

Interactive comment on “Attribution of the Australian bushfire risk to anthropogenic climate change” by Geert Jan van Oldenborgh et al.

Geert Jan van Oldenborgh et al.

oldenborgh@knmi.nl

Received and published: 16 July 2020

This paper describes an attribution analysis of a number of factors that are known to either contribute to or reflect wildfire risk using observational data, observationally constrained data products (reanalyses), and a collection of CMIP5 model simulations. It also considers the impacts of internal climate variability as reflected in the large-scale modes of variability that influence the Australian climate, and it includes a discussion of vulnerability and exposure factors associated with the impacts of the summer 2019/2020 wildfires.

I found the paper frustrating to read and evaluate. One clear impression is that the authors were in a terrible hurry, producing text that often appears not to have been

C1

carefully proofread, not thinking carefully about how to describe their methods in a clear way, not always justifying methodological choices, not justifying choices of data products or evaluating those products with a sufficiently critical eye, and attempting to be overly comprehensive. Reading the paper is a bit like being forced to “drink from a firehose” — there are so many details and so many small aspects that can be criticized, that is difficult to know exactly how and what to criticize in a review. The fact that all code is being made available doesn’t really reassure me very much. Readers who want to understand what was done, sufficiently so that the work can be replicated, shouldn’t be placed in a position of having to read code but rather, should be provided with explanations in the paper that are clear enough so that they can develop and implement their own code.

We thank the reviewer for their comprehensive feedback on our paper. Indeed, the article was written in a hurry in order to make the results available in a timely fashion during the aftermath of the fires. For example, the study’s findings were used one week after the discussion paper was published in discussions of the various state commissions and Royal Commission at the Bureau of Meteorology on the link between climate change and bush fire risk. During the analysis we focused on making sure the results were correct, which took time away from creating a more carefully written text. We would like to take the opportunity of the revisions to make the article more readable and appreciate the reviewer’s detailed comments to that end.

Some specific comments:

1-14 *The abstract does not mention the long section on vulnerability and exposure factors, and there is no reference to vulnerability and exposure in the title. Does that section really belong in the paper?*

Thank you for highlighting this. We have incorporated a reference to the vulnerability and exposure context in the abstract. Generally, many studies conducted by the World Weather Attribution group include a section on vulnerability and

C2

exposure, as some of the co-authors work in this field and help us put the physical results (often focused on physical processes) into context of impacts on the ground.

- 16 *The very first sentence of the paper starts by being sloppy in the way in which Australian station data are characterized. The word “homogeneous” has a very clear and well understood meaning in the context of observational data products (i.e., meaning that observations have been carefully evaluated and adjusted to ensure that they are free of artefacts resulting from changes in instrumentation, instrument siting, instrument housing, observing and reporting practices, etc., etc.), and surely the claim here is not that Australian station data is homogeneous in that sense. Clearly, avoiding the obvious inhomogeneity due to the lack of proper instrument shielding early in Australian instrumental record is necessary, but we shouldn’t just accept that all of the subsequent record is homogeneous.*

This is a good point. We intended to refer to the introduction of Stevenson screens as standardized measurement equipment, so we have revised this to ‘standardized’.

- 26-27 *What is the source of this estimate? Is it possible to have any confidence in that number or the range that is given?*

The source is a Red Cross report, accessible via hyperlink. We have changed this to an explicit footnote, so that it becomes clear how to access the source. This report information is collected by the Red Cross, based on official state reports of damaged infrastructure. The loss of wildlife number is an estimate from Prof. Chris Dickman at the University of Sydney (<https://www.sydney.edu.au/news-opinion/news/2020/01/08/australian-bushfires-more-than-one-billion-animals-impacted.html>), communicated as a ‘conservative’ estimate based on taking the typical wildlife loss after habitat destruction based on case studies and scaling it to the burned area of this year’s bushfires.

C3

- 27-29 *Again, what are the sources?*

We are not aware of peer-reviewed studies yet quantifying economic and health impacts from these fires, but we have added a reference to insurance claims and a general reference to health impacts from wild fires.

- Figure 1** *Is there a URL and a date for where this image was obtained?*

We made the image ourselves using MODIS satellite data (obtained from <https://firms.modaps.eosdis.nasa.gov/>) for the fire radiative power and Dinerstein et al. (2017) data for the forest cover. References to the data sources have been added.

- 93 *I imagine daily maximum temperatures are meant. There are many instances in the paper where a second reading of the words, just to see if they connect logically, would have helped enormously. There are also a large number of run-on sentences in the paper that are difficult for readers to parse and understand.*

We clarified that this refers to daily maximum temperatures. We will revise the entire paper for readability.

- 102 *This subsection is entitled “Event definition”, but it doesn’t talk specifically about event definition at all. I think what is needed is a clear statement that the event of interest will be defined using the FWI. This section gives some justification for doing that by considering the relationship between FWI and area burned, but event definitions per se are not discussed in this subsection.*

We have substantially revised this subsection to clarify that it mainly deals with general parameters of the event definition (such as fire season, spatial domain, way of aggregating), while specific event definition details for each driver variable (temperature, precipitation, fire indices) are given at the beginning of their respective sections later on in the paper.

C4

Figure 2 caption *Please tell me what is meant by a “one-sided confidence interval about zero”. I assume you mean the interval from -1 to the expected 95th quantile for the correlation coefficient under the null hypothesis that the correlation is zero. If this is correct, then it would be better to call this the 5% significance critical value for a one-sided test of the null hypothesis that the correlation is zero against the alternative hypothesis that the correlation is positive.*

We have revised the caption, including the wording suggested by the reviewer, which was helpful and greatly appreciated.

129 *Often, acronyms like ASF20C appear before they are defined.*

We thank the reviewer for flagging this. We have scanned the whole text to remove instances of this problem.

137-152 *Some careful justification for the distributional choices would seem to be in order. These distributions emerge in statistical extreme value theory as limiting distributions under idealized conditions, where the limit is taken either as block length increases without bound in the case of the GEV, or as the exceedance threshold increases without bound in the case of the GPD. Given the way the data are processed, we are likely a long way from being able to be satisfied that the actual distributions are well approximated by these limiting distributions. Indeed, it seems likely that the relative quality of the fit will diminish as you go deeper into the tail, even if quantile plots look to be ok. In particular, one should be worried about extrapolating beyond the available data. Some aspects of this are discussed later in the paper, but those limitations don't really seem to prevent the authors from referring to values that appear to correspond to very long return periods in some instances. In the case of precipitation deficit, any of a number of possible candidate distributions could presumably be considered if using as much as 30% of sample values. These would have different deep tail characteristics, affecting calculations of probability ratios, but might not be discernably better or*

C5

worse than the GPD based on standard diagnostics of the fit. So how does one proceed in a careful way take this source of structural uncertainty into account? It might be as important as the structural uncertainty represented by the spread between models.

We agree that this is an important point and have spent quite some time in the past investigating which distributions fit the data well enough for common extremes. In this case we decided to use the following for the three types of extremes we study.

Heat extremes The highest weekly average of the year is a block maximum. Even though the number of independent blocks in a summer is small, as the reviewer mentions, a GEV distribution fits the observational data well, see Fig. 4. The uncertainties are indeed rather large. More convincing evidence is the agreement in models with large numbers of simulated years, all of which show no deviations from the GEV fits in the return time plots up to the return times of a few hundred years that we use. CESM (4000 years of data) shows a thinner tail for the highest values, whereas CanESM (3500 years of data) shows a fatter tail above 1000 yr return times, EC-Earth and GFDL-ESM2M (2000 years each) are described well by the GEV. We assume that the very hottest events involve nonlinear physics, but apparently the models do not agree which sign the deviation from the GEV should be. However, we evaluate the distribution at 100 years for the current climate, so we use the distribution to interpolate, not extrapolate, and the differences in the PDFs are mainly due to the (considerable) differences in the parameters of the function and not to deviations from the assumed form of the GEV. Because of the large uncertainties in the tails of these distributions with negative shape parameters we take care to never claim much accuracy of the numbers coming out of the fits anyway.

Drought extremes We verified that these are not described well by either a normal

C6

distribution nor by a gamma distribution (indeed, we have not found a single case where a gamma distribution described precipitation data well enough). The GPD appears to fit the data well enough, partly due to the extra fit parameter that gives more flexibility. The low number of data points and hence strong dependence on the threshold in the observational case does limit its usefulness. We have attempted to take this uncertainty into account by using both the 20% and 30% threshold results. In this case, model experiments with larger numbers of data points show that the GPD fits the data well up to the return times of thousands of years, which are actually sampled in the climate model cases. Note that we evaluate the distribution at a return time of only around ten years.

The functional form of the dependence of the covariate is just a convenience that makes sure we never get negative values, there is no theoretical justification for assuming Clausius-Clapeyron scaling for droughts. However, as the changes with the covariate are small (in fact, mostly compatible with zero), the exponent is close to a linear function anyway and this choice allowed us to use a ready-made routine. The main problem was technically fitting the function in a parameter space where only a very small area satisfied the constraints of no negative precipitation and we had to change to a different version of the simplex minimisation routine we used to do the log-likelihood maximisation in order for the fit routine to find the maximum (the GSL version could not do it).

Fire Weather Index We again use block maxima, like in the heat extremes analysis. The GEV fits all models well in a return time plot (equivalent to a Q-Q plot), but with very different parameters. We again conclude that the differences between the distributions of the different models are dominated by the parameter differences and not by deviations from the GEV.

As for droughts, there is no theoretical justification for the functional form of the dependence on the covariate (smoothed global mean temperature)

C7

beyond the requirement of a positive-definite distribution (which would be violated with a simple linear dependence that is used in many other articles).

We have added a summary of these arguments into the article and supplementary material.

160 *Why 4-years and not some other degree of smoothing? Exactly how is the smoothing done, and how is time referenced to the smoothed values? For example, if using a 4-year running mean, which year is the value associated with in covariate dependent functions?*

The 4 years are a previously developed compromise (King et al., 2015) between the typical time scales of ENSO and more decadal-scale variability and a well-defined value at the end of the record. The smoothing is a running mean. Year 3 of 4 is used as covariate. Critically, uncertainties from these choices are typically very small. We have added clarifications on all of these points.

165-172 *Choices for how the GEV and GPD distributions are parameterized should be justified and carefully argued, not just stated. For precipitation, exponential scaling might make sense at the upper end of the precipitation distribution, but why would I consider that to be reasonable at the lower end of the distribution, and why, in that case, should the scale parameter be linked to GMST? Building in something that scales like Clausius-Clapeyron might not be the best idea for the dry end of the precipitation distribution.*

See the answer to the comment on lines 137–152 above.

187-189 *Is it obvious that this is the best way to proceed? If the analysis was literally performed as described here, the effective block size for the models would be 5- or 10-times the block size used for the observations. That means that for the models, the block maxima used to fit GEV distributions would sample a much deeper part of the tail than is possible with the observations since the distributions*

C8

for the model output would have been fitted to what are effectively 5-year or 10-year blocks rather than 1-year blocks as for the observations. How then, can I make sense of differences in parameter estimates between fits to observed and simulated parameter estimates.

Indeed it would have been better if both fits would have been done with the same block size. However, this is not possible due to the paucity of data in the reanalysis and the functional form of the model data not agreeing with a GEV function all the way down to very low return times in one climate model that is used in the final synthesis (CanESM2, the others have distributions that are compatible with a GEV for all return times). The description in the methods section has been updated to mention that it only refers to one model. The method described is the one that gave a good fit in the region of the return time of interest, 31 years. The main goal of the method, to establish lower bounds on the probability ratio, can therefore be achieved with it.

Regarding the test on the fit parameters that the reviewer comments on, indeed the return times for which these are defined differs between observations and models when excluding the shortest ones (below 5 yr). However, we are interested in return times, around 31 years in the current climate, which correspond to longer return times in the climate of 1900. The assumption we make is that the data in this range are described well by the GEV we fitted to. We do not use the range of return times below 5 years that are excluded in the CanESM2 fit. The parameter values can therefore still be compared to the 'observational' ones.

191-192 *I think this is all that is said about bias correction in the paper except for another brief mention at line 442, but surely this is important and should be discussed (and defended) in some detail. Exactly what was done, and how does this avoid overusing the observational data?*

All the bias correction we do is to evaluate the extreme value function at the same return time as the return time of the observational analysis. As the reviewer

C9

comments, it is very important not to overuse the observational data so we do not attempt to correct the PDF further than this one-parameter correction (see also the reply to Tim Palmer's comment). The usual minimal bias correction in climate model analysis is a correction of the mean, which is implicit in all IPCC change plots. We found that the effective inclusion of some effects of the biases in variability and skewness by evaluating at the same return time rather than just shifting the mean gives more realistic results, especially in heat extremes. These have distributions with a negative shape parameter in which the return time is very sensitive to the return value for which it is evaluated. Evaluating at the same return time as the observations removes this sensitivity. We have updated the bias correction description to better reflect this.

200 *Exactly what do you mean by the χ^2/dof statistic (what is calculated, and what is the basis for the interpretation given to this statistic)?*

The goal is to determine whether the apparent uncertainty across models is just due to internal variability or whether it is indeed indicative of actual model structural differences. To that end, we simply compare the spread of the model results with the spread expected from their estimate of uncertainty due to internal variability Δx , $\chi^2 = \sum [(x - \bar{x})/(\Delta x)]^2$. This should be roughly equal to the number of degrees of freedom, $N - 1$. If it is larger, we interpret this as evidence that we have to take into account another source of uncertainty, the model spread (which is part of the model uncertainty). The same method is used by Aurélien Ribes et al. (2020), although he agrees with us that it is a rough estimate as any discrepancy could still be due to chance. However, it is the best we can do given the information available. We have extended the discussion in the methods section with the explicit formula and more details.

240 *For each observational product, the paper should draw attention to the key limitations that would affect the analysis in this paper. For example, although Stevenson screens begin to be used in 1910, there could be many other reasons to*

C10

be concerned about the homogeneity of temperature observations, such as variations in station coverage over time (e.g., spatial sampling in 1910 would undoubtedly have been different than in the 1970s). Also, the paper should make a clear distinction between observational products on the one hand, and observationally constrained products (re-analyses) on the other. The latter are clearly non-homogeneous, with inhomogeneity due to changes in data sources, quality and quantity over time being of particular concern in the southern hemisphere where the observational constraint is much weaker. Ensemble reanalysis products, such as the 20th century reanalysis may be able to provide information about the strength of the observational constraint and how it varies in space and time (if the spread between ensemble members is large, the constraint is obviously weak or non-existent; if the spread is small, one has further work to do to determine if it is small because the analysis is being effectively constrained by the observations or whether this is coming about for another reason). Further, it should be noted that surface variables are often not very well constrained in reanalyses. The classification of variables by strength of observational constraint that is given in Appendix A of the Kalnay et al paper describing the original NCEP 40-year reanalysis (BAMS, 1996, [https://doi.org/10.1175/1520-0477\(1996\)077<0437:TNYRP>2.0.CO;2](https://doi.org/10.1175/1520-0477(1996)077<0437:TNYRP>2.0.CO;2)) still largely holds and should be considered.

This is certainly a valid point and we attempted to provide more justification for our dataset choices in response to this and other reviews received.

We have added the number of stations entering the analyses to the information already given. Beyond that, the only information on the inhomogeneities would come from the daily gridded ACORN dataset that has been corrected for this, but unfortunately it is not yet available.

We have indeed attempted to distinguish between observational datasets and reanalyses. JRA-55 is the only dataset with adequate coverage in the southern

C11

hemisphere before 1979. We have added a section discussing the quality of the longer reanalyses, showing that the ECMWF products ERA-20C and CERA-20C indeed are unusable in this region. The 20CRv3 reanalysis does surprisingly well despite the limited constraints in the early part of the record, but is not used in the analysis either given the availability of good observational datasets.

We disagree that the classification in Kalnay et al. (1996) is still useful. Modern models are good enough to assimilate near-surface temperature observations and hence reanalyses like JRA-55 and ERA5 have become much better at simulating these variables than NCEP/NCAR R1. This obviously does not hold for long-term reanalyses (ERA-20C, CERA-20C, 20CRv3), of which only 20CRv3 approaches a reasonable simulation of the weekly maximum temperature extremes in southeastern Australia.

Figure 3 Use the same vertical scale on both panels (or better yet, plot the two time-series on the same graph).

At the request of Antje Weisheimer we have expanded Fig. 3 to five panels so showing all lines in one plot is not an option anymore. We have replotted all these panels on the same vertical scale (attached).

251-255 A number of reanalysis products are mentioned here, but the paper also uses others (e.g., ERA-5).

This section only refers to the reanalyses that were used in the heat attribution. ERA5 is not used there due to its brevity and availability of better (longer) alternative datasets. For the FWI, where more variables are required, we do rely on ERA5, despite its brevity, as explained in the relevant section. The results from ERA5 for the hottest week of the year are compatible with the other observational and reliable reanalysis datasets. There is a higher trend over the period since 1979, but the difference with the long-term trends from the century-scale datasets is not statistically significant due to the larger uncertainties resulting

C12

from the lower number of data points. We have included this in text. We do not show this in the summary plots because the period over which the trend is estimated is so different from the observational datasets and models.

240-269 *An overview of the strategy for using the different observational and reanalysis products would be useful. This would demonstrate that there is some overarching reasoning that knits the selection of products together and that has informed the choice of products. I have to say that the choices are really confusing, both for reanalyses and for the observational products. For example, GMST is apparently from GISTEMP (mentioned at line 123, but not in this observational data section), but the gridded global surface temperature dataset that is used is Berkeley Earth (line 242), and other well studied and documented global gridded temperature data products such as HadCRUT4 are not mentioned at all. Why these particular choices? For the gridded products, the infilling strategy and error models, which vary between choices, are presumably important considerations, particularly in the southern hemisphere and especially when considering a relatively small land area in the southern hemisphere that is sandwiched between ocean to the east and a very dry, sparsely observed continent to the west.*

We will add more structured text to explain the dataset choices, as this has been a consistent concern across reviews.

The heat events are defined as the maximum of 7-day maximum temperature, so these can only be extracted from a dataset with daily resolution. Neither GISTEMP nor HadCRUT4 has daily resolution, these are monthly datasets with very coarse resolution designed to study global and large-scale temperature changes. The only place where these are useful is as the covariate in the extreme value function that acts as a proxy for the radiative forcing of global warming. For this we take the (smoothed) GISTEMP global mean temperature, as HadCRUT4 underestimates the warming trend due to its undersampling of the polar regions (although the differences are so small that in practice it makes very little differ-

C13

ence).

As we mention in the text, there is only one global long-term daily maximum temperature dataset that we are aware of, Berkeley Earth. We checked this dataset against local Australian datasets of daily maximum temperature, of which the AWAP dataset is a full analysis, which interpolates the station data in space. The ACORN station dataset does not, but provides a useful qualitative check. Reanalyses of maximum temperature are somewhat less reliable but again serve as useful checks for the observational datasets. As the reviewer mentioned, some are not reliable in the southern hemisphere before the satellite era.

We revert to monthly and daily mean data to check whether the January 1939 event is real or not, but prefer to use datasets with high spatial resolution for this. The $5 \times 5^\circ$ resolution of HadCRUT4 is not sufficient to resolve this relatively small-scale event.

We have clarified this discussion in the text.

270 *I find it very surprising that the entire observational discussion for TX7x, including results from AWAP and mention of one of the reanalyses, is limited to only 6 lines of text. Statistical model fitting results are shown in Figure 4, but are really not discussed in any meaningful way — and Figure 4 itself is not explained in a way that most readers would be able to understand. Specifically, cumulative frequency distributions for 1900 and 2019 are shown, but there is no explanation in the text or in the figure caption explaining how the points that are shown are derived from the observations. Evidently observations are adjusted to particular years using smoothed GMST values for those years to make adjustments via the fitted distribution. Shouldn't one be concerned that this could induce some circularity, particularly if one of the intents of the figure is to illustrate the fit of the statistical model to the observations? Results from one reanalysis are mentioned, but silence concerning other reanalyses begs a question about whether they did not “tell a similar story” — do they tell a similar story?*

C14

The same method has been used numerous times in previously published literature going back to van Oldenborgh et al. (2015) and King et al. (2015). We might have assumed that most readers would be familiar with this approach and have thus omitted some of the details. Indeed the observations have been shifted to the climate of 2019 and 1900 using the functional dependence on the smoothed global mean surface temperature. We have added references to earlier work using these methods to the methods section and a more complete explanation of the figure to the caption.

We are not sure how this can induce some circularity: the 4-year smoothed annual mean global temperature is not affected in any meaningful way by the local highest one-week maximum temperature in southeastern Australia.

We did not want to discuss the 20CRv3 results quantitatively as it does not assimilate near-surface temperature observations and therefore cannot be expected to reproduce these reliably. We added a sentence discussing this: ‘As the 20CRv3 reanalysis does not assimilate near-surface temperatures we do not expect the quantitative results for TX7x to reflect reality, but note that qualitatively they agree well with the other estimates of the observations.’ The other long reanalyses have been rejected as unreliable at this stage.

- 281 *See my comment concerning lines 187-189. What explains the apparently much narrower uncertainty bounds on the climate model-based parameter estimates as compared to the model-based estimates? Is the explanation that the model-based analysis actually uses annual blocks rather than blocks constructed by pooling data for a particular simulated “year” across ensemble members (which is literally what lines 187-189 appear to say)? In this case, samples of annual maxima are 5- or 10-times as large as from observations, which, all else being equal, should result in confidence intervals that are about $5^{-0.5}$ or $10^{-0.5}$ as wide as for observations (i.e., $\sim 45\%$ or $\sim 32\%$ as wide, respectively). But this interpretation also doesn’t seem quite right because the model confidence intervals*

C15

seem narrower than these expectations.

First, we assume the reviewer means ‘observation-based’ instead of ‘model-based’ in his second sentence.

In this case, all fits are based on annual maxima, larger blocks are only used for the Fire Weather Index (FWI) and Monthly Severity Rating (MSR). Indeed, the models have more ensemble members than the one realisation of the real world, leading to smaller uncertainties due to natural variability. The number of ensemble members per climate model is listed in Table 1 and varies from 15 to 50 for the models with daily maximum temperatures. They also have differing record lengths, from 55 to 170 years compared to the 110 years of the AWAP dataset. For models with reasonable variability, which we have checked in the model validation section, this gives uncertainties due to natural variability of about $\sqrt{110/65}/\sqrt{15} \approx 34\%$ (HadGEM3-A) to $\sqrt{110/70}/\sqrt{50} \approx 18\%$ (CanESM2) of the observational estimates, which agrees by eye with the length of the bars in Figure 5.

We have added a clarifying sentence to this end.

I have many largely similar comments about sections 4 and 5 that I won’t repeat here. Hopefully the message that the paper needs to document the work and justify choices and interpretations much more carefully has come across.

We have updated these sections with similar clarifications as the heat event attribution section.

Regarding Section 6 — a very strong conclusion is drawn on lines 565-566, but it is not obvious to me that the strong quantitative evidence and supporting modelling experiments that would be required for such statement has really been presented. Quantitative evidence seems to be restricted largely to estimates of correlation coefficients which, if considered as simple regression diagnostics (i.e, focusing on r^2 rather than r), would correspond to explained variance amounts of the order of 5-15%.

C16

We estimated the contribution of the IOD and SAM from the scatter plots in figures 19 and 20 as the fraction of the anomaly of 2019 described by the (purple) linear trend line from zero DMI-ENSO or SAM to the value observed in 2019, which in both cases is about one third. Mathematically the explained fraction is

$$\frac{dP}{dI} \Delta I / \Delta P \approx \frac{1}{3}, \quad (1)$$

with P the July–December precipitation and I the DMI-ENSO or SAM index. However, we are deliberately unspecific about exact numbers given the large approximations in the value of the relationship and the linearity of it as the reviewer points out. In response, we weakened the statement in the paper to ‘More quantitative estimates will require further analysis and dedicated model experiments, as the linearity of the relationship between these indices and the regional climate is not verifiable from observations alone.’

References

- Dinerstein, E., Olson, D., Joshi, A., Vynne, C., Burgess, N. D., Wikramanayake, E., Hahn, N., Palminteri, S., Hedao, P., Noss, R., Hansen, M., Locke, H., Ellis, E. C., Jones, B., Barber, C. V., Hayes, R., Kormos, C., Martin, V., Crist, E., Sechrest, W., Price, L., Baillie, J. E. M., Weeden, D., Suckling, K., Davis, C., Sizer, N., Moore, R., Thau, D., Birch, T., Potapov, P., Turubanova, S., Tyukavina, A., de Souza, N., Pinteá, L., Brito, J. C., Llewellyn, O. A., Miller, A. G., Patzelt, A., Ghazanfar, S. A., Timberlake, J., Klöser, H., Shennan-Farpón, Y., Kindt, R., Lillesø, J.-P. B., van Breugel, P., Graudal, L., Voge, M., Al-Shammari, K. F., and Saleem, M.: An Ecoregion-Based Approach to Protecting Half the Terrestrial Realm, *BioScience*, 67, 534–545, <https://doi.org/10.1093/biosci/bix014>, 2017.
- Kalnay, E., Kanamitsu, M., Kistler, R., Collins, W., Deaver, D., Gandin, L., Iredell, M., Saha, S., White, G., Woollen, J., Zhu, Y., Leetma, A., Reynolds, R., Chelliah, M., Ebisuzaki, W., Higgins, W., Janowiak, J., Mo, K. C., Ropelewski, C., Wang, J., and Jenne, R.: The

C17

- NCEP/NCAR 40-year reanalysis project, *Bull. Amer. Met. Soc.*, 77, 437–471, [https://doi.org/10.1175/1520-0477\(1996\)077<0437:TNYRP>2.0.CO;2](https://doi.org/10.1175/1520-0477(1996)077<0437:TNYRP>2.0.CO;2), 1996.
- King, A. D., van Oldenborgh, G. J., Karoly, D. J., Lewis, S. C., and Cullen, H.: Attribution of the record high Central England temperature of 2014 to anthropogenic influences, *Environmental Research Letters*, 10, 054 002, <https://doi.org/10.1088/1748-9326/10/5/054002>, 2015.
- Ribes, A., Thao, S., and Cattiaux, J.: Describing the Relationship between a Weather Event and Climate Change: A New Statistical Approach, *Journal of Climate*, 33, 6297–6314, <https://doi.org/10.1175/JCLI-D-19-0217.1>, 2020.
- van Oldenborgh, G. J., Haarsma, R., De Vries, H., and Allen, M. R.: Cold Extremes in North America vs. Mild Weather in Europe: The Winter of 2013–14 in the Context of a Warming World, *Bulletin of the American Meteorological Society*, 96, 707–714, <https://doi.org/10.1175/BAMS-D-14-00036.1>, 2015.

Interactive comment on *Nat. Hazards Earth Syst. Sci. Discuss.*, <https://doi.org/10.5194/nhess-2020-69>, 2020.

C18

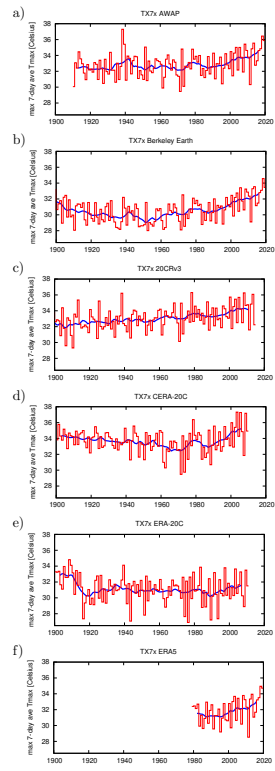


Fig. 1.