves or others looking to develop new DP solvers and/or DRL solution algorithms so as to scale the problem (consider more actions, longer time horizons, etc.), the current work provides the much-needed optimal solutions to a general climate adaptation problem that can serve as a benchmark in the future. A thorough theoretical and numerical comparison with several adaptive solutions and modeling approaches, elucidating their advantages and disadvantages, strengths and weaknesses, can be investigated in the future by us and others.

Related to the above, added/revised text in the manuscript can be now found in both the introduction and discussion sections.

Lines 255-267: To summarize, this study makes the following significant contributions to the planning of climate-change-related risk mitigation:

1. Presents the foundational elements of a general framework for the life-cycle adaptation under climate change based upon stochastic control methodologies of MDPs and POMDPs. The framework is sufficiently general to be readily extended to consider a wide range of actions and is compatible with state-of-the-art point-based MDP/POMDP solvers with global optimality guarantees.

2. Formulates a general POMDP problem, by considering climate model uncertainties updated in time through relevant observations.

3. Considers environmental impacts, the social cost of carbon, and related carbon emissions and uptake in the decision-making process.

4. Accounts for various nature-based infrastructures as potential flood risk mitigation measures and demonstrates their effects in mitigating carbon emissions.

Lines 524-533: In this work, the life-cycle adaptation of infrastructure under climate change has been framed as a stochastic-control MDP/POMDP problem, which is generalized to accommodate a wide variety of action types, even system-altering ones. Within this framework, optimal actions are identified to be taken adaptively in time based on observed data and realized conditions, reducing life-cycle costs and carbon emissions and responding to the important uncertainties that characterize the problem. The efficacy of the MDP-based adaptive framework is guaranteed over other methods for the same optimization problems posed due to better-informed policies, which are chosen according to the actual evolving climate conditions and global optimality guarantees, mathematically defined.

Lines 604-636: Overall, this work aims to present a generalizable stochastic control MDP/POMDP framework for life-long optimal infrastructure adaptation based on the observed evolving climate conditions. The developed framework is general in that it can incorporate a wide variety of possible action types, also considering nature-based infrastructure options along with grey infrastructure choices without making any structural changes to the problem setting. This work also provides globally optimal solutions while accommodating various other features, such as model uncertainty aspects and the social cost of carbon considerations in the analysis. This framework provides globally optimal solutions while accommodating various features, including, among others, a wide variety of possible action types, also considering nature-based infrastructure options, model uncertainty aspects, and social cost of carbon considerations in the analysis. As with all MDP/POMDP-based frameworks, the computational complexity increases with an increase in the number of states, possible dynamic actions, and time horizon length, which at some point cannot be practically handled with MDP/POMDP solvers that provide guarantees of global optimality. This issue can be overcome, at the expense though of optimality guarantees, by pursuing deep reinforcement learning-based solution techniques, allowing consideration of a larger number of adaptive actions and longer time horizons. Given this higher flexibility and fidelity, the identified adaptive policies may be expected to result in significantly higher cost savings in comparison to static cost-benefit analysis solutions, something we are aiming to explore in the future. Other adaptive approximate methods, such as those based on DPS, for example, can also offer benefits in reducing the dimensionality of the decision variables of the problem and, therefore, may offer solutions and viable alternatives in situations where state-action relation mappings are available, and optimality guarantees are not required or cannot be achieved. Moreover, the various considered objectives were aggregated in this work into a single scalarized objective function. In future work, we also aim to extend our framework to consider relevant multi-objective formulations to the problem by balancing various tradeoffs among the constituent objectives. Finally, the presented framework here can be further compared in the future with other adaptive solutions in the literature. Since it can provide globally optimal solutions for related optimization problems, as developed in this work, it can also serve as an additional point of reference for assessing the accuracy and performance of newly suggested methodologies.

Response to the reviewers

Reviewer 1

The authors' considerable effort to address the comments is highly appreciated. I am also happy with most responses and modifications to the manuscript. I have only two remaining points, which the authors should consider for the final manuscript:

We thank the reviewer for the insightful comments and suggestions, and the appreciation for our work. We are responding to the remaining comments in detail below.

Reviewer Comment 1.1 — I am not entirely satisfied with the response to my comment 1.3. What the authors write is correct, but my comment was directed towards the fact that the current practice is indeed a static policy within an open-loop planning framework (as the authors nicely put it). The authors comparison however assumes that the static policy is never changed, which makes it unnecessarily pessimistic. To be more specific, I refer to Figure 5: If my understanding is correct, both the MDP and the static policy optimization recommend the construction of F1 initially. If one only looks at the decisions taken within the nearer future, they are identical for the two approaches. The fact that the optimal static policy recommends not to build F2 later is irrelevant, as this decision will be revised later. But since this revision is not accounted for in the current calculations done in the paper, the estimated life-cycle costs for this policy are pessimistic. (If one were to compute the expected cost of the open-loop planning with static policy, that would be closer to the one of the MDP model). (The fact that also the MDP model can also be updated a later point is irrelevant here, as the comparisons assume that dynamics underlying the MDP model are the true one.) I would urge the authors to state this clearly in the manuscript, as this is an important point that is often misunderstood.

We thank the reviewer for their further clarification. The reviewer's comment correctly highlights that an open-loop policy, when executed anew at a subsequent time, may identify a novel action to be implemented in the future. However, this type of open-loop optimization comes with the risk of yielding less-informed and potentially suboptimal solutions, as it still operates on an averaged forecast over the remaining future without incorporating real-time feedback, apart from the initial conditions used in every loop. In contrast, closed-loop policies are derived using dynamic programming principles, which provide strong mathematical guarantees. These policies do not require rerunning the optimization problem with updated conditions, since the obtained policies can continuously suggest actions based on real-time feedback. A more detailed comparison with the approach suggested by the reviewer and studying its implementation details warrants future study, given also the availability of benchmark solutions provided by our method. We have also now added the following statement related to this in the manuscript.

Lines 546-554: A final note about the cost-benefit-based static policy considered in this work is that it is only evaluated here based on the initial conditions of the studied horizon, as typically done. Future work can, however, further study these practical approaches based on open-loop control formulations, i.e., recomputing the cost-benefit policy in time with updated conditions as things evolve. However, such static optimization and resulting open-loop policies come with the risk of yielding less-informed and potentially suboptimal solutions, as the identified actions are based upon an averaged forecast over the remaining future without incorporating real-time feedback.

Reviewer Comment 1.2 — This is a minor point, but in line 261ff the authors write that the study "presents a general life-cycle adaptation framework under climate change that extends the MDP-based problem formulation …". This makes it sound like the MDP formulation was the current state of the art or practice. However, the current state of practice is the use of static policies in an open-loop planning framework, as discussed above. The current state of the art in science are POMDPs, which were considered in some of the references now included in the paper.

We thank the reviewer for this comment. We have revised the statement to now emphasize that our work builds on the stochastic control-based frameworks of MDPs and POMDPs. It extends these formulations to incorporate various types of actions and makes them compatible with state-of-the-art dynamic programming solvers that provide global optimality guarantees within the problem's settings. The following text is now added in the paper:

Lines 257-261: Presents the foundational elements of a general framework for life-cycle adaptation under climate change based upon the stochastic control methodologies of MDPs and POMDPs. The framework is sufficiently general to be readily extended to consider a wide range of action types and is fully compatible with state-of-the-art point-based MDP/POMDP solvers with global optimality guarantees.

Reviewer Comment 1.3 — Some of the new text that the authors included involves very long sentences that are a bit hard to read (E.g., lines 183 – 193). It is recommended to check the manuscript for readability.

Thank you for this comment, we have revised the manuscript to eliminate the long sentences to enhance readability.

Reviewer 2

I have spent some time (and a few rounds) reading the revision of "Optimal Life-Cycle Adaptation of Coastal Infrastructure under Climate Change," and I unfortunately cannot convince myself that the authors have quite responded to a core critique from the initial review (both from myself and from Reviewer 1), which I will discuss more below. I unfortunately cannot recommend publication in Nature Communications (or a similar broad-audience journal) for this reason, though a major revision might be in order if this was a miscommunication or misinterpretation, rather than an attempt to avoid the critique.

Let me be more direct to avoid any potential miscommunication. The fundamental problem with this manuscript (for this journal; this may not apply to more discipline-specific journals) is that, for an article of this sort (I might frame this more in terms of decision analyses, but I think this applies more generally) to be useful to a broad audience, it should contain at least one of the following elements:

1) Development of a new method which addresses limitations of existing methods

2) Application of an existing method to a problem which highlights new insights about the problem or its setting

3) Application of an existing method to a problem which is a demonstrated advance on the state of the art (or the state of the practice, if it can be articulated why this method may be more implementable by practitioners than other methods which advance on the state of the practice but are not commonly used).

That this article fails to do any of these three was at the root of my comments about the framing of the paper (including the large gaps in the literature review) and the comment Reviewer 1 and I made about the overly stylized example (in terms of the decision space) and the corresponding lack of model fidelity to a "real" decision problem. The paper is applying an existing method (FRTDP) in what is hopefully a new context, and so the contribution is not in method development. The paper's framing does not make any claims to new insights (with the exception of the inclusion of the social cost of carbon in changing the cost-benefit analysis), but rather that the FRTDP algorithm improves on the state of the art through its adaptativity and by having optimality guarantees.

The problem with these claims is that there are many other adaptive methods; while several (such as DPS) do not have optimality guarantees, these are of limited practical application since one's model is never "correct," and the "optimal" solution to the modeled problem may be suboptimal for the "true" problem. In their revision, the authors have tried to address the issue of model error through hand-waving away the issue by claiming that, in principle, as the FRTDP algorithm has optimality guarantees, it would find the optimal solution for any given problem framing, so anyone could just select the method they prefer. The issue with this claim is that the benefits of the method are presented as an advance on other methods, but only compared directly to a static method (which is a strawman compared to the literature, though perhaps not compared to the practice; more on this below). It might even be that for this simplified problem (limited decision levers; independent [versus autoregressive] noise around the future projections; see also the response to my comment about sampling states of the world, which might be required for problems which are more complex), the FRTDP algorithm does indeed outperform other algorithms which have been used in the literature, but the paper does not make this comparison, instead relying on its characterization of other methods as a reason to prefer FRTDP. But the characterization of DPS, for example, is a bit glib: it claims (on lines 99-103) that DPS relies on heuristics (true, but may not practically matter), that the actions are "arbitrarily defined" and "explore only a limited subset of the policy space." But this is both true of the example application and is true of pretty much every model-based decision problem, not unique to DPS (one could set up a DPS problem to capture the same set of policy options as the FRTDP example). The problem with the stylization of the problem in this manuscript is that it is unclear whether, as a practical matter, the theoretical guarantees of the FRTDP algorithm make a difference when the problem becomes sufficiently complex, and this is impossible to understand from the manuscript.

This would not be a problem if the FRTDP algorithm addressed practical implementation problems with other adaptation methods, which are why many of these methods have not caught on in application; these range from framing complexity to computational complexity to transparency and communication complexity. But none of this is addressed by the manuscript other than comparing the number of samples the authors needed (which is conditional on the problem framing) to the number of all possible combinations, which is not the fair comparison. The best way to address this issue would be to include other analogous methods in the comparison to show that FRTDP improves on them in this case, and argue why this ought to generalize; it would also be appropriate to make the argument on the basis of the properties of the other methods, though the authors do not have a comprehensive enough literature review to do this, and I am afraid based on their characterization of DPS discussed above that this may not be a useful exercise in this case.

As a result, I feel that the issue of model fidelity and model error is a central problem for publication in a broadly read journal like Nature Communications, and one which the authors failed to respond to in their revision. This may be less of a problem for other journals which are more focused on disciplinary audiences, but I would not speak for them.

Rather than emphasizing the methods, the authors do have a route towards new insights, namely the role of monetizing future carbon emissions in changing decisions; in this case, the MDP formultion and the FRTDP algorithm do provide a route towards identifying these effects, and the issue of "optimality" is less important. But the authors did actively choose to not reframe their manuscript in this fashion in the previous revision, so I am unsure if they are interested in doing so. However, that would in my view obviate the entire argument against publication I raised above (though others might emerge).

We would like to thank the reviewer for their time with our revised manuscript, and the general comments and discussion. We generally agree with the criteria listed by the reviewer for the basis of acceptance. We have thus provided an overview of our work's contributions below based on the categories mentioned by the reviewer. Owing to the length of the review, we have broken down our response into distinct sections, as below:

1) Development of a new method which addresses limitations of existing methods

This work develops a generalizable life-cycle adaptation framework that incorporates various sources of uncertainties and a variety of action types to mitigate climate-related risks. We argue that this framework is a significant contribution to the field, as no such comprehensive methodology currently exists. Existing related literature typically addresses climate risk mitigation with simplified MDP formulations dependent on the state defined by a single variable type/metric. For example, [1–4] have typically addressed levee-heightening problems, where the current levee height fully determines the state of the MDP problem, maintaining its Markovian properties, as the future states and costs are defined by the current state-action pair.

However, when considering a multitude of different types of risk-mitigation actions, the problem becomes more complex. In a more general setting, the current infrastructure state and level of protection are influenced by the entire history of previously executed adaptation actions up to the current time step, given that these actions can completely alter the system configuration and dictate how the system possibly evolves in the future. This characteristic makes the problem non-Markovian, as the history of actions becomes relevant now. Our work addresses this issue by augmenting the state space in a novel way to enforce a Markovian formulation, also importantly ensuring compatibility with state-of-the-art dynamic programming solvers with optimality guarantees and offering also valuable benchmark solutions for comparison with other approaches and solvers that can offer heuristic solutions, when such methods may offer practical or other advantages, for any reason possible.

In this work, we solve the developed framework with a state-of-the-art dynamic programming solver called FRTDP, which allows the problem to be solved with global optimality guarantees. Solving the climate adaptation problem with global optimality guarantees is important as the solution can serve as a benchmark for the development of new DP solvers and/or solutions by deep reinforcement learning in future work by ourselves and others. This solution method, while providing optimality guarantees, imposes certain limits on the number of possible actions that can be considered for decision-making. However, this limitation is not inherent to the developed climate adaptation framework presented here. Instead, this work represents a significant advance in establishing the problem formulation while ensuring compatibility with other advanced offline dynamic programming (DP) solvers (e.g., LRTDP, HSVI, RTDP, SARSOP). It also provides the groundwork for developing more sophisticated and scalable DP solvers for problems of this type, as presented in the manuscript.

2) Application of an existing method to a problem which highlights new insights about the problem or its setting

In this work, we extend the MDP/POMDP-based formulation for life-cycle adaptation to address climate change-induced sea level rise. Within this developed framework, we assess the performance of various nature-based solutions, examine the impact of model uncertainty, and consider the inclusion of carbon emissions in the decision-making costs. To the best of the authors' knowledge, there exists no generalized control-based framework that can be used to incorporate a wide variety of action types, such as nature-based solutions and grey infrastructure choices, together in the decision-making problem. This work has also provided new insights by incorporating the social costs of carbon emissions as part of the costs in the control-based decision-making problem and its influence on the resulting policies. We show that not considering the full costs of climate change delays actions. This highlights the importance of not only understanding the consequences of climate change but also considering the full costs, including carbon costs, when seeking to mitigate its effects.

The generality of the framework also allows us to consider various sources of uncertainty (we provide a detailed taxonomy of the sources of uncertainty that can/have been incorporated into our framework in the manuscript lines **58-70**). Under model uncertainty, we see that the framework can infer the underlying sea-level-rise models with the incoming observations in time. This is also a new observation; previous work, e.g., [2], has shown that the models cannot be inferred in the case of river discharges. However, this is not the case in the context of sea level rise scenarios, where higher sea level rise values are more likely to be observed under more extreme climate models. Additionally, we have demonstrated this study considering actual IPCC SSP scenario models, unlike most other similar studies, such as in [3], for example.

3) Application of an existing method to a problem which is a demonstrated advance on the state of the art (or the state of the practice, if it can be articulated why this method may be more implementable by practitioners than other methods which advance on the state of the practice but are not commonly used).

This work builds upon the stochastic-control-based formulation for the life-cycle adaptation of infrastructure systems and extends the state-of-the-art formulation to incorporate a variety of action types, state metrics, and sources of uncertainties and cost objectives. The state-of-the-art MDP/POMDP-based framework typically only addresses MDP problems concerning a single state variable, for example, the particular setting of levee heightening, which vastly simplifies the problem formulation as explained above. Furthermore, only a few works have systematically studied the effect of model uncertainty within this framework subjected to different sea level rise models. We have also studied the effect of incorporating carbon emissions within this framework and show the effect of the choice of the social cost of carbon on the learned policies.

This work thus establishes a general climate adaptation framework that provides essential benchmark solutions for comparison with existing adaptive approaches while also paving the way for the development of more sophisticated dynamic programming solvers and reinforcement learning-based solutions, which can be applied to the framework without requiring any structural changes.

Overall, we believe that our suggested methodology and this paper offer contributions in all three categories above and provide useful insights about modeling and solving climate-change-related adaptive policy problems, as also hopefully the reviewer can attest, based on the information that it seems from our discussion here to have acquired from our paper.

The problem with these claims is that there are many other adaptive methods; while several (such as DPS) do not have optimality guarantees, these are of limited practical application since one's model is never "correct," and the "optimal" solution to the modeled problem may be suboptimal for the "true" problem. In their revision, the authors have tried to address the issue of model error through hand-waving away the issue by claiming that, in principle, as the FRTDP algorithm has optimality guarantees, it would find the optimal solution for any given problem framing, so anyone could just select the method they prefer. The issue with this claim is that the benefits of the method are presented as an advance on other methods, but only compared directly to a static method (which is a strawman compared to the literature, though perhaps not compared to the practice; more on this below). It might even be that for this simplified problem (limited decision levers; independent [versus autoregressive] noise around the future projections; see also the response to my comment about sampling states of the world, which might be required for problems which are more complex), the FRTDP algorithm does indeed outperform other algorithms which have been used in the literature, but the paper does not make this comparison, instead relying on its characterization of other methods as a reason to prefer FRTDP. But the characterization of DPS, for example, is a bit glib: it claims (on lines 99-103) that DPS relies on heuristics (true, but may not practically matter), that the actions are "arbitrarily defined" and "explore only a limited subset of the policy space." But this is both true of the example application and is true of pretty much every model-based decision problem, not unique to DPS (one could set up a DPS problem to capture the same set of policy options as the FRTDP example). The problem with the stylization of the problem in this manuscript is that it is unclear whether, as a practical matter, the theoretical guarantees of the FRTDP algorithm make a difference when the problem becomes sufficiently complex, and this is impossible to understand from the manuscript.

To the best of the authors' knowledge, the current literature has applied DPS primarily again to dike/levee heightening type of problems with their characteristics as explained above, where the relationships between current water levels and trigger actions are defined by equations based on exploratory analyses or expert judgment rather than systematic mathematical derivations. This is what we meant as "arbitrary". As such, DPS explores the limited policy space captured by these state-trigger relations, which may not always include the optimal policy. Moreover, these state-action relationships must be derived separately for each action and application, making these methods less generalizable and harder to extend to different contexts and settings. This issue is also related to the 'communication' and 'transparency' aspects, also mentioned by the reviewer.

On the other hand, methods based on stochastic control formulations have sound mathematical guarantees and do not require an explicit definition of state-action relationships. These methods can be readily extended to incorporate any type of action as long as the effect of the action on the environment is adequately modeled in the framework, as has been carefully and originally handled in this work. Importantly, the MDP-based framework is also flexible to be extended to incorporate various sources of uncertainty in a principled manner, as we explain in the paper, together with other reported relevant innovations. Therefore, we chose a closed-loop control-based formulation to solve the problem at hand and extend the current literature, which also has typically applied these methods in levee heightening problems and thus retaining a simplified Markov decision process (as previously explained). We have added the following statements in the manuscript to clarify this point.

Lines 101-107: With DPS, the proposed actions are condition/state-dependent, as observed in time; however, the expressions relating the observable state variables to the triggering actions are action-specific and derived based on exploratory analyses or expert judgment, complicating somewhat the extension of DPS to several different action types and applications. DPS solutions are based on heuristic searches and thus lack optimality guarantees while also exploring only a relatively limited subset of the policy space, which is captured by the state-action expressions.

This would not be a problem if the FRTDP algorithm addressed practical implementation problems with other adaptation methods, which are why many of these methods have not caught on in application; these range from framing complexity to computational complexity to transparency and communication complexity. But none of this is addressed by the manuscript other than comparing the number of samples the authors needed (which is conditional on the problem framing) to the number of all possible combinations, which is not the fair comparison. The best way to address this issue would be to include other analogous methods in the comparison to show that FRTDP improves on them in this case, and argue why this ought to generalize; it would also be appropriate to make the argument on the basis of the properties of the other methods, though the authors do not have a comprehensive enough literature review to do this, and I am afraid based on their characterization of DPS discussed above that this may not be a useful exercise in this case.

We have included a comprehensive review of the different dynamic adaptive approaches that are present in the literature in the manuscript lines **96-121, 122-147**, and we have described the characteristics of DPS in detail in response to the previous comment and in the manuscript line **101-107**. The comparison has been made based on the characteristics of the other methods since it is not straightforward to apply different approaches, including DPS, to the current scenarios, especially including floodwalls and nature-based solutions in the same case study. The authors did not find a single paper in the literature that makes a comprehensive comparison of the currently existing adaptive approaches in a single case study, and the authors believe that such a comparative study can be a strong contribution to the community in itself. This work puts forward a novel MDP/POMDP framework and uniquely provides exact adaptive optimal solutions for a specific modeling environment used. In the future, we can direct our efforts toward a thorough comparative study, using also our derived methodology here as a benchmark solution, since it is distinct from other approaches in being able to provide globally optimal policies.

As a result, I feel that the issue of model fidelity and model error is a central problem for publication in a broadly read journal like Nature Communications, and one which the authors failed to respond to in their revision. This may be less of a problem for other journals which are more focused on disciplinary audiences, but I would not speak for them.

We agree that models can only capture the true reality with some level of approximation. Notwithstanding, to assess the performance of various algorithms, we always use the same models as benchmarking environments, and solutions built on stochastic control principles have nice mathematical guarantees in suggesting optimal sequential decisions. Additionally, there can be multiple sources of model error, such as errors in assuming the wrong climate evolution model, errors in simulating realistic storms/ hurricanes, errors in simulating realistic floods over a city map and estimation of losses, and so forth. The developed POMDP framework can adeptly address the issue of uncertainty associated with underlying climate models, as shown in the paper, and can suggest related optimal actions. We have demonstrated the POMDP framework by considering two possible SSP scenarios, SSP2-4.5 and SSP5-8.5, and show the evolution of the model beliefs with the incoming sea level rise observations in time. It is not evident how such a principled analysis on model uncertainty can be performed with other adaptive approaches, such as DPS, for example, and we claim that this is still an open topic with such methods. Thus, this work proposes a new and general framework that can adeptly accommodate various forms of uncertainty allowing other solution techniques, with or without providing optimality guarantees.

Rather than emphasizing the methods, the authors do have a route towards new insights, namely the role of monetizing future carbon emissions in changing decisions; in this case, the MDP formultion and the FRTDP algorithm do provide a route towards identifying these effects, and the issue of "optimality" is less important. But the authors did actively choose to not reframe their manuscript in this fashion in the previous revision, so I am unsure if they are interested in doing so. However, that would in my view obviate the entire argument against publication I raised above (though others might emerge).

We highly appreciate this comment in recognition of the new insights on the effect of the incorporation of carbon emissions on the learned policies. The authors feel that this is a significant development that showcases the effect of including carbon emissions in the cost objective of the developed framework but is not the core idea of the work. The central theme of this paper is, first and foremost, to illustrate an MDP-based control framework that can be generalized to include several action types, taking into account various sources of uncertainty and different cost objectives and making it amenable to be solved with state-of-the-art dynamic programming solvers that can offer mathematical guarantees. We understand the point raised by the reviewer that perhaps under some scenarios, the optimality guarantees might not be relevant or necessary. However, we would like to point out that the developed framework is original and generalizable irrespective of the choice of the specific solver used to solve it. For example, reinforcement-learning-based solutions can potentially solve the framework in applications with more complexity and fidelity without offering guarantees, and we are exploring extending this work in that direction in the future, together with comparisons with appropriate FRTDP benchmark solutions.

Specific Comments (these are somewhat limited because I flagged some before coming to the conclusion that the paper does require a much more thorough reframing for publication in Nature Communications):

Lines 14-17: This initial line in the abstract is the core argument about the focus on a static counterfactual, with very little evidence in the introduction to support it.

Thank you for this comment. This line in the abstract highlights the fact that most decision-making and decision-support approaches today still largely rely on cost-benefit analysis, particularly when federal agencies and authorities are concerned. For example, Implementation of Cost Benefit Analysis of Coastal Risk Management Prevention Measures against Natural Hazards (2016) [5], USACE and FEMA [6], and [7], among others. Many works in the literature also use cost-benefit analysis to assess the performance of different action alternatives; see [8–14] for example. The following sentence has now be added in the manuscript:

Lines 71-73: A conventional approach for climate change planning has been primarily driven by cost-benefit analysis (CBA), particularly when federal and national agencies and authorities are concerned [5–7].

Adaptive approaches to climate-related risk management have also been explored, and we provide a detailed review of the current literature in the manuscript **96-121, 122-147**. However, there remains a notable gap in the comparative analysis of different adaptive approaches applied to real climate change-related problems. This deficiency can be attributed to various factors, including the availability of source codes and the flexibility of methodologies to extend to diverse application scenarios, among others. A thorough survey would indeed be very useful to the research community. Though this was not the aim of this paper, which focused on developing a new methodology, this paper also puts forward an important step in this direction by providing a detailed taxonomy of solution approaches and different sources of uncertainty associated with climate change adaptation as well as offering solutions with optimality guarantees that can also assist in establishing benchmarks.

The primary objective of this work is thus to develop a generalizable framework underpinned by the rigorous mathematical foundations of stochastic control. We extend the existing MDP/POMDP-based problem formulations to accommodate a broad spectrum of action types. This extension also allows for the evaluation of various nature-based solutions in conjunction with grey infrastructure choices within a unified MDP-based framework, a demonstration that has not been previously achieved. Furthermore, we provide a novel analysis of model uncertainty, specifically associated with the IPCC AR6 SSP models, on the derived policies. We also elucidate how the probabilities of underlying models evolve in response to new sea level rise observations. Additionally, we examine how changes in the social cost of carbon influence optimal policies, another aspect that had not been previously published. We chose to compare with static baselines since they represent the state of practice. At this time, comparison with different adaptive approaches, such as DPS, ROA, and others, is not possible, mainly due to lack of flexibility and other reasons mentioned above. A comprehensive survey comparing various adaptive approaches across different problems would be a significant contribution, but in our opinion, this warrants a separate and dedicated publication. The following related sentence is now added to the paper:

Lines 632-636: Finally, the presented framework here can be further compared in the future with other adaptive solutions in the literature. Since it can provide globally optimal solutions for related optimization problems, as developed in this work, it can also serve as an additional point of reference for assessing the accuracy and performance of newly suggested methodologies.

Lines 49-56: This sentence makes it seem as though the analysis would address a more complex problem formulation involving multiple interacting decision levers, which would actually be a relevant level of complexity for comparison to other methods, and one which is rarely done in the literature, but the actual examples do not do this.

We thank the reviewer for pointing this out and would like to take this opportunity to reiterate the originality of our work. The developed MDP/POMDP-based problem formulation is original in that it is the first study with general action types, provided that their effects on the environment are appropriately modeled, dealing also with the non-Markovian characteristic of the general climate adaptation problem. This level of generality is not present in any previous MDP/POMDP-based formulations. Our framework also integrates uncertainty related to underlying climate models and includes future carbon emissions as an additional cost objective in decision-making. Additionally, our framework has been designed to ensure that existing state-of-the-art MDP/POMDP solvers with optimality guarantees can be readily applied, which, as recognized, is not something easily implemented [2].

Further research is needed to develop even more sophisticated solvers that can handle increased complexity and fidelity within this framework. Reinforcement-learning-based solution methods also offer the potential for addressing higher-fidelity environments, and we are actively exploring this avenue, but they do not provide the strong optimality guarantees as the DP solvers used here, which we feel is an appropriate first step for later scaling and complexity.

Fig 1 Caption: Aren't there 21 climate projections used? Which SSP5-8.5 realization did you use in panel A?

The SLR model in this work is derived based on regional SLR projections using CMIP6 models from the AR6 of the IPCC [15], for the Battery tidal gauge [16]. The SLR projections for each SSP are obtained from [17], and for each SSP model an ensemble of 21 Global Climate Models (GCM) is considered. In Fig 1 A, the uncertainty range obtained from the entire ensemble of 21 GCMs is plotted, as obtained from [17] for NYC.

Lines 99-103: Again, this is not quite a fair characterization of DPS, and I worry a bit about whether it is hyperbolic to make the method in the manuscript look better. DPS does not have to explore any more of a ”limited” subset of the policy space than any other approach (though it may get more complex to formulate, implement, and communicate as the policy space expands, particularly to multiple levers), and often the state space mapping is formulated using exploratory analyses of relevant sensitivities and options, so is not any more arbitrary than the limited policy options used here.

We thank the reviewer for this comment and would like to clarify the challenges associated with the implementation of DPS in different applications. As the reviewer correctly pointed out, state space to trigger action mapping is derived from exploratory analyses and must be carefully defined for each type of action, which is typically not a straightforward task for climate change-related applications. For instance, in this work, we considered various nature-based solutions, and to integrate these into our framework, we only had to model their effects on incoming flood levels through their respective wave attenuation properties, as well as account for their implementation and maintenance costs (including carbon emissions).

Our suggested approach and solution optimally learn to implement these options based on their impacts on reducing flood losses and emissions. However, there is no existing literature that we are aware of on generally establishing state mapping equations for these nature-based solutions or other actions within a DPS-based framework. This lack of guidance presents a significant challenge in adapting DPS to accommodate a broader range of action types and, hence, in offering a direct comparison. Moreover, these methods heuristically search in the limited policy space captured by the specific state-action mappings and hence do not guarantee convergence to the optimal policy.

We would also like to point out, however, that DPS does offer benefits in reducing the dimensionality of the decision variables of the problem and, therefore, should be preferred in situations where action-specific state mappings are well understood and can somehow be adequately modeled, and optimality guarantees are not required within the problem settings. We have now added a statement to highlight this in our Discussions section:

Lines 624-628: Other adaptive approximate methods, such as those based on DPS, for example, can also offer benefits in reducing the dimensionality of the decision variables of the problem and, therefore, may offer solutions and viable alternatives in situations where state-action relation mappings are available, and optimality guarantees are not required or cannot be achieved.

Lines 108-112: Are you claiming that adaptive pathways necessitate a “few” forward simulations? There are no citations to justify this, or specifics given on what ”a few” means.

The tipping points or useful life of the actions are determined over an ensemble of future trajectory scenarios which allow for a wide variety of uncertainties about future developments to be taken into account in the planning process [18]. Distribution of the tipping points is usually generated, and the median or quartile values are used to generate the adaptation pathways [18]. We have cited this paper [18] in our literature review on adaptive pathways, and we have added/modified the following related sentences in the manuscript:

Lines 113-118: Alternative approaches have also been suggested through adaptive pathway solutions [18–20], identified through explorative modeling of adaptation tipping points in time. These methods rely on forward simulations/trajectories of possible futures in time to identify the tipping points [18], sharing many of the limitations of the cost-benefit and DPS methods and resulting in the adaptive pathway solutions likely to be sub-optimal for the actual evolving conditions in time.

Line 113: From a real-world perspective, there is no guarantee that even a theoretically "optimal" solution is optimal, hence the importance of model fidelity and the salience of the problem framing, which is totally missing from this manuscript despite its grandiose claims.

Thank you for the comment. We would again like to point out here that the examples considered in this work are in no way a representation of the limitations of the developed MDP-based framework. The framework is general and can accommodate various actions with diverse effects on the environment without changing the problem settings by adequately modeling their effects on the environment. The solution of the framework has been done exactly with guarantees in this work as a foundation and benchmark that can be further scaled in the future. It is mentioned that reinforcement learning-based solutions can potentially solve a more realistic problem with higher fidelity, most importantly, which can be developed using the same framework as developed in this work.

Lines 615-624: As with all MDP/POMDP-based frameworks, the computational complexity increases with an increase in the number of states, possible dynamic actions, and time horizon length, which at some point cannot be practically handled with MDP/POMDP solvers that provide guarantees of global optimality. This issue can be overcome, at the expense though of optimality guarantees, by pursuing deep reinforcement learning-based solution techniques, allowing consideration of a larger number of adaptive actions and longer time horizons. Given this higher flexibility and fidelity, the identified adaptive policies may be expected to result in significantly higher cost savings in comparison to static cost-benefit analysis solutions, something we are aiming to explore in the future.

Lines 129-130: There are other examples of approaches in the literature for sea-level rise adaptation which also use MDP and closed loop controls (I pointed out some in the previous review) but the paper still reads as though it is claiming to be the first to have done so. This shows the limits to the improvement of the literature review in this revision. I get the sense from reading the paper that the authors are trying to overstate the novelty of the paper and the strength of its contribution, which is not needed (or helpful!).

We appreciate the reviewer's comment and would like to clarify how our MDP-based problem formulation differs from those in the existing literature. As noted in the literature review section, current MDP-based approaches typically focus on problems defined by a single state variable type. For instance, in levee heightening scenarios [1–4], the state of the environment is fully described by the levee height. This simplification arises because knowing the current levee or dike height ensures that the environment transitions remain Markovian, with future states and costs determined solely by the current state and action pair.

However, in a more general environment setting where we can have any set of actions such as seawalls, floodgates, and green infrastructure, among others, the environment state definition needs considerable modification. Since an action taken in the past changes the behavior of the system and how flooding will incur costs in the future, the current flood-induced damages and costs will depend on the available protections, as determined by the history of past executed actions. Thus, the information of already implemented actions needs to be maintained over time, and this information is referred to as the "system" information in this work. Given the current flood level and the "system" information, the flood-induced costs can now be calculated in a way that preserves the Markovian property. Markovian manner. Thus, the state space in a general environment setting, with a broad range of actions that can change the system, needs to be originally augmented, as shown in this work. To be also compatible with state-of-the-art dynamic programming solvers, the problem is formulated as an infinite horizon MDP/POMDP problem by appropriately adding a final terminal state which indicates the end of the decision horizon, as explained in the paper.

Lines 138-139: I don't understand what you mean by "quite specific to the characteristics of their particular application scenarios." Do you mean that the decision problems are formulated to be specific to e.g. levee heightening or sea-level rise generally? Depending on what you mean, I am not sure how this is avoidable in any situation. If you mean that the example is specific to a particular case study, this seems like a more tenuous claim: details will always be specific to the particular case study, but it's rare that methods would only work in one location, isn't it?

The existing literature on adaptive decision-making typically addresses MDP problems that can be defined by a single state variable type, as just explained in our response to the previous comment. Previous works [1–4] have typically implemented an MDP-based formulation for levee heightening problems where the current height of the levee constitutes the state definition of the problem. However, our adaptive framework can consider a diverse set of action types, such as floodwalls, green infrastructure, and so forth. This requires an original formulation of the MDP problem by augmenting the current state with the information of previously executed actions, as previously explained in detail. Our devised framework is thus general in the sense that it lays out an environment formulation that allows a wide variety of action types without requiring any structural changes to the problem setting. To incorporate any action within our framework, one only needs to model its effect on flood-related losses and carbon emissions and estimate the costs (and carbon emissions) of implementation and maintenance of the action.

Other adaptive methods, such as DPS, also typically formulate the decision problems to be specific to levee heightening-related settings, where the mapping between the states and the trigger actions are often established based on experience. Thus, it is not straightforward to systematically derive and justify the state mapping equations for other actions, such as nature-based solutions, withdrawal, and so forth. On the other hand, the effects of any flood protection asset on incoming flood levels are usually well established and can be readily incorporated within our framework. Finally, our suggested methodological formulation also allows the problem to be solved exactly based on various dynamic programming-based solvers. We have revised the following line in the manuscript to address the reviewer's comment

Lines 606-611: The developed framework is general in that it can incorporate a wide variety of possible action types, also considering nature-based infrastructure options along with grey infrastructure choices without making any structural changes to the problem setting. This work also provides globally optimal solutions while accommodating various other features, such as model uncertainty aspects and the social cost of carbon considerations in the analysis.

Lines 149-153: This again makes it seem as though the paper will incorporate multiple levers in the same example, which would (I believe) be a strong contribution, but this is not done.

We would like to thank the reviewer for acknowledging the limitations of the currently existing literature in consideration of a variety of decision-levers, especially in a control-based framework. We would like to further clarify that our work sets forth an original MDP/POMDP formulation in order to incorporate a wide variety of action types, including nature-based solutions along with grey infrastructure choices in the same case study, while retaining the Markovian properties of the framework. The framework also allows solutions with state-of-the-art MDP/POMDP solvers with guarantees that act as a foundation and benchmark that can be further scaled in the future. We have modified the following line in the manuscript to clarify that our framework can address a variety of action types without requiring any structural changes to the problem setting.

Lines 606-611: The developed framework is general in that it can incorporate a wide variety of possible action types, also considering nature-based infrastructure options along with grey infrastructure choices without making any structural changes to the problem setting. This work also provides globally optimal solutions while accommodating various other features, such as model uncertainty aspects and the social cost of carbon considerations in the analysis.

Line 290-296: The assumption that historical variability would be stationary in the future (despite possible changes to land-water storage or rapid ice sheet changes, which may be captured in a very limited way in the sampled climate projections) is actually quite a strong claim and deserves some justification.

Thank you for this comment. In this work, the historical variability is captured by fitting a polynomial regression model to past historical observations, ideally accounting for small-scale natural variability. We have modeled the historical variability to be stationary since the long-term non-stationary effects are assumed to be captured by the climate projections, as obtained from the IPCC models. It should be noted that the sea level rise model used in this work offers a higher fidelity than most existing standard approaches in the literature, which typically assume a storm surge model represented by a GEV distribution with a linearly increasing mean to account for the sea level rise process [2, 9, 21]. We would also like to highlight the work [8], which incorporates abrupt events through additional deeply uncertain parameters. This model can also be integrated into our framework, with probabilistic information collected from the model used to generate the sea level rise transitions, without any loss of generality. Finally, non-stationarity modeling effects should further help in identifying changing patterns in time, thus better informing optimal adaptive actions and increasing the benefits of sequential/adaptive decision-making. In response to this comment the following text has been added in the manuscript:

Lines 296-299: Without any loss of generality, any sea level rise transition model can also be similarly utilized in the presented framework, including those accounting for large-scale uncertain events like rapid ice sheet collapse as considered in [8].

Line 377 (and others): Is the improvement over the static baselines actually meaningful? The authors report decreased costs of several thousands to tens of thousands of US dollars (depending on the specific example) per meter or coastline, and it is not entirely clear what this means in the context of how stylized the problem is and the barriers to implementing adaptive solutions. Would these savings justify this approach? It might be better to calculate the overall savings for the case study (which could then be made into a rate), because right now the savings don't seem like much.

The case studies are conducted on a per meter basis of coastline, assuming a uniform distribution of city/community characteristics, such as value and building density from the shoreline to the highest elevation under study, similar to the work in [9]. The city/community is also presumed to expand uniformly across the width in this illustrative application, and strategies are implemented uniformly across the width parallel to the coastline. Consequently, losses and savings are reported in dollars per meter. For Manhattan, which has an overall coastline length of approximately 21 kilometers, the potential savings could be in the order of millions of dollars, assuming the distribution of losses remains uniform along the coastline. The shown case studies in the paper demonstrate in detail the application and implementation of our framework based on realistic settings and real-world numbers. It was not our intention to identify and showcase where this framework necessarily offers very significant financial benefits, although such cases can certainly be found and exist under different settings. In the presented applications, the need for immediate actions is dire, hence the small difference with the static baselines. However, from optimization and methodological perspectives, which are core contributions of this work, identifying globally optimal policies that surpass suboptimal ones by 2%-5% is a very difficult task, showing the power and effectiveness of our suggested approach. In addition, as we also mentioned before, our approach can serve as a baseline solution in comparison to other suggested policies, and a decision-maker can then quantify the sub-optimality of a plan and make informed decisions about it, as needed and is practically feasible. We have the following statements in the manuscript related to this:

Lines 533-543: As expected, the MDP-based policies thus lead to improvements over the considered static policies, albeit marginal to moderate ones for the analyzed settings. This is because the adaptive policies involve actions at the beginning of the considered horizon, or else an existing seawall already in place, which are also identified by the static policies. However, the fact that the static and dynamic policies differ by a single action since only one action was deemed optimal for the static policies in all cases, a non-trivial $2\% - 5\%$ life-cycle cost savings is already realized, that shows the potential of sequential decision-making and timely actions based on observed conditions. It is expected that with even higher fidelity and consideration of more available actions, a greater reduction in life cycle cost, relative to the static policy, will be seen.

Lines 621-624: Given this higher flexibility and fidelity, the identified adaptive policies may be expected to result in significantly higher cost savings in comparison to static cost-benefit analysis solutions, something we are aiming to explore in the future.

Lines 386-389: As mentioned above, the issue about the role of cost of carbon in e.g. incentivizing green infrastructure vs. grey infrastructure is an interesting one, and one which your framework is suitable to address. But it reads as somewhat of a throwaway in the current framing.

We thank the reviewer for recognizing this aspect of our study as a novel contribution to the community. While we have emphasized in the paper that our framework can adeptly incorporate future (and suitably discounted) carbon emission costs in decision-making, we do not consider this to be the most central contribution of our work, but rather, it is one of the significant features and that our framework, which has been suitably formulated to accommodate general actions while preserving the Markovian property, supports and can provide insights for. In future work, we aim to further investigate the value and role of the social cost of carbon within our framework and computed/learned policies.

I don't have MATLAB (and so did not replicate), but the workflow for reproducibility looks reasonable. While the README provides the steps for the reproducability workflow, it is not clear from the README how to use the code for a different problem, which is ok, but limits the code as a resource for the community beyond the reproduction.

We would like to thank the reviewer for examining the codebase. The workflow remains valid for any MDP/POMDP environment developed within the framework presented in this work. It aligns with typical MDP/POMDP environments, where the main functional modules include the state transition, observation and reward models. The state transitions and observation models are driven by the climate models, as explained in detail in the paper, and can be utilized as-is when employing the same climate models.

The reward model must encompass the costs of implementation, maintenance, and the flood protection impacts of an action. We have detailed the rewards for actions, such as floodwalls, and specific nature-based solutions, such as salt marshes and oyster reefs. The exact same methodology can be applied to design other flood protection actions within this framework.

We have structured the problem formulation to ensure compatibility with any off-the-shelf, state-of-the-art offline dynamic programming solvers (e.g., LRTDP, HSVI, RTDP, SARSOP). This approach is analogous to gym-based reinforcement learning environments, where general and diverse applications can be solved by suitably modifying the transitions and rewards function methods while maintaining compatibility with any reinforcement learning algorithm. Yet, based on the reviewer's comment, we will also further enhance the README file and instructions when we upload our code on GitHub.

References

[1] M. J. Lickley, N. Lin, and H. D. Jacoby, "Analysis of coastal protection under rising flood risk," *Climate Risk Management*, vol. 6, pp. 18–26, 2014.

[2] O. Špačková and D. Straub, "Long-term adaption decisions via fully and partially observable Markov decision processes," *Sustainable and Resilient Infrastructure*, vol. 2, no. 1, pp. 37–58, 2017.

[3] R. Hui, J. Herman, J. Lund, and K. Madani, "Adaptive water infrastructure planning for nonstationary hydrology," *Advances in Water Resources*, vol. 118, pp. 83–94, 2018.

[4] M. Pozzi, M. Memarzadeh, and K. Klima, "Hidden-model processes for adaptive management under uncertain climate change," *Journal of Infrastructure Systems*, vol. 23, no. 4, p. 04017022, 2017.

[5] "Implementation of Cost Benefit Analysis of Coastal Risk Management Prevention Measures against Natural Hazards," 2016, accessed: 2024-08-08. [Online]. Available: https://civil-protection-humanitarian-aid.ec.europa.eu/system/files/2017-11/ecoshaz_guidelines_1.pdf

[6] Federal Emergency Management Agency, "Guidelines for implementing executive order 11988, floodplain management, and executive order 13690, establishing a federal flood risk management standard and a process for further soliciting and considering stakeholder input," 2015, accessed: 2024-08-08. [Online]. Available: https://www.fema.gov/sites/default/files/documents/fema_implementing-guidelines-EO11988-13690_10082015.pdf

[7] ENWinfo, "Grondslagen en Richtlijnen voor Waterveiligheid," 2023, accessed: 2024-08-08. [Online]. Available: https://www.enwinfo.nl/publish/pages/183541/grondslagenen-lowresspread3-v_3.pdf

[8] R. L. Sriver, R. J. Lempert, P. Wikman-Svahn, and K. Keller, “Characterizing uncertain sea-level rise projections to support investment decisions,” *Plos One*, vol. 13, no. 2, p. e0190641, 2018.

[9] R. L. Ceres, C. E. Forest, and K. Keller, “Optimization of multiple storm surge risk mitigation strategies for an island city on a wedge,” *Environmental Modelling & Software*, vol. 119, pp. 341–353, 2019.

[10] Y. Miura, H. Qureshi, C. Ryoo, P. C. Dinenis, J. Li, K. T. Mandli, G. Deodatis, D. Bienstock, H. Lazrus, and R. Morss, “A methodological framework for determining an optimal coastal protection strategy against storm surges and sea level rise,” *Natural Hazards*, vol. 107, pp. 1821–1843, 2021.

[11] Ward, Philip J and Jongman, Brenden and Aerts, Jeroen CJH and Bates, Paul D and Botzen, Wouter JW and Diaz Loaiza, Andres and Hallegatte, Stephane and Kind, Jarl M and Kwadijk, Jaap and Scussolini, Paolo and others, “A global framework for future costs and benefits of river-flood protection in urban areas,” *Nature Climate Change*, vol. 7, no. 9, pp. 642–646, 2017.

[12] J. H. Kwakkel, “The exploratory modeling workbench: An open source toolkit for exploratory modeling, scenario discovery, and (multi-objective) robust decision making,” *Environmental Modelling & Software*, vol. 96, pp. 239–250, 2017.

[13] M. Gong, R. Lempert, A. Parker, L. A. Mayer, J. Fischbach, M. Sisco, Z. Mao, D. H. Krantz, and H. Kunreuther, “Testing the scenario hypothesis: An experimental comparison of scenarios and forecasts for decision support in a complex decision environment,” *Environmental Modelling & Software*, vol. 91, pp. 135–155, 2017.

[14] R. Lempert, “Scenarios that illuminate vulnerabilities and robust responses,” *Climatic Change*, vol. 117, pp. 627–646, 2013.

[15] Arias, Paola and Bellouin, Nicolas and Coppola, Erika and Jones, Richard and Krinner, Gerhard and Marotzke, Jochem and Naik, Vaishali and Palmer, Matthew and Plattner, G-K and Rogelj, Joeri and others, “Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change; Technical Summary,” 2021.

[16] “Station Home Page - NOAA Tides amp; Currents — tidesandcurrents.noaa.gov,” https://tidesandcurrents.noaa.gov/stationhome.html?id=8518750, accessed 08-08-2024.

[17] G. Garner, T. Hermans, R. Kopp, A. Slangen, T. Edwards, A. Levermann, S. Nowicki, M. Palmer, C. Smith, B. Fox-Kemper, H. Hewitt, C. Xiao, G. Aalgeirsdóttir, S. Drijfhout, N. Golledge, M. Hemer, G. Krinner, A. Mix, D. Notz, I. S. Nurhati, L. Ruiz, J.-B. Sallée, Y. Yu, L. Hua, T. Palmer, and B. Pearson, “IPCC AR6 Sea Level Projections,” Aug. 2021. [Online]. Available: https://podaac.jpl.nasa.gov/announcements/2021-08-09-Sea-level-projections-from-the-IPCC-6th-Assessment-Report

[18] M. Haasnoot, J. H. Kwakkel, W. E. Walker, and J. Ter Maat, “Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world,” *Global Environmental Change*, vol. 23, no. 2, pp. 485–498, 2013.

[19] T. D. Ramm, C. S. Watson, and C. J. White, "Strategic adaptation pathway planning to manage sea-level rise and changing coastal flood risk," *Environmental Science & Policy*, vol. 87, pp. 92–101, 2018.

[20] J. H. Kwakkel, M. Haasnoot, and W. E. Walker, "Comparing robust decision-making and dynamic adaptive policy pathways for model-based decision support under deep uncertainty," *Environmental Modelling & Software*, vol. 86, pp. 168–183, 2016.

[21] R. L. Ceres, C. E. Forest, and K. Keller, "Understanding the detectability of potential changes to the 100-year peak storm surge," *Climatic Change*, vol. 145, pp. 221–235, 2017.