



Project no. GOCE-CT-2003-505539

Project acronym: ENSEMBLES

Project title: ENSEMBLE-based Predictions of Climate Changes and their Impacts

Instrument: Integrated Project

Thematic Priority: Global Change and Ecosystems

D2B.8. Working paper on model weighting for the construction of probabilistic scenarios in ENSEMBLES

Due date of deliverable: December 2005

Actual submission date: June 2006

Start date of project: 1 September 2004

Duration: 60 Months

Organisation name of lead contractor for this deliverable: UEA (Clare Goodess)

Revision [draft, 1, 2, ..]

Project co-funded by the European Commission within the Sixth Framework Programme (2002-2006)		
Dissemination Level		
PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the Consortium (including the Commission Services)	

1. Introduction

RT2B provides a major interface between the global and regional climate modelling and the impacts modelling activities in ENSEMBLES (Figure 1). Its first primary objective, for example, is to construct probabilistic high-resolution regional climate scenarios and seasonal-to-decadal hindcasts using dynamical and statistical downscaling methods in order to add value to the earth system model output from RT1 and RT2A and to exploit the full potential of the Regional Climate Model (RCM) ensemble system developed in RT3. A range of techniques for constructing these scenarios will be explored, including Monte Carlo sampling, Bayesian methods, ensemble averaging, pattern scaling and weighting (Figure 1). The latter issue (i.e., weighting model output on the basis of performance, typically evaluated by comparing observed and simulated data) is the subject of this working paper.

Section 2 reports on a cross-cutting ENSEMBLES session on weighting (organised by Clare Goodess) which was held in Athens prior to the 2005 General Assembly. Subsequent discussions and plans for setting up some demonstration case studies to explore how user-relevant probabilistic outputs can be developed and applied in an end-to-end approach are described in Section 3. ENSEMBLES deliverables dealing with weighting issues from the perspective of other research themes are identified and outlined in Section 4, while Section 5 lists the ENSEMBLES data sets that may be useful for constructing climate model weights. Finally, Section 6 makes recommendations on the construction and use of weights in RT2B.

2. Summary of the cross-cutting ENSEMBLES session, 5 September 2005, Athens

2.1 Discussion questions

The aim of this session was to address a set of seven questions put together following a small, informal meeting held in Reading in July 2005 to discuss a number of ENSEMBLES cross-cutting issues (attended by Paco Doblas-Reyes, Clare Goodess, Chris Hewitt, Andy Morse and Antje Weisheimer).

1. Is weighting a necessary and appropriate technique?
2. How should weights be calculated?
3. How should weights be used to construct PDFs and other forms of probabilistic scenarios?
4. Can weights from global and regional climate models and statistical downscaling, and for climate change and seasonal-to-decadal timescales, be combined?
5. Can weights from impacts models also be combined?
6. At what stage(s) should the weighting be applied in an integrated (from the coupled model, through the downscaling to the application model) prediction system be carried out?
7. How can the performance of a weighted prediction be compared with an unweighted one?

2.2 Summary of presentations

The session started with brief presentations outlining key issues, methods and questions by the following people:

- Clare Goodess: General overview
- James Murphy: Global climate change simulations
- Paco Doblas-Reyes: Global seasonal-to-decadal hindcasts
- Filippo Giorgi: Regional climate change simulations and Reality Ensemble Averaging
- Andy Morse: Seasonal-to-decadal impacts modelling
- Tim Carter: Climate change impacts modelling

Their PowerPoint presentations can be downloaded from the RT2B ENSEMBLES web site. The summaries below were produced by Lars Barring as part of his rapporteur's report to the General Assembly.

Clare Goodess: General overview of discussion questions and other issues

Weighting is not the only issue in developing probabilistic techniques. Broader issues and questions include:

- What uncertainties should/can be included?
- Are PDFs the best way of representation?
- How might (less/non-technical) users make use of probabilistic information?

Statistical downscaling has been used to produce probabilistic scenarios, e.g., using a weather generator as part of the CRANIUM project (<http://www.ncl.ac.uk/cranium/>): but how could weighting be incorporated in such an approach?

An example of weighting applied at several different stages of an end-to-end approach to assessing uncertainties in impacts scenarios was presented: Wilby, R.L. and Harris, I., 2006: A framework for assessing uncertainties in climate change impacts: low flow scenarios for the River Thames, *Water Resources Research*, **42**, W02419, doi:10.1029/2005 WR004065. The question is: how can we improve on this kind of study in ENSEMBLES?

James Murphy: Towards PDFs of regional climate change from perturbed physics ensembles in RT1:

Three broad methodological categories exist in the literature at present:

- Direct use of GCM ensembles without formal application of observational constraints
- Methods designed to be constrained by observations, and, as far as possible, model-independent
- Methods designed to give results dependent on both observational constraints and GCM ensemble distributions

Summary:

- Significant progress has been made since the IPCC TAR, but there is much still to do.
- Results depend on choices such as:
 - what uncertainties are sampled in the chosen ensemble of models
 - the constraining observations
 - prior distributions for uncertain inputs
 - whether or not, and how, to weight ensemble members
 - what should control the range of outcomes (e.g. observational uncertainties, ensemble spread)
- No obvious basis exists for choosing a “best” method yet, but the sensitivity of results to experimental choices should be quantified as far as possible

Paco Doblas-Reyes: Some answers from seasonal-to-decadal (s2d) forecasting

- Is weighting appropriate?
Yes, when robust
- How should the weights be computed?
Using hindcasts and in a robust way
- How do weights relate to PDFs?
The combination method should be probabilistic
- Can weights from different systems be combined?
Good question
- Can weights from impact models also be combined?
Combination/calibration can be carried out at any stage of the forecast process
- The performance should be compared with an unweighted/uncalibrated prediction

Filippo Giorgi: The reality ensemble averaging (REA) method

- different components:
 - individual model biases to present day observed climate
 - natural variability in 30 year observational series
 - model convergence to ensemble mean of future climate
- planned developments:
 - now only temperature, also include precipitation, and ...

Andy Morse: Case study – s2d context

The conclusion from the example presented, based on DEMETER rainfall data for Southern Africa and a malaria application model, is that, in this particular case, weighting seems most appropriate after rather than before the application model stage.

Tim Carter: Climate change impacts modelling

Weighting: some issues relevant to impact modelling:

1. Climate inputs
 2. Non-climate inputs
 3. Impact model uncertainties
 4. Interactions and feedbacks
- Impact modellers can handle weighted climate inputs in the same way as unweighted climate inputs, but there may be need for:
 - consistent weighting of non-climate scenarios

- refinement of impact evaluation methods
- consideration of feedbacks
- In principle, impact outcomes could be weighted based on impact model performance. In practice, it is uncertain if this will occur in ENSEMBLES

2.3: Summary of open discussion

- i. no alternative to weighting, i.e. cannot avoid the issue - either deliberate weighting or equal weights
- ii. any weighting scheme should account for structural deficiencies in the models
- iii. suggestion: a weighting scheme should not be tailored to any specific application - flexibility is needed
- iv. suggestion: a weighting scheme should include *many* variables, ('every' variable in the ENSEMBLES standard set?) not 'only' monthly/annual mean temperature/precipitation
- v. suggestion: to have a 'straw man' developing an example as the basis for further more focussed discussions
- vi. can weighting be used to root out unbelievable model runs etc.? - if so, such a process needs to be transparent
- vii. different impact models will maybe use different sub-ensembles (based on performance evaluation)
- viii. different impact models will have different requirements regarding performance of driving models in case study regions, etc ...
- ix. ... and during specific calibration periods when other calibration data is available
- x. suggestion for development of guidelines for a standard ensemble to be used
- xi. difference between end-users driven by s2d time-scale (agriculture, electricity/power, ...) and those driven by longer time-scales (infra-structure construction, forestry, ...)
- xii. discussion on the issue that different end-user/stakeholder categories may have widely different needs - and competence to express their needs - and to ingest different types and time-scales of climate projections
- xiii. suggestion to continue the discussion in a specific group, based on the present ad hoc group, also with representatives from
 - RT4: fundamental process and model uncertainties, uncertainties in extremes
 - RT5: observations, model validation
 - RT7: scenario uncertainties and feedback

Figure 2 (produced by Paco Doblas-Reyes) indicates the different stages in the end-to-end approach at which weighting could be undertaken in the s2d context. A modified version is shown in Figure 3 for the anthropogenic climate change context.

Figure 1: Schematic diagram of RT2B activities

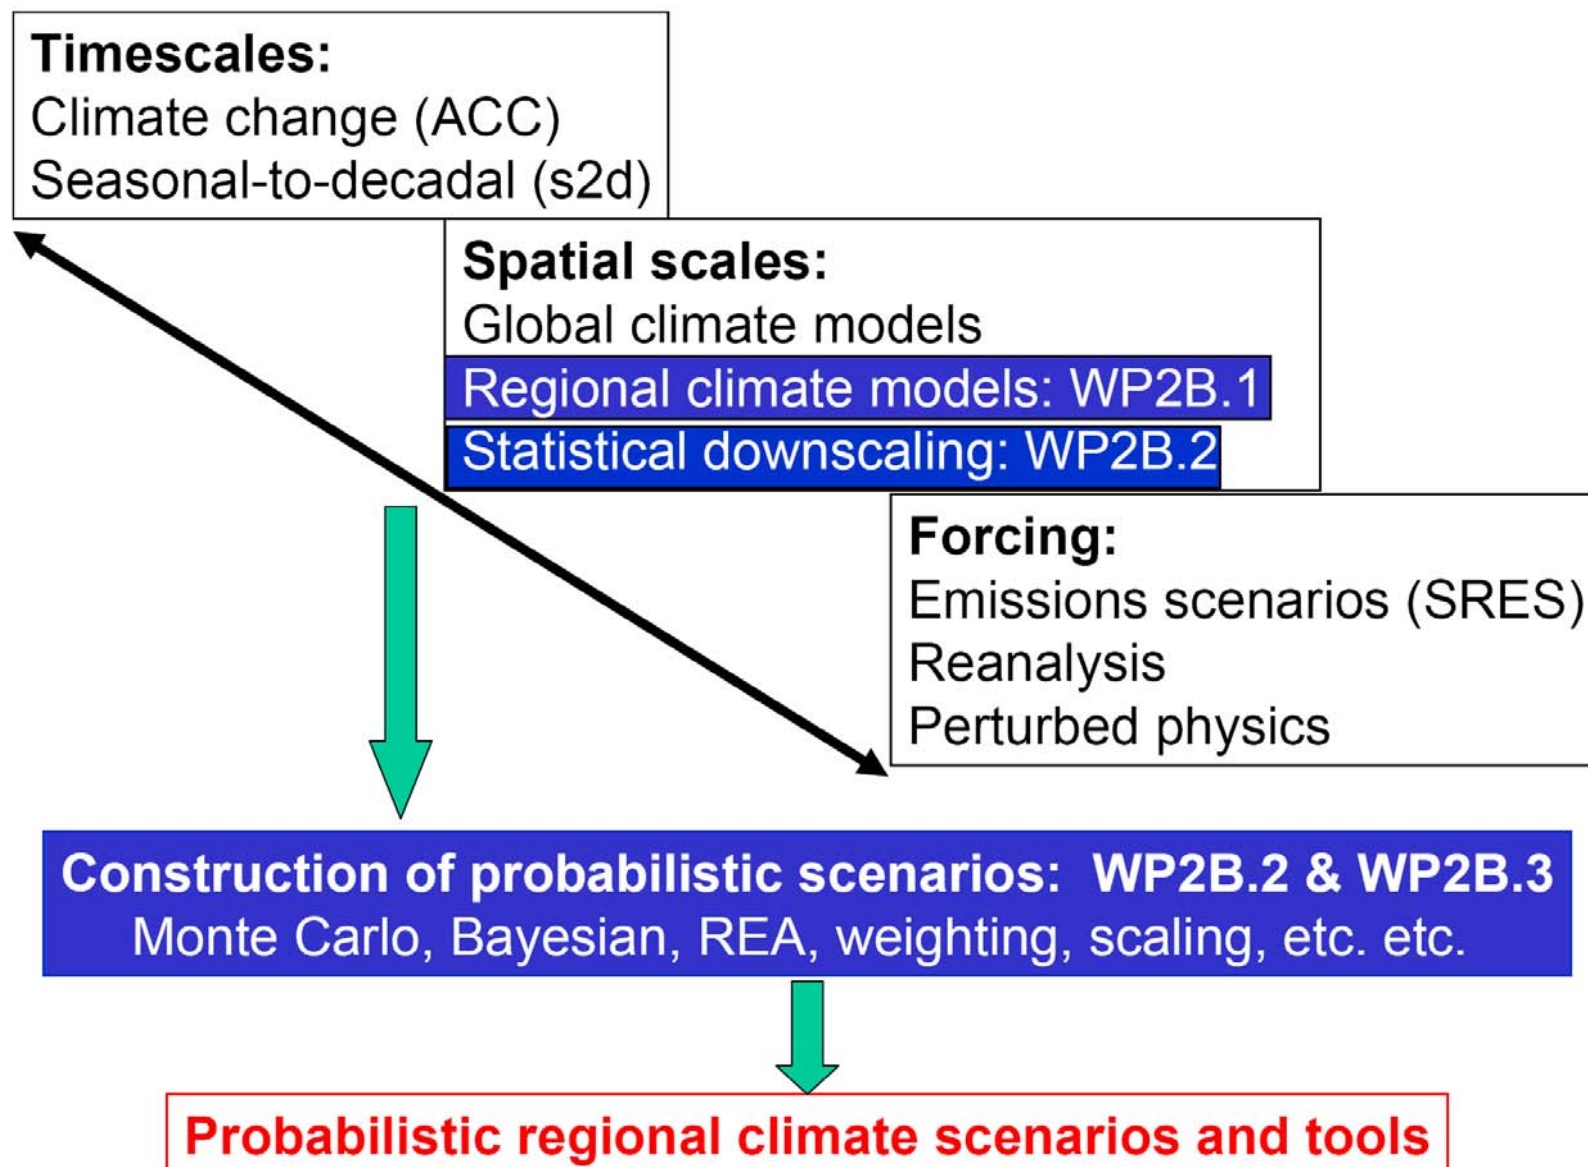


Figure 2: Weighting in the end-to-end approach: seasonal-to-decadal context

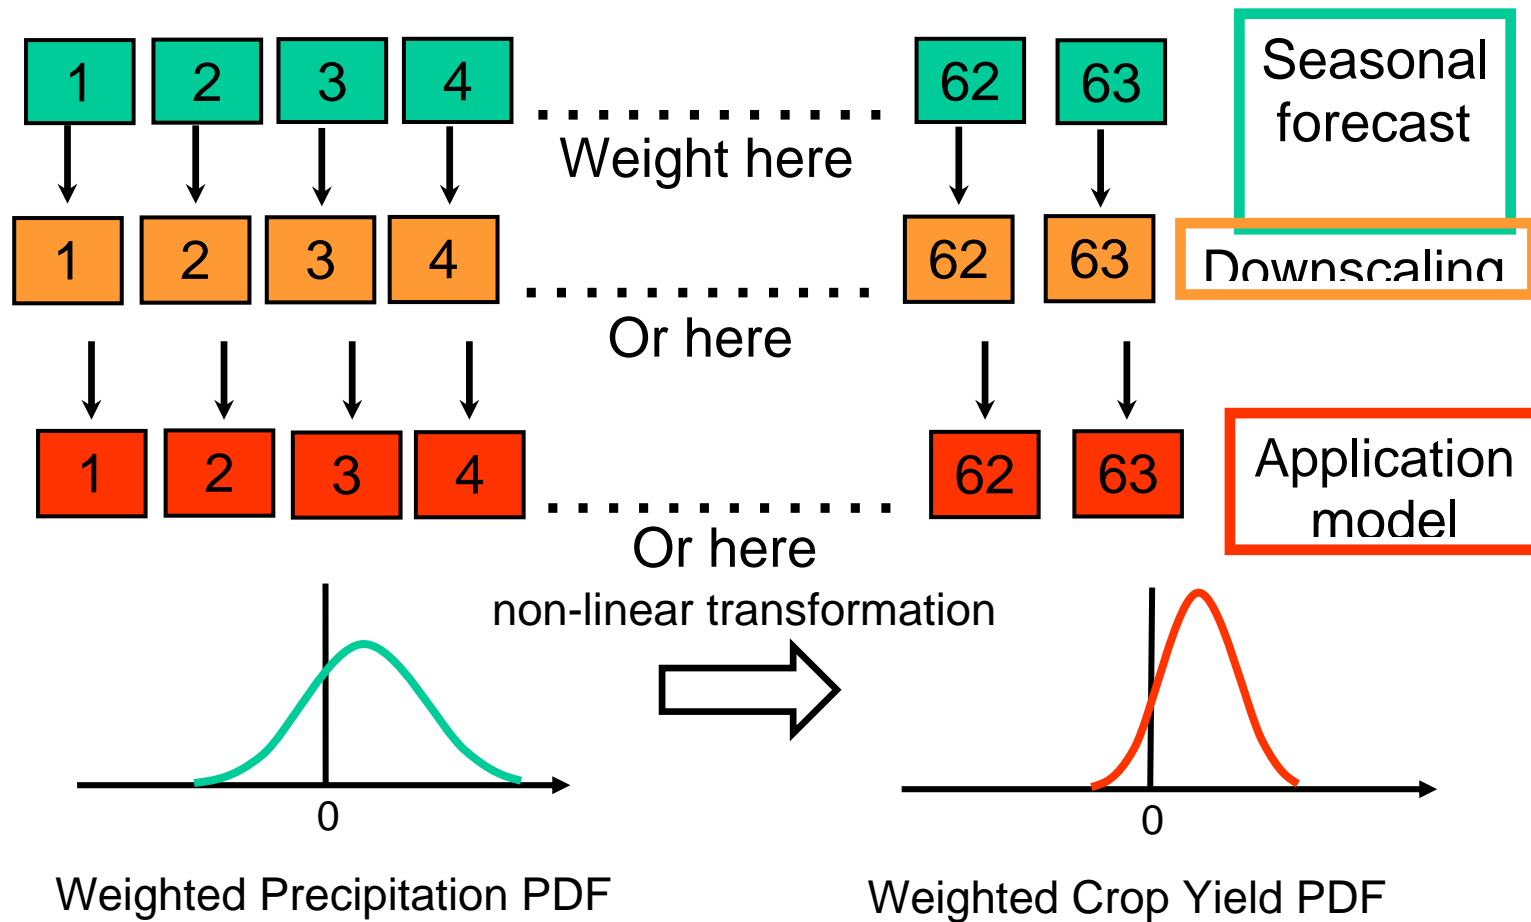
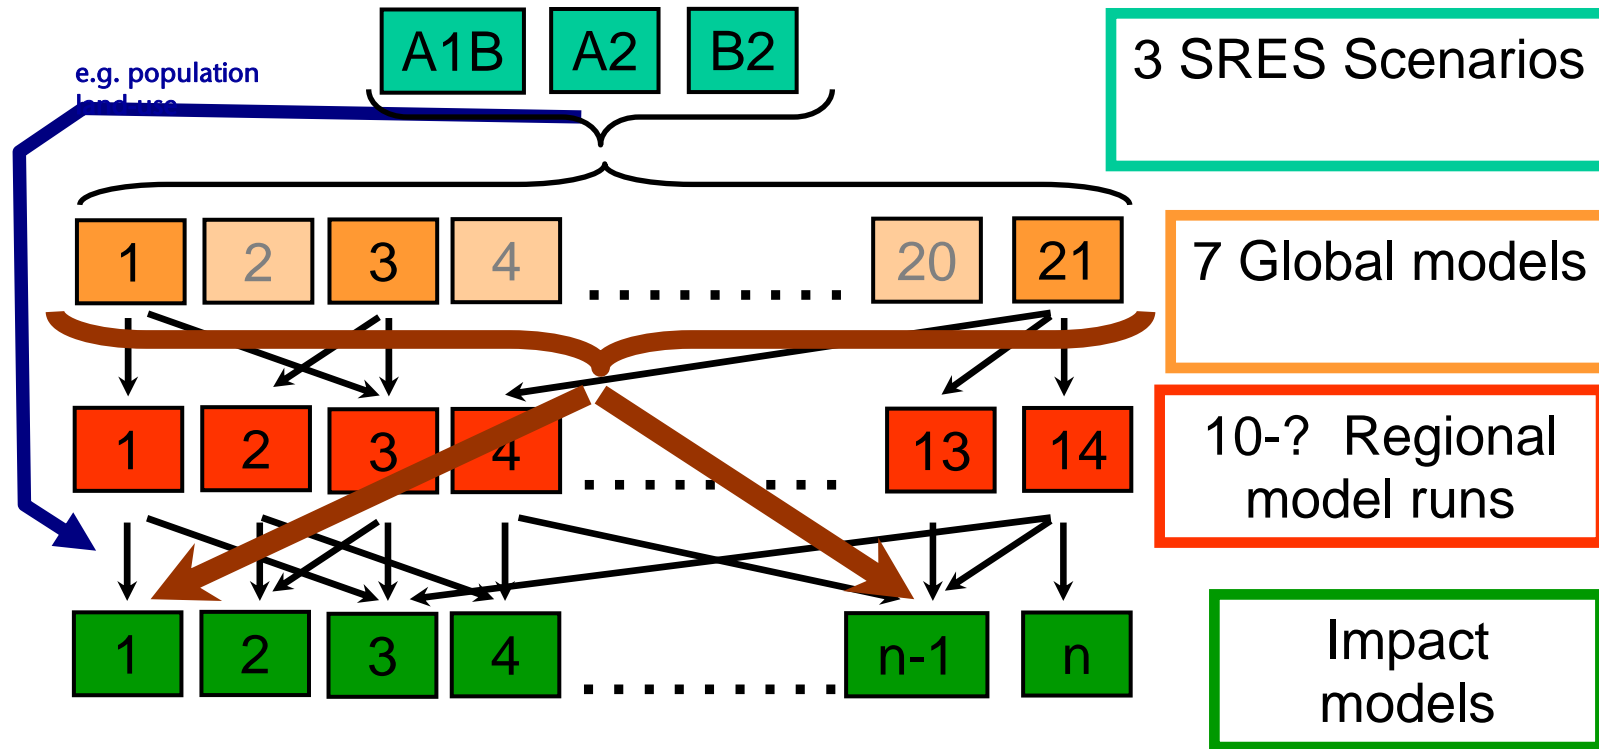


Figure 3: Weighting in the end-to-end approach: the anthropogenic climate change context.



3: Demonstration of end-to-end applications

Based on discussion during the Athens cross-cutting session on weighting (see Section 2) and follow-up discussions during the 2005 General Assembly, it was agreed to identify a few demonstrations of end-to-end applications, with the aim of stimulating discussion on how to handle the construction of methods (including weighting) for provision of user-relevant probabilistic outputs. The summary below of the proposed demonstrations is taken from subsequent email discussion instigated by James Murphy and replies from Tim Carter, Paco Doblas-Reyes and Andy Morse.

In terms of climate model outputs, there were two suggestions for provision of data direct from 'global' model ensembles:

- (1) Use of output from the IPCC multi-model archive at PCMDI (suggested by Filippo Giorgi). This is a set of coupled AOGCM simulations from different modelling centres, with observed historical forcing to present day, then future forcing based on one or more SRES scenarios.
- (2) Use of results from the Hadley Centre's "perturbed physics" ensemble (James Murphy). This consists of 17 variants of HadCM3 with multiple parameter perturbations, run from 1860 to present day with observed historical forcing, and on to 2100 with SRES A1B forcing.

It was suggested to focus on the A1B scenario for case (1), in order to allow a clean comparison of how the outputs from the impact models depend on variations in the inputs arising from alternative designs for the driving AOGCM ensemble.

A third suggestion (Clare Goodess) was to test applications in which input data is provided from a regional climate model ensemble (i.e., PRUDENCE data). This would not be directly comparable with (1) or (2), as the ensemble design is very different (a partly-filled matrix in which alternative AOGCMs are used to drive alternative RCMs). Hence discussion has focused on use of the PCMDI archives and the perturbed physics ensembles. The planned outputs from the latter are outlined below:

- a. There is an RT1 ENSEMBLES deliverable based on the perturbed physics experiments, due in September 2006: Interim probability distributions of transient climate change over Europe will be produced, for use by other RTs in testing methodologies for prediction of climate change impacts. These will be "interim" because they will not (yet) account for variations in the likelihood of alternative model versions, and will not (yet) take account of "structural" model uncertainties. They will be based on 17 member GCM ensemble simulations of the transient climate response, augmented by pattern scaling from the equilibrium responses of a larger (~160 member) ensemble plus statistical emulation of regions of parameter space not sampled by actual GCM simulations. They will also sample the impacts of internal climate variability. They will be specific to the SRES A1B scenario.
- b. It is hoped to include both marginal PDFs of single variables and joint PDFs of selected multiple variable combinations (for example, surface temperature and precipitation changes) as part of (a). However, it will only be feasible to provide PDFs of annual or seasonal means on this time scale. PDFs of monthly changes will eventually be produced, but not until 2007 due to the time needed to extract all the monthly data. For

seasonal/annual variables, the methodology to develop joint PDFs will not be available until summer 2006. Marginal PDFs may be available earlier, hopefully by early spring 2006. See Section 4 for an outline of a recent deliverable on the approaches that will be used.

In terms of linking to impact models, it has been suggested that an example be tried where the input data required by the impact model comes in the form of a PDF derived from a climate model ensemble, plus an example where the impact model requires more detailed daily, monthly or seasonal time series inputs from the actual climate model runs. It would be good, if possible, to determine the sensitivity to the nature of the input, if examples could be found of impacts models which can take either a PDF or time series input. In cases where time series input is definitely required, it would be good to test the sensitivity of the results to the time resolution of the time series (e.g. daily cf monthly). For example, Lenny Smith has suggested using a model which predicts energy requirements, using either monthly or daily surface temperatures.

Another aim should be to identify an impacts model which can be run for both seasonal-to-decadal (s2d) predictions, and for longer term anthropogenic climate change predictions.

By the end of 2005, the following four impacts models were proposed for these end-to-end demonstrations.

1: Malaria incidence statistical model developed by IRI (proposed collaborative work by IRI and ECMWF). Requires PDFs of mean precipitation change for 2046-65 minus 1981-2000 averaged over December to February, for the tropical band 30S to 30N. Can be applied to perturbed-physics output, IPCC data and DEMETER output – using both weighted and unweighted ensembles.

2: Dynamical malaria model (proposed work by University of Liverpool). Requires time series of daily precipitation and daily diurnal maximum and diurnal mean surface air temperature for selected parts of Africa. Regions, time periods and seasons need to be defined. This is the model which was used to assess different approaches to weighting (see Section 2.2). The main scientific interest here would be to look for shifts in seasonal cycles.

3: Statistical ‘palsa mire’ model (i.e., of permafrost hummocks found in northern Sweden), developed by SYKE (proposed by Tim Carter). Requires joint PDFs of annual mean surface air temperature and precipitation change for northern Scandinavia. Monthly mean temperature and annual precipitation from the 17 ensemble members (see above) would be required to compare the approach of inputting PDFs to create impact response surfaces against an approach of running multiple scenarios based on inputs from different GCM runs.

4: DIAS winter wheat and DAISY nitrogen leaching model (proposed by Tim Carter). Requires joint PDFs of surface air temperature and precipitation for a set of European locations, plus daily time series of temperature and precipitation from the 17 ensemble members. This would allow the impact response surface method (i.e., feeding the impact model by adding time mean changes sampled from a PDF onto observed daily time series) to be compared against the ‘brute force’ method (i.e., feeding the impact model with full simulated daily time series from multiple ensemble runs of GCMs). This proposed demonstration is more intensive in terms of data requirements and might realistically be tackled after the other three proposed case studies described above.

At the time of completing this working paper (early June 2006), data extraction from the perturbed physics simulations should have started. Once this is well underway, the specific requirements for the proposed demonstrations will be finalised. It is noted that we need to be realistic about the data volumes that can be extracted for this exercise.

Although weighting issues will be addressed as part of some of the proposed demonstrations, the main focus will be on a broader exploration of how user-relevant probabilistic outputs can be developed and applied in an end-to-end approach.

4. Other ENSEMBLES deliverables dealing with weighting issues

This section identifies a number of completed ENSEMBLES deliverables (all available from the deliverables section of the ENSEMBLES web site) which deal with weighting and relevant model validation issues.

D1.2: Systematic documentation and intercomparison of ensemble perturbation and weighting methods

Outlines a suite of papers by the Oxford and Hadley Centre teams detailing new developments in the methodology of ensemble climate predictions of 21st century climate. The Bayesian approach being developed at the Hadley Centre differs somewhat in philosophy from the STABLE Inferences from Data (STAID) approach followed by Oxford. Relevant work on seasonal to decadal prediction by ECMWF is also noted.

D2B.6: Refinement of the Reality Ensemble Averaging (REA) framework

This brief report by Filippo Giorgi outlines a refined REA scheme in which a model weight depends on a combined functional form that includes model mean bias, interannual variability and multidecadal trend performance for both temperature and precipitation. The model convergence criterion has been removed.

D3.2.1: Definition of measures of reliability based on ability to simulate observed climate in hind-cast mode

This ENSEMBLES RT3 concept paper by Filippo Giorgi discusses (1) how to calculate model weights and (2) how to produce PDFs based on the ENSEMBLES multi-model set. With respect to (1), the framework discussion raises many issues which are also considered in this working paper. A work plan for WP3.2 work on weighting is proposed, using the Reality Ensemble Averaging (REA) method of Giorgi and Mearns (2002) and the Climate Prediction Index (CPI) based method of Murphy *et al.*, 2004 as starting points to evaluate RCM weighting procedures.

D5.3: Scientific article/report and R software on optimal statistical methods for combining multi-model forecasts to make probabilistic forecasts of rare extreme events

This discussion paper by ECMWF and the University of Reading discusses the need for clear definitions of relevant extremes on s2d timescales and the importance of appropriate calibration and verification methods.

D5.4: Verification methods for forecasts of extreme events

This discussion by David Stephenson proposes some possible scores (such as extreme dependency scores and odds ratios) for the verification of extreme event forecasts. Problems with conventional approaches are discussed, and some possible new approaches suggested.

Deliverable D5.5 (due to be completed shortly) should provide a ‘preliminary report on systematic errors in the ENSEMBLES models’ and hence may also be relevant.

5. Data sets available in ENSEMBLES for constructing weights

The following observed datasets (see also Section 3.2 of deliverable D2B.2) may be useful for constructing model weights:

- The CRU gridded data sets - <http://www.cru.uea.ac.uk/cru/data/hrg.htm>
- ERA-40 (though caution is needed with some variables, e.g., a method for bias correction of Alpine precipitation has been developed – deliverable D5.2).
- Daily station data from the European Climate Assessment & Dataset - <http://eca.knmi.nl/>
- The STARDEX seasonal indices of extremes (six for precipitation and four for temperature) – available for 491 European-wide stations for the period 1958-2000 - http://www.cru.uea.ac.uk/cru/projects/stardex/euro_indices.tar.gz
- The ENSEMBLES 25 km gridded dataset of maximum and minimum temperature, precipitation, sea level pressure and snow depth. RT5 plan to have an early release in Autumn 2006, with the final dataset due in Spring 2007. Deliverables D5.8 and D5.9 describe the station density/homogenisation methods and proposed gridding methods respectively. RT5 also plan to release simulations (about 1000 depending on computer requirements) to address uncertainty rather than a single data set. They also plan to release different resolutions to match the different models – as they would prefer that users don't interpolate to match observations to their results as this adds uncertainty.
- The observed case-study data sets for scenario construction and assessment to be developed in RT2B (deliverable D2B.15, due in December 2006).

6. Recommendations for the construction and use of weights in RT2B

From the ENSEMBLES discussions and work on weighting issues to date, the following recommendations for the construction and use of weights in RT2B have emerged:

- Weighting schemes should be robust, i.e., demonstrated to be relatively insensitive to choice of variables, time periods used, etc.
- Weighting schemes should be informed by the underlying physical processes and expert knowledge
- Transparency is needed in terms of how the weights are calculated and applied
- Weighting schemes should be seasonally stratified and based on a broad range of variables, encompassing the reproduction of extremes, inter-annual and longer-term variability, as well as mean climate

- It would be desirable to have a ‘common, comprehensive and flexible’ weighting scheme, although some users may also want schemes which are more tailored to their particular application
- Weighting schemes should be developed in consultation with end users
- Need to avoid ‘double-counting’ of errors/biases in the end-to-end approach
- Need to be able to provide and compare weighted and unweighted scenarios/outputs

The first three bullet points can be thought of as guidelines for the construction of ‘appropriate’ weighting schemes and hence are relevant to Question 1 in Section 2.1. Most of the other bullet points relate to Question 2. It is clear that a lot of work, including extensive sensitivity analysis, is needed before all seven questions can be fully addressed.

Questions 4 and 6 in Section 2.1 are those which have been given least consideration to date, but are particularly relevant to RT2B.

For the climate change context (Figure 3), it is recommended that the regional scenarios constructed in RT2B should be conditional on emissions scenario and will focus on the A1B SRES scenario. This is consistent with the approach being taken in RT1 and removes the need, in RT2B at least, to weight emissions scenarios.

The next step in the end-to-end approach at which weighting could be applied, is with respect to the global models (Figure 3). For the climate-change context, RT2B will have access to the PDFs of regional climate from the RT1 perturbed physics ensembles (see Sections 2 and 3). These will be constructed using the Hadley Centre’s Climate Prediction Index as a weighting tool. It is not currently clear whether common weighting schemes will be developed for the RT2A stream 1 (i.e., AOGCM IPCC AR4) climate change simulations, or to what extent model evaluation work undertaken in RTs 4 and 5 could feed into the development of such schemes. RT2B will need to determine whether any weighting schemes developed for the AOGCMs are at the appropriate spatial and temporal scale for regional scenario construction.

At the next stage in the end-to-end approach, i.e., downscaling (Figure 3), RT3 will provide a weighting scheme for the RCMs based on validation of the ERA-40 forced simulations (initially based on the 50-km simulations, with the final set (D3.2.2 due February 2007) based on the 25-km simulations). An issue for RT2B is, whether or not the same set of weights should also be calculated using the WP2B.1 RCM simulations (which would encompass both GCM and RCM performance)? Another issue for RT2B is, whether the same weighting scheme can be applied to statistically downscaled output? Some of the statistical downscaling methods will use what is essentially a delta change approach, e.g., using RCM change fields to perturb weather generator parameters – the question arises as to whether it is appropriate to weight these change fields? As with the above stage, the extent to which RT4 and RT5 RCM evaluation work will/could input directly to the calculation of weights is not clear. RT2B also needs to consider the implications of the RT5 decision (which is consistent with the ENSEMBLES philosophy) to produce 1000 iterations of the gridded observed data set (see Section 5).

The final stage in the end-to-end approach concerns the impacts model applications (Figure 3) and hence is of primary concern to RT6. In order to facilitate relevant sensitivity analyses in RT6 it will be important for RT2B to produce both weighted and unweighted scenarios and outputs, as recommended above. The impacts models being used for RT6 work on climate change and s2d timescales are outlined in deliverables D6.3 and D6.4 respectively.

The above discussion focuses on climate-change applications, but many of the issues and questions raised are also relevant to the s2d-context (Figure 2). Issues relating to model weighting and debiasing (bias correction) are being explored as part of the work on deliverable D2B.12 on ‘Recommendations for the application of statistical downscaling methods to seasonal-to-decadal hindcasts in ENSEMBLES’ (see the outline of this deliverable at <http://www.cru.uea.ac.uk/projects/ensembles/pmwiki/pmwiki.php?n=D2B.12>). Here, it is noted that, if the same dynamical and statistical downscaling models are used for s2d and climate-change applications, then it should in theory be possible to use the same regional weighting schemes. It is also noted that whilst bias correction is often used by the s2d community (and has the potential to affect the need for weighting as well as the value of any weights used), it is currently rarely used by the climate-change community.

The simplest weighting scheme that could be used in RT2B would give equal weight to all models, i.e., all models would have a weight 1. This is our current starting point. A slightly more sophisticated scheme might give a weight 0 to any models considered ‘too poor’ and weight 1 to all others. RT2B work over the coming months will determine whether or not more complex weighting schemes are necessary and appropriate.

This working paper is intended to set the background for this work, to make a preliminary set of recommendations and to identify issues and questions requiring further work. Additional questions as well as comments on this document are therefore encouraged and should be sent to the author, Clare Goodess (c.goodess@uea.ac.uk).

Bibliography

- Desai, S.X.R., 2005: *Robust adaptation decisions amid climate change uncertainties*, PhD Thesis, University of East Anglia - <http://www.uea.ac.uk/~e120782/thesis.pdf>.
- Dessai, S., Lu, X. and Hulme, M., 2005: Limited sensitivity analysis of regional climate change probabilities for the 21st century, *Journal of Geophysical Research*, **10**, D19109, doi:10.1029/2005JD005919.
- Giorgi, F., and L.O. Mearns, 2002: Calculation of average, uncertainty range and reliability of regional climate changes from AOGCM simulations via the “Reliability Ensemble Averaging” (REA) method, *Journal of Climate*, **15**, 1141-1158.
- Giorgi, F. and Mearns, L.O., 2003: Probability of regional climate change based on the Reality Ensemble Averaging (REA) method, *Geophysical Research Letters*, **30**, doi:10.1029/2003GL0171310.
- Murphy, J.M., Sexton, D.M.H., Barnett, D.N., Jones, G.S., Webb, M.J. and Collins, M., 2004: Quantification of modelling uncertainties in a large ensemble of climate change simulations, *Nature*, **430**, 768-772.

Tebaldi, C., Smith, R.L., Nychka, D. and Mearns, L.O., 2005: Quantifying uncertainty in projections of regional climate change: A Bayesian approach to the analysis of multimodel ensembles, *Journal of Climate*, **18**, 1524-1540.

Van Ulden, A.P. and van Oldenborgh, G.J., 2005: Large-scale atmospheric circulation biases and changes in global climate model simulations and their importance for regional climate scenarios: a case study for West-Central Europe, *Atmospheric Chemistry and Physics Discussions*, **5**, 7415-1455.

Wilby, R.L. and Harris, I., 2006: A framework for assessing uncertainties in climate change impacts: low flow scenarios for the River Thames, UK, *Water Resources Research*, **42**, W02419, doi:10.1029/2005 WR004065.

Acknowledgements

The following ENSEMBLES colleagues are thanked for their contributions to discussions on these issues at meetings in Reading and Athens, and by email: Lars Bärring, Tim Carter, Paco Doblas-Reyes, Filippo Giorgi, Chris Hewitt, Andy Morse, James Murphy and Antje Weisheimer. Thanks also to everybody who attended the cross-cutting session on weighting in Athens.