

Interactive comment on “An investigation of carbon cycle dynamics since the Last Glacial Maximum: Complex interactions between the terrestrial biosphere, weathering, ocean alkalinity, and CO₂ radiative warming in an Earth system model of intermediate complexity” by C. T. Simmons et al.

C. T. Simmons et al.

christophesimmons@gmail.com

Received and published: 10 June 2016

Please find below our detailed responses to the reviewers' very helpful comments, which we believe will vastly improve the manuscript. The original text of the reviewers' comments are included in this response, with dashes ("———") used to separate our response from the original text of each review.

[Printer-friendly version](#)

[Discussion paper](#)



Responses to Reviewer 1 Anonymous Referee #1 Received and published: 12 April 2016
General comments: This paper investigates how ocean and land carbon cycle respond to the deglacial climate changes using an Earth system model of intermediate complexity. The novelty of this paper is to evaluate the effects of the climate variability and the associated carbonate compensation on the atmospheric CO₂ from multiple sets of deglacial transient simulation. They also show the importance of CO₂-driven warming for the reconstruction of ocean circulation and the oceanic carbon release on the deglaciation. The paper argues that deglacial warming accelerates Atlantic meridional overturning through the sea-ice retreat at both hemispheres and then the replacement of low-alkalinity deep water to the surface contributes to the CO₂ release to the atmosphere. This result provides new information for the deglacial climate system. Overall, this paper helps understanding the question, how the climate system drives the glacial-interglacial atmospheric pCO₂ change, although the results of multiple sensitivity experiments cannot reproduce the full deglacial rise in atmospheric CO₂. I recommend that it would be accepted with minor revisions addressing the concern below.

Specific comments:

Specific comments: The authors argue that the deglacial warming increases zonal wind speeds in the subtropics and thus dynamical upwelling (page 16, l. 11). Why does the warming enhance zonal winds at the tropics in this model? Since this wind-upwelling response is a key in this paper to explain the CO₂ release from the tropical ocean to the atmosphere, it would be helpful to show the spatial pattern of physical anomaly (winds or vertical upwelling). —————

We can add a panel to Fig. 4 to show the increased wind stress in the subtropics. We had previously included a qualifying sentence that indicated that a better atmosphere model is needed to confirm this effect (this sentence was trimmed from the manuscript before submission to reduce the word count, but we can include it again). Basically, the warming of the tropics increases the pole-to-equator temperature gradient, which then

[Printer-friendly version](#)[Discussion paper](#)

affects the thermal wind perturbation (the “wind feedback” described in Weaver et al. 2001), which is added to the NCEP reanalysis winds to stimulate a change in the wind field as the model climate evolves. In our simulations, the perturbation provided by the thermal wind field between the LGM and the interglacial is greatest in the subtropics, but a better atmosphere model with more sophisticated internal wind dynamics would certainly be better for evaluating how a denser, more strongly CO₂-fertilized biosphere modifies the global wind field.

———— Furthermore, it may be valuable to show the delta 14C signal to confirm the ventilation changes. For example, the difference in delta 14 C between surface and deep ocean changes (ventilation age) is used as an index of ocean stabilization in paleo proxy field. Does the response of modeled ventilation age also support the enhanced ventilation? —————

Reviewer 2 has also requested to see these fields, so we plan to show the delta C-14 signal in different ocean basins, both vertical and horizontal profiles and hovmöller diagrams (where appropriate), so that the ocean ventilation changes are more apparent.

———— The configuration of sedimentation outflow/riverine inflow and total carbon conservation is a little unclear. The model is assumed to accumulate (sediment) all of calcite at depth upper 1450m. In this case, does the model also treat the burial of organic carbon in sediments at these depths? If so, does the model conserve total nutrient in the ocean in the long-term experiments? In addition, I suppose that the sediment model needs the deposition flux of detrital material as well as organic carbon and CaCO₃ flux for the calculation of early diagenesis. Did you give a spatially uniformed flux of detritus or the spatiotemporal pattern to the model? The detail description of deposition fluxes is helpful to follow your experiments and understand how to treat the sedimentation rate in the long-term transient simulations. —————

The NPZD biogeochemistry (Schmittner et al. 2003) and sediments of the UVic 2.9 model is rather simple. The evolution of the biogeochemistry in the oceans is spa-

[Printer-friendly version](#)[Discussion paper](#)

tiotemporal, as is the falling detritus (much of the organic carbon content of which is respired according to a temperature-dependent parameterization). With regard to the hard-tissue pump, as the calcite falls as detritus through the water column, it dissolves with an e-folding depth of 3500 m (as described in Section 2.1, paragraph 2). Thus, while the dissolution of falling shells is prescribed, the dissolution of sediments themselves are a function of the respired CO₂ content of overlying waters (influenced by alkalinity and CCD changes) and the particulate organic carbon in the sediments themselves (i.e., the fraction which does not respire in the water column). The sediments only exist below 1450 m depth, so there is no sediment accumulation in shallower regions. We are willing to include these and more details about the sediment model in the discussion of the results if this would help clarify the important contributors to the long-term sediment evolution in the model's transient simulations.

———— At the end of Conclusions (page 34, l. 16), the authors argue the effect of the tropical temperature increase on ocean dynamics and vegetation carbon. Furthermore, I think that the change in soil carbon reservoir also affects the atmospheric CO₂ concentration. In fact, the soil carbon is decreased by around 30 GtC in CO₂rad CA than in FC CA (Figure 1c and d). Is this because warming enhances decomposition of soil carbon? How large does the soil carbon change contribute to the atmospheric CO₂ rise? —————

While the soil carbon can be diagnosed in the current figures by subtracting the vegetation carbon (Fig 1d) from the total terrestrial carbon (Fig. 1c), it is difficult for the reader to make these distinctions, so we will add a separate soil-carbon time series. This is a very good point that needs to be emphasized more in both Section 3.2 (on terrestrial carbon) and in the conclusions, that ocean warming and greater ventilation are only indirect contributors to the atmospheric carbon budget. Indeed, part of the difference between the CO₂rad C and FC CA simulations is due to a decrease in soil carbon, as respiration rates are greater globally in the CO₂rad case. Furthermore, by the early interglacial period, some biomes are destabilized by the warmer world where

[Printer-friendly version](#)[Discussion paper](#)

CO₂ fertilization remains low. The vegetation die-back leads to the respiration of vegetation and soil carbon, the latter being reduced by less litter input in grid cells with less vegetation, and this contributes to some of the difference between FC CA and CO₂rad CA simulations during the early Holocene. However, because total global vegetation carbon is greater in CO₂rad CA than in FC CA by the end of the simulation period, greater soil respiration is the most important contributor to the net change in terrestrial carbon storage. Furthermore, as the solubility of CO₂ in the oceans is markedly lower in CO₂rad CA compared to FC CA, and the oceans in CO₂rad CA are better-ventilated, the ocean in CO₂rad CA absorbs less atmospheric carbon and thus has slightly less ocean carbon storage than FC CA (Fig. 1b). This basically translates to much of the excess respired soil carbon remaining in the atmosphere.

Responses to Reviewer 2

The manuscript by Simmons et al. is devoted to an important issue of climate-carbon cycle interactions during the last 23 thousand years. Mechanisms responsible for the 100-ppm CO₂ increase during deglaciation have been identified during the last two decades. Still, while we know the main processes behind the CO₂ increase, their particular role and the strength differ from model to model. In particular, differences among the box, 2-D and 3-D models of the ocean biogeochemistry are particularly striking.

The experiments done by Simmons et al. are of especial importance as the UVIC model has 3-dimensional ocean model. The experimental design of their study is quite complex, and it is not that easy to follow the logic of experiments in the discussion section. My main suggestion is to restructure the results part, especially the marine section. Another major remark is to include more 2-D vertical plots of marine biogeochemistry, which would justify the results described in the text. In the current version, most of plots are timeseries or surface plots, which do not fully exploit an advantage of 3-d models over the box model.

[Printer-friendly version](#)[Discussion paper](#)

Specific comments The title is unusually long. In fact, it is more an abstract than a title: “An investigation of carbon cycle dynamics since the Last Glacial Maximum: Complex interactions between the terrestrial biosphere, weathering, ocean alkalinity, and CO2 radiative warming in an Earth system model of intermediate complexity”. I strongly recommend reducing it, for example, to “An investigation of carbon cycle dynamics since the Last Glacial Maximum using an Earth system model of intermediate complexity”. Besides, the interactive terms in the current title are subjects of different categories: terrestrial biosphere (model component), weathering (process), alkalinity (variable), CO2 warming (process/forcing). It is awkward to mix them into one list/group. ———

We agree that the title is a bit long and confusing, and we will change it to the suggested title, “An investigation of carbon cycle dynamics since the Last Glacial Maximum using an Earth system model of intermediate complexity,” as this was indeed the title for earlier conference presentations which concentrated on some of these simulations.

————— The structure suffers from the same mixture of different categories: 3.1 Atmospheric CO2 - component 3.2 Terrestrial carbon - component 3.3 Physical and Dynamics ocean changes - processes 3.4 Alkalinity response to ocean ventilation – variable 3.5 Sensitivity to Weathering and Carbonate Compensation - process 3.6 The Alkalinity Response to Holocene Terrestrial Uptake - variable Sections 3.1 and 3.2 are relatively easy to read and perceive, while sections 3.3 – 3.5 are very difficult to read. The main reason, in my view, is that the authors try to focus on particular components/processes based on sensitivity studies, which are too different in terms of processes and their effect on the carbon cycle and climate. For example, all FC experiments have rather low CO2 and, consequently, smaller radiative warming. All PC experiments have much warmer climate, which affects all biogeochemical fluxes. CO2rad CA is the third group of experiments, which differ from the previous two categories of runs.

My suggestion is to restructure the results in accordance the experimental setup in the Table 1. The line of result presentation could be as follows:

[Printer-friendly version](#)[Discussion paper](#)

Section 3.1. FC experiments: no radiative warming. - 1st: Changes in the CO₂/carbon dynamics in the CA experiment - 2nd: Sensitivity experiments: effects of higher/lower weathering on CO₂, changes in the ocean DIC, terrestrial carbon (if any)

Section 3.2 CO₂ rad experiments: prescribed radiative warming, but interactive atmospheric CO₂. - 1st: Effect of radiative change on the ocean/atmospheric circulation, surface climate - 2nd: Effects of weathering changes on CO₂, ocean carbon, terrestrial carbon

Section 3.3. PC experiments: prescribed radiative warming and atmospheric CO₂. - 1st: Changes in the carbon dynamics in the CA experiment - 2nd: Effects of high weathering changes on the ocean carbon —————

These are helpful suggestions regarding the layout of the manuscript, where presently each subsection of the result section focuses on different processes and variables (each dominated by a significant finding) rather than a more systematic analysis of the simulations themselves. However, at the same time, we do see value in placing the results of simulations with different transient forcing together in the same figure and discussing them together. For instance, it is more interesting to compare the changes in meridional overturning circulation between the CO₂rad and FC simulations than it is to compare the MOC differences between the each FC simulation. Similarly, the difference in terrestrial carbon between the CO₂rad simulations and the PC simulations are more interesting than comparing the minor differences in terrestrial carbon between the different CO₂rad simulations. As the experimental design for the simulations in Table 1 intends an inter-comparison of simulations with different forcings, the paper's most significant findings may be more difficult to identify in the text if each set of simulations is described individually, and the paper would likely need to be expanded to include more discussion. This would be difficult, as it is already rather lengthy.

As the reviewer indicates that Section 3.1 (Atmosphere) and Section 3.2 (Terrestrial biosphere) are straightforward and easy to read, we suggest maintaining these sec-

[Printer-friendly version](#)[Discussion paper](#)

tions as-is, but with the addition of a paragraph at the beginning of these two subsections describing the main results of the groups of simulations individually (FC, CO2rad, PC) before launching into the inter-comparison of the different simulations. Also, to improve readability, we propose inserting the material from current Section 3.6 (terrestrial-alkalinity feedbacks) into Section 3.2, and thus having three subsections to Section 3.2: (1) terrestrial biosphere results description and inter-comparison, (2) terrestrial-alkalinity feedback, and (3) sensitivity to the distribution of Antarctic ice shelves (the reviewer suggests making this a separate subsection in a later comment). Then, following the logic of describing first the atmosphere (component) and then the terrestrial biosphere (component), we would combine Section 3.3-3.5 into one subsection: the oceans (component), with subsections based on processes (1) physical, (2) chemical (including sediment processes). Each of these subsections could then be in the format that the reviewer suggest (first presenting FC simulations, then CO2rad, and finally PC where relevant, followed by an inter-comparison). We propose these changes to the layout as a way to rationalize the presentation of the results in a more logical and cohesive manner, while at the same time maintaining the emphasis on the significant results that derive from the inter-comparison of the simulations.

———— Changes in the climate and ocean circulation - very important outcomes of the study —should be discussed in more details and supported by better figures. The meridional overturning in Atlantic should be also shown as a 2D plot (depth/lat) as it is the main advantage of the 3D ocean model over box models. What is a vertical distribution of water masses at LGM vs 15.5 kyr BP? The D14C data could be also shown at the 2D plot (depth/lat). It is an added value to the time series of changes in these quantities. —————

We would certainly be willing to include additional 2D figures that provide more information about the ocean state, including the proposed depth/latitude plots at the LGM vs 15.5 kyr BP for certain simulations. Reviewer 1 also requested to see vertical Delta C-14 changes on an x-y grid, which could also be included in the new figure.

[Printer-friendly version](#)[Discussion paper](#)

———— Sections 2.2-2.3: The ocean volume change at LGM due to sea level drop (roughly 3%)— was it accounted in the LGM and transient simulations? If not, what is the possible effect of sea level change on the atmospheric CO₂? —————

The volume and bathymetry of the ocean is equivalent to the present-day configuration, and to make this clearer, in the first paragraph of Section 2.1, we will add volume to the list of qualifiers, i.e., “In its present configuration, the model’s ocean regime is defined by unchanging present-day bathymetry, volume and sea level, and thus some important features of LGM and deglacial topography, such as continental shelves above sea level, are not featured in the simulations discussed here.” In addition to specifying “volume” in the results section as one of the deglacial changes not captured in our simulations, we will make qualifying statements in the text that reinforce the fact that sea-level processes (e.g., shallow water sedimentation, coral reefs, shelf carbonate weathering, tidal mixing, terrestrial storage, etc.) are not included. That said, sea level change processes are probably most important starting in the early Holocene, whereas many of our most significant results are obtained for the first ten thousand years of deglaciation.

———— p. 9, l. 7-8: if the ice sheet decay does not produce a freshwater flux into the ocean, what is a reason for the saw-tooth output of the model (eg Fig. 1a,c,d)? It is not discussed in the paper. —————

The saw-tooth pattern of the results (particularly apparent in fast-response variables such as the terrestrial carbon and atmospheric CO₂) is largely due to the rapid retreat of ice sheets. The model only sees the ice retreat once every thousand years, so each thousand years there is an abrupt transition (for ex., 17500 B.C., 16500 B.C., etc.). This causes, for example, abrupt environmental changes in formerly ice-covered regions (and, occasionally, ocean-adjacent regions) every thousand years. Perhaps the most obvious by-product of this feature of the model is the rapid expansion of terrestrial vegetation (mostly grasses initially) over newly-exposed terrain once every thousand years. This leads to a jump in terrestrial carbon variables and a correspond-

[Printer-friendly version](#)[Discussion paper](#)

ing decrease in atmospheric CO₂. We shall add a sentence in the results section (Section 3.2 concerning terrestrial carbon) and the following sentence to the last paragraph of Section 2.1: “In addition, the abrupt transition every 1000 years of the model’s prescribed ice sheets leads to a saw-tooth output for certain fast-response variables.”

——— p.16, l. 17: could you comment on the effect of an absence of peat/permafrost module in the terrestrial biosphere model? ———

As mentioned in the methods section (Section 2.1), the version of the UVic model used in this study neither includes glacial permafrost and peatlands or other forms of passive carbon storage. The thawing of glacial permafrost and replacement of glacial and deglacial peatlands by other biomes should lead to a net release of carbon to the atmosphere, whereas the expansion of peatlands from the early and mid-Holocene would lead to a net uptake of carbon by the terrestrial biosphere. While we explored passive high-latitude carbon storage in Simmons et al. (2015), our model does not have a peatland component. For this reason, in Section 3.6, we used the PC simulation, which simulates the effects of increasing (deglacial) then decreasing (early Holocene) atmospheric CO₂ and oceanic pCO₂ without reference to the sources of these changes (whether they be caused by peatlands, permafrost, volcanoes, or a combination of these changing features). Then, upon freeing the carbon cycle at the mid-Holocene, we can investigate the evolution of the ocean chemistry in response to these changes. For this reason, we think that moving Section 3.6 to the terrestrial biosphere discussion would be more appropriate. We will also reemphasize the drawbacks of not having permafrost and peatlands in the model into Section 3.2, toward the end of the section.

——— Section 3.3: This section is very important part of the result discussion, but it needs to be reorganized, otherwise it is extremely difficult to read it. See major comments above. ———

We will be happy to reorganize this section as a new subsection (3.3.1) under dynamical ocean changes (also see the above comments).

[Printer-friendly version](#)[Discussion paper](#)

———— p. 18, l. 6: “vegetation” should be “vegetation biomass” —————

We will make the suggested change.

———— p. 19, l. 9: “these simulations lack of freshwater fluxes. . .” – perhaps, without “of”? Section 3.4 (p.22): l.18-21: Discussion of the effect of circulation changes on DIC and carbonate ion concentration (FC HW and CO₂ rad HW) could move into the new section 3.2 (see comment above). 2-D plots of changes in DIC and carbonate ion (Atlantic vs Pacific) are needed to justify discussion of spatial differences in distribution of these species (eg “reduced DIC storage in the deep Atlantic ocean” – how could readers see it without explicit map of 2-D vertical profiles of DIC storage in the Atlantic?). The most interesting is to see at what depth the DIC concentration changed. —————

As per the reviewers request, we will create a new figure with DIC changes (latitudinal with depth) for each ocean basin. It may also be advantageous to show a hovmöller diagram for DIC changes in individual ocean basins at specified depths for certain simulations (i.e., FC HW and Co₂rad HW shown together).

———— Why the authors picked up the HW and not LW experiments (eg p. 24, 1st para)? Any rationale for this decision? The LW experiments lead to higher CO₂ at the end because the ocean alkalinity is reduced. —————

In the history of writing this paper, the HW experiments were completed first, with the LW experiments being added later as sensitivity simulations to illustrate certain ideas. Thus, while we chose to focus on the HW experiments, a similar analysis could also be done between FC LW and CO₂rad LW. The contrast between the FC HW and CO₂rad HW simulations is particularly interesting, however, with respect to ocean chemistry, as the carbonate compensation in the FC HW simulation would support a deepening of the CCD through the deglacial period, whereas in the CO₂rad HW simulation, carbonate compensation would shoal the CCD after ~15 kyr BP. However, carbonate compensation leading to CCD shoaling is predominant in both the FC LW and CO₂rad LW simulations, so there is less of an alkalinity change contrast between

[Printer-friendly version](#)

[Discussion paper](#)



the two simulations as there is between CO2rad HW and FC HW. Furthermore, both the exposure of shelf carbonates and higher silicate weathering of glacial moraines during the mid-deglacial is expected to support a higher net weathering rate during this period, hence why the CO2rad HW and FC HW simulations were contrasted in the discussion of deglacial alkalinity changes (Page 24, 1st paragraph). The LW rate would be more characteristic of the Holocene or early deglacial. We will add a qualifier to indicate our logic to the paragraph contrasting the CO2rad HW and FC HW simulations.

———— P.25, l. 9: “sensitivity to weathering”: What does it mean “another important factor”? Weathering experiment was already discussed in the previous section. l.20. “increase in calcifiers” – could you add a 2-d plot of calcification rate to justify this statement? —————

This sentence makes reference to the difference in the atmospheric CO2 concentrations between the FC LW and FC HW simulations by the end of both simulations (~15 ppm). We agree that this discussion is probably better to be integrated in the alkalinity description, and thus we plan to reorganize this discussion as part with the new subsection on ocean chemistry. With regard to the calcification rate, we can include this in a new figure, but in order to obtain this output (which is not part of the model output variables currently available), we will have to redo these simulations.

———— p. 26, l. 7: what are “the other runs”? Be explicit. —————

“Each of the other simulations” will be added to the text.

———— p. 27, l. 10: What is “this simulation” – FC CA (the last mentioned above) or FC HW? —————

The FC HW simulation.

———— p.28, l.16-23: “Greater ocean ventilation” etc. – again, this statement should be justified by 2-D plot (depth/lat) of circulation changes Section 3.6: p.29, l.20-23, p.30

—————

We can add a streamfunction figure in addition to the previously-mentioned D14C cross section of the Atlantic to address this concern, but ocean ventilation can be understood by referencing a time series Fig. 5 c-d (which I suggest referencing in this location to make that clearer).

———— Section 3.6: p.29, l.20-23, p.30, l.1-13: I do not understand why prescribed concentration simulation is the best illustration of the carbonate compensation effect proposed by Broecker et al. In my view, a discussion of effect of carbonate compensation on atmospheric CO₂ requires interactive (FC-type) CO₂ simulation. Prescribed atmospheric CO₂ changes (first down and then up) are mirrored in the carbonate ion concentration, so I have a trouble with understanding causality in this experiment. The FC experiments in the next para are more insightful, and they basically show no desired effect of CO₂ growth during the late Holocene (l. 18-20). Experiments with more extensive ice shelves might be useful, but they also do not show a sustained growth of CO₂ from 6 to 0 ka. I miss this point in the conclusions section. —————

The prescribed carbon (PC) simulation is useful here because it forces atmospheric CO₂ and the pCO₂ of the oceans to decrease during the early Holocene. None of the other free carbon simulations produced the same decrease in atmospheric CO₂ during this critical time period (the early Holocene), so a more realistic simulation is needed in order to determine the influence of terrestrial uptake on long-term ocean chemistry changes. Freeing the PC simulation's carbon cycle at 8 kyr BP allows us to model how carbonate compensation evolves in response to an extraction of CO₂ from the atmosphere and ocean between 11 kyr BP and 8 kyr BP. What we essentially found is that the effect on alkalinity during the mid-late Holocene is rather small, but the influence on ocean chemistry is greater when the distribution of Antarctic ice shelves are more realistic. We think that a better introduction to the separate experimental setup in this section (including a new table overviewing the subset of simulations discussed here rather than just describing them) will help the reader understand the experimental setup in this section.

[Printer-friendly version](#)[Discussion paper](#)

———— The last para on the sensitivity of alkalinity to the terrestrial carbon uptake (p.31, l. 14- 23) seems to be very different from the scope of the section. It is interesting to know an effect of ice shelves on ocean alkalinity, but this should be a separate subsection with a clear title and message. ————

As suggested by the reviewer, the discussion of Antarctic ice shelf extent is essentially a sensitivity study within a sensitivity study and merits its own subsection, which will be done for the revised version of the paper.

———— Conclusions: p.34, l. 13-15: I do not understand the point 1. If the sedimentation rate is higher than the weathering rate, the alkalinity is decreasing and CO2 is increasing, independently on the scale of the weathering rate. ————

This point 1 will be rephrased as “A lower early-deglacial weathering rate leads to an earlier and larger increase in atmospheric CO2 in our simulations, as suggested by Rickaby et al (2010).”

———— Figures: General – plots of time series are sub-optimal. The label font should be increased, and a grid added to quantify time series values in the middle of the plot. ————

We are happy to both increase the font and add subgrid reference lines on the x and y-axes to make values at different dates in the simulation clearer.

———— Figure captions are very long. Some of them contain rationale of experiments or details of experimental description, which belongs to the main text. ————

We will try to reduce some of the figure captions. However, we recognize that often readers will prefer to look at figures and figure captions rather than read all the details provided in the main text, so we believe that it is helpful to provide some basic details of the simulations in the figure captions to help make the figures easier to interpret by themselves.

———— Figs. 2-3. What is shown on these maps – terrestrial carbon densities? I

guess, they should be in units of kgC/m².

Figs. 6-7. The same unit question as above. _____

These are not terrestrial carbon densities but rather the total quantity (in Tg C) of vegetation/terrestrial carbon. As the physical area of each grid cell is different (with grid cells near the equator having the largest area, and grid cells near the poles having the smallest area), simply reporting carbon density does not give a sense of the total carbon storage in each grid cell. For example, polar grid cells often have a very high carbon density, but much less total carbon storage compared to mid-latitude or tropical grid cells. These figures (total carbon storage per grid cell) helps clarify which regions actually hold the most carbon. We did not specify that these are carbon storage per grid cell and should make that clearer in the figure captions. However, we are willing to convert the data back to carbon densities if that would be preferable.

Interactive comment on Clim. Past Discuss., doi:10.5194/cp-2016-24, 2016.

Printer-friendly version

Discussion paper

