ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
1	2	11	An Executive Summary (a few pages) upfront prior to the Introduction chapter would be useful. Terminology in climate-literature is crucial to avoid confusion and conflict. As such, it would be useful to have a Glossary with all key terms in the report defined (preferably from authoritative sources/literature).	We appreciate this comment and the utility of an Executive Summary and glossary. Although we are unable to add these at this time, we will keep these structural components in mind for future versions of the documentation.
2	3	2	The general level of the report is reasonable, and the graphics are generally clear and intelligible. The appendices could use more detail in places – especially on the representation of extreme behavior as a function of general warming level (see question 3). Sectoral calculations should better summarize the methodologies of the underlying literature.	We have attempted to provide additional information in Appendix A to address the reviewers concerns.
3	3	5	As noted in the response to question 3, the report lacks detail on the execution of the sectoral applications of the framework in some cases (wildfire, extreme precipitation).	We have attempted to provide additional information in Appendix A to address the reviewers concerns.
4	3	9	As noted in the response to question 3 – I am concerned that many of the methodological details for the sectoral cost analysis are not adequately described in the report itself, sometimes referring the reader to closed-access (Shane Underwood et al, Price et al,,) or in-review (Neidell et al, Neumann et al, Fant et al) publications and reporting only highly processed results of sectoral analyses.	As noted elsewhere in response to more detailed comments, all three previously "in review" papers are now published as open access. Updated citations and DOIs are now provided for those papers. EPA has made the Price et al. paper open access as well, here: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8477377/. The Underwood et al. paper was not sponsored by EPA and is unfortunately not available open access.
5	3	10	2) As noted in previous sections, the sectoral analysis could be better expanded upon in places. The working logic and implicit statistical assumptions in mapping mean changes to extreme distributions (impacting urban drainage and extreme temperature, labor and wildfire) are not always clear from the report itself – and referenced papers are not universally open access.	The urban drainage paper (Price et al.) are now available open access, here: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8477377/. The Underwood paper was not sponsred by EPA and so unfortunately we cannot make open access. EPA 2017a does provide some additional information about the Price et al. paper and methods, and is available at www.epa.gov/cira. Overall we have attempted to clarify where possible the intermediate steps of processing in Appendix A.
6	2	2	“11-year moving average” – justify this is appropriate for GMSL considering tidal epochs are 19 years.	The focus of the coastal sectors is on long-term mean changes in sea level (localized sea level projections are decadal) as opposed to variations in levels contained in the 19-year epoch. While sea level variations in the 19-year period are considered in the high tide flooding and traffic sector, it is modeled such that hourly anomalies from the 19-year mean are equally likely each year. In this way, the SLR impacts have the 19-year variations "smoothed out," i.e. each year contains the statistics this lunar nodal cycle.
7	2	3	“11-year windows” – this is another crucial issue that applies to the entire report: The choice of a 11-year window (why not 5? Why not 10?) needs to be adequately justified.	We have added a sensitivity analysis in Appendix B that shows the impact of short and longer window sizes on the arrival year. The analysis shows that the window size does not have a significant impact on the arrival year, ranging from 0 to 3 years.
8	2	4	(1) The flipping of axes in subfigure A starts the reader off with some confusion. Such a diagram needs some axes flipped, but since this is the most important information, it should be one of the un-flipped ones. (2) The lack of units for impacts is conf	The reference appears to be to Figure 1 rather than Figure 4. The original Figure 1 has been removed from the revised documentation, as several reviewers found it confusing
9	3	2	Better labeling required – what is the baseline, is it annual mean precipitation? What is the confidence?	Better labeling has been added.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
10	3	9	There are also some analytical assumptions in the general approach which are not fully tested, and the report would be strengthened by publishing sensitivity tests. The report notes (line 720) some limitations in the use of a single scenario to establish damages at different warming levels, but it does not demonstrate that the impacts of this assumption are trivial. The report uses only RCP8.5 simulations to provide the temperature reference, but it has been shown that this scenario will exhibit much larger land-sea warming differences than an aggressive mitigation scenario such as RCP2.6 (https://agupubs.onlinelibrary.wiley.com/doi/full/10.1002/2015GL063569) . This problem is noted in the report, with a defense of computational constraints (line 295) and that the problem is mitigated by using US temperatures as a reference (line 317). However, for many scenarios this will simply transfer the uncertainty to the mapping function from Global to US temperatures (line 331), which should itself be scenario dependent. As such, a sensitivity study assessing low warming impacts from RCP2.6 simulations would be a valuable addition to address this concern.	This is a good point. We had done some internal testing of these assumptions but those were not part of the documentation. We have added a new Appendix B that addresses this and other concerns. We were not able to compare RCP2.6, unfortunately, because the impact categories with strong precipitation-induced impacts. Instead we use RCP4.5 for the roads sector and find that although there are differences in the resulting impacts at integer degrees, there is not a clear bias (high or low) and differences are well within the differences across GCMs.
11	3	9	In addition, using US temperatures is not by itself sufficient to demonstrate that the results of impacts at different warming levels are scenario independent – when one would expect rather different synoptic conditions at different warming levels in different scenarios due to different radiative forcings and warming patterns, which impact mean precipitation (https://agupubs.onlinelibrary.wiley.com/doi/full/10.1002/2015GL065854) and cold event frequency (https://link.springer.com/article/10.1007/s00704-019-02850-8). The report would be greatly strengthened by showing the sensitivity of impacts to scenario choice at lower warming levels.	See response to Comment #10.
12	3	10	3)Some sensitivity analysis of key assumptions (notably the use of only the RCP8.5 scenario as a source of information for temperature threshold impacts) would be a valuable addition.	As mentioned, these are included in a new appendix to this report.
13	4	2	The term “global temperature change” needs to be some standard definition. Is it GMST? Or just land? Or something else?	True. It is GMST and figure caption text has been updated to mention this explicitly.
14	4	3	The basic analytical approach of the TBF is clear, but there are a number of supporting frameworks which are not explained to readers. (1)The TBF uses an 11-year averaging window for identifying GCM years. This is based on the challenge of internal variability in the GCMs, which is not discussed. However, it is frequently alluded to, with phrases like the “first arrival year” of 11-year averaged temperatures. This analytical choice should also be related to the standard practice in climate science of using 30-year averages, and the difference explained. It is also not clear if damages are to be averaged over 11-year periods, or what the analytical basis for that is. (Clearly some averaging is beneficial, so how should averaging decisions be made?)	This is a good point and we have added a sensitivity analysis in Appendix B that shows the impact of short and longer window sizes on the arrival year. The analysis shows that the window size does not have a significant impact on the arrival year, ranging from 0 to 3 years. The larger variations are for a shorter window. Increasing the size of the window has less of an effect. We did not try 30 years, however, since there are some relatively minor mathematical issues with larger windows when temperatures thresholds are crossed late in the simulation that ends in 2100.
15	4	6	Why are values adjusted to 2017? I thought the socioeconomic baseline was 2010.	This was confusing and text has been rewritten here. 2017 is the year of the base property values based on recorded data. In the model, values are adjusted to reflect appreciation over time dependent on GDP per capita.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
16	4	7	The use of 11-year averages is likely to raise eyebrows. Standard practice for defining smooth climate is 30-years, and this is because sub-decadal (e.g. ENSO) and multi-decadal (e.g. NAO) cycles will make shorter averages an unreliable estimate. Specifically, some models will produce 11-year averages that breach a temperature threshold, but then fall below this level.	See response to Comment #7.
17	4	7	On the other hand, the 11-year average for GMSL seems unnecessary. In the absence of variations, won't this just push the recorded year earlier than the actual year that GMSL reaches a level, because of the convex rise in sea-level?	This is true. The 11-year average has little effect and is probably unnecessary but we've kept it in for the sake of consistency with temperature and impacts, both of which use 11-year windows.
18	4	7	Since the input results include the effect of changing population, it is important to isolate the effect of climate change. Water quality will degrade over time even under a no-climate-change scenario.	You are right. And we do control for this but it was not included in the text. Added the following: "These climate change impacts are relative to a "control" scenario (one with socioeconomic growth and historical climate) to isolate the climate change impacts from the impacts of socioeconomic growth."
19	4	10	How are the 2010 and 2090 constant socioeconomics used to model time-varying socioeconomics? See comment in the main question response.	We have rewritten the section in the main report that describes how socioeconomic trends are incorporated into the Framework, and specifically how the 2010 and 2090 constant runs are used for this purpose.
20	4	10	Are these costs constant or time-varying?	The impact per delay hour per passenger do not change over time. Added as limitation in Appendix A.
21	5	3	There are two main issues on which I remain confused after carefully reading the full report and Appendix, which I therefore think could use some improved clarity: 1.I am very confused by the description of the SLR scenarios (lines 349-367 and Figure 4). I find Figure 4 quite mystifying – it seems to be analogous to Figure 3, but the rows are not different models and I'm not clear what the different colors represent or how they map on to the SLR amounts across the different rows. I am also quite unclear how figures like Figure 7 are produced for sectors that are a function of SLR – what does it mean to plot the same SLR change (e.g. 150cm) for multiple levels of warming? How is this mapping done? I think this just requires some more detail on the Sweet et al study that supports these SLR scenarios and more detail on how these are mapped to levels of warming.	Some text has been added to help clarify. We've also made significant differences to the SLR is handled in FrEDI, particularly the relationship between SLR and temperature.
22	5	7	-In the coastal property sector it is not clear whether coastal property values evolve based on socio-economics over the 21st century, or are fixed at the 2006 values mentioned in line 1644. If they are assumed not to change, that should be explicitly mentioned.	Text has been adjusted as it was confusing. We draw from data sources ranging from 2006 to 2017 for the property value dataset. In the model, values are adjusted to reflect appreciation over time, dependent on GDP per capita.
23	4	5	In two cases (Valley fever and Asphalt Roads), a new baseline estimate is constructed using the LOCA weather dataset (replacing PRISM and USHCN weather, respectively). However, this makes the baseline inconsistent with the rest of the results. The original baseline, modeled with the weather that the climate data is bias-corrected to, should be used, because any bias introduced by that not matching the TBF downscaling would be second-order since only differences from the baseline are being reported. On the other hand, subtracting off a LOCA-based baseline introduces a first-order bias. In the case of Asphalt Roads, this would also allow all 19 originally modeled climate models to be used.	It is a good point. Baseline inconsistency is occasionally an issue for non-CIRA studies where we are not able to rerun the analysis with the CIRA baseline. That was the case for Asphalt roads, a non-CIRA study (Underwood et al. 2017). Note that this study did not use the LOCA climate data--a separate set of bias corrected projections were used. Underwood et al. (2017) establishes a baseline commonly used in asphalt design so this alternative baseline has value for that particular study. Our approach, while not perfect, is an attempt to bring these impacts into the CIRA (and FrEDI) framework as much as possible as it is valuable comparison with the CIRA roads sector, in spite of the differences in baseline and projection. We have added a bullet on this under limitations and assumptions. Valley Fever, however, is a CIRA study. In this case, the model was trained with weather that matched the record of Valley Fever incidence, but the when we estimate the marginal impact of climate change, we use the CIRA baseline so it is consistent with the other sectors. We've added text to make this clearer.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
24	4	7	Why is it not possible in some cases to determine the difference in temperature between a study's baseline period and 1986-2005? And what does it mean that binning windows (is this the width of the temperature bins or the year periods?) are developed (how?) based on the available baseline (what is this baseline information?)?	This was confusing as written. Text has been adjusted with a reference to Appendix A, which explains the individual sectors in more detail.
25	1	3	A similar concern arises with respect to the authors' relationship between global temperature and GMSL (note 19); however, while it is possible to intuit what the authors did to derive the relationship between global temperature and CONUS temperature, even this is not possible for GMSL. The authors state that this is based on "data from Sweet et al. (2017)", but Sweet et al. (2017) did not present their own relationships between forcing and sea-level change, simply situate their GMSL scenarios in the context provided by the probabilistic, RCP-based projections of Kopp et al. (2014). Further, as the authors elsewhere note, there is no physical basis for a time-independent relationship between global temperature change and GMSL, which the authors specify (to a high degree of precision) in note 19; GMSL can be mapped to temperature change only at a specific point of time (and, even then, with substantial uncertainty). Arguably (e.g., Hermans et al 2021, 10.1029/2020GL092064), a time-independent relationship might be specifiable in terms of integrated temperature; but in any case, the authors provide no clear basis for the derivation of the relationship they are suing. Recommendation: The authors should clearly show in an Appendix the data they are using to calibrate the relationship between global temperature and GMSL and their fitted model. They should explicitly incorporate time-dependence and uncertainty.	We have provided an appendix in which we show the updated method of calculating global mean sea level rise from temperature using a semi-empirical method developed by Kopp et al., 2016.
26	3	3	The overall framework is clearly described in the main text, and nicely summarized in Figure 6. The provision of the R toolset is commendable with clear documentation, but the EPA may want to consider expanding the toolset to other languages to increase general usability (Python in particular).	Conversion of the tool to other languages would broaden the potential interest and usage of the tool. At the moment we do not have plans to recode into another language, however a Python wrapper could allow the user to call this R tool in Python.
27	3	3	The assumed interplay between assumed temperature pathways and sea level in projected scenarios is unclear in the report as it stands. An additional figure showing the relationships between temperature scenarios and SLR scenarios would be helpful	We have provided an appendix in which we show the updated method of calculating global mean sea level rise from temperature using a semi-empirical method developed by Kopp et al., 2016.
28	3	10	4)The relationship between assumed temperature pathways and SLR pathways in the aggregated cost estimates presented in Figures 3.6 and 3.7 is insufficiently explained.	This is addressed through the revised SLR binning and Global temperature to SLR methods.
29	4	4	(1) What are the units of GMSL here (cm?)? (2) This relationship suggests an immediate response of GMSL to changes in temperature. A more normal approach (DICE, FUND, and PAGE, I believe) would include a relaxation rate parameter, to incorporate a delay, which would describe changes in SLR, so that SLR effectively integrates over temperature.	We have provided an appendix in which we show the updated method of calculating global mean sea level rise from temperature using a semi-empirical method developed by Kopp et al., 2016.
30	4	4	The title of this figure is confusing.	We updated the title of this figure. "FIGURE C-1. GLOBAL MEAN RADIATIVE FORCING AND TEMPERATURE FOR THE GCAM REFERENCE SCENARIO"
31	4	5	Also, it is not clear why the Hector-GCAM scenarios are produced and used. While they do represent a custom scenarios, the same would be true of the RCPs. Nothing seems gained in constructing a new scenario in this context.	We have added the Hector-GCAM scenarios to the appendix as a case study. The tool is built to be able to run any scenario and not be bounded by the RCPs or SSPs. The initial analysis was done using RCP8.5 to estiblish the temperature to damages relationships.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
32	4	7	There is an inconsistency between how impacts temperature data is put into the system and how reduced complexity temperatures are used. For impact temperature data, CONUS temperature is related to impacts separately for each GCM, and then the impacts are averaged. For FaIR or Hector, GMST is used to produce an average CONUS temperature. So, one averages of Y and the other over X.	We have updated the text throughout and hope to have clarified this issue. Hector is used in the case study appendix to highlight how any scenario can be used within the tool. Emissions are passed to Hector, which calculates a global mean temperature. This global mean temperature is then converted to CONUS temperature within the tool.
33	4	7	Given that these alternative models are available, and either would be preferable, why are they not used? In particular, even if the full BRICK model cannot be embedded into the R tool, the Kopp et al. model should be possible.	We agree. We have since added in the semi-empirical sea level estimates from Kopp et al 2016. The text throughout the manuscript has been changed to reflect this as well as an appendix explaining the details of the semi-empirical sea level rise module.
34	4	10	Which information is drawn from each of these sources?	The sentence has been rearranged to reflect which information is coming from each source.
35	4	10	Is this a 90% reduction in total emissions, or just industry and fossil fuel emissions?	We have updated the text to reflect that the change in emissions is from industry and fossil fuel CO2 emissions.
36	4	General	It would be useful at this point to mention that these models are only used to produce an example trajectory, and that existing scenarios (i.e., the RCPs) may better conform to standard projections used in research.	We have updated the text, "This appendix includes additional information on the data and underlying models used in the example climate impact analyses presented in Section 3. Section C.1 provides more information on these climate scenarios and Section C.2 describes the socioeconomic assumptions."
37	5	4	I was very confused here about whether Hector, FaIR or the semi-empirical models were included as part of the tool. On the one had it says “The framework begins with an emissions trajectory” but then says the Framework “accepts global mean temperature”. If there is not functionality in the tool to convert emissions to temperature, I would recommend simply saying “users have to provide trajectories of either global or local temperature”	The reduced complexity climate models are not a part of the tool. We agree with the reviewer and have updated the documention thoughout.
38	1	3	Broadly, yes. However, it is not clear about how the relationships among climate variables (e.g., CONUS temperature, global mean surface temperature, and global mean sea level) are derived, and also – to the extent I am able to trace these relationships – does not appear to incorporate the (substantial) uncertainty that exists in them.	More details on each of these relationships is now provided in the appendix on Method Details.
39	1	3	The authors do not describe the derivation of their parametric relationships in footnotes 18 and 19 in a traceable or reproducible manner, and appear to assume perfect certainty.	Appendix D goes into detail on how the CONUS temperatures and global mean sea level are calculated from a global mean temperature change.
40	4	7	Are the authors using time-varying valuation approach? If so, this will conflate the effects of climate change and socioeconomics. In other words, even in the absence of a climate change effect, valued impacts will be higher in 2050 and 2090, and this would be incorrectly attributed to climate change.	A time varying approach for valuation is used, but only after the climate effect is estimated without reference to time. Time variation in valuation is re-introduced once the impacts by degree are identified, and then the two are combined to develop a custom impacts trajectory. The case study in new Appendix C should make this more clear.
41	4	7	As with 1558, the value of lost productivity increases across the century with socioeconomics.	The reviewer is correct, but this effect in the underlying study method is first controlled for, then time-varying valuation is only introduced once a new impacts trajectory is developed, as in the Appendix C case study.
42	4	9	Is it true that the hot cities have minimal extreme cold damages? How much do they have?	Yes some cities have no estimated extreme cold damage - cities where the maximum temperature for the 1 percentile coldest day is greater than or equal to 10°C (7 cities - 4 in South Florida, 2 in Southern California, and Phoenix) are excluded from analysis of cold exceedance.
43	4	9	What portion of the US population is contained in the 49 cities?	In the 2010 Census, the 49 cities accounted for 91.3 million of the total US population of 309.3 million, or nearly 30 percent. We added this information to the relevant Appendix text.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
44	1	1	The introductory chapter clearly explains the purpose of the report. However, there is an important omission regarding the context of the tool. Namely, tool is clearly an application of what is known in the pattern-scaling/climate-pattern emulation literature as the time-shift approach (e.g., Herger et al., 2015, 10.1002/2015GL063569; Tebaldi and Knutti, 2018, 10.1088/1748-9326/aabef2). This is used extensively in AR6 to present rules in terms of warming levels as well as scenarios. It is surprising to see this literature unreferenced. Recommendation: The authors should ensure their introduction is grounded in the peer-reviewed literature related to the time-shift approach for climate pattern emulation.	We have added several citations to make the literature review more comprehensive - Herger et al. was particularly relevant. However, we do note that few of the studies in this area take the output all the way to monetized impacts.
45	2	1	Martinich and Crimmins (2019) not in reference list	We added to the reference list.
46	2	1	Meinhausen 2011 should be appropriately referenced as Meinhausen et al., 2011	We made this change.
47	2	1	“Climate Impact Lab” – provide weblink or reference	We added a weblink in a new footnote
48	2	1	Houser et al. not in reference list	Reviewer is correct - need to add: We added to the reference list.
49	2	2	X-reference to Sweet et al. is inconsistent with other x-referencing.	We made this change.
50	2	2	Sweet et al. (2017) not in reference list	We added to the list
51	2	2	Replace 2015 with 2014.	We made this change.
52	2	2	Inconsistent x-referencing format	We made this change.
53	2	2	And also compare adaptation options/measures.	We edited the text to add this point.
54	2	2	Sweet et al. (2017) not in ref list	We added to the list
55	2	2	Bierwagen et al. 2010 not in reference list	We added to the reference list.
56	2	2	Chen et al. 2015 not in reference list	We added to the reference list.
57	2	2	Replace “O’Neal” with “O’Neill”	We made this change.
58	2	3	Replace Nadja, P. with “Popovich, N.”	We made this change.
59	2	3	Wobus et al. (2017b) not in reference list	The intended reference was to Wobus et al. (2017) - we made this change.
60	4	6	An extra “are”.	Thank you, we have fixed this error.
61	4	7	Missing a close parenthesis.	Thank you, we have fixed this error.
62	4	General	These pages are numbered as C-##.	Thank you, we have fixed this error.
63	1	2	Broadly, yes. The overall clarity of the presentation would benefit from a Figure 1 schematic that clearly shows the process of calibrating the Temperature Binning Framework to represent a particular sector. I believe the flow charts in Appendix A may try to do the former, but it would be good to have a broader representation upfront to help guide the discussion. Recommendation: The authors should incorporate a Figure 1 schematic.	Thank you for this helpful suggestion. We have added a new figure to the beginning of the methods chapter (Ch 2) that provides a simplified outline of the process. The figure focuses more on the full framework than the sectoral data processing as suggested in some comments. It is difficult to generalize the sectoral data processing in a flow chart due to the nuances in the incoming data sector by sector. We have described the general process narratively in the text of this new figure.
64	1	2	In addition, the authors frequently quote numbers with a number of degrees of precision totally unjustified by the underlying methods (e.g., Table 1, where they specify temperature changes to the nearest mK.) Recommendation: The authors should not use degrees of precision unsupported by the underlying methods. Temperatures, for example, should be specified in general to the nearest 0.1 K, though for trends the nearest 0.01 K might be justified.	We now show temperatures to 0.1 degree in Table 1 and have made several other edits to the precision presented throughout the report.
65	1	2	In describing the benefits of mitigation, the results would be more cleanly communicated if the mitigation scenario results, as well as the difference between baseline and mitigations scenario, were shown.Recommendation: The authors should add absolute values for the mitigation scenario in Figures 16-17, and consider rearranging Figures 11-17 so easy comparisons can be made between baseline and mitigation scenarios.	The case study, now presented as an appendix, includes total impacts for all scenarios prior to showing the benefits achieved in each emission reduction scenario (we now include multiple emission reduction scenarios to de-emphasize the 90% reduction scenario).

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
66	1	3	With respect to the relationship between CONUS and global temperature, the authors employ a one-to-one mapping (specified, absurdly, to 5 significant figures) between global mean temperature and CONUS temperature, even though large ensemble results clearly show that there is substantial variability in the relationship between global temperature and CONUS temperature, even within a single model (consider Figure 5 from Kay et al., 2015); further structural uncertainty is introduced when considering projections across models. Recommendation: The authors should, in an appendix, clearly show the relationships between CONUS and global temperature over time from multiple GCMs they are using to calibrate the relationship in note 18, and should explicitly incorporate uncertainty into that relationship. Recommendation: Given that the relationship between global and CONUS temperatures varies across model realizations, the authors should unambiguously document the GCMs they are using, including model replicates.	Thank you, we have added a new appendix that includes the temperature translation function derivations as well as the global temperature to sea level rise derivation. The ability to evaluate uncertainty in this relationship is not currently built into the Framework but is something we will consider adding in future versions.
67	1	4	As noted above, the report would substantially different from a Figure 1 flowchart showing – at a broader level of generality than the sector-specific flow charts in appendix A – the entire process of calibration and application. Recommendation: The authors should incorporate a Figure 1 schematic.	We have added a new Figure 1 that provides a roadmap for the framework. Please see the response to comment #63.
69	1	11	The layout is, frankly, not very attractive and would benefit from attention. The report looks like it makes heavy use of Excel defaults in the figures, and has quirks like writing out “DeltaT” instead of a more typical and attractive “ΔT”. Figures would be easier to read if their captions appeared below rather than above figures. The header is applied irregularly on different pages. Some figures (e.g., Figure 3, 5) appear to have no caption. Table 6 is split unnecessarily across pages. Figures in the Appendix split across pages are hard to read.	We have made an effort to revise the Figures and Tables in this final version. Specifically we have standardized many of the Figure outputs direct from the model so they are cleaner and more consistent (particularly in Appendix A); we have added captions to all tables (above the table) and figures (below the figure); we have changed the header structure from print format (where certain headers only appeared on odd pages) to an electronic editions (with headers appearing on all pages).
70	2	1	Use acronym for Temperature Binning Framework and use acronym through the rest of the report	In this version we have moved from the title of "Temperature Binning Framework" to the acronym "FrEDI"--the Framework for Evaluating Damages and Impacts. We sometimes refer to "the Framework" and in a limited set of references specific to the implementation of the Framework in R, we use "the Tool".
71	2	1	“Sectors” should be clearly defined here, with examples, at first use	We have added examples to clarify how we use the term "sectors" in this report.
72	2	1	Typo in “an overview how”	Thank you, this typo has been corrected.
73	2	2	“impacts per capita” should be defined and explained clearly.	We have clarified that this is intended to be a generalized example and therefore the "impacts per capita" measure could be thought of as any sectoral impact.
74	2	2	It appears as though the “arrival windows” are centered within each 11-year bin: clarify this in text.	This is a correct interpretation. We have edited this caption to make this point clear.
75	2	2	“PM 2.5”, “VSL” should be defined and explained	We have spelled out these acronyms in the table and pointed the reader to Appendix A for more details on the
76	2	2	“Non-CIRA study” – Is it only that CIRA studies are considered credible for the TBF? How about peer-reviewed literature outside of CIRA? Clarify why “non-CIRA” needs to be highlighted.	We have expanded the note in this this table to further clarify what non-CIRA studies mean and why we make that distinction. All studies in the Framework are considered credible, however the distinction is made to highlight that while the underlying CIRA studies all use consistent assumptions about socioeconomics and use the same set of GCMs (and same downscaled dataset), non-CIRA studies often do not follow the same assumptions.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
77	2	2	Even when some sort of normalization can be performed per capita? Explain.	We have expanded the section on socioeconomic scalars considerably in this revision, partially to respond to this question. Differences between impacts in 2010 and 2090 are driven, in the first-order, by factors such as total population and GDP). In this case, as suggested in the comment, which can be calculated and controlled for either on the front end of the estimation process, or on the back end in results presentation. For several sectors, there are second-order effects in play that further differentiates impacts over time--factors such as changing population distributions or energy demands per capita. These are accounted for in the year-specific adjustment factors.
78	2	2	Clarify why not, including specifically for urban drainage especially as precipitation is used in the framework directly tied to temperature degrees - is the underlying model/data not available?	Correct, FrEDI does not include differences in impacts for Urban Drainage, Wind Damage, and parts of the Wildfire analysis due to the absense of the linkage in the underlying studies. We have clarified this in the text.
79	2	2	Clarify what “states of the world” mean	We have clarified in the text that by states of the world, we mean options for future societal responses to climate change.
80	2	3	Clarify how the interpolation can be impacted by potential non-linearities/accelerations/decelerations (e.g. GMSL) as opposed to a constrained linearity assumption.	In terms of temperature-driven sectors, we have added further discussion on this point to the uncertainties section and the methods sections, as well as developed a test of non-linearities in these scaling factors in a new appendix. For SLR-driven sectors, the new method no longer requires these factors and therefore non-linearities are incorporated in the annual estimates.
81	2	4	Explain what the negative numbers mean. If these are not adverse impacts, and benefits instead, explain.	Thank you, we have clarified that these numbers represent a reduction from the baseline.
82	2	6	This statement appears to be a bit difficult to comprehend as there clearly are socioeconomic influences to Urban Drainage. If this is an assumption, justify and if a conclusion, state how it was reached.	The underlying studies used for these two sectors (Asphalt Roads and Urban Drainage) do not model the impact of changing socioeconomic conditions. The text now clarifies that this is specific to the underlying studies and not a general assertion.
83	2	7	This should be explained in more detail for each sector, impact, model etc. as to how exactly "aggregation", "disaggregation" and "summing" are performed, whether or not any weighting is performed etc.	We have clarified this in the text. The aggregation is done via simple summing except in cases where the original unit of analysis does not fall within one NCA region (e.g. HUC8-based analyses). For the latter cases, we use spatial weighting to assign results to the NCA region. To date, no cases use dissaggregation (e.g., National results that need to be downscaled to regional results). We have therefore removed this discussion from the text.
84	2	7	The “separation” and “reassembly” needs to be explained in more detail as the two – climate effects and socioeconomic trends - are inter-twined for a variety of sectors and impacts.	We have restructured Chapter 2 to simplify the presentation of the way population, GDP, and other socioeconomic-related trends are incorporated in the Framework. Discussion about the Framework's limited ability to account for feedbacks between socioeconomics and climate is included in the limitations section.
85	2	7	Not clear why socio-economic driver for Urban Drainage is “none” when clearly there is in reality. Is it that no authoritative literature exists from which sectoral impacts data can be used? Clarify.	The socioeconmics listed in Table 2 represent the drivers included in the model from the underlying study. The underlying study for urban drainage (a previously completed CIRA sectoral study) does not include socioeconomic drivers as they relate to changes over time.
86	2	10	Include “analysis” underlying the equation in an appendix to this report.	We have added this information in as an appendix to the report.
87	2	10	Provide further details as to how this – non-integer warming - is addressed compared to the integer approach described earlier, including any interpolation, arrival times, bin centering etc.	We have clarified that impacts for non-integer degrees of warming are interpolated between known integer-based results.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
88	3	1	The introduction makes the point well that the temperature-binned framing avoids some of the more problematic communication issues (climate sensitivity uncertainty, complex scenario definitions), but this rather undermined by the fact that the report ultimately reverts to a scenario framing to present the key results and the temperature binning approach primarily as an impact emulation tool in the headline applied results of the report (Figures 11-17). This is something of a missed opportunity – creating a large dependency on one scenario, and one idealized mitigation pathway in the results.	Thank you for this feedback. We have moved the exploration of damages associated with particular temperature trajectories to an appendix. The intention behind this example is to show how the Framework can be used to estimate impacts for any climate trajectory, however we do not intend to endorse any particular scenario. Also to this point, we have added additional emissions reductions scenarios to the appendix to reduce focus on the 90% emissions reduction scenario.
89	3	1	The report also confuses the primary purpose (the documentation of the temperature binning approach), which is quite adequate, with a secondary purpose which is a case study of a mitigation cost-benefit assessment, which is in itself incomplete – omitting major sectors (agriculture, biodiversity, insect-borne disease transmission, immigration pressure) and using very idealized mitigation pathway assumptions. My concern is therefore that the applied mitigation cost assessment will be cited as a major outcome of this report – and this aspect of the report is incomplete, and liable to significantly underestimate the financial impact of mitigation. This concern is somewhat supported by some of the public comments to the report, which interpreted the aggregate and sectoral results as the primary products of the report, rather than as case studies.	We have moved the case study to an appendix to de-emphasize impact results for this particular scenario (see response to comment #88). We have also added caveats to the new appendix regarding the aggregation of sectoral impacts. FrEDI does not provide a complete accounting of the benefits of mitigation, however it does provide an informative partial accounting in that impacts are presented by sector.
90	3	1	The lack of formal scenarios also removes dependencies between mitigation action and infrastructure development. A world which mitigated to 1.5C would not have the same balance of road and rail infrastructure as an unmitigated world, for example, but such complexities are explicitly ignored in this analysis. My recommendation is that the report is more tightly focused on the core question – impacts at temperature levels, and that the national level cost-benefit impact assessment is left to a future report where the activity can be conducted more completely and comprehensively.	We have moved the case study to an appendix (see response to comment #88). In response to the feedbacks between road and rail infrastructure under various mitigation scenarios, we include a note in the uncertainy section that these feedbacks are currently not modeled in FrEDI.
91	3	2	This figure is a little confusing – the table should make more clear that the rows correspond to different long term sea level rise scenarios.	This table no longer appears in the report now that the SLR binning method has been revised.
92	3	3	It would be useful to at least summarize the details of the intermediate working for each sectoral impact. At present, it is often unclear whether the results arise from a GCM or from a metamodel (e.g. for wildfire occurrence). Statistical assumptions used in the mapping of temperature and precipitation extremes are first order important, but are not clear in the current text.	The Framework itself uses no statistical mapping of temperature and precipitation extremes but rather reflects the variation in temperature and preciptation extremes implicit in the GCM results. Among the underlying studies, to our knowledge, and with the exception of the newly added inland flooding results, there is no statistical mapping of temperature or preciptiation either - for inland flooding that process is explained in the revised Appendix A. All underlying sectoral studies estimated impacts using specific temporal and spatial resolution outputs from GCM results, as processed by LOCA for bias-correction and spatial downsclaing, and as appropriate for the sector. The wildfire study does use some intermediate processing of temperature and precipitation, which is complex but is described in the supplementary material for that ERL published study - for space reasons it is not included here, but is referenced in Appendix A.
93	3	4	The text is clearly written throughout. Efforts should be made throughout the report to label and define axes and units in Figures, which are in places ambiguous.	Thank you, we have updated the graphics to have more consistent formatting and labeling.
94	3	4	Units and/or labels missing on plots	Thank you, we have updated the graphics to have more consistent formatting and labeling.
95	3	4	Labels are too small, axes are unlabeled	Thank you, we have updated the graphics to have more consistent formatting and labeling.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
96	3	4	Explicitly write “climate sensitivity”	We have clarified the figure title to explain this graphic shows six climate sensitivities.
97	3	7	Finally, in the case study presented in section 3 - the aggregate economic impact is calculated with an incomplete list of potential climate impacts – notably excluding the effects of biodiversity loss, agricultural impacts, a comprehensive assessment of disease transmission, immigration pressure, water availability, direct effects of temperature on renewable and non-renewable energy capacity and conflict. As such, the subset of impacts considered here is insufficient to be considered as an aggregate cost of climate impacts. As recommended in the general conclusions, there is an argument to move this section to the appendix to avoid confusing the messaging of the report. I would recommend a focus on individual sectors, rather than on the aggregate economic impacts presented in section 3.	Thank you. Please see response to comment #89.
98	3	9	Having made a convincing case for the benefits of temperature framing at the start of the report, this is potentially undermined by reverting to a scenario framing at the end of the report. The treatment of the benefits of mitigation in sections 3.6 and 3.7 is intended as a case study, but risks being interpreted generally. There is no guidance as to whether the 90 percent reduction in emissions relative to the reference scenario is ambitious or not, nor is there sufficient guidance on the parameters and assumptions of the reference scenario (notably, the emissions reduction considered is significantly short of a Paris-compatible scenario). As a result, the estimated cost benefits of climate mitigation are largely arbitrary, which is a problem if the case study is interpreted as a product of the report.	Thank you. Please see response to comment #88.
99	3	9	I would recommend either expanding the range of reference and mitigation scenarios considered in Section 3 to make clear the scenario-based assumptions of the computed costs, or (perhaps preferably for a technical report) moving this section to the appendix and focusing on the impacts at different warming levels in the main text. The warming levels could usefully be put into a general emissions context through the use of global mean temperature RCP pathways or cumulative emissions relationships (e.g. IPCC AR5 SPM Figure SPM5).	Thank you. Please see response to comment #88. We have added additional mitigation scenarios to the case study in response to this comment.
100	3	11	Section 3 is a case study application of the binning methodology to an example scenario. As acknowledged in the report (line 831), it is not intended as a general assessment of future climate costs or mitigation benefits: only one arbitrary reference and mitigation scenario are considered, and only a subset of sectors are considered. However, there is a risk that this case study may be interpreted more generally than intended – and this is underlined by some the public responses to the report – even from expert reviewers, which focused largely on the choice of sectors considered, rather than the binning methodology.	Thank you. Please see response to comment #89.
101	3	11	Given this, there may be some logic in avoiding this confusion by making the case study itself an appendix. This would allow the key message of the report to be unambiguously the communication of best estimates of climate impacts as a time-independent function of temperature, which should be the focus of the report.	Thank you. Please see response to comment #89.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
102	4	2	In general, the writing and graphics are well-designed for this audience. However, the report frequently introduces terms that are either not standard, or would not be known by a wide audience, without defining them. I have some cases of this listed in the specific comments below. For example, the term “era” is commonly used, but only explained through examples (e.g., “a baseline era (e.g., 1986-2005)”). Upon the first introduction of the term, it should be explained: “Within this report, the term ‘era’ is used for a range of years across which temperatures, impacts, or socioeconomics are averaged.” I discuss uncertainty in question 8, but I would also note here that uncertainty means different things to different semi-technical audiences. A clear definition is needed.	Thank you, we have made an attempt to better define these terms upon first usage througout the text.
103	4	2	The figures and tables are generally good, but one significant issue is the pervasive use of total dollar losses. These will not be intuitive to many readers, and certainly not useful when described for 2090 socioeconomics. Where possible, I suggest reporting percent losses or dollars per capita loss.	consider for future. not what EPA Guidelines for Economic Analyses say
104	4	2	(1) The flipping of axes in subfigure A starts the reader off with some confusion. Such a diagram needs some axes flipped, but since this is the most important information, it should be one of the un-flipped ones. (2) The lack of units for impacts is confusing. (3) The y-coordinates of subfigure C are not made clear graphically. Either add a fourth panel with a diagonal line for mirroring A, or drop the L-organization and just show three graphs as normal.	removing figure
105	4	2	Is this an incomplete sentence or a run-on? The framework “assigns” what “to” what? Neither of the “to” words used seems to produce any grammatical interpretation.	Thank you, this sentence as been re-written to clarify that the Framework draws on the regional and temporal variation in climate outcomes from the underlying GCMs but cannot run custom regional climate scenarios at the region level.
106	4	2	The use of a “baseline year” for temperature and GMSL is not a standard concept. Explicitly say that temperatures are reported as anomalies, or as differences from a given year.	Thank you, we have clarified this terminology in the report.
107	4	2	The term “adjusted” here is not clear, nor how the physical and economic scalars are determined.	We have simplified this explanation. Further details on scalars used can be found in the main text.
108	4	2	Is a “model” here a GCM or an impact model? Is an “impact type” here an impact or a group of impacts? When interpolation is done between impact years, is the result subannual (since previously we are told that all values already interpolated to annual timeseries)?	"Models" refers to GCMs in this context. Impact types are the smallest groupings of impacts available from the underlying studies within each sector (e.g. mortality and morbidity estimates for health sectors, or ozone mortality versus PM2.5 mortality for the Air Quality sector). Impact interpolation is done at the annual level, where interpolation occurs between the 2010 and 2090 impact years. We have clarified this in the text.
109	4	3	(2) It is assumed that socioeconomics influences damages, but not explained how the readers or the TBF understand that relationship. Socioeconomic growth simultaneously increases the financial and population exposed to risks while increasing their capacity to adapt. I believe that only the first of these effects is incorporated (since adaptation is described in static scenarios). Most socioeconomic change occurs through growth processes, but linear interpolations are used. What is the analytical justification for this? Within economics, most socioeconomic relationships are described through elasticities and social welfare functions. These concepts are not introduced or used, and this puts the analytical foundation of the TBF on uncertain footing.	Your interpretation is correct that the current Framework only accounts for changing exposure due to socioeconomic changes, not changing adaptive capacity. We have added a discussion of this limitation to the text.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
110	4	3	(3) Appendix A spends an appropriate amount of time discussing the process of isolating the impacts of climate change from the results (except for cases I mention below), but the whole concept is missing from the analytical framework. Since this is a framework for taking pre-existing results, it is important to help readers understand how to think about the climate change component of those results. In some cases, this is easy and can be done by just subtracting off the outcome of a baseline period. In other cases, there may be a non-stationary effect that influences future years, like the rise of income causing electricity demand to increase. Things can get more complicated if climate change and socioeconomics interact, so that neither alone produces the impacts in the baseline period. A clear definition and process description should be provided for how climate change impacts can be isolated from study results.	The reviewer raises a good point about the complications of non-stationary effects and the potential for climate change and socioeconomics to interact. To address the comments, we added clarifying language in two locations in the document. First, we expanded the paragraph at the end of section 2.3 about considerations in identifying and adjusting for baseline effects when incorporating studies in the framework. In the case of electricity supply & demand, as the reviewer mentions, the rising of income causes electricity demand to increase. The outputs of FrEDI are designed to incorporate the change in demand such that the results for, for example, 2 degrees warming in 2050, reflect both the impact of 2 degrees of warming AND the demand expected given GDP/capita in 2050. A second scenario that reaches 2 degrees in 2070 will have different results because of the change in demand. See Table 5 for an example of how this is incorporated in FrEDI. Second, we added specific language in a new Table 7 discussing uncertainties in the framework and their estimated influence on results. The entry for the row on "Socioeconomic and demographic change over time" directly addresses the reviewers comments and acknowledges that the Framework is limited in abilty to address these complications in cases where the underlying work may not provide a means to isolate the climate impact from baseline effects and processes.
111	4	3	I do not think this sentence supports any understanding of the socioeconomic scaling factors (or whatever it is meant to address). What is the time dimension transformed to?	We have clarified the language in this section to better explain the sectoral data pre-processing steps.
112	4	3	Only results for these six models are displayed, and results for other GCMs from studies used are ignored. It would be useful to understand the expected plan to incorporate more GCMs, or if that will not be a priority.	We have added a footnote to clarify that the Figure only shows the GCMs used in CIRA studies, however the same process is employed for the GCMs used in non-CIRA studies processed for FrEDI.
113	4	4	In multiple cases, the units of the table are only reported in the caption, presenting a barrier to interpretation and use. For example, this is the case with Table 7 and Figure 9.	Thank you for this feedback we have added additional labels and notations with the graphics and tables.
114	4	4	The sets of Figures A and B in the sector appendices can be improved. In almost all cases, the 2010 and 2090 socioeconomics results are exact duplicates, with a different y-axis. This relationship would be clearer if either one set of graphs is shown with two y-axes, or if the unscaled graph is shown and the scaling is just reported for 2010 and 2090. If this is done, a second set of figures (new Figures B) could show a time-varying socioeconomic result. A baseline scenario (2010) and a time-varying scenario (SSP2?) would be much more informative than a separate constant-2090 scenario.	We have updated all graphics in Appendix A for more consistency and readability. We only show impacts by degree, not time-varying scenarios to keep the focus on the pre-processed data in the Framework for each sector, which is the focus of Appendix A. Time-varying results are shown in the case study appendix.
115	4	4	In many cases, the damage functions for MSLR impacts are reported along an x-axis that reads “Degrees of Warming” (Figure 7, A-13, A-14). I was under the impression that MSLR impacts were binned according to 25 cm MSLR ranges. Where are these temperatures coming from, and how are they used?	SLR-driven impacts are now shown against cm of GMSLR with the new SLR binning method. Previously we showed them by temperature for easier comparison to other sectors (using the GMSL and global temperature conversion method described elsewhere in the report).
116	4	4	The graphs are fairly difficult to read. One suggestion would be to reduce the horizontal and vertical whitespace, and to report y-axis labels in billions. Also, why are SLR impacts displayed against degrees of warming?	We have regenerated these graphics to be more readable (including adding grids and reporting impacts in $billions). We also now show SLR-driven sectors in relation to cm of GMSL. We initially showed impacts by degree for easy comparison across all sectors but agree the revised presentation is clearer.
117	4	4	Why are the seven impacts on the left given more space than the nine on the right?	This was done to emphasize the difference in y-axis scale for the two plots. We have added more explanation of what the graph is showing in the caption to further clarify.
118	4	4	What are the units of premature mortality? Individuals? And are these results annual?	Premature mortality is measured as the number of deaths. We have clarified this in the table text. The results are annual, as stated in the title but we have added an additional reference to this in the caption.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
119	4	4	Do not include GCMs in legend that are not included in the analysis. I initially assumed that these other GCMs must somehow be perfectly overlaid with the displayed GCMs.	We generate graphics for all sectors together, as an output of the tool and the processing script automatically outputs the same legend in order to keep consistent formatting by GCM across sectors. In the future we will consider generating custom legends by sector. In the meantime we have noted which GCMs are not included in each sector.
120	4	5	First, although the report explains the principles of how sectors can be included, it does not explain what the authors of non-CIRA studies should do. Should they be emailing the TBF team with data files structured a certain way? What is the quality-management process and what is the bar for validation? How will overlapping estimates (double-counting) and different modeling assumptions (model uncertainty) be handled as the collection of impacts grows? Without this information, the crucial potential of the TBF to act as a framework for coordination and communication will be underdeveloped.	We have added a discussion in Section 2.2 on the process by which EPA monitors the literature for additional studies for inclusion. For the sectors included thus far, we have been able to work with a wide variety of incoming data formats, so there is not much that needs to be done on the front end, besides meeting the criteria described in Section 2. The case of Asphalt Roads and the CIRA Roads sector provide an example of how overlapping estimates will be treated. We utilize the "primary" sector run designation to develop summaries, etc. that do not include overlaps. A similar method can also be used for runs from the same model with various assumptions. This is the case with the Air Quality sector where we currently include two runs that contain different assumptions about emissions levels.
121	4	5	Below are some concerns with how the existing set of studies is incorporated. According to appendix A, results for sectors are “adjusted” to isolate climate change impacts by subtracting off the baseline damages. This approach assumes that damages are produced by a separably additive process. In particular, it would often be reasonable instead to assume that climate change amplifies risks, rather than adds on new ones. Then, you would want to compute something like Damage / Baseline – 1, rather than Damage - Baseline. This distinction matters in any case where socioeconomics scales the results, because it changes the order of multiplications and subtractions.	We agree that the details of adjusting for baseline damages can be complex and sector specific - this concern occupied great attention in the pre-processing of sectoral results, as well as in the underlying sectoral studies, most of which were sponsored by EPA. While it is difficult to generalize, we reviewed the processing to confirm that each of the relevant sector studies, some of which employ a "no climate-change" baseline trajectory rather than a single baseline period estimate, employed best practices to estimate only climate change attributable impacts, which are in turn reflected in the results of the Framework.
122	4	5	For two of the sectors (Asphalt Roads and Urban Drainage), the costs are calculated relative to an unquantified adaptation scenario (Asphalt Roads damages are relative to the damages under proactive adaptation; Urban Drainage is relative to no-adaptation). While the exclusion of unquantified damage channels is common, it should be made more clear (1) that there are climate damages from the alternative adaptation scenarios that are being excluded, and (2) when each of these channels is included. On (2), I presume that Asphalt Roads should only be included in the no-adaptation and reactive-adaptation setups, and that Urban Drainage should only be included in the proactive-adaptation setup.	Move some text from uncertainty to adaptation about default being best estimate, not no adaptation
123	4	5	This “loss” is actually the shutting down of a transfer from consumers to purveyors of winter activities. There is no actual loss to society.	The valuation measured used in the underlying winter recreation study is lost revenues. We have added discussion to the case study application appendix on the caveats of comparing this type of loss to the welfare losses more commonly used in the sector studies. Other recreation loss studies measure lost welfare value, as consumers move to their second-choice activity during ski area shut downs. EPA is currently working to add new recreation impacts (as measures of lost welfare) to the CIRA study set.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
124	4	6	I think there are ways to improve both the description of the use of socioeconomic data and how it conforms to economic principles. The way that socioeconomic drivers are described, through a “scaling” process, is unnecessarily confusing for simple cases. That is, in the simple case of per-capita impacts and linear valuation, the very general terms used for the scaling approach added more confusion than elucidation. It was not clear to me through most of the report that “scaling” in these simple cases just meant that impacts are described in per-capita physical units.	Thank you--we have restructured Chapter 2 to simplify the presentation of the way population, GDP, and other socioeconomic-related trends are incorporated in the framework. Specific to this point, we have pulled out the discussion of per-capita physical units and valuation seperate from the remaining adjustments made to account for more complex relationships between impacts and changing socioeconomic conditions.
125	4	6	The entire socioeconomic scaling system seems to assume an elasticity of 1, and specifically that the degree of adaptation is not sensitive to the future wealth of the United States. This is clearly incorrect. As some examples: (1) The effect of extreme temperatures on urban mortality will have an elasticity with population below 1, since denser areas will be more able to protect their people; (2) Impacts due to Air Quality, Extreme Temperatures, and Lost Labor all decrease with income, as people are more able to protect themselves, so that even results reported in dollar terms would have an elasticity below 1.	See response to comment #109.
126	4	6	The way that population numbers are used across sectors is not always consistent, and no explanation is given for these differences. Mortality from air quality and from extreme temperatures are both normalized by population from ICLUSv2, but in the former case total regional population is used (because other datasets might not have ages), while in the latter city populations are used, with conversion factors. Why can the latter approach not be used in the former sector? For labor, the estimates are first divided by high-risk worker population and then multiplied by total regional population. That sounds like it assumes that all workers are high-risk.	We have attempted to clarify this in Table 4. Sectoral Impacts Linked to Custom Socioeconoimc Scenarios, and the surrounding text. Not all of the underlying scenarios study the full population for each region. For example, the Extreme Temperature study only looks at populations living in cities. We still want to index these values to regional populations so that they can scale with the user input regional populations, hence the conversion factors. Other underlying sector studies, like Air quality, estimate impacts across all population and therefore do not require a conversion factor. For the labor sector, the number of high-risk workers is held constant over the century (following the underlying study) but it is only a portion of total population and total workers.
127	4	6	In some cases, it is mentioned that a linear interpolation is used for socioeconomics between years/eras (like 2010 and 2090). For population, this is a plausible assumption when the years used in the interpolation are 20-or-fewer years apart, but not more than that. For income growth, it is never a reasonable assumption, since income grows exponentially. It can be interpolated using a linear-in-logs approach for years 20-or-fewer years apart.	Please see the new analyses presented in Appendix E. Here we show the differences in various methods of interpolation. Also, for sectors with direct links to GDP/capita we are able to introduce elasticities and other non-linearities as appropriate.
128	4	6	The report does not distinguish between the scalars being used in the normalization process (line 467) and the re-scaling process (line 469). Presumably one is the inverse of the other.	Yes, this is the correct interpretation. We have clarified that in the text.
129	4	6	What does it mean to “separate” multipliers “outside” of a model? Based on the examples given in Table 5, I believe that the authors mean that sector-specific time-series are used. These appear to be extracted from the sector-specific models (so, certainly separated to be outside of those models).	See response to comment #109. We have rewritten this section and no longer use this phrasing. Your interpreation of the process is correct.
130	4	7	How are impacts for non-integer temperatures estimated? Linear interpolation is only discussed in the context of “era”-based impact studies, but I suspect this is the intended approach.	Yes, this assumption is correct. We have added to this discussion to clarify this process.
131	4	7	The approach here is certainly not using the best available data and methods. The best available method would be to apply the methods of the underlying studies directly to new scenarios.	Thank you, we've clarified that this is not the best available data and methods, but rather a reliable method for rapidly developing estimates.
132	4	7	The socioeconomic scenario is not specified.	The case study application, now in an appendix, uses the default socioeconomics scenarios in the tool (i.e. the ICLUSv2 population projections and the CIRA GDP projection, based on EPPA v6).

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
133	4	7	Does the wage rate change according to economic scenario? Wages should rise with productivity, or at least at the rate of GDP per capita.	Yes, wages, and therefore the labor impacts, scale with GDP per capita. This has been clarified in Section 2 of the main report.
134	4	7	Why are the miles per region divided out and then immediately multiplied in?	Although miles per region are not modeled to change over time in the underlying study, modeling impacts per mile instead of total impacts makes it easier to compare across regions. This could be done in pre-processing or in post-processing, but we made the decision to calculate per-mile impacts in the tool and re-scale on the back end.
135	4	7	Is there an interaction between population and warming in determining damages? For example, would 2 °C warming on the current population result in damages equal to 2 °C warming on a larger, future population, times the increase in population? This is how the sector is modeled, but given the potential for capacity expansion, it seems plausible that fully modeled impacts would be less or more than the linearly scaled estimate.	The scaling between Electricity Transmission impacts and temperature is not a linear function of population. Therefore, we use a year-specific adjustment factor, the "Electricity demand growth factor" to account for the difference in impact per capita over time. This demand growth factor is derived from the underlying study. The section in Chapter 2 pertaining to these adjustment factors has been udpated for clarity.
136	4	8	The default socioeconomic scenario appears to be constant 2010 socioeconomics (although this isn't specified in the projection results discussions starting in section 3.6). Socioeconomic scenarios are both highly uncertain, and constant 2010 values are (obviously) implausible. That does not mean that they are not a sensible communications tool, but much more needs to be done to help readers know "how results should be interpreted and used". (1) Explain that these losses are not what would be experienced, but rather meant to provide a relatable set of numbers. (2) Report the numbers as %GDP, in addition to any dollar terms. (3) Show how the projected values change for projected socioeconomics, both in constant $ and %GDP terms.	The case studies use projected socioeconomic scenarios, not constant 2010 values. We have clarified that in the report.
137	4	8	This claim of "computational constraints" sounds implausible, because the framework does not appear to be conceptually capable of handling multiple RCPs: these would include a kind of variation (and socioeconomic-time dependence) which is not discussed.	Yes, this was a simplifying decision, not the result of computational constraints. We have clarified this in the report. We have also added an appendix that compares RCP8.5 binned results to RCP4.5 binned results for two example sectors.
138	4	11	Section 3 first presents dollar losses, and then (in section 3.4) explains that these previous results are based on the information provided there. It would help with understanding for the physical impacts to come before the economic impacts.	We appreciate this comment, however we chose to show economic impacts before physical impacts because economic impacts are the main output of the Tool, and are available for all sectors where physical impacts are only available for a subset of sectors.
139	4	General	Earlier in the report (lines 921-922), we were told that economic impacts are computed from physical impacts (which would be the normal process). Here it is said that physical impacts are computed from economic impacts. This seems backwards and an opportunity to introduce error.	While the underlying studies produce economic impacts as a function of physical impacts, the data processed for FrEDI comes in as economic impacts. Therefore we calculate physical impacts as a function of economic impacts and known multipliers for reporting purposes.
140	5	3	2. I found the discussion of the time-dependent socio-economic scalars fairly confusing, i.e. the second type of socio-economic scaling described on page 17. It was only after reading the sector details in the appendix that I felt I had some grasp on what is meant here. It helped that the sector details specifically described model runs with and without population growth to isolate the socio-economic component. I am still unclear though whether these sectors can accept alternate socio-economic trajectories. At some points the text seems to suggest they can (e.g. line 501 "custom combinations") but then that they cannot (i.e. the "can not be separated outside the underlying model"). If these sectors effectively have fixed socio-economics determining results, that is an important thing to flag for users, as it means there may be some (admittedly small) discrepancies with any user-input scenarios they provide.	See response to comment #124. We have rewritten this section for better clarity. For the sectors that also use the second type of scalar (now termed "year-specific adjustment factors"), custom socioeconomic scenarios can be used but the year-specific adjustment factors will not change. This allows for first order scaling (i.e., for total population and GDP) but as you note, it can cause second-order discrepancies between assumptions in the inputs and the year-specific adjustment factors. We have noted this in the limitation section and added more analysis of the adjustment factors in the appendices.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
141	5	4	I think that pretty much all the information that users would need is somewhere in the report – I found the Appendices A and B particularly helpful. But I am not sure it is all optimally communicated for a reader looking to use the tool. In particular, reading Chapter 2, where the meat of the methods is communicated, for the first time, I found myself confused between several things: 1) the analysis that was done to pre-process impact studies in a standardized way; 2) things users could do (or have to do) in order to use the tool; 3) things users might want to do prior to using the tool (e.g. the discussion of Hector and FaIR, line 607); 4) things that might be done in the future with the framework (e.g. incorporating more damage sectors). Part of the problem is that each section (e.g. "economic impacts", "socio-economic scaling") is both describing the processing of impact results that was done to produce the tool and how the user engages with the tool. The result is confusion regarding what is part of the tool vs the overall framework. I would suggest splitting these out more clearly into two sections, something like 1. Methodology to Create Temperature Binned Impacts and 2. Using the Tool; or within each section within the chapter more clearly distinguish between what was done by EPA to create the binned impacts vs what the user is required or can do to use the tool.	Chapter 2 has been rewritten to address these comments and clarify the socioeconomic scalars (in response to other comments received). The new Figure 1 also attempts to clarify which parts of the Framework are pre-processing versus the parts that are necessary for each model run.
142	5	4	Relatedly, I found Figure 6 very confusing. I think there is an opportunity here to much more clearly distinguish the climate impacts modeling already completed (i.e. the sector-level models), the results processing done for the tool and the user inputs / tool outputs. It partly seems from this diagram that the user has to run the sectoral models or the SLR models etc to use the tool, when that is not the case. I think the figure could be substantially simplified – right now it is describing not the actual framework, but a particular use case of the framework (i.e. using it to compare damages under a baseline vs policy emissions scenario). But the caption makes it seem as though the whole flow chart is a necessary part of the framework.	We have removed this figure and instead include a new Figure 1, that essentially covers the same information regarding the parts of the Framework, but does so in a simplified manner. The Framework is inclusive of all steps shown in the former Figure 6, however, only a subset of those steps are required for each impact evaluation. We have attempted to clarify this distinction in the new Figure 1.
143	5	4	This is a more general confusion throughout the chapter of what constitutes the "Framework" as opposed to the "Tool". For instance, caption to Figure 6 states "The TBF starts with an emissions trajectory…", but the user actually has to supply temperature trajectories. I would suggest improved clarity throughout the chapter around this issue of what is the framework (which I think includes the underlying sectoral studies and the pre-processing to create damage functions, not the emissions-temperature link, which is what any climate model will do) vs the tool (which provides regional damages given user-defined temperature trajectories, adaption options etc).	We have added a new Figure 1 that provides a roadmap for the framework. We separate the framework into pre-processing, economic impact calculations, and post-processing. The emission scenario to temperature conversion is considered a part of the framework, but takes place during pre-processing.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
144	5	5	I do think an issue that should be dealt with earlier in the text and more explicitly is the question of overlap of damages between sectors. Unless I am missing something, this really seemed to be only mentioned in passing in parts of Chapter 3 (eg. Line 1025). But this is something users should be aware of earlier on – it is not clear in the discussion in Chapter 2 that adding up impacts from all the sectors will double count some damages. This will become a more important issue as more sectoral impact models are added (e.g. having the Hsiang et al 2017 mortality included). I think there could be a more explicit discussion of this issue in Chapter 2 (including how this overlap is handled in the aggregation functionality in the tool). This could also address other areas where it would seem there might be overlap but may not be given the details of the underlying studies (most notably PM2.5 mortality, Southwest Dust, and Wildfire morbidity and mortality).	We added more discussion in the methods section about the Framework's ability to summarize across sectors and account for potential duplication of impacts. For example, the summaries in the case study, now presented in an appendix instead of Chapter 3, do not include Asphalt Roads as it overlaps with the "All Roads" sector. We also added some text about the potential for double-counting PM2.5 benefits, which we estimate would be quite small for the relevant concentration-response function used and PM2.5 concentrations typically encountered in CONUS.
145	5	5	One question around non-CIRA studies is whether there is specific inclusion criteria for studies that could or would be included in the tool? The section describing the non-CIRA studies already included and the ones being prepared for inclusion felt a little ad-hoc. Why these studies and not others? Are these the only ones available at sub-national resolution? Are is there some methodological quality criteria EPA is using to filter potential studies? More information on criteria for including non-CIRA studies might help researchers with potentially relevant impact studies to figure out whether they could be incorporated into CIRA.	See response to Comment #63. The discussion on criteria for inclusion has been expanded in Section 2.2
146	5	6	As mentioned in response to Chapter 3, I found the discussion of the “time-dependent” scalars quite confusing. In particular, it is not clear to me whether this time-dependent scaling can be adjusted based on user-defined socio-economic scenarios. I feel like the terminology time dependent scalar is quite confusing as well and could be difficult to parse for a semi-technical audience.	See response to comment #109. We have rewritten this section for clarity.
147	5	6	After reading the information on page 19 several times, I remain confused about what exactly these scalars are needed for. Unlike the other type of dealing with socio-economic trends (ie. type 1 described in lines 502-503), which is accounting for how changing population and income will affect climate damages, it *seems* to me that that this second type is doing something quite different, which is backing out the climate change effect from sectoral model runs that include changing costs from both climate and socio-economics. If that is the case, then I don’t think they should be described as similar things (e.g. “both types of scalars”, line 550), though I realize I might well be wrong in this interpretation as it is not clear to me.	See response to comment #109. We have rewritten this section for clarity.
148	5	7	I would suggest adding the references for the underlying sectoral studies to the main text (e. Table 2). These studies are essential for understanding the values produced by the tool and so having the references easily accessible to readers, not just in the Appendix, seems important.	Thank you, we have added these to Table 2.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
149	5	9	2. I think the translation in the tool from global temperature to CONUS temperature change needs better justification. Footnote 16 gives this as 1.30764*T + 0.34057. But the inclusion of an intercept term here implies that there is warming in the US even with no anthropogenic climate change. This could give users non-sensical results. For instance, if a user put in a global temperature trajectory of zero (i.e. no global warming), damages in the US from global warming would be non-zero. This is an important part of the tool as many users might not have access to CONUS temperature trajectories and might be relying on it. I think it needs better documentation (and ideally no intercept term). It could also potentially have a more complex functional form – it is not obvious the scaling from global to continental temperature change will remain constant as a function of warming, for instance as various albedo feedbacks change at higher levels of warming.	Thank you for this thinking. We have added the Global to CONUS formula and underlying data points to a new appendix. As shown in that appendix, the relationship between Global and CONUS temperature changes remain fairly consistent across GCMs and over time. We acknowledge we only have six GCMs to rely upon. In addition, the point about inconsistent conceptual logic of an intercept term is well taken - we re-estimated the translation with no intercept - see Appendix D, Table D-2 for the updated result.
150	5	9	I think there is a mistake here in the translation to CONUS temperatures. For instance, 1.5 degrees relative to preindustrial might be about 0.7 degrees relative to 1986-2005. Plugging that into the formula given in footnote 16 gives ~ 1.2 degrees, less than half the 2.9 degrees reported. Similar issue for 2 degrees.	Thank you for this catch. There was an order of operations error that has been corrected for this version. The resulting CONUS temperatures are now 1.7 and 2.4 degrees, respectively (average global warming from pre-Industrial to the 1986-2005 baseline is about 0.4 degrees, accounting for the difference between the final numbers and your estimations).
151	5	11	The title and heading in the table make it seem as though damages are calculated in 2020 and 2090, instead of just using socio-economics from those years. This is confusing since 3 degrees of warming in 2020 is clearly not possible.	Thank you for this comment. We have reworded the column titles to specify these results represent certain socioeconomic conditions.
152	5	11	Units of premature mortality are not clear – total lives lost? Life years?	Premature mortality is measured as the number of deaths. We have clarified this in the table text.
153	5	11	Title “2-Degrees in 2020” is misleading to the casual reader. This should changed to specify 2 degrees warming using 2020 population and GDP.	We have updated the figure title to clarify that the results reflect 2020 conditions but do not suggest 2-degree warming by 2020
154	5	11	The figure caption (line 1087) and the y axis describes what is shown here as “Impacts”. This is misleading given how that term is used throughout the report. Instead it should be “Benefits” or “Reduction in Damages” (for instance y axis of Figure 17). The caption of Figure 17 has the same issue.	Thank you, we have adjusted the figure axes and captions to read as "Annual Reduction in Impacts" for similar graphics in the new appendix.
155	4	4	The inclusion of the ECS 2 line just adds confusion, since it is a single, unbalanced arbitrary point within the confidence intervals. Also, show the shading as done for the IPCC likely range, as done in fugure figures.	We have simplified this graphic and updated the results to show the AR6 ECS.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
156	1	6	In evaluating how the framework handles this issue, I was trying to imagine how the framework would integrate the Climate Impact Lab mortality working paper (Carleton et al., 2020, http://www.nber.org/papers/w27599). This paper explicitly varies the temperature-mortality relationship as a function of recent-past climate and income per capita. The former introduces some time dependence (e.g., mortality will be reduced after a couple decades at 2°C of warming vs upon first arrival at 2°C warming), and the latter introduces substantial time and socioeconomic scenario dependence. I think this would put it in the same category as the studies in Table 5 with time-dependent scalar, but how would the authors handle the fact that these impacts vary *greatly* with choice of socioeconomic scenario? (E.g., per-capita mortality projections differ substantially among the different SSPs.) It's not clear to me how this would work. Recommendation: The authors should clearly show the data they are using to calibrate non-trivial socioeconomic scalars, and if these scalars do not admit a precise representation of the underlying model data, should explicitly incorporate uncertainty into the analysis. They should also be careful not to claim the framework can generally handle custom socioeconomic scenarios if their ability to do so is limited to impact models of relatively low sophistication; and to clarify how they would handle a model like that of Carleton et al., 2020, if they can.	We have provided additional information and analysis about the adjustment factors in the appendix on sensitivities and rewritten substantial sections of section 2 to better explain the treatment of "custom socioeconomic scenarios". We would need to spend more time with the Climate Impact Lab mortality data, however if impacts are provided per capita, that first order scaling can be applied to custom scenarios in the Framework. The second-order adjustment factors for the underlying SSPs would likely be stored in the tool as alternative runs (similar to the two emissions assumptions in the air quality sector). It would then be up to the user to select the SSP run that best matches their specified scenario. The Tool would likely have a default selection that best matched the default population and GDP trajectories.
157	4	6	Many future RCP 8.5 impact projections will be tied to SSP5, which is an outlier socioeconomically. Using SSP5-based impacts to calibrate other socioeconomic scenarios is sure to produce odd results, even with the "re-scaling" process.	While we used RCP8.5 in data processing steps described in detail in Appendix A, the socioeconomics are not from SSP5 but are custom to CIRA, and closer to SSP2. Detailed analyses comparing the socioeconomic inputs to SSPs are included in EPA (2017a), which can be found here: www.epa.gov/CIRA
158	1	6	I do not believe the relationships that allow the use of custom scenarios are clearly described. While for some examples – e.g., valuation with a VSL, or scaling per-capita mortality to total mortality – this is straightforward and can be done precisely, in other cases the authors are using scale factors that must be approximations calibrated via regression against the model output. I found nothing that allowed me to assess the extent to which these approximations accurately describe the response surface of the underlying models.	See response to comment #109. We have rewritten this section for clarity.
159	2	4	Clarify why there is no coastal property impact at 30cm	This is addressed in the revised SLR binning method.
160	1	5	More clarity would be useful with respect to the experiments that were run into CIRA to enable the current application. For example, my understanding from Appendix A is that the CIRA experiments were run with time-varying climate and fixed 2010 or 2090 socioeconomics. Is this correct, or were the future CIRA experiments run with time-varying socioeconomics? If the latter, how were the impact vs. temperature curves extracted from the time-varying socioeconomics? Recommendation: The authors should clarify the specific experiments run by the CIRA models, and whether fixing socioeconomics helps with integration into the framework.	The CIRA experiments included both fixed 2010 and 2090 socioeconomics at a minumum for sectoral analyses, effectively as control runs. Fixed socioeconomics provide a useful standardization, as explained further in Sarofim et al. (2021). Some experiments also allowed for time-varying socioeconomics. We added some clarifying text in Appendix A.
161	3	2	There is a fairly high level of assumed knowledge on ongoing EPA projects (climate impacts lab, CIRA) – the report should make efforts to define and introduce activities, their primary objectives and products.	Footnotes have been added to the main report where the CIRA and Climate Impact Lab projects are first introduced. Each footnote describes the objective of the project, makeup of contributors, and where readers can go to learn more information.
162	3	2	Some introduction of the CMIP models used, why they are used and why they represent climate uncertainty would be appropriate	An extensive footnote has been added describing the source of GCM (CMIP5) and downscaled data products (LOCA) used in the framework, including how/why the subset of 6 GCMs were chosen and where the reader can go for additional information about this selection.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
163	3	10	6) There is no discussion of the degree to which the use of 6 CIRA GCMs as a sample of future uncertainty may fail to represent tail risks due to common structural approximations in current generation GCMs.	A description of this limitation has been added to the report, including the potential implications on the damage functions and results.
164	4	3	Standard CMIP5 projections are also provided with “model historical” results for years prior to 2006. Readers may wonder why these are not used.	A footnote has been inserted describing that "the underlying downscaled dataset also contains hindcast results (‘model historical’) for the years 1950-2006. Since the purpose of the Framework, and its underlying studies, is projecting future damages, this hindcast dataset is less relevant."
165	4	4	Most of the specific points below highlight problems issues with the figures and tables, and other points of missing information are discussed in response to question 10. An important discussion missing on the requirements of the input data is the definition of “impacts”. The report states that impacts can “be integrated into the Framework if the necessary information is available” (line 403). However, this seems insufficient. The following questions seem crucial: (1) Is an impact limited to consequences of climate change? For example biodiversity loss projections are often described out to 2100, as driven by land use changes, and not accounting for the effect of the temperature changes. (2) What are impacts measured relative to? The implicit answer seems to be 1986 – 2005, but in most cases this is inadequate. For example, death rates are likely to decrease even as temperature increases, due to improvements in health care. If a timeseries of death rates were added to the system as described, it would produce the result that climate changes saves lives. One possible answer to this would be compared to a counterfactual world with changes in health care but no climate change. (3) Should the transfer of economic value from one sector to another be included? Currently, the winter recreation sector is included, and transfers are certainly a valid form of impact. However, transfer do not represent a loss to society, so any analysis that aims to quantify the social cost of carbon would want to exclude impacts like these.	These are all good questions. The remainder of the referenced paragraph provides some basic information on what is required from new studies to be included in the Framework, also refers the reader to Section 2.3, where the question of baselines is addressed (lines 596 to 600 in April draft report), and the possibility of adjusting for differing baselines is introduced. Certainly a study must also be well-conducted - and issues of projected baseline death rates might, for example, be adjusted in data processing with simple re-analysis if necessary. We added text about a study having "a clearly defined and transparent baseline scenario." On the winter recreation question, we added to footnote 8 (to the referenced paragraph at line 403) that new work that values multiple forms of recreation using a willingness to pay measure (building on work by Chan and Wichman of RFF) - which could reduce the need to include studies that use transfers as potential proxy values for societal losses.
166	4	7	Starting in figure 12, damages across sectors are just added together. Since welfare valuation methods are used, a compensating variation would need to be used to describe the combined welfare losses. Alternatively, the non-welfare damages can be put into a CGE to allow for redistribution. In either case, the total loss is not equal to the sum of the sector-specific losses, and a simple sum is not a reasonable assumption.	Currently the Framework is not capable of using a compensating variation to describe combined welfare losses, or employing a CGE to aggregate expenditure denominated damages. We have added caveats to the main text, figures, and relevant Appendices to note these limitations.
167	4	11	As a tool for coordination, however, the report might be able to do more. There are different audiences for this kind of framework: Researchers producing the underlying studies, policy-makers using the outputs, and academics interested in the framework itself. The last group is the smallest of the three, but it seems like much of the report is written for them. A document that is targeted at encouraging coordination would do more to provide a clear process for new sectors to be included (see question 5).	Moving forward, EPA intends to carefully monitor the literature to identify appropriate sectoral studies for inclusion in the Framework. In order to advance the utility of the Framework, EPA encourages researchers and practioners to develop additional sectoral impact studies that can be considered for use in FrEDI. A description to this effect has been added to the documentation.
168	2	3	“Baseline climate era (1986-2005)” – this is a crucial issue that applies to the entire report: Climate normal are typically defined over a 30-year period, and is the international standard. The justification for a 20-year period for the baseline needs to be thoroughly addressed and the potential for picking up variability due to the smaller window should be adequately explained, in addition to explaining how this choice may deviate from – or be comparable to - conclusions drawn, or data used, by the rest of the community when using the standard 30-year period.	We added the sentence "The size of the binning window is a balance between smoothing out interannual variability and the inclusion of years at the beginning and end of the window that would not be representative of the window’s average temperature: the smooth behavior of the damage curves for most sectors and GCMs indicates that 11 years is likely sufficient" to the document.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
169	1	8	No. As noted previously, it is not at all clear how uncertainty in the relationship between climate variables, or in the socioeconomic dependence of impact functions, is handled. Throughout the report, there is a love of excess precision that is inappropriate given the uncertainties involved. The discussion in 2.6 is fine, but I am simply not comfortable with some of the uncertainty-free approximations made earlier.	Discussion of uncertainty and limitiations of the tool has been significantly clarified and expanded in revised Sections 2.6 and 2.7, including explanation of the deterministic underpinnings of the approach. The report and particularly the tables and figures have also been edited to remove a suggestion of excess precision. Nonetheless there remain limits to the ability to quantitatively assess uncertainty in this framework designed for rapid assessment.
170	1	9	EPA has objectively used, applies, and documented the underlying data, though as noted above, in question 6, more detail would be appreciated. As noted above, I do not think uncertainties related to climate variables and the response to socioeconomic scenarios approrpaitely tested.	Discussion of uncertainty and limitiations of the tool has been significantly clarified and expanded in revised Sections 2.6 and 2.7, including explanation of the deterministic underpinnings of the approach. The report and particularly the tables and figures have also been edited to remove a suggestion of excess precision. Nonetheless there remain limits to the ability to quantitatively assess uncertainty in this framework designed for rapid assessment.
171	2	1	“risks” – Sensitivity and vulnerabilities are part of risk. Do the authors mean “impacts” as opposed to “risks”?	We agree and made the suggested change.
172	2	1	The arguments presented by the authors as to how the TBF addresses “challenges” of scenario-based modeling, and how it is different, is not convincing and needs better justification or should be removed. The two - scenario-based modeling and what is attempted here, are fundamentally different modeling paradigms. Indeed, the TBF is just using data from the scenario-based modeling. The two are not alternatives – the TBF is just a “post-processing” algorithm.	We agree and edited the text to clarify that FrEDI relies on scenario-based analyses but post-processes the results to provide standardized impacts by degree.
173	2	1	“more reliable fit” – provide justification or reference.	This text was deleted in response to another comment.
174	2	1	It is not clear from this report how it would accommodate different “study methodologies” as opposed to different impacts data or socio-economic pathways. Reword or explain.	We are referring here to ability to pre-preocess the results of both process-based and econometric impact studies, as described elsewhere in this section.
175	2	2	“Coastal Properties and High Tide Flooding” – Is this an EPA report? Provide reference or explain.	The SLR component of the framework has been substantially revised, and so this portion of the report has been substantially rewritten and the reference has been deleted - the Coastal Property and High Tide Flooding sector studies are, however, newly published and are referenced in Appendix A.
176	2	3	Clarify why 2015$ and not 2020$ especially as 2020 socioeconomic conditions are used	The use of 2015$ is consistent with prior work in this series, particularly Sarofim et al. 2021a and EPA 2017a - they are used here to enhance traceability of results. In addition, at the time of initial publication, final data to adjust to 2020$ were not available. Users can of course adjust all values to any years dollars as a post-processing step.
177	2	3	“but at higher cost” - Provide justification and/or a reference for this statement that "protecting" is higher cost than "elevation.	We provided the reference.
178	2	6	Explain why armoring is not considered as a reactive adaptation measure when it is routinely practiced nationally after individual and specific storm damage. Structural elevation, e.g. in building codes, for new development, but also changes in building codes for existing developments, is also often a pro-active adaptation measure, not a reactive one. This section needs to be explained clearly with adequate justification as the model output and results are not intuitive.	The underlying study (Neumann et al. 2021) outlines the logic for classifying measures as reactive or proactive. The general concept is that reactive measures are either responsive to events (without foresight about future events) or can be undertaken without coordinated action between individuals and governments. Elevation, for example, is modeled at the individual property level in response to highly localized hazards, not as a collective action of municipal governments to modify building codes. While it is true that small scale armoring does occur as a reactive action, more substantial armoring requires coordinated action or, in several states, waivers from existing coastal policies. We have added the recommended explanations to Section 3.5 and the appendix.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
179	3	3	Similarly, future wildfire projections are hugely uncertain – and it is not clear that most of the current generation GCMs are fit for purpose for the regional projection of wildfire activity (see Sanderson 2020 https://www.nature.com/articles/s41558-020-0707-2). Wildfire projections are also a first order function of land use assumptions which are not independent of the mitigation pathway.	We agree that wildfire projections are uncertain, and that existing literature likely underestimates impacts, but the cited Sanderson and Fisher paper appears to focus on events rather than regional and era scale trends in wildfire activity. While the GCMs are likely not suitable to define the specific temporal pathway of fire events, the overall trends, at the scale of multi-state US regions, are the focus of the underlying wildfire paper. We have added additional text to the Appendix A discussion on this sector study.
180	3	8	There is very little quantitative representation of uncertainty or confidence in costs, though uncertainties would be expected to vary significantly across the sectors considered. For example, we have a relativity high confidence in the strength of the relationship between US-level warming and regional heat extremes (https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2018EF000943), but the confidence in precipitation extremes is larger and dependent on scenario (https://agupubs.onlinelibrary.wiley.com/doi/full/10.1002/2015GL065854) and region (https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2019GL082362). Meanwhile, our confidence in current GCM's ability to represent future wildfire evolution is low (https://www.nature.com/articles/s41558-020-0707-2).	Discussion of uncertainty and limitiations of the tool has been significantly clarified and expanded in revised Sections 2.6 and 2.7, including explanation of the deterministic underpinnings of the approach. The issue of sector-level comparative uncertainty is now addressed, and it is acknowledged that there remain important limits to the ability to quantitatively assess uncertainty in this framework designed for rapid assessment, especially dependent on the extent to which uncertainty analysis has been completed in the underlying sectoral literature.
181	3	8	The framework would benefit from representing the confidence of cost projections for different sectors – at least qualitatively, but ideally quantitatively such that the framework has the capacity to represent the impacts of relative uncertainty in projections – both at a global and regional level. This is currently done only for global climate sensitivity, and implicitly across the range of CIRA GCMs, but it would be useful to clearly communicate that some sectors are (1) fundamentally more uncertain in common GCM errors which may influence impact-temperature relationships (e.g. common simplistic wildfire parameterizations) and (2) subject to uncertainties arising from exogenous factors which are not considered in this analysis (e.g. non-CO2 emission pathways). (3) subject to scenario dependencies which are not well represented in the temperature binning framework (e.g. mean precipitation change which is a direct function of CO2 concentrations as well as mean temperatures)	Discussion of uncertainty and limitiations of the tool has been significantly clarified and expanded in revised Sections 2.6 and 2.7, including explanation of the deterministic underpinnings of the approach. In particular, a table outlining the main sources of uncertainty and a qualitative characterization of the influence of each factor is now included in Section 2.6. Nonetheless there remain important limits to the ability to quantitatively assess uncertainty in this framework designed for rapid assessment, especially dependent on the extent to which uncertainty analysis has been completed in the underlying sectoral literature.
182	3	9	The largest single factor in the future cost evaluation, air quality change, is highly influenced by non-CO2 emissions – which are not sampled in this report. This is a noted limitation (line 756), justified because non-CO2 emissions impacts are challenging to emulate. But the methodology would benefit from some quantitative estimate of the role of non-CO2 emissions in air quality in particular (also, to some extent extreme temperatures and precipitation). This could be done using two different non-CO2 emission rates, rather than the single 2040 rate used in the current analysis.	We agree. To be clear, the analysis presented in Section 3 of the report uses a 2011 non-CO2 emissions scenario. In footnotes in this revised section,we included a short summary of the impact of using alternative non-CO2 emissions, and also references to Appendix A where both sets of non-CO2 emissions results are presented by degree and for alternative socioeconomic scenarios.
183	3	10	1) At least some communication of comparative uncertainties in costs arising from missing externalities or uncertain temperature response would be very valuable, establishing that projections of climate related costs in different sectors are not equally confident. This could be a qualitative or quantitative assessment, and could make use of IPCC style confidence language.	Discussion of uncertainty and limitiations of the tool has been significantly clarified and expanded in revised Sections 2.6 and 2.7, including explanation of the deterministic underpinnings of the approach. In particular, a table outlining the main sources of uncertainty and a qualitative characterization of the influence of each factor is now included in Section 2.6. Nonetheless there remain important limits to the ability to quantitatively assess uncertainty in this framework designed for rapid assessment, especially dependent on the extent to which uncertainty analysis has been completed in the underlying sectoral literature.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
184	3	10	5)There is no discussion of the indirect effects of climate mitigation action on infrastructure (energy, transport, agriculture), and how this might modulate costs incurred due to climate change.	Some discussion of feedback effects, including indirect effects of mitigation policy on impacts, has been added to the limitations section of the revised report (Section 2.7) to address this comment.
185	3	General	My biggest single critique is an insufficient of representation of uncertainty (beyond the inter-GCM spread) in impact cost projections. Even if a quantitative assessment is beyond scope for this report, I would encourage the EPA to consider at least qualitative guidance for the confidence in reported costs in different sectors, and to consider how to objectively represent end-to-end uncertainties in future iterations of the framework – due to the structure of the framing itself (pattern scaling assumptions in mapping global to CONUS temperatures), common structural errors in GCMs (relevant for e.g. wildfire projections) or non-CO2 related risk factors (e.g. aerosol emission impacts on air quality).	Discussion of uncertainty and limitiations of the tool has been significantly clarified and expanded in revised Sections 2.6 and 2.7, including explanation of the deterministic underpinnings of the approach. In particular, a table outlining the main sources of uncertainty and a qualitative characterization of the influence of each factor is now included in Section 2.6. Nonetheless there remain important limits to the ability to quantitatively assess uncertainty in this framework designed for rapid assessment, especially dependent on the extent to which uncertainty analysis has been completed in the underlying sectoral literature.
186	4	1	The discussion of alternative techniques for developing damage function raises some issues. (1) The result of the Binning framework is a damage function, but the concept of a damage function is not introduced. (2) The critique of econometric techniques (line 205) seems out of place, since these techniques are used to do something quite different, and are potentially an input into the Binning framework. (3) It is not clear what the point of the IAM discussion is. IAMs are hugely relevant, since they use damage functions of the form the Binning framework produces (although this is not made clear), or to produce scenarios which could be used with the Binning framework. But somehow IAMs and the Binning framework are set up as something to be contrasted. Also, one guesses that the authors are referring to CB-IAMs, and not PB-IAMs.	Added a clarifying phrase "designed for damage estimation" after "integrated assessment model" [Notes from Jim - need to clarify that in fact we were referring to Benefit-Cost IAMs used to support SCC, and not process based IAMs which use detailed process-based models to characterize physical and economic damages.
187	4	1	The use of GCM outputs in analyzing impacts is better described as the gold-standard, rather than the “default”.	Agreed - we made the suggested change.
188	4	2	What does “normalized by temperature” mean? What are opportunitistic ensembles, what is an ensemble of climate scenarios, and what are 21st century eras? Also, on line 200, I believe the authors mean “therefore”, not “thus”.	The "opportunistic ensemble" point, as well as other points made in this comment, were made in Sarofim et al. (2021) in a different context, which is not applicable here. These statements have been removed.
189	4	3	In the introduction, it’s suggested that the goal is to provide reliable fits at low warming levels (199-200)—mostly historical warming levels if one is only concerned with 1 degree warmer than 1986-2005. Is that goal further developed?	This point was made in Sarofim et al. (2021) in a different context - we removed the point as the context is not relevant here.
190	4	4	The report states that co-emitted air pollutants are not included in the Framework (line 150), so why is ozone and PM2.5 included?	The air quality study considers two scenarios for emissions of PM and ozone precursor emissions, but these scenarios are not tied to GHG mitigation scenarios because of data and model limitations. The air quality study is focused on the impact of future meteorology on ozone and PM2.5 concentrations, rather than changes in co-emitted pollutants. We added text at line 150 and clarification in the notes section of Table 2 to address this comment.
191	4	5	Does a “no wildfires” baseline assume that there are no wildfires without climate change? This is clearly not true.	The "no-wildfires" scenario is not a baseline but a tool to estimate the incremental impact of wildfire emissions in the baseline. We agree that is not a realistic scenario, but it is used in the referenced study to identify excess risk of wildfire emissions. We added a clarification in the Appendix.
192	4	5	Are changing storms (e.g. hurricanes) and the interaction with non-climate-change driven changes to coastlines included? If not, these are additional limitations.	We agree - while these factors are discussed in the referenced literature, we added a new bullet in the Appendix as well.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
193	4	5	Here sea walls and beach nourishment are included in the “reasonably anticipated adaptation”, but in the coastal properties sector, that requires the “proactive” adaptation scenario. Does that make the “reactive” adaptation scenario combined across these two sectors inconsistent? Also, how is it determined where coastal protection will be performed?	The reviewer is correct, this an inconsistency regarding seawalls, though beach nourishment is part of the reactive and proactive coastal properties adaptation scenarios. As outlined in the Fant et al. (2021) high-tide flooding paper, however, the impact of this inconsistency is slight - Figure 3 and accompanying text in that paper notes that alternative routing reduces the no adaptation impacts by 77%, while the marginal additional impact of ancillary sea wall protection increases the total to an 80% reduction. We nonetheless agree that the inconsistency is worth noting and added text to Appendix A to further explain the logic.
194	4	5	Does this assume that all roads get equal usage?	No, current traffic is an important input to the model and affects estimation of delay costs, and traffic is adjusted for future local population growth. We added clarification to the Appendix text.
195	4	5	Do I understand correctly that there is a total budget for repairs, after which the state becomes bankrupt? Or is it a fixed budget per year? If the former, this is a serious problem. If the model is designed with a single budget over multiple years, this can be resolved just by running the model multiple times, ending in different years, with extra money each time.	The no-adaptation scenario considers a fixed budget per year, to illustrate the potential implications on road users (in terms of delays and vehicle repairs and operating costs) of failing to adjust repair budgets in response to climate change - for this reason it is not the central case in the FrEDI tool. Other response scenarios do not include fixed budgets. These issues are discussed further in the underlying study (now published with Open Access availability).
196	4	6	One of the characteristics of socioeconomic trajectories is the extent of inequality, including sub-national inequality. It is not clear why national GDP used, when the impact relationships are available at the regional level.	This is an important point - CIRA work as described in EPA (2017a) has used national-scale GDP driver data, in part on the recommendation of the MIT modelers who conducted that portion of the analysis. We are aware that other tools produce spatially differentiated GDP estimates, and some produce long-term forecasts with spatial differentiation (at least through mid-century). We can consider adding a user capability to use regionally differentiated GDP drivers in future revisions of the Framework.
197	4	6	In many cases, I believe that costs-rates are assumed to be constant in time. For example, rail costs are in damages per mile. However, labor is used to repair rails, and the cost of labor will increase with income. This is mentioned as a limitation for the rail impact (though often left off for other impacts), and it is not clear why labor costs are not allowed to scale with socioeconomcs. The Asphalt Road impact is another case where labor costs are important, and this would make the damages sensitive to the socioeconomic scenario.	Labor costs are an important component of repair costs, but the underlying sectoral studies do not provide breakouts of repair costs by labor, capital/materials, and equipment. Some of the public health studies, however, do report these breakouts. A refinement of the treatment of labor for the infrastructure categories could be considered for a future enhancement of the approach.
198	4	7	Are medical costs constant across time? They should scale with income.	Consistent with EPA's practice for health risk valuation in economic analyses, and the "Guidelines for Economic Analysis," the direct component of medical costs are held constant over time. As noted in the "Guidelines," real medical costs could increase or decrease over time based on technological change and other factors. The indirect costs (lost work time) are adjusted over time consistent with income projections.
199	4	7	Does the model include changing costs for repairs to baseline risk? In this case, changes in the costs will be mis-attributed to climate change.	The model does not include changing costs for repairs to baseline risks.
200	4	8	It is not clear what “uncertainty” as used in the report means. Is it related to the expected error, the confidence interval around estimates, a lack of knowledge, or (as appears to be the case in many instances of the use of the term) a mismatch between the assumptions of the Framework and reality?	Discussion of uncertainty and limitiations of the tool has been significantly clarified and expanded in revised Sections 2.6 and 2.7, including explanation of the deterministic underpinnings of the approach. In particular, a table outlining the main sources of uncertainty and a qualitative characterization of the influence of each factor is now included in Section 2.6. Nonetheless there remain important limits to the ability to quantitatively assess uncertainty in this framework designed for rapid assessment, especially dependent on the extent to which uncertainty analysis has been completed in the underlying sectoral literature.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
201	4	8	The proper quantitative representation of uncertainty appears to be missing in this framework. Impacts by bin are averaged across climate models, dropping crucial policy-relevant information. While uncertainty does not always need to be quantified, it would appear to be the responsibility of the Framework to justify not using information available to quantify it (e.g., GCM-specific results are directed to be averaged, although this information is displayed in section 3), and to provide guidance on how uncertainty as it is "provided" by the framework should be used or engaged with.	Discussion of uncertainty and limitiations of the tool has been significantly clarified and expanded in revised Sections 2.6 and 2.7, including explanation of the deterministic underpinnings of the approach. In particular, a table outlining the main sources of uncertainty and a qualitative characterization of the influence of each factor is now included in Section 2.6. Nonetheless there remain important limits to the ability to quantitatively assess uncertainty in this framework designed for rapid assessment, especially dependent on the extent to which uncertainty analysis has been completed in the underlying sectoral literature.
202	4	8	Also, avoiding quantification of uncertainty (or only quantifying the uncertainty across GCMs) is likely not justifiable. Uncertainty comes from ECS assumptions, GCM models, GCM runs, region-specific weather outcomes, socioeconomic models and their uncertainty, impact models and their uncertainty, and uncertainty in valuation and equilibrium adjustments. These sources of uncertainty tend to compound, and are important for describing risk.	Discussion of uncertainty and limitiations of the tool has been significantly clarified and expanded in revised Sections 2.6 and 2.7, including explanation of the deterministic underpinnings of the approach. In particular, a table outlining the main sources of uncertainty and a qualitative characterization of the influence of each factor is now included in Section 2.6. Nonetheless there remain important limits to the ability to quantitatively assess uncertainty in this framework designed for rapid assessment, especially dependent on the extent to which uncertainty analysis has been completed in the underlying sectoral literature.
203	4	8	It is not clear how the binning framework highlights/exposes structural differences between GCMs.	We clarified that this is a point made in Schleussner in support of time-slice/temperature binning type approaches, rather than a feature of this Framework.
204	4	8	How much uncertainty is added? I assume that readers are meant to take the stated "understanding of the added uncertainty" to mean something fairly vague, but isn't it the role of this framework document to provide any such "understanding" (or at least a process for reaching that understanding)?	We modified the language to clarify that, while the Framework has an extrapolation capability, that when analyzing user scenarios that go beyond the scope of the underlying sectoral studies, a flag is generated for the user. The Framework itself does not provide an estimate of the added uncertainty in extrapolating beyond the range of temperatures for which the underlying sectoral studies were exercised.
205	4	9	The document does a good job of documenting the data that has been used in the framework and in applying the sector-specific damages. The document does a less thorough job of considering the sensitivity to these assumptions. Several sectors have to be interpolated to get integer warming points, and then are interpolated again in the projection process. How much fidelity to the original results is lost in this process? The role of socioeconomics is generally assumed to be linear. In cases where there are multiple socioeconomic years in the underlying studies, how well does the linear approximation reproduce the results for years that are not used to create those scalings? Impacts are mostly driven by local weather, not CONUS-wide climate. Under cross-validation, how well can results predicted from CONUS-wide climate reproduce local impacts? All three of these robustness checks should be performed.	Some of these points are now addressed in the revised uncertainty and limitations sections of the report (Sections 2.6 and 2.7). We also now include new sensitivity analyses of the impact of interpolated socioeconomic scenarios, which are referenced in Section 2.6.
206	4	9	It is important to note that different kinds of impact models represent different processes, and that process-based simulation models are not generally commensurate with econometric models (Franziska et al. 2021). Process-based models often do not capture the reactive effect that humans and the environment have to impacts.	Thank you for the reference - we added a new footnote to make this point. Our main point is that both process-based and econometric models have insights, if differing insights, to add regarding economic impact estimation.
207	4	9	Is the size of the 50-year storm allowed to change, or is the frequency of storms above a present-day 50-year storm threshold allowed to change?	The size is allowed to change - a new footnote clarifies this at the location noted by the reviewer.
208	4	9	What then is included and what is not included? Not all US cities, not all conveyance systems?	Not all U.S. cities are included in the underlying study. This is now clarified in the limitations noted by the reviewer.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
209	4	10	The description of the population and GDP projections (appendix C.2) is incomplete. How is information from the UN projected values and from the ICLUSv2 model combined? What assumptions are made about region-specific and time-varying rates of urbanization and fertility rates (a table is needed)? What does it mean that EPPA-6 and EIA projections are “harmonized”? It also sounds like there is a discontinuity in GDP growth—from 2% in 2015 to 2.5% in the projection. What explains this, and is it fixable?	The technical documentation for the ICLUSv2 model provides information on the demographic and other modeling assumptions, and is now more clearly referenced in the main report (see EPA 2017b). Full technical documentation for the EPPA-6 GDP projections is provided in EPA (2017a) and that report's technical appendices. Any discontinuity cannot be fixed at this stage, but the Framework was designed to accept custom population and GDP assumptions precisely because we anticipate that some users would prefer to use alternative inputs to those commonly used for EPA's Climate Impacts and Risk Analysis (CIRA) framework.
210	5	1	I think the introductory chapter mostly does a good job of outlining the framework, putting it in the context of previous literature and providing the context for the main conceptual issues addressed in the next methodology chapter. I do think the over-arching need for the overall framework could be more clearly articulated (i.e. why is there a need for the processing and standardization done by the framework?). The need is clear to me but I worry it would be lost on the majority of readers without “hands on” experience of trying to create damage functions from the supporting climate impacts literature. I have in mind something that motivates how the tool makes climate impact results more accessible such as “A large number of studies have simulated the effects of climate change on various socio-economic outcomes within the US. But many use distinct climate or socio-economic scenarios that are incompatible with each other, or report outcomes in units that require further processing to be comparable across sectors. Underlying impact models often require specialized, sector-specific knowledge to run or, in some cases, may require substantial computational resources, making them inaccessible for a typical user. This framework and tool bridges this gap: by processing climate impact modeling results, users can explore impacts across multiple sectors in a standardized way as well as exploring the effects of temperature, socio-economic, and adaptation scenarios of interest.”	This suggestion is useful, and we have incorporated the language nearly verbatim in Section 1.3 of the documentation.
211	5	1	I did not understand the point here about why the framework would be better for small levels of warming	This point was made in Sarofim et al. (2021) in a different context - we removed the point as the context is not relevant here.
212	5	1	I am not clear the distinction being drawn here between the framework and IAM damage function (other than IAM damage functions are typically global in scope). To me it seems that the framework is essentially providing sector-specific damage functions for the US, of the same kind that might well be used in a cost-benefit IAM.	We agree with the reviewer - the topic sentence of this paragraph notes that the two have some similarities. The main distinction we make is that the Framework is not an optimization approach, but it does provide flexibility to adopt temperature trajectories from a range of sources.
213	5	7	I found the level of detail given in the sector-specific section in the Appendix mostly very good. I noted a few issues: -The Health Sectors (Appendix A.2) seemed to be missing information on how adaptation costs and benefits were modeled. This discussion is good for the infrastructure sectors but is notably missing for the earlier sectors, even though they have adaptation options associated with them.	The Health sector where adaptation is considered is Extreme Temperature Mortality, but the scenario does not reflect a benefit-cost calculation but an assumption that cities will gradually adapt to a hotter environment through physical acclimatization of their residents, infrastructure replacement with more heat suitable shading and air conditioning, and behavioral changes so that the stressor-response will look like that of the current Dallas context. More refined adaptation analyses for this sector are currently the subject of ongoing research. We have clarified these points in Appendix A.
214	5	7	A “no wildfire” mortality baseline is mentioned. It seems like this should be a “no climate change” baseline given the context – these are not the same thing	The "no-wildfires" scenario is not a baseline but a tool to estimate the incremental impact of wildfire emissions in the baseline. We agree that is not a realistic scenario, but it is used in the referenced study to identify excess risk of wildfire emissions. We added a clarification in the Appendix.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
215	5	8	In my opinion this is currently a weakness of the framework developed here, though possibly an unavoidable one given the climate impacts literature that is being relied on. The uncertainty most clearly and extensively discussed in the report is that of climate models – but given the "temperature binning" approach used, this is a relatively minor component of uncertainty in most sectors. A large source of uncertainty is the uncertainty involved in translating temperature or SLR trajectories into damages, but this is unaddressed, other than a brief mention in Figure 6. I am surprised that the supporting sectoral studies did not provide some confidence interval or uncertainty ranges that could be propagated through to the tool in order to better communicate / address this source of uncertainty in the output. To me this is a key limitation of the tool as currently designed and should be addressed explicitly, for instance with a paragraph in Section 2.6	Discussion of uncertainty and limitiations of the tool has been significantly clarified and expanded in revised Sections 2.6 and 2.7, including explanation of the deterministic underpinnings of the approach. In particular, a table outlining the main sources of uncertainty and a qualitative characterization of the influence of each factor is now included in Section 2.6. Nonetheless there remain important limits to the ability to quantitatively assess uncertainty in this framework designed for rapid assessment, especially dependent on the extent to which uncertainty analysis has been completed in the underlying sectoral literature.
216	5	8	I think the discussion of uncertainty here needs more context to make it useful to the reader. In particular, while the listing of sources of uncertainty is quite comprehensive, they are dealt with in completely different ways within the tool. For instance, sources of uncertainty included in the tool (Adaptation Options, GCMs), completely un-dealt with (Structural Uncertainty), and uncertainty that could be investigated by the user if desired (Climate Sensitivity, Baseline Uncertainty) all appear in the Figure in exactly the same way. Note that population and GDP inputs also have uncertainty associated with them (a similar type of input uncertainty to the "Baseline Uncertainty" associated with emissions).	Discussion of uncertainty and limitiations of the tool has been significantly clarified and expanded in revised Sections 2.6 and 2.7, including explanation of the deterministic underpinnings of the approach. In particular, a table outlining the main sources of uncertainty and a qualitative characterization of the influence of each factor is now included in Section 2.6. Nonetheless there remain important limits to the ability to quantitatively assess uncertainty in this framework designed for rapid assessment, especially dependent on the extent to which uncertainty analysis has been completed in the underlying sectoral literature.
217	5	8	This is a critical discussion for properly interpreting the output from the tool as quantification of a select set of sectors and impacts. I would add that even within modeled impacts, there are reasons costs would be underestimated because of missing valuations. Most obviously would be willingness to pay to avoid sickness (as opposed to cost of illness) and similarly welfare losses from lost winter recreation (as opposed to simply ticket revenues).	We agree and added the reviewers point to the referenced text.
218	5	9	Mostly I think the analysis done for the tool and the key assumptions in processing the impact studies are well documented. There are two points I think need better documentation: 1. The VSL used to value mortality damage scales sub-linearly with income (line 1359). This is an important assumption since the mortality damage are a major component of total damages and this parameter will have a big effect on how they scale with GDP per capita. But this elasticity of 0.4 is not referenced and is not mentioned in the main text. The value of 0.4 seems to be quite far on the low end of ranges reported in the Masterman and Viscusi paper referenced in footnote 26 so this is a key assumption determining damages that should be justified.	The current default EPA policy for use of an income elasticity adjustment to VSL, based on the most recent Science Advisory Board review of this parameter, uses the 0.4 value, as described in the referenced documentation for the BenMAP-CE model in Appendix A. We added a footnote to the main text to clarify the basis for this assumption. For the same reasons outlined by the reviewer, however, many of the underlying studies that rely on VSL test the sensitivity of impact results to this assumption and include consideration of an elasticity of 1.0 that is more consistent with current literature.
219	5	10	Other than the points identified in response to previous questions, I do not think there is essential information missing from the technical documentation. The processing of individual sectors and the key elements of each sector's modeling, including important assumptions are well described in Appendix A. The one important issue is that several papers describing the underlying studies are still listed as "in press" – ideally these will either be published by the time the tool is released, or perhaps made available as working papers.	Three papers previously indicated as "in review" or "in press" have now been published and are available with open access (Neidell et al. 2021; Neumann et al. 2021; and Fant et al. 2021). The Price et al. urban drainage paper has been added to PubMed as open access, here: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8477377/. In the revision, all cited literature is now available, and only one paper (Underwood et al. 2017, on asphalt roads), not sponsored by EPA, requires subscription access.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
220	2	1	"changes in sea level" – sea level changes are driven by warming. Explain any dependencies or "double-counting" here.	Since the April draft, the Framework's treatment of SLR has been substantially revised to better tie SLR directly to levels of warming - in the process we have eliminated certain potential inconsistencies. There is no double counting that we are aware of.
221	4	4	The reduction in benefits in the high tide flooding and traffic in the 2060s needs to be explained.	Revisions to the methodology used to estimate the influence of a global temperature trajectory on GMSL, and of GMSL (mapped to local SLR) on regional-scale damages in response to reviewer comments also eliminated this discontinuity in the high-tide flooding sector results.
222	4	10	The role of adaptation scenarios is also unclear. For example, line 442 claims, "The Framework estimates results for all available adaptation scenarios." But how it does that is not stated. In the appendix, it is clarified that the adaptation estimates are provided by the underlying studies. The central contribution of the framework in this regard is then to develop an ontology of adaptation, grouping the adaptation information for communication. But it is not clear, (1) if any cross-study standardization of adaptation terminology is being done; (2) if adaptation scenarios not included in the list are to be excluded from the framework (and if not, what the point of the list is at all), or (3) if the framework somehow adds any additional information, other than allowing itself to be applied to multiple sets of data.	The Framework provides a rapid means to assess the incremental benefit of adaptation for these scenarios. As outlined in NCA4, the nomenclature of reactive versus proactive scenarios was designed to acknowledge that additional effort may be needed to reach the full potential of adaptive capacity, and that effort could be cost-effective in the medium-term, if more costly upfront. We added some text to clarify this point.
223	4	10	What is done with these two sectors? Are they excluded, included by not scaled by population, or scaled by population even though they should not be? Given that this is listed in the Limitations section, I suspect the first or third, but I do not know which.	These two infrastructure types are included in the overall results for this sector, but are not adjusted for population. We added a note to clarify in the Appendix
224	4	7	It is not clear what the modeled results for non-endemic counties actually are. It sounds like nothing is removed from them in the baseline, and then any initial Valley fever rates appear in the 1 °C and higher bins, even though these are not caused by climate change.	For currently non-endemic counties, the approach assumes no cases occur, in either the baseline or the future climate. A county must first pass the endemicity screen before climate-induced cases are estimated. In the case of counties that, in the historic baseline, do not meet the endimicity screen, but in future climates do the meet the endimicity screen, all cases are correctly attributed to climate change. Note further that excess cases attributed to climate change are estimated by comparing a modeled future to a modeled baseline, to ensure that we do not incorrectly attribute cases to climate change. There are a few cases where counties do not pass endemicity screens but historic cases have been recorded - those historic cases are effectively ignored in the method, a conservative assumption that follows the reasonable inference that the case would have been contracted in a neighboring county but diagnosed in the non-endemic county. We added a footnote to clarify.
225	4	4	I do not understand what these heat maps are of (and saying that the colors are "determined from the underlying data" is useless). A graphic would be helpful.	We added example figures to illustrate the outputs.
226	4	7	The last box mentions a "weighted average of valuations by impact type", but this weighting and averaging is not discussed in the text.	The weighted average reflects a composite of mortality and morbidity case incidence, which are incident in the same ratios per unit of climate stress. The weighted average valuation is described in the text that immediately precedes the figure (Figure B-10 in the revision). We clarified the text somewhat to connect it to the figure.
227	4	7	Between values of what, for all of what?	We will clarify this wording. tempBin generates a series of annual outputs for all user inputs if there are missing values or if they are not already on an annual basis.
228	4	8	10 °C and 150 cm are below the 99-%ile of the uncertainty in temperature and SLR. These are plausible warming/SLR levels, would be returned by a model like FaIR, and should be able to be handled by the projection.	Thank you. We updated the model to reflect results up to 250cm. Future versions of the model will use a different, continuous approach to estimating SLR impacts and will more directly address this comment.
229	4	10	tempBin is described elsewhere more sensibly as the projection function. I do not think that it performs any "binning".	We edited the text to clarify and describe as "projection."

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
230	4	10	If present_values, cumulative_impacts, and get_reports are not documented/exposed to users, they should not be shown here.	These are deprecated functions and were removed from the tool documentation.
231	4	10	How is the list of valid values for sectorsList retrieved?	The reviewer's question prompted us to add a function that users can call to view the list of sectors.
232	4	10	How are adaptation scenarios requested, and is there a place in the tool where they are listed?	The reviewer's question prompted us to add a function to view the list of sectors and available adapptations. Adaptations are specific to each sector, and also can be determined from the output dataframe. Users do not need to specify adaptation scenarios, which we clarified in the text.
233	4	10	Column names have not been discussed in the documentation. How should users know which they want, or what they mean?	Users can visualize any numeric results. We also added a list of column names in that users can visualize, and added an example in the text.
234	4	11	Although the format of the R package documentation is standard, the documentation approach is not. Function documentation should be descriptive, not procedural. In programming, functions are used to build abstraction barriers, where the implementation details are not the concern of users. When we are told that “the population scenario is joined with” a data frame called “drivers” (lines 2305-2306), it seems like we are being given details that have no bearing on our ability to understand and use the function, and thereby are obfuscating it. We do not care how the function goes about adding rows to the result data frame (lines 2310 – 2314), just what the results data frame is ultimately going to contain.	We appreciate the reviewer's comment and provided extra detail to be as transparent as possible for the review process. We attempted to reduce procedural explanations where they could be more confusing than illuminating.
235	4	General	Why are 2010 and 2090 socioeconomic runs used at all, since this is the process for generating arbitrary socioeconomic runs? We already have all of the drivers defined elsewhere, and any relationship between 2090 and 2010 would presumably be represented as a linear scalar.	We interpolate between 2010 and 2090 socioeconomic runs for certain sectors to account for time-dependent aspects of climate impacts. This level of detail may be too procedural for this documentation, so we simplifed the explanation in this Appendix (now Appendix F).
236	4	General	Use file.path rather than paste to avoid platform-dependence.	The reviewer is correct - we updated both the code and the text.
237	4	General	There are quite a few typos in this section: “reads in and format”, “column contents of”, “(defaults to (`column=NULL`)” (and by the way, I think it defaults to column=“annual_impacts”).	We corrected the typos and incorrect references.
238	4	General	The standard way to describe the data formats in popform are as “wide” and “long”.	We incorporated this phrasing in the text.
239	4	General	A vector graphics form of the plots (e.g., PDF) is generally more useful than PNG.	We agree and made this change.
240	5	4	get_reports() is listed as a function but is not documented in this section with the other three main functions	Thank you. get_reports() is now a deprecated function and we removed it from the documentation.
241	5	4	Possible values for the “sectorsList” argument should be specified in this section. In addition, it is not clear from this function documentation how the user specifies particular adaptation scenarios?	We added a function that users can call to view the list of sectors called "get_sectorInfo()" and added adaptation options to that list. Adaptations are specific to each sector and users do not need to specify them. We clarified in the documentation.
242	5	4	Give temperature year baseline in the “tempfile” argument information. “Baseline scenario” for slrfile argument should be “baseline period”?	We corrected the wording to describe the baseline period.
243	5	4	I am confused why the slrInput table has 0 SLR in 1970 when I thought SLR values were relative to a 2000 baseline?	We provided this as an example of a custom SLR input, but replaced it in the revision with a scenario in line with the reference temperature scenario for clarity.
244	5	4	Give possible value for “column” argument	Users can visualize any numeric results. We added a list of column names that users can visualize.
245	5	4	Specify what unique cell on the heat map represents e.g. sector * model? Sector by year?	We clarified this in the text and added example figures to illustrate the outputs.
246	5	General	I think Chapter 3 is very useful in providing case studies in demonstrating the tool capabilities. If at all possible, I would suggest releasing the code used to do this analysis as a “vignette” included with the R package. This would be extremely valuable in terms of getting users quickly to a point where they can conduct interesting analysis using the tool.	Thank you. Good suggestion. We will consider developing a vignette for the case study for the next version.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
247	4	4	Information on output/result summaries is largely absent. The R tool produces heat maps, but I do not think these are shown in the report. The get_plots function also mentions reporting annual results, but how are these plotted? All of the projection results are reported as tables or line plots. Does the tool help produce those? How about other forms of outputs, like bar and box-and-whisker plots, choropleth maps, histograms, and stacked ribbons?	Good suggestion. We will add a couple of example figures to illustrate the outputs. The get_plots function only creates the two outputs described here. We are considering adding more plotting functions for the future. Users can visualize the dataframe of outputs from tempBin in R using various common R packages (ggplot2, etc.) or import the results into another program (Excel, MATLAB, etc.) We can clarify this here.
248	1	5	As noted above, the linkage between the authors' approach and the peer-reviewed literature on the time-shift method is not clearly conveyed. Recommendation: The authors should ensure their introduction is grounded in the peer-reviewed literature related to the time-shift approach for climate pattern emulation.	We have added several citations to make the literature review more comprehensive
248.1	2	3	"without consideration of when that warming occurred" – this is crucial and needs to be explained better as the results of the framework, but not the framework itself, are directly dependent on the timelines/durations considered during input.	The text has been modified to "Temperature binning aids comparability of independent analyses by using estimates of physical impacts without consideration of when that warming occurred or which scenario was used (other authors have used the nomenclature "time-slice" or "time-shift" for similar analyses) to drive estimations of economic impacts" to make it clearer that the physical climate patterns are analyzed independent of the date from the original GCM calculation, and then used to drive the monetized impact calculations.
249	3	3	Most sectors present only the impact costs as a function of warming level, and the reader is referred to a reference study for the methodology. Many of these reference studies are not open access – my institution did not have access to a number of key references. This makes the calculations relatively opaque – for the urban drainage cost estimate, it is unstated in the entire report what statistical assumptions are made on the evolution of future extreme precipitation, subject to significant uncertainty (see Pendergrass et al 2015 https://agupubs.onlinelibrary.wiley.com/doi/full/10.1002/2015GL065854, Sanderson et al 2019 https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2019GL082362), which is not communicated in the report.	Thank you for pointing this out. The urban drainage paper is now available open access, through PubMed, here: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8477377/ , and all papers used in the underlying sector work, except for the Underwood/asphalt roads paper (which was not sponsored by EPA), are also now published and open access, including those which were indicated to be "in review" in the prior April 2021 draft. We have also added new language on uncertainty with particular focus on difficulties in estimating exteme events.
250	3	5	The report makes it fairly clear how sectoral studies need to be framed to be incorporated into the framework, laying out the need to express damages as a function of warming level. However – I am not convinced that the temperature binning framework is appropriate for all sectors where there are significant non-warming related aspects to future damages which cannot be trivially classed as 'adaptation'. In such cases, non-temperature mitigated factors may be more dominant and where a comprehensive scenario-based assessment might be more appropriate. This is true, for example, for the air quality impacts – where pollutant emissions are held fixed at a single projected estimate of 2040 levels, and temperatures vary throughout the century. This leads to a highly unrealistic scenario – given that pollutant emissions are not independent from CO2 based mitigation. It is particularly concerning because this is the most significant component of the damage function.	The following text was added to address this comment: "For climate impacts on air quality, the tool includes the two future relationships between climate and air quality derived in Nolte et al (2021), one based on a 2011 US emissions inventory and the other based on a 2040 US emissions inventory. The climate mortality penalty for the latter scenario is about half of the penalty for the former scenario. If precursor emissions were to be reduced further, that might further decrease the climate penalty. Nolte et al. (2021) did not consider elevated methane concentrations or changes in transboundary air pollution transport which could also influence the climate penalty." and "• Co-benefits and ancillary benefits and costs of climate policy – This tool only examines the direct impacts of climate change. It does not, for example, estimate the benefits of reducing co-emitted air pollutants such as nitrogen oxides, volatile organic carbons, or particulate matter due to climate policy."
251	3	6	The air quality assessment requires – at least – an additional degree of freedom to represent the level of ambition in clean air policies at the city and regional scales. Co-benefits of greenhouse gas mitigation in terms of reduced pollutant emission should be at least discussed.	See response 250

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
252	3	6	Electrical demand could allow for time-dependent evolution not just of the electricity demand, but also of the generation portfolio and interconnectivity.	The reviewer is correct - the energy sector is particularly difficult to address because climate change and other factors affect demand, while climate change and climate change policy, both, affect supply. We added a new section on uncertainty and called this out specifically as an important uncertainty in these estimates. We also added text to the Appendix discussion of this study.
253	3	6	Labor impacts could allow for a degree of freedom relating to changing working patterns – not just GDP and population.	While this is an interesting idea, until a methodological study is done to examine such a sensitivity analysis, we can't include it in this framework. Language has also been added to highlight this kind of challenge: "Changes in other drivers – Some sectors in this analysis have significant non-climate drivers. For example, changes in land use and forest management could have substantial implications for the climate response of impacts such as wildfires or dust. If the underlying study did not consider such sensitivity analyses, the framework cannot yet consider them."
254	3	6	Dust and wildfire are heavily dependent on land-use assumptions, but there are currently no external drivers which represent land-use (agriculture patterns, forestry and fire management practices).	The following text was added to address this comment: "Changes in other drivers – Some sectors in this analysis have significant non-climate drivers. For example, changes in land use and forest management could have substantial implications for the climate response of impacts such as wildfires or dust. If the underlying study did not consider such sensitivity analyses, the framework cannot yet consider them."
255	3	7	The temperature-binning framework risks systematically ignoring the direct impacts of greenhouse gas mitigation ambition on the sectoral vulnerabilities e.g. renewable energy availability will modulate the electricity robustness parameters, and are themselves a direct function of mitigation ambition – but this linkage is missed from the cost-benefit analysis because only temperature-mitigated impacts are considered. This is, of course, a function of the temperature binning methodological choices, but the limitations should be discussed and the impacts on aggregate results should be potentially considered.	The reviewer is correct that, particularly within the energy sector, mitigation ambition is a key factor in sectoral vulnerabilities. We added language to a new undercertainties section that directly addresses this limitation. The framework is not yet equipped to address the degree of change in enegy supply infrastructure associated with changing mitigation policy ambition.
256	3	7	Similarly, for the air quality cost assessment (the largest single component of the aggregate climate costs), the temperature-binning approach combined with very idealized non-CO2 emissions assumptions means that the assessment omits changes due to transient pollutant emissions in a transient scenario. Given that there is a high correlation between CO2 and non-CO2 pollutants – this is potentially missing a major factor when assessing the benefits of CO2 mitigation (https://www.nature.com/articles/nclimate2009). A sensitivity analysis on the 2040 assumptions for non-CO2 emissions would be a valuable addition.	The direct effects of co-reductions in air pollutants are outside the scope of this kind of analysis: that can be addressed by traditional air quality cost-benefit analysis of a given policy. There will be a 2nd order term where reducing air pollutants reduces the sensitivity to climate: the following text is now included in the documentation: "For climate impacts on air quality, the tool includes the two future relationships between climate and air quality derived in Nolte et al (2021), one based on a 2011 US emissions inventory and the other based on a 2040 US emissions inventory. The climate mortality penalty for the latter scenario is about half of the penalty for the former scenario. If precursor emissions were to be reduced further, that might further decrease the climate penalty. Nolte et al. (2021) did not consider elevated methane concentrations or changes in transboundary air pollution transport which could also influence the climate penalty." and "• Co-benefits and ancillary benefits and costs of climate policy – This tool only examines the direct impacts of climate change. It does not, for example, estimate the benefits of reducing co-emitted air pollutants such as nitrogen oxides, volatile organic carbons, or particulate matter due to climate policy."
257	3	7	The sectoral analysis of extreme temperature related mortality, based on historical relationships, is insufficient for future projections. The inability to represent mortality due to heat extremes in regions which have not experienced heat extremes in the past should be seen as a severe limitation (line 1418), notably because such regions may be expected to have the fewest pre-existing adaptations to extreme heat.	The limitation for the Pacific Northwest is an issue (as has been seen with the recent heatwave), as is the incomplete geographic coverage more generally. However, this documentation has limited scope for comments on the limitations of the underlying studies.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
258	4	1	It would be beneficial in the introduction to better place the TBF within the experience/literature of both damage functions and past attempts to combine multiple sectors into a common assessment. On damage functions, readers can be pointed to Nordhaus & Moffat (2017) (and the Tol papers that they improve upon), since this work has long indexed damages (sometimes sector-specific damages: see, for example, the Stern Review) to climatic temperature. The innovation two differences here is that the focus is entirely on sector-specific damages (although that is only made clear in the examples and it is not clear why top-down estimates would be out-of-scope), and that temperatures are standardized rather than coding exactly what is reported in the underlying studies (although the benefits of this are unclear, given that the damage functions are turned into linearly splines eventually anyway).	We added a sentence citing Nordhaus & Moffat (2017): "These IAMs contain relationships between temperature and damages, with a range of geographic and sectoral resolutions – Nordhaus and Moffat (2017) and Diaz and Moore (2017) recently assessed the damage function representation in these models in the context of the broader literature."
259	4	1	The process of finding where GCMs reach particular temperature thresholds and using these for the purpose of communicating results is also not new. This is the process used by the EU's PESETA to determine their scenarios. Two useful papers for understanding how other researchers have navigated the translation from GCMs to economic impacts are Diaz & Moore (2017) and Ciscar et al. (2019).	We added a sentence citing Nordhaus & Moffat (2017): "These IAMs contain relationships between temperature and damages, with a range of geographic and sectoral resolutions – Nordhaus and Moffat (2017) and Diaz and Moore (2017) recently assessed the damage function representation in these models in the context of the broader literature."
260	4	1	"by degree of warming" and by "custom temperature trajectory" reads like inconsistent goals, and if you believe that temperature trajectories contain any information as trajectories (rather than unordered points), they are inconsistent.	While the impact sectors considered in this framework do not have "memory", the trajectories themselves do have temporal information: while 3 degrees in 2050 and 3 degrees in 2100 will have identical climate characteristics in this framework, the population, GDP, and other aspects of the impacts can differ.
261	4	1	The introduction of socioeconomics and timing is not adequately described—even for its inclusion in a brief overview. As written, it just introduces confusion.	We substantially re-wrote and clarified the discussion of socioeconomics and timing in the revision, including the specific text referenced by the reviewer.
262	4	2	"Binning", as used in the title of the report, is also problematic. Binning means collecting multiple similar things into bins, and has a specific definition in statistical analysis. Specifically, it is not the intention that different temperatures are being grouped together (how the term "temperature binning" is used in climate econometrics work, for example). In this case, I think that the implication is that results from different GCMs are collecting into common bins, indexed by temperatures. Although the term is defined (lines 91-93), I imagine that a better job can be done. That definition describes "synthesizing results" (what does that mean?) of "sectoral impact results" (what is a result of a result?) "by temperature change" (the word 'by' here does not explain how temperature changes are used). It is not difficult to understand the concept, but as the first introduction of the term this leaves much to be desired.	We substantially re-wrote this section to address these comments.
263	4	2	The terms "scenario-based analysis" and "by degree" are not defined. One can guess that scenario-based analysis is meant to refer to the process in lines 154-157, and "by degree" refers to the temperature-binning framework, but not without some confusion.	The text has been modified to "Temperature binning aids comparability of independent analyses by using estimates of physical impacts without consideration of when that warming occurred or which scenario was used (other authors have used the nomenclature "time-slice" or "time-shift" for similar analyses) to drive estimations of economic impacts" to make it clearer that the physical climate patterns are analyzed independent of the date from the original GCM calculation, and then used to drive the monetized impact calculations.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
264	4	5	Second, it is worth noting that the techniques described here are appropriate for only a subset of climate impact studies, and that the frontier of research is moving in a direction that is not conducive to the simplifications made here. The past decade has made clear that (1) speed of approach to temperatures are of primary importance, for adaptation, economic growth effects, and mitigation implications; and (2) global average climate change are not reflective of local heterogeneity, and inequality of local impacts is essential to understand welfare loss. A growing number of studies will incorporate these principles into their analyses, and the TBF will need to make progressively less palatable approximations to incorporate this work.	We disagree. As an analogy: GCMs have become progressively more complex, with higher resolution and representation of more atmospheric, ocean, and carbon cycle processes. And yet, the utility of reduced complexity climate models such as FaIR has not disappeared. The world can use both the complex models that produce rich, complex results, and the simpler models that are tuned or fit to reproduce some bulk properties of these complex models. Similarly, there will continue to be a use for frameworks which can produce rapid estimates of damages, even if the framework does not have all the depth of analyses which can have the luxury of concentrating on a single sector or region and so explore local heterogeneities. Seperately, we would be interested in specific examples of the literature that has demonstrated that damages are more related to the speed of temperature change rather than the absolute magnitude (outside of ecosystem studies looking at the speed of species migration relative to the rate at which their ideal temperature is moving poleward or upward): while it makes sense that adaptation is easier when the rate of change is lower, quantitative studies that produce monetized estimates of these effects are to our knowledge still rare.
265	4	General	The temperature-based approach is highly relevant for policy-makers and the general public, and the TBF formalizes well-established approaches used by the IPCC and European Commission. There is discussion in the report that it will be updated as new sectors are added, but there is a wide gap between these comments and the necessary planning for this to become a living document. In addition to how new sectors can be added (see question 5), there should be a process for improving the basic assumptions of the framework can be improved as new research is developed. What aspects of the framework are “core” immutable assumptions, and which are subject to improvement? For example, the specific baseline years and the current GCMs should be made as easy to change in the future as possible. The the use of 1-degree and 25-cm bins (rather than un-binned sector-specific points) and the particular process of including socioeconomics (as simple multipliers), while discussed at length here, seems like a simplification worth dropping in the future. Identifying these as worthy of reconsideration, and setting up a process to revise them, would help this framework stay relevant as the research evolves.	We agree that integer degree bins are not necessary for the incorporation of new sectors, and have made minor edits to make it clear that while FrEDI generally does use integer degree bins for consistency and ease of incorporation, it is not a requirement.
266	1	5	I have some concerns about the integration of studies with more sophisticated socioeconomic dependencies, which I discuss in my answer to Question 6, below.	See response to comment #156.
267	1	7	This seems reasonable and suitable, provided the uncertainties in the relationships among climate variables (and, potentially, socioeconomics) are clearly addressed, as noted above.	Thank you for this comment. We hope the revisions in this draft address your initial concerns.
268	2	General	The authors should be congratulated for attempting to simplify what are clearly extremely complicated, and still poorly understood aspects of impacts from climate change and the inter-linked socioeconomic drivers. However, simplification comes with its own challenges. With the rather large number of caveats included in the report, and on a complicated subject such as being addressed, and with significant scope for misinterpretation of results that come out of the TBF, the value of the TBF and the report remain quite limited especially when considering dependencies and cascading impacts which are not addressed within the framework.	The FrEDI framework is designed primarily for rapid analysis of economic impacts for custom climate trajectories, removing the need to run the underlying studies. In the revision, we attempted to clarify both the purpose and the key uncertainties that are presented by the approach. We hope that in future revisions we can both quantify the uncertainties and, potentially, say more about cascading impacts.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
269	3	1	The introduction clearly lays out the logic for the framing of climate impact assessments in a temperature-binned framework, and how physical and economic impacts can be represented as functions of national mean warming levels – allowing for expansion in the future. It also lays out the conceptual framework for the value of adaptation across sectors in a logical fashion. It notes the differences between an emissions-scenario based approach and a temperature binned approach.	Thank you for these comments.
270	3	4	Tables are generally quite clear, and the report does a good job of defining the dependencies for each sectoral impact.	Thank you for these comments.
271	3	6	The methodology for the incorporation of socioeconomic data into the framework is clear, but the degrees of freedom allowed to represent socioeconomic drivers could be expanded in some sectors:	Thank you for these comments. We've addressed your specific comments in subsequent responses.
272	3	General	The technical report is well written, and provides a valuable construct in which to integrate sectoral climate impact analyses into a common framework. Though some of my responses have been critical in places, I do find the report as a whole to be a valuable product which will be useful for the framing and communication of climate cost projections.	Thank you for these comments. We appreciate your review of the work and your constructive feedback.
273	4	1	The introduction provides an accessible overview of the goals of the framework on the contents of the report. This framework has the potential to provide a much-needed basis for collaboration and communication around climate risks, and makes a number of strong assumptions and simplifications to do so. Some of these are already highlighted elsewhere in the report, and others I mention below, and they include assuming linear socioeconomic relationships, ignoring the speed of warming, and making very limited use of uncertainty quantification. These simplifications should be clearly stated, but they can also be justified. By simplifying the representation of climate damages, the TBF can easily incorporate a wide range of studies. In particular, assuming linear socioeconomic relationships allows studies with only static 2010 and 2090 results to be used; the speed of warming allows all temperature results to come from a single emissions scenario; and uncertainty information remains under-reported and incompletely understood.	Thank you for these comments. We have responded to the individual points made here in subsequent specific comments.
274	4	5	My responses on the conceptual basis for the TBF are discussed in question 3, so I will focus here on specific steps taken. These are clear in general, but in some cases the decision process is opaque (discussed below and in question 6).	Thank you, we address specific issues raised in subsequent responses.
275	4	7	The procedural approach of applying the damage functions to temperature scenarios and applying valuation or scalings to the results is very standard (e.g., throughout the IAMs). It is reasonable and reasonably well-described. I will focus on some more specific decisions made in this process.	Thank you for this comment.
276	4	10	While the information on the “temperature binning” part of the framework is sufficient, this is unlikely to be a source of confusion and this part of the approach is already pervasive amongst existing climate economics work, both within academia and as directed at policy-makers. However, the codification of “scaling” as a generic approach is new, and except for simple cases, not well explained (see question 6).	We have significantly re-worked the section on "scaling" to better describe the process. See revised section 2.
277	4	11	I hope that my comments above do not suggest that I want the perfect to be the enemy of the good (particularly since this is an area where knowledge is growing rapidly, so no one can say what the perfect framework should look like). The TBF is an excellent basis for a wide range of impacts to be incorporated into CIRA and into other policy-communication.	We appreciate your support for this project and thank you for your constructive feedback.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
278	5	2	Yes – the report is very well written and describes a substantial amount of necessary technical detail in an accessible way. I think Chapter 3 is particularly valuable in providing examples of exactly how the methodology can be used to investigate a variety of questions that might be of interest. I also think the provision of the important technical details on individual sectors in an Appendix is a good idea – these details are documented for users that need to refer to them, but more casual readers can get a sense of the overall tool without being overwhelmed with this detail.	Thank you, we have added substantially to the appendices to provide additional details.
279	5	5	Mostly this is well described. I particularly liked the systematic and detailed sector-specific information provided in Appendix A. I thought the flow diagrams of the pre-processing steps and what is incorporated into the tool in this section were quite effective.	Thank you for the feedback.
280	5	General	This is an accessible and thorough documentation of a fairly complex analysis of a number of climate change damage sectors to produce a useful tool for users to understand the effects of changing emissions pathways across the US. My main suggestions for improvement involve better organization of Chapter 2, distinguishing between the processing of impact studies done for the tool and how users might use it (more details in response to Question 4), and better defining what exactly is included in "the framework".	Thank you for this comment. We have attempted to rework Chapter 2 and better frame the process in response to your comments. See specific comments for more detailed responses.
281	5	11	I think the figures and table used through the report are mostly clear and effective, other than some exceptions noted above.	Thank you for this comment.
282	ReRev 1	1	Yes, by and large, the new section adequately explains the analytic approach. One point that required some thinking to figure out is that the bars in Figure 1 are simply relating the projections presented in Kopp et al 2016 for comparison – they are not reflecting any uncertainty in the analysis system. This confusion is created because the authors refer to the 'Mann et al 2009 data set' as opposed to the 'Kopp et al 2016 GMSL projections driven by the Mann et al 2009 temperture data set.' I would suggest clarifying this.	We have updated the text to reflect this.
283	ReRev 1	1	In addition, the example SLR scenario shown in Figure 2 – while by no means implausible – is not one that could be produced using the semi-empirical model specified. Are the authors allowing exogenous specification of sea-level rise scenarios, such as the one shown in Figure 2?	The tool allows the user to supply an exogenous sea level rise scenario. We have updated the text to reflect this.
284	ReRev 1	1	State "posterior distribution calculated with the Mann et al 2009 temperature data set" rather than "Mann et al 2009 posterior distributions"	We have updated the text to reflect this.
285	ReRev 1	1	State "sea level rise calculated with the Mann et al 2009 calibration" instead of "sea level rise from the Mann et al 2009 data set"	We have updated the text to reflect this.
286	ReRev 1	1	In the figure caption, state "median projections from Kopp et al 2016, using the calibration to the Mann et al 2009 temperature data set".	We have updated the text to reflect this.
287	ReRev 1	2	This section has largely addressed my comments. I do note that, given the absence of treatment of uncertainty, the approach for custom scenarios will in general yield sea level projections that range from just below the RCP 2.6 scenario to about the RCP 8.5 scenario – i.e., always bounded between Sweet et al 2017 Low and Sweet et al 2017 Intermediate scenarios. This excludes scenarios of potential economic relevance, such as the low confidence SSP5-8.5 projections presented in the AR6.	We agree with the reviewer and added text to describe future plans for exploring uncertainty and calibrating to AR6 values.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
288	ReRev 2	1	The new section mostly well describes the new approach to estimating sea-level rise damages. I find it much easier to follow the logic, both of the semi-empirical sea-level model and the interpolation of damages. A few small suggestions would be to:	Thank you, we address specific issues raised in subsequent responses.
289	ReRev 2	1	1. Mention why Mann et al. 2009 values from the Kopp et al paper were used instead of other possibilities from that paper	Either set of parameters can be used in this model. We added text on addressing uncertainty in the text.
290	ReRev 2	1	2. Mention the distinction between global sea-level rise (which is being calculated in the semi-empirical model here) and US national or regional SLR, which might be quite different	We added text to the new method description that describes our method to translating GMSL to local SLR, which reflects multiple site-specific factors including land subsidence/uplift rates.
291	ReRev 2	2	Mostly yes – the approach described here is much more intuitive and easier to understand how damages are being calculated. I am still somewhat confused though by the continued use of the 250cm, 200cm etc SLR scenarios in Figure 2. How do these relate to the semi-empirical SLR model described on the previous page? I figure these are the scenarios used to drive the impact model, but in that case what exact role is the semi-empirical model playing in the impact estimates? Spelling out these connections more explicitly might be helpful.	The six SLR scenarios from Sweet et al. are used in the underlying sector models for the two SLR-driven sectors. They define six pairs of GMSL and impact trajectories. The new method compares the custom SLR scenario (in the example here, that is the GCAM reference scenario as estimated semi-empirically in FrEDI) to the six SLR trajectories and finds the two scenarios from Sweet et al. that bracket the custom scenario in each year, in terms of sea level rise heights. Using that information, and where exactly the custom scenario falls in between the two bracketing Sweet et al. scenarios (e.g. is it directly in between, 25% of the distance between, etc.) to then interpolate damages for the custom scenario.
292	ReRev3	1	First, this is a sensible approach to model sea level rise, and it is a significant improvement over the instantaneous response used previously. It is worth noting that this model is more complicated than used in standard IAMs, which I believe use a first-order differential equation, rather than the second-order approach used here. But it is not very complicated, and Dr. Kopp's work on SLR is certain to be preferable from a scientific standpoint. However, there are a few mistakes in the equation given here, as I understand it. I have put these into the specific comments list, but they significantly change the dynamics of the model. I suspect that they are incorrect only in the text and not as simulated. The sea level damage method left me confused, both about what changed and the exactly the implementation of the new method. I think that a simple equation would be useful here. I suspect that the process is something like: Damage(t) = (SLR(t) - SLR_below(t)) / (SLR_above(t) – SLR_below(t)) Damage_above(t) + (SLR_above(t) – SLR(t)) / (SLR_above(t) – SLR_below(t)) Damage_below(t) Are all of these values (SLR(t), SLR_*(t), and Damage_*(t)) 11-year averages, or only some of them? Are all of these values available in every year, or are they further linear interpolations (as it looks from the graph)?	Thank you, we have addressed the issues with the translation function in specific comments. We have also added a formula to the new binning description. The values are available annually.
293	ReRev3	1	Just a typesetting issue, I believe that the (t) in dTe(t)/dt and the (2000) in the Te(2000) are not meant to be subscripts. They are just time indexes, as in the top equation.	We have updated the text to reflect this.
294	ReRev3	1	This expression for dTe(t)/dt will cause Te to increase indefinitely, as long as T > Te(2000). This is not the case with Kopp et al.'s original expression. They would have the expression: dTe(t)/dt = (T(t) – Te(t)) / tau1, which is a normal relaxation expression.	We have updated the text to reflect this.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
295	ReRev3	1	This would have phi decrease linearly forever (e.g., going negative). In Kopp et al., it's written: d phi(t)/dt = -phi / tau2. That is, phi drops toward 0 asymptotically.	We have updated the text to reflect this.
296	ReRev3	1	The text says that the previous approach did not control for the year when the SLR is reached, but it was not clear how the new method does this until I wrote the equation above. I believe that the new method calculates damages off by only comparing it with the damages under the scenarios for the same year. This should be made clearer, and I am not certain that it is right to say that this is a way to "control" for the year.	Your interpretation is correct--the new method uses impact estimates from the specified year to interpolate impacts. We have softened the language in this description to say that previously we did not "account" for the timing of SLR so it is not suggesting we now "control" for the year.
297	ReRev3	1	What does the text mean when it says that under the method you "identify damages associated with the input sea level rise in a given year"? How are they identified? (Is it the result of the 11-year rolling average, or is it the result of the process in the next sentences?) What is the "input" sea level rise? (Is that the custom scenario's SLR?)	Thank you for this feedback--we have clarified this section to better describe the new SLR binning method. We calculate the damages associated with the custom sea level rise trajectory by interpolating between the damages in the two underlying sea level rise scenarios that have sea level rise heights closest to the input custom scenario in the given year.
298	ReRev3	2	These changes do address my concern about the instantaneous response of SLR to temperature. I did have other comments related to sea level rise and damages (what should be included in each adaptation scenario, the consequences of an 11-year average for SLR, and questions about units in some of the graphs). However, this is certainly the most crucial issue, and a welcome change. Sea level rise is notoriously fat-tailed, because of the unknown timing and consequences of ice shelf collapse and the poverty of data on current changes in the ice mass. In this situation, the median may not be an appropriate summary statistic to use for damage projections. In particular, using the median rather than the mean seems likely to underestimate mean damages.	We updated the text with a paragraph on addressing the uncertainty that surrounds sea level rise as well as updating the figure to reflect the range in parameteric uncertainty. We use the median values from the data sets provided in the supplement of Kopp et al., 2016.
299	ReRev4	1	The explanation of the empirical SLR estimates are quite clear. The empirical sea level rise approach is based on existing literature and is reasonable for the application. Limitations to this empirical approach should however be noted - primarily that long term sea-level rise associated ice-cap melt with future emissions is not parameterised. For pre-2100 applications, this is probably reasonable. It would not be accurate for multi-century applications though, which should be noted. That said, the uncertainty ranges for the parameters and SLR projections should be presented - central estimates are of limited use, and imply unjustified confidence in the projected range.	We updated the text with a paragraph on addressing uncertainty as well as updating the figure to reflect the range in parameteric uncertainty.

ID	Reviewer	Charge Question	Reviewer Comment	Response to Comments
300	ReRev4	1	The interpolated cost estimate seems a little less justified - it is not grounded in existing literature, and the justification for interpolating between existing scenarios is not self evident. In particular - in the illustration in Figure 2, the sample SLR scenario has more rapid sea level rise than any of the scenarios considered in the original studies. In reality, this would require a more rapid adaptation of infrastructure, and adaptation timescales would cause greater damages than in a more slowly rising scenario. However -in the interpolation scheme presented, the end-of-century damages are identical in the orange scenario and in the dotted black scenario in Figure 2, despite the fact that the latter has undergone a much more rapid sea level rise. As such, I am not convinced that this cost mapping is appropriate - even as a first order estimate.	We believe that concerns about the interpolation are likely attributable to the sample SLR scenario, which was meant to be illustrative only, is not a realistic outcome, and is not based on the new reduced complexity SLR estimation method. We believe the interpolation ought not be a general concern, for two reasons. First, the scenarios used are not necessarily independent, they differ only in terms of the rate of SLR. Second, careful examination of Figure 2 indicates that the damages per unit of SLR are closely consistent across scenarios, meaning that interpolation does not have a large effect on the results. It is reasonable to conclude, however, that future work could better establish a relationship between damages and adaptation costs and the rate of SLR than the literature currently relied on in the Framework. Refinements of the underlying model to better capture delays in planned/proactive adaptation, however, which can take time to implement after a flooding risk is recognized for a variety of reasons (see the Chambwera citation in main text, for example), could change this conclusion - so this is an area we will continue to watch as new models and results are introduced in the Framework.
301	ReRev4	2	The quantitative basis for the empirical SLR estimates is sound, but the lack of any uncertainty quantifification is not justified. The range of plausible SLR associated with any given temperature pathway should be presented in a public comminication document in order to meaningfully assess future risk.	We updated the text with a paragraph on addressing uncertainty as well as updating the figure to reflect the range in parameteric uncertainty.
302	ReRev4	2	The basis for the interpolation between scenarios is less clear. It is not self-evident that interpolating between costs in independent scenarios is reasonable when there are rate-dependencies in the cost function (i.e. faster sea-level rise will result in higher damages).	We believe that concerns about the interpolation are likely attributable to the sample SLR scenario, which was meant to be illustrative only, is not a realistic outcome, and is not based on the new reduced complexity SLR estimation method. We believe the interpolation ought not be a general concern, for two reasons. First, the scenarios used are not necessarily independent, they differ only in terms of the rate of SLR. Second, careful examination of Figure 2 indicates that the damages per unit of SLR are closely consistent across scenarios, meaning that interpolation does not have a large effect on the results. It is reasonable to conclude, however, that future work could better establish a relationship between damages and adaptation costs and the rate of SLR than the literature currently relied on in the Framework. Refinements of the underlying model to better capture delays in planned/proactive adaptation, however, which can take time to implement after a flooding risk is recognized for a variety of reasons (see the Chambwera citation in main text, for example), could change this conclusion - so this is an area we will continue to watch as new models and results are introduced in the Framework.