Interactive comment on "A crop yield change emulator for use in GCAM and similar models: Persephone v1.0" *by* Abigail Snyder et al.

Anonymous Referee #1

Received and published: 12 November 2018

General comments: The paper by Snyder et al. presents a crop yield change emulator for use in multi-sector economic models. The emulator is intended to support studies of the feedback loop among socioeconomics, Earth system changes, and crop yield changes, by providing multi-sector economic models with rapidly generated yield responses to Earth system changes, including some quantification of crop response uncertainty. Compared to using computationally expensive global gridded crop models to estimate yield changes, the emulator has the advantage of being more computationally efficient, allowing to run large ensembles of simulations necessary to explore future response options.

The work presented in the paper is highly relevant because it promises to allow new agricultural impact studies that do not only consider the impact of Earth system

changes on the agricultural system, but also how responses by society to cope with these changes may feed back to the Earth system. This bidirectional feedback has often been missing in previous impact assessments, or could only be analyzed for a small number of simulations due to its high computational requirements. While the emulator is developed for a specific multi-sector economic model, GCAM, it can also be used by other similar-in-scale models. The software implementing the emulator is published alongside the paper.

The paper is well written. The authors clearly describe the data processing and derivation of the emulator response functions (section 2) and their evaluation (section 3), finishing with a short example of an application of the emulator and some final conclusions and discussion. Although I have to admit that I am not very familiar with Bayesian statistics and found some of the corresponding terminology confusing I was still able to follow what was done.

I have two main concerns with the paper: the first refers to the C3MP dataset used to derive the emulator response functions, and the second refers to how these response functions are going to be used. The C3MP dataset is a large set of 99 CTW sensitivity tests carried out by a number of site-based crop models covering a range of different crops at a total of 1135 sites, of which data from 575 sites are actually used in this study. The climate change signal used in these sensitivity tests is completely synthetic since it consists of applying a temporally uniform temperature offset or precipitation multiplier to a historical baseline weather timeseries. However, in reality, climate change is not constant over time. For example, precipitation might increase during part of the year, while decreasing during other times. I am not convinced that the constant CTW perturbation experiments are equivalent to using a more realistic climate timeseries, and I would suggest that the authors either show in the paper or at least point the reader to other literature showing that this is a valid approach. Otherwise, the authors risk that there is already a bias in the data used to train the emulator which would propagate to the emulated yields.

My second concern is that response functions for some of the 25 production groups are based on a very small number of sites, in the most extreme case only two sites. I cannot help but wonder how representative these results really are, keeping in mind that each production group represents one crop-irrigation-latitude combination, with latitudes only distinguished into extended tropics and mid-latitudes. In addition, these sites are not only used to derive a mean yield response, but also an estimate of response uncertainty. The high and low response functions are supposed to represent site responses at the mean plus/minus one standard deviation level but how meaningful are these estimates based on such a small sample size? The emulator response functions are only derived for two regions (extended tropics and mid-latitudes) but I wonder at what level of spatial detail they will be used later. In Figure 8 and 9 of the paper, the authors take spatial patterns of climate change from the HadGEM2-ES model and use their response functions to derive corresponding patterns of yield change. These patterns are shown, but not evaluated in any way. I would suggest that the authors compare their derived patterns of yield change to simulations of yield change from global gridded crop models for the same climate data. After all, the intent of the emulator is to replace simulations by global gridded crop models. Such crop yield simulations for the HadGEM2-ES RCP8.5 scenario used in Figure 8 and 9 of the paper are, for example, available from a number of crop models and for a number of different crops within the ISIMIP data archive at https://esg.pik-potsdam.de/projects/isimip/ Such a comparison would help to address both of my concerns voiced above.

Specific comments:

- Page 2, ll. 25 − 28: You state that previous emulators were restricted to emulating yield change under RCP scenarios. I am aware of at least two other global crop yield change emulators derived from crop model simulations that are applicable to any future climate scenario: Oyebamiji et al. (2015) and Ostberg et al. (2018). It might be useful to contrast them to the work presented in this paper.

- Page 4, starting with line 11, and Figure 1: You outline three use cases for the Perse-

phone yield emulator, but none of this is actually done in this paper. So I'm not sure if the Methods section is the right place for this.

- Page 6, section 2.2: The Copernicus guidelines request that datasets should be cited with a reference in the reference list. Is there a reference for the C3MP dataset?

- Page 6, ll. 8 − 14: This part does not refer to the setup of C3MP or to your processing of the C3MP data. Instead, it refers to how climate data is pre-processed before use with the finished emulator. It should probably be moved to section 4.

- Page 7, ll. 7 − 9: Unless I am mistaken, the global gridded crop models within AgMIP also conducted the 99 CTW sensitivity tests. They should offer a much better global coverage. Would it be worthwhile adding some discussion of why this paper used the site-based results instead of the global simulations?

- Page 8, ll. 19 − 20: Are the 8 different functional forms documented anywhere? Are these the same as are used in the emulator?

- Page 11, ll. 9 − 10: Is the value of b0 constrained between -0.02 and 0.02 or is it 0.02%? The following sentence suggests that it is 0.02%.

- Page 13, ll. 1 − 8: Given the very small sample sizes of about 1/3 of the production groups the 84.135th and 15.865th percentiles do not seem very meaningful

- Section 3.2: It seems to me that this is rather a test whether the crop models used in training feature these important relationships. Or would you say it is possible that these relationships are present in the models, but missing in the emulator? Since the evaluation is positive I guess it means that the relationships are present in the models and retained in the emulator.

- Page 20, l. 4 − page 21, l. 1: Climate change can affect both the start and the length of the growing season. Is this accounted for or is the same growing season used under climate change as during the reference period?

- Page 21, ll. 8 – 10: On the one hand you talk about passing CTW changes for regions into the emulator, on the other hand you mention a gridded map of yield changes. So are the yield changes at the same spatial resolution as the climate data or is there a difference between resolutions (region versus grid)? Please clarify.

- Page 21, ll. 17 – 18, and Figure 9: Looking at Figure 9, it seems to me that most regions show a positive yield change under the high response function, not just "a few regions". Also, I think that in the bottom row of Figure 9 the maps for mean and high response are swapped. The bottom right map (which should be the high response) looks identical to the map for Maize in Figure 8 (which shows the mean response).

- Page 23, ll. 6 – 8: Here, you emphasize the rapid evaluation time of the response functions relative to a global gridded crop model, but I think you should really try to show that the emulator response is actually comparable to what you would get using a global gridded crop model. This step is missing in the paper.

- Page 24, ll. 8 – 10: Given that the response functions are only derived for two latitudinal bands I would say that they cannot really be used to characterize the range of uncertainty within national or multi-national units (unless the respective unit is covered by C3MP sites).

- Page 24: In the paragraph on caveats, I would suggest to add discussion of potential biases arising out of way the CTW experiments are set up. Another source of uncertainty results from the fact that crop model simulations generally omit adaptation options such as changing sowing dates or switching to different cultivars. This is done for simplicity and comparability but is not realistic considering that agriculture is a highly managed system, and potentially creates another bias in the crop model simulations of yield change that are used to train the emulator.

Technical corrections:

- There are a number of typing errors throughout the paper (mostly missing characters,

sometimes missing words). Please check during the revision.

- The greyscale lines in Figure 6, 8 and 9 are hard to see at all, let alone distinguish the different shades of grey.

References:

Ostberg, Sebastian, Jacob Schewe, Katelin Childers, and Katja Frieler. 2018. "Changes in Crop Yields and Their Variability at Different Levels of Global Warming." Earth System Dynamics 9 (2): 479–96. doi:10.5194/esd-9-479-2018.

Oyebamiji, Oluwole K, Neil R Edwards, Philip B Holden, Paul H Garthwaite, Sibyll Schaphoff, and Dieter Gerten. 2015. "Emulating Global Climate Change Impacts on Crop Yields." Statistical Modelling 15 (6): 499–525. doi:10.1177/1471082X14568248.