Robustness of climate response in HadCM3 to various perturbations of the Atlantic meridional overturning circulation.

Hadley Centre technical note 48

Michael Vellinga

6 May 2004

Met Office

Robustness of climate response in HadCM3 to various perturbations of the Atlantic meridional overturning circulation.

Michael Vellinga
Met Office, Hadley Centre for Climate Prediction and Research, FitzRoy Road, Exeter, Devon Ex1 3PB, United Kingdom
email: michael.vellinga@metoffice.com

May 6, 2004

Executive summary

Uncertainty in stability of the meridional overturning circulation of the ocean ('MOC'*) in climate models contributes to uncertainties in climate projections. Disruptions to the Atlantic MOC, and its associated northward heat transport, have the capability of causing both local and global climate anomalies. No rapid, full collapses of the MOC in HadCM3 have been observed in simulations under pre-industrial conditions, or a range of greenhouse gas forcing scenarios. In those simulations, negative (stabilising) feedbacks dominate over positive (destabilising) feedbacks and maintain the MOC. In previous work we have identified the dominant stabilising feedback that leads to the recovery of the MOC after a particular external perturbation. To identify if, in fact, this recovery mechanism is a robust property of the model under pre-industrial conditions, we have applied a series of external freshwater perturbations to the Atlantic MOC.

1. In the perturbation experiments there is generally a rapid weakening of the MOC, followed by a gradual recovery. (Figs. 1- 2). No permanent 'MOC-off' state was found.
2. The degree of MOC weakening depends near-linearly on the magnitude of the perturbation. Perturbations are most effective if applied as instantaneous freshening of northern high latitudes (Fig. 3a). Rate of MOC recovery varies little between the experiments (Fig. 1), and shows no clear correlation with the magnitude of the perturbation (Figs. 3b).
3. The MOC recovers in most experiments in a generic fashion. This suggests that the stabilising process identified previously in HadCM3 is robust. (Figs. 7– 9). An important aspect of this process is the near-invariance of the ocean freshwater transport in the subtropical Atlantic. This invariance is not understood, and is a model property that needs to be considered more carefully.
4. Surface climate anomalies after MOC disruption can be characterised by a low number of generic spatial patterns. The amplitude of these patterns depends linearly on the degree of MOC weakening, and estimating the climate response to a particular MOC disruption under pre-industrial conditions would be relatively straightforward. (Figs. 10-11).
5. Climate response after an artificially forced MOC disruption by 2049 under an IS92a greenhouse gas ('GHG') scenario was compared with a similar disruption under pre-industrial conditions. Global cooling due to MOC collapse under GHG is about 0.5^oC stronger than under pre-industrial conditions (Fig. 12). Over the North Atlantic and western fringes of continental Europe, cooling caused by an MOC collapse is stronger than the warming caused by the increased GHG forcing. (Fig. 13).

We have identified a robust mechanism for MOC recovery in HadCM3 (item 3 above). This knowledge provides a focus for our attempts to reduce uncertainty in projections of future MOC changes. Ideally, we would like to do this by constraining dynamics of oceanic freshwater transport in the model by the observational record. This will require us to explore what useful information is currently available from observations, and what may become available in the future. It will be equally important to compare this specific aspect of HadCM3 to that in other climate models.

* 'Meridional overturning circulation' is a more accurate description of the nature of the flow than the traditional 'thermohaline circulation.'

1 Introduction

One of the factors contributing to uncertainties in climate simulations is the stability of the meridional overturning circulation ('MOC'). Differences in stability properties of the MOC among various climate models lead to considerable spread in projected weakening of the MOC in simulations by these models, all forced by the same IS92a emission scenario (Cubasch *et al* (2001)). Changes in the strength of the MOC can have an important effect on global climate conditions (Vellinga and Wood (2002)), and it is therefore important to identify and understand the processes that control stability of the MOC in climate models in order to increase the level of confidence in climate projections.

Vellinga *et al* (2002) describe how the MOC in HadCM3 recovers after being disrupted by one particular external freshwater perturbation. The recovery of the MOC is accomplished through a combination of oceanic and atmospheric feedbacks.[†] In this report we will describe what we have learned from a series of perturbation experiments that have been carried out by HadCM3. These experiments were designed so as to determine the dependence of climate response on magnitude and nature of the applied perturbation.

Firstly, the sensitivity of the MOC response itself on the perturbation size is discussed (Section 2). Then we will assess if the recovery mechanism described by Vellinga *et al* (2002) is generic, in that it is typical for the recovery the MOC in a wide range of perturbation experiments. This issue is addressed in Section 3 of this report. If the recovery process is indeed generic, and does not depend critically on the nature of the applied perturbation it will be important that we have a high level of confidence in its accurate representation in climate models.

Even if to the best of our current knowledge we do not believe that a sudden collapse of the Atlantic MOC is likely to occur over the next 100 years (Wood *et al* (1999); Thorpe *et al* (2001)), there remain important questions that need to be answered. Firstly, is an MOC collapse more likely in the future, when increased greenhouse gas forcing may have shifted the mean climate state (briefly addressed in Section 2 c)? Secondly, are the climatic anomalies caused by a sudden MOC collapse similar in a greenhouse-gas enriched climate and in a pre-industrial climate (Section 4)?

2 Model experiments and MOC response

We have analysed the evolution of the MOC in HadCM3 in a series of perturbation experiments. The perturbations were chosen to reduce high latitude salinity in the upper layers of the North Atlantic. The motivation for this approach is twofold: firstly, freshening of the upper ocean lowers density which suppresses oceanic convection, surface buoyancy loss reduces, and hence water mass transformation. These processes are important in driving a large part of the zonal mean overturning circulation in the Atlantic Ocean. This circulation is sometimes referred to as the North Atlantic Deep Water ('NADW') cell. Secondly, freshening of the upper North Atlantic Ocean reduces the meridional steric height gradient, a quantity to which the quasi-equilibrium strength of the MOC in HadCM3 is closely related (Thorpe *et al* (2001)).

[†] The barotropic circulation in the ocean (which is partly wind-driven) is maintained, even in the absence of the MOC. It exports the freshwater that was put in the North Atlantic, and replaces it with more saline waters from the subtropics. Due to a change in rainfall patterns, these subtropical waters are significantly more saline than they are in the unperturbed climate. As a result the waters of the North Atlantic gradually turn more and more saline, and become dense enough to support deep water formation. The MOC can therefore re-start.

The perturbation experiments fall into two categories, namely initial perturbation experiments and anomalous flux experiments. Both are described in the following two subsections.

a Initial-perturbation experiments

The first category contains experiments in which an initial perturbation to the salinity field was applied (generally in the upper 500m). Each experiment is designated 'Pnn', with nn the experiment number. The perturbations are generally negative, but differ in magnitude, and area over which they were applied. Only in P08 did we apply a positive salinity perturbation, in the South Atlantic, between 1200-3000m. This was done to test the response to a reduction of the north-south steric height gradient by increasing water density at the southern end of the NADW cell. A summary of the different experiments is given in Table 1.

name	area	freshening (psu)	size (Sv*year)	length of exp. (yrs)
P01	20N-90N	-3.0	65.6	110
P02	0-20N ∪ 55N-90N	-0.8	60.0	270
P03°	55-90N	-3.0	19.8	160
P04	55-90N	-0.5	2.7	90
P05	55-90N	-1.5	9.9	50
P06	55-90N	-0.75	5.0	50
P07	38-50N	-2.8	9.8	60
P08	20-30S	+0.74	-8.8	50
P09*	55-90N	-2.2	11.7	90
P10†	50-90N	-3.4	22.1	40

Table 1: MOC perturbation experiments in HadCM3. Unless stated otherwise, experiments were started from year 100 of the HadCM3 control integration. (°): experiment described in Vellinga *et al.* (2002) and Vellinga and Wood (2002). (*) started from year 1000 in HadCM3 control. (†): started from year 2049 in an IS92a emission scenario run. The area where the perturbation was applied is shown in the second column. The volume-averaged change in salinity in that area (third column) can be expressed as resulting from an equivalent freshwater pulse (fourth column, a volume expressed as $Sv * year$, $1\,Sv \equiv 10^6 m^3 s^{-1}$; $1\,Sv * year \approx 31.5 \cdot 10^{12} m^3$.). The last column shows the length in years the experiment was run for.

The evolution of the Atlantic MOC at $30^o N$ [‡] in all perturbation experiments is shown in Fig. 1. In none of the experiments did the MOC evolve into a permanent off-state. Generally, the MOC weakens strongly in the first decade after the perturbation is applied, then gradually recovers. Exception to this behaviour is experiment P07, in which the strongest weakening is delayed by three decades. In P07 the salinity perturbation is applied in the northern subtropics and mid latitudes, while in most other experiments salinity was perturbed at high latitudes. In P07 the the weakening of the NADW cell is initially strongest between $30^o N - 30^0 S$; the anticipated weakening of the northern (downwelling) limb of the NADW cell occurs in years 20-30 (as in Fig. 1). In P08, although there is a considerable change in the deep circulation in the South Atlantic, there is no change in the NADW cell at 30^o that exceeds the natural, internal variability.

[‡] This latitude was chosen because it is close to the latitude of maximum northward ocean heat transport.

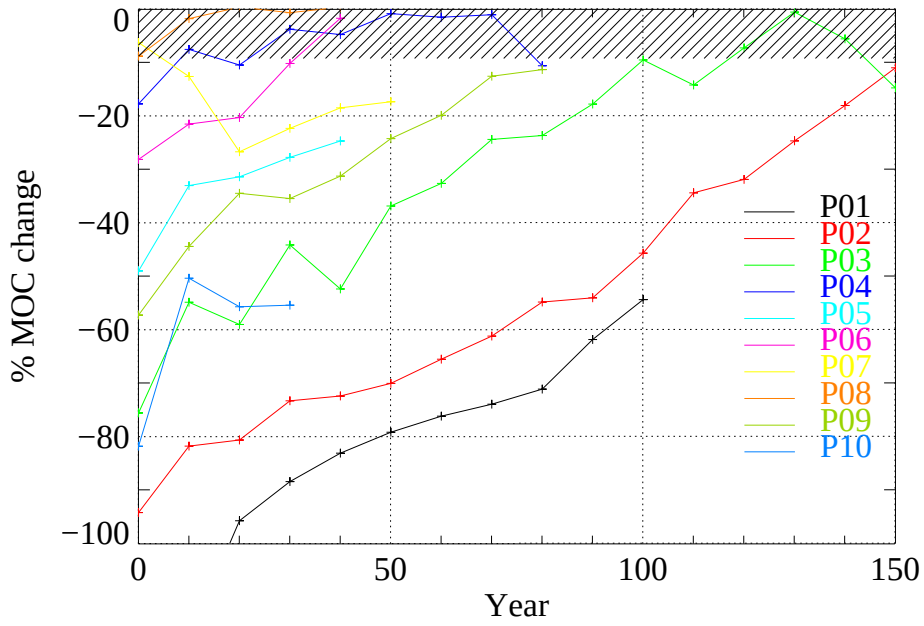

Figure 1: P experiments: decadally averaged MOC at $30^o N$, expressed as fractional change relative to the average value over a (parallel) 100-year period of the corresponding control experiment. Hatching indicates the range of two standard deviations of decadally averaged anomalies in the control run, and provides a measure of internal variability in the unperturbed climate.

b Anomalous surface flux experiments

In the second set of experiments (each of which is designated 'Fnn') an anomalous surface flux of fresh water was applied over the northern North Atlantic for a prolonged period. Anomalous surface fluxes of different strengths were applied. In particular, F01 was applied over the Nordic and Labrador Seas around Greenland only. A summary of the different experiments is given in Table 2.

In all flux anomaly experiments the MOC slows down gradually. In F04 the flux anomaly was removed after 100 years. By then the MOC had reduced to about 85% of its original strength. Thirty years after the flux anomaly was removed, the MOC started to recover slowly. Even in this experiment where over 100 Sv*year of freshwater was applied to the North Atlantic, the state with a very weak MOC is not stable.

name	area	size (Sv)	length of exp. (yrs)
F01	40N-70N (around Greenland)	0.04	190
F02	45N-90N	1.1	240
F03	45N-90N	0.27	300
F04‡	45N-90N	1.1	200

Table 2: MOC surface freshwater flux perturbations in HadCM3. All experiments were started from year 300 of the HadCM3 control integration. (‡): parallel to F02, but with anomalous surface flux switched off after 100 years.

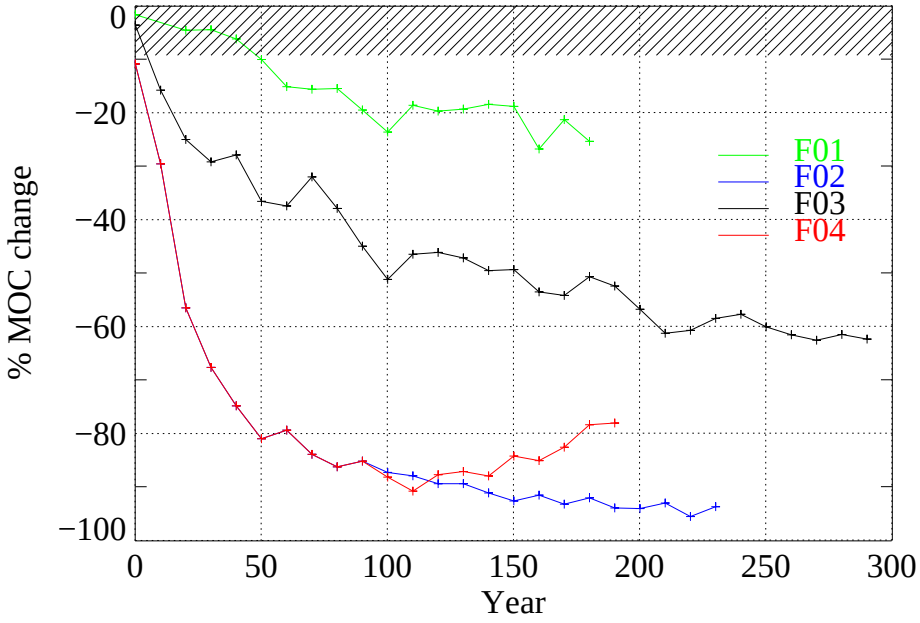

Figure 2: F experiments: decadally averaged MOC at 30^oN, expressed as fractional change relative to the average value over a (parallel) 100-year period of the corresponding control experiment. Hatching indicates the range of two standard deviations of decadally averaged anomalies in the control run, and provides a measure of internal variability in the unperturbed climate.

c Sensitivity of MOC response to perturbation size

The amount of initial weakening of the MOC is proportional to the size of the perturbation (Fig. 3a). The sensitivity is approximately linear for perturbations up to about 20 Sv*year. Perturbations stronger than that do not cause additional weakening. While the maximum amount of induced weakening of the NADW cell could be 100% this is not reached, and the sensitivity saturates. The strongest weakening in P07 is less than in most of the other experiments. This confirms the view that it is the northern high latitudes where the MOC is most sensitive to salinity changes (Rahmstorf (1995); Manabe and Stouffer (1997)).

The positioning of P10 in the near-linear regime indicates that the MOC response to the given amount of freshening is not sensitive to increased GHG conditions. Based on the single experiment carried out under GHG conditions one might conjecture that the likelihood of an MOC collapse in a GHG and in a pre-industrial climate are similar, assuming the properties of external perturbations to the MOC (e.g. through surface freshwater flux anomalies) does not change.

The amount of weakening that occurs in F02 and F03 is less than might have been expected from an initial perturbation of the same magnitude. Part of the anomalous surface flux in F02 and F03 has probably been 'wasted'. Given that in these experiments the anomalous fluxes were applied over a large surface area, ocean transports will have carried part of the fresh anomalies away from the sensitive deep water formation areas in the Nordic Seas.

The estimated linear rate of recovery of the MOC after the perturbation is shown in Fig.3b, again against the magnitude of the perturbation. For a wide range of perturbations, the MOC recovers at similar rates. Again, there is no fundamental difference between the response in P10 (the GHG run) and the other experiments. This suggests that, whatever the recovery

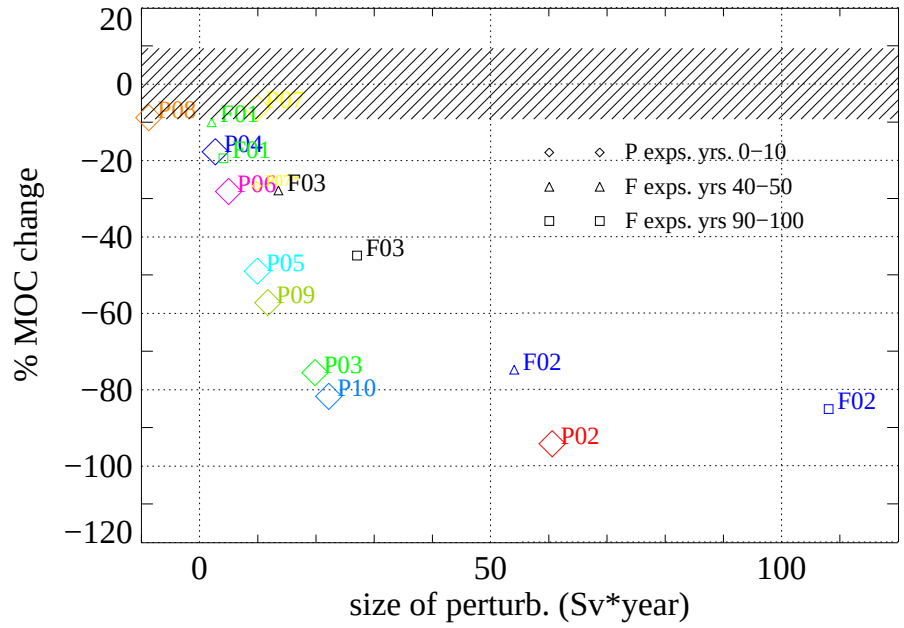

(a) MOC change (vertical) vs. North Atlantic freshening (horizontal) during: yrs 0-10 P exps. (diamonds), yrs. 40-50 of F exps. (triangles) and yrs. 90-100 of F exps. (squares). For P07 a second (smaller) marker with asterisk shows weakening in decade three.

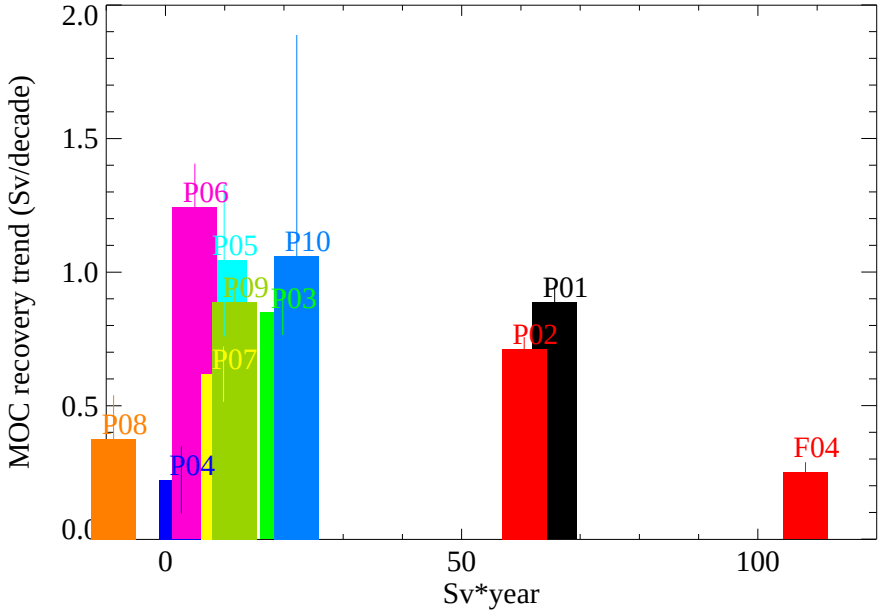

(b) Linear trends in MOC (vertical, in Sv/decade) vs. North Atlantic freshening (horizontal, in $Sv * year$). Average (median) value of the trend is 0.74 (0.85) Sv/decade). Thin lines indicate the error in each estimate.

Figure 3: Sensitivity of MOC at $30^o N$ in the Atlantic to the amount of freshening in P and F experiments.

mechanism, it operates at about the same effi ciency for perturbations of up to 60 Sv*years. Only in F04, where over 100 Sv*year of freshwater has been released over the ocean, is there

no significant change over the 30 years after the anomalous surface flux was removed. It may be the sheer amount of relatively fresh water that has accumulated in the North Atlantic why recovery in F04 is so slow. Possibly, differences in climate feedbacks between P and F experiments are responsible for the difference in recovery rate. Mechanisms by which the MOC recovers in the experiments will be discussed in the next section.

3 Recovery mechanisms

Experiment P03 has been analysed extensively in Vellinga *et al* (2002), and most aspects of the evolution of climate anomalies are well understood. We will therefore consider it to be the 'baseline' experiment to which the other experiments can be compared. In particular, we want to determine to what extent response and processes in the other experiments can be characterised by those found to be important in P03 (as documented in Vellinga *et al* (2002)). A useful view of the analysis that will be presented in this section is that it will determine to what degree the response in the other experiments lies 'in the direction' of that of P03.

a Characterisation of response patterns

To enable a comparison between anomaly patterns $\vec{\mathcal{A}}(t)$ (e.g. precipitation) in the various experiments, we define an orthonormal basis of M spatial patterns $\hat{e}_i$ onto which $\vec{\mathcal{A}}(t)$ is expanded:

$$\vec{\mathcal{A}}(t) = \sum_{i=1}^{M} \alpha_i(t)\hat{e}_i \tag{1}$$

($\vec{}$ further omitted). For each pattern or mode $\hat{e}_i$ there is an expansion coefficient α_i, that evolves in time (t). We choose a basis in which the first basis element $\hat{e}_1$ is the anomaly pattern from baseline experiment P03 (or 'canonical pattern') that we want to identify in $\mathcal{A}$ (see the appendix for more details.) To identify to what extent the response in (say) P07 contains that pattern from P03, we first calculate for each t the projection $\Pi_1(t)$ of $\mathcal{A}$ from P07 along $\hat{e}_1$: $\Pi_1(t) \equiv \alpha_1(t)\hat{e}_1$. Next, the time evolution of the Euclidian norm $\|\Pi_1(t)\|$ of the projection is compared with that of the total anomaly $\|\mathcal{A}(t)\|$, and with that of the residue $\|\mathcal{R}(t)\|$ (where $\mathcal{R}(t) = \mathcal{A}(t) - \alpha_1(t)\hat{e}_1$).

Because of the orthonormality of the basis, the Euclidian norms have the property that $\|\mathcal{A}(t)\|^2 = \|\Pi_1(t)\|^2 + \|\mathcal{R}(t)\|^2$. A useful measure of the importance of pattern $\hat{e}_1$ in determining the total anomaly $\mathcal{A}$ is the 'squared fractional norm' (or SFN): $\|\Pi_1(t)\|^2/\|\mathcal{A}\|^2$. If the total anomaly $\mathcal{A}(t)$ is 'close' to the canonical pattern (i.e. $\|\Pi_1(t)\| \approx \|\mathcal{A}(t)\|$ and $\|\mathcal{R}(t)\| \ll \|\mathcal{A}(t)\|$), then SFN ≈ 1. If the anomaly is significantly different form the canonical pattern: (i.e. $\|\Pi_1(t)\| \ll \|\mathcal{A}(t)\|$ and $\|\mathcal{R}(t)\| \approx \|\mathcal{A}(t)\|$), SFN will be near-zero.

It is straightforward to extend this analysis to more canonical patterns to test $\mathcal{A}$ against (see Appendix). In that case we have a projection $\Pi_i(t) \equiv \alpha_i(t)\hat{e}_i$ for each additional canonical pattern, and its norm should be considered in the manner described above for Π_1.

b Adjustment of Atlantic Ocean freshwater budget

Vellinga *et al* (2002) show how the recovery of the MOC in P03 is linked to the evolution of the salinity field in the North Atlantic. We will therefore analyse anomalies in the Atlantic

freshwater budget $\Phi(y,t)$:

$$\Phi(y,t) = \int_{y}^{y_0}\int SF\,\mathrm{dxdy} + FWT(y) \tag{2}$$

where $\Phi(y,t)$ is the sum of the net surface flux of freshwater between latitudes y and y_0 and the northward freshwater transport across latitude y.[§] If $\Phi(y,t) = 0$ it means the volume of Atlantic Ocean between latitudes y and y_0 is in freshwater balance, and experiences no change in volume-average salinity at given time t. If Φ is negative then there is net removal of freshwater from the volume, and salinity will increase. We will choose $y_0 = 75^o N$, and consider the North Atlantic excluding the Arctic.

From P03 we take the anomaly of Φ after 6 decades as our first mode $\hat{e}_1$ (Fig. 4). It represents a re-distribution of freshwater within the North Atlantic: from being small in the north, there is an increasing imbalance in the freshwater budget at more southerly latitudes, reaching a maximum near $20^o N$. It tapers off further south and is relatively small in the Southern Hemisphere.

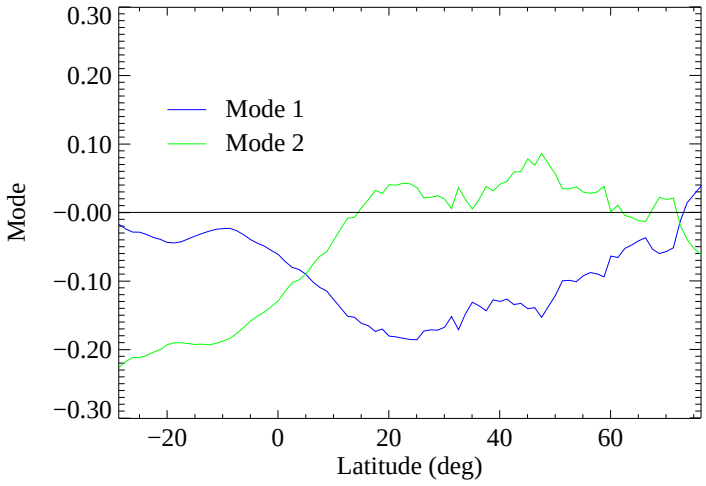

Figure 4: Latitudinal profiles of $\hat{e}_1$ (blue) and $\hat{e}_2$ (green) of anomalies of the Atlantic freshwater budget Φ

In the first few decades of initial perturbation experiments the evolution of Φ depends strongly upon the details of the initial condition (i.e. the applied perturbation). Consequently, projection of $\mathcal{A}$ along $\hat{e}_1$ is small, and $\|\mathcal{R}\|$ is relatively large. After the first few decades, the effect of initial conditions weakens. The time evolution of the norms of projection, $\|\Pi_1\|$, and of the total anomaly, $\|\mathcal{A}\|$, are quite close in all initial perturbation experiments, except in P07 and P08 (Fig. 5). This means that $\hat{e}_1$ describes much of how anomalies of Φ evolve, once the influence of the initial conditions has tapered off. (This early stage is marked by the relative importance of the residual $\mathcal{R}$.)

[§] In the terminology of Pardaens *et al* (2003) Φ is the sum of the implied and the actual freshwater flux.

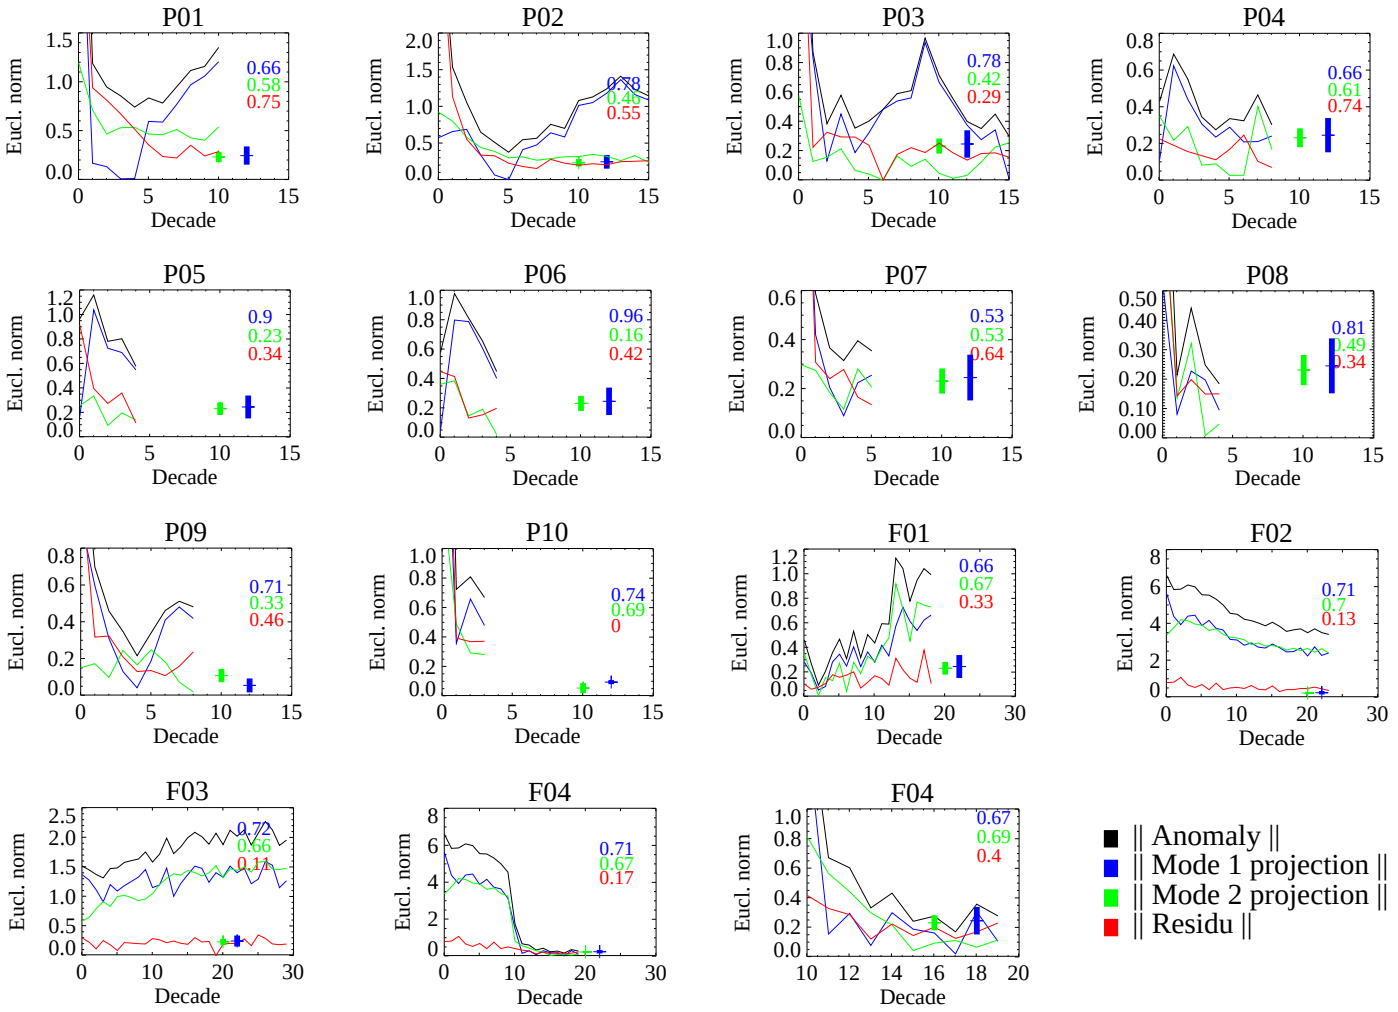

Figure 5: Time series of $||A||$ (black), $||\Pi_1||$ (blue), $||\Pi_2||$ (green) and $||R||$ (red). The coloured crosses in the lower right-hand corner are an indication of typical norms consistent with internal variability: shown are the median value (horizontal line) and deviation from the median (vertical range) of the norms of modes 1 (blue) and 2 (green) in a 100 year slice of the corresponding control run. The coloured numbers in the right-hand side of each panel are the regression coeffi cients for each of the norms in the multiple linear regression: $||A|| = a_0||\Pi_1|| + a_1||\Pi_2|| + a_2||R||$ and indicate the relative contribution of each of the norms to $||A||$.

Version 2.0, May 6, 2004

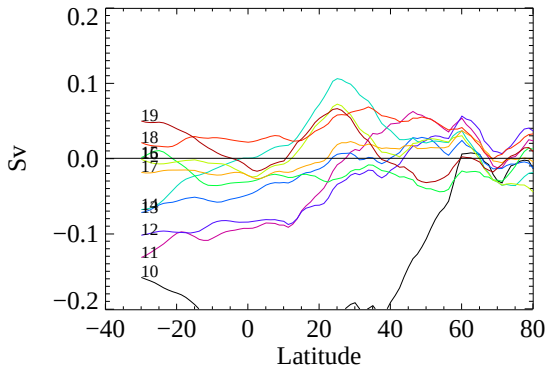

Figure 6: Latitudinal profiles of Φ for the last 10 decades of F04. Decade nos. are shown at the left of each curve. Decade 10 is the first decade without anomalous flux forcing.

We conclude that in HadCM3 the generic response of the freshwater budget to an instantaneous salinity perturbation that covers the northern deep-water formation regions, and that is strong enough to seriously disrupt the MOC consists, to leading order, of a hemispheric redistribution of the initial anomaly, and of a small export of freshwater out of the Atlantic basin across $30^o S$. This mechanism of 'flushing' out the initial salinity perturbation from the north Atlantic is responsible for the long-term recovery of the MOC.

In the case of the flux perturbation experiments (panels marked 'Fnn', Fig. 5), mode 1 alone is not capable of describing changes in Φ. A second mode $\hat{e}_2$ (Fig. 4, green profile) is needed to adequately describe the evolution of the freshwater budget. Mode 2 is mainly a Southern Hemisphere signal, complementing the Northern Hemisphere freshwater imbalance described by mode 1. In flux perturbation experiments Φ the expansion coefficients $\alpha_{1,2} < 0$. This means that, in contrast with initial perturbation experiments, the Atlantic becomes fresher in time. Note how in the forced case F01-03 $\|\Pi_1\| \approx \|\Pi_2\|$, i.e. modes 1 and 2 are excited to a similar extent. The sum of $\hat{e}_1$ and $\hat{e}_2$ adequately captures the phase of the flux perturbation runs with forcing applied, leaving a relatively small residual (panels F01-F04, Fig. 5).

In years 100-200 of F04 (i.e. after the freshwater flux has been removed) the projection of Φ onto $\hat{e}_1$ and $\hat{e}_2$ neither exceeds the values for the control simulation, nor the residual. The weak projection of $\mathcal{A}$ in this case onto recovery modes $\hat{e}_{1,2}$ is consistent with the slow recovery of the MOC in F04 after year 100 (cf. Fig. 2). The temporal evolution of $\Phi(y)$ during this phase of F04 is shown in Fig. 6. Initially, there is a tendency for the Atlantic to export freshwater at a rate of about $-0.15\,Sv$. In the North Atlantic Φ changes little, however, and by decade 16 $\Phi > 0$ almost everywhere. The salinity and velocity structure in the Atlantic in F04 are fundamentally different from the control run (not shown). This altered 'mean state' of the ocean is why the adjustment process is so different from those in the initial perturbation experiments. The generic MOC recovery scenario by way of mode $\hat{e}_1$, as identified in most Pnn experiments, is not applicable in those situations where the state of the Atlantic Ocean in some sense is 'too dissimilar' from that of the control integration.

In the 6 decades of P07 anomalies of Φ are very erratic in nature, and no single pattern can be identified to describe Φ. In P08 neither changes in the North Atlantic MOC, nor of the two modes $\|\Pi_{1,2}\|$ exceed those of the control simulation (Figs. 1 and 5).

The time evolution of $\hat{e}_1$ of Φ is described by $\alpha_1(t)$ (cf. Eq. 1). Surprisingly, there is no clear correlation between $\alpha_1(t)$ and concurrent changes in the MOC (not shown). This

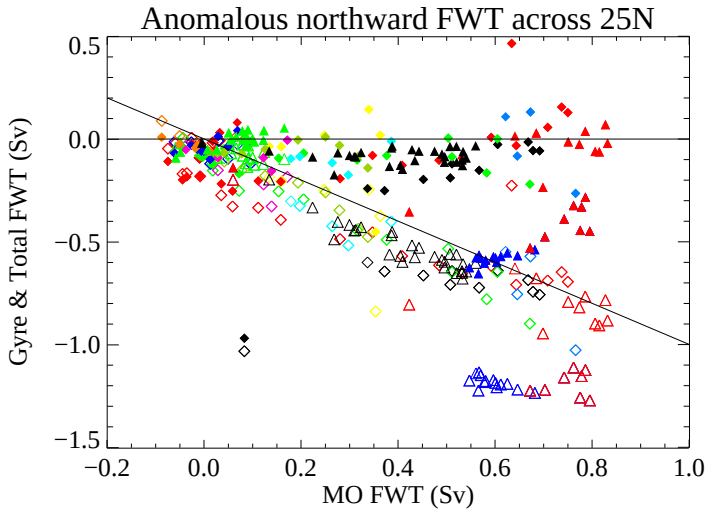

Figure 7: Changes in the northward freshwater transport at 25^oN (in Sv). Horizontal axis: changes in the freshwater transport carried by the overturning circulation; vertical axis: changes in the freshwater transport carried by the gyre circulation (open symbols), and the total circulation (closed symbols). Diamonds (triangles) refer to P (F) experiments. Colours are as in Fig. 3a. Also shown is the line $y = -x$

means that the main mechanism for the recovery of salinity operates at a rate that is largely independent of the degree of disruption to the MOC. It turns out that this is because anomalies in freshwater transport at 25^oN ('FWT' in Eqn. 2) are uncorrelated with changes in the MOC¶. It is found that the net amount of excess southward freshwater transport across 25^oN is rather similar for most of the perturbation runs (filled symbols, Fig. 7). This is consistent with the near-constant rate of MOC recovery shown in Fig. 3b. The anomalies in the amount of freshwater transported by the overturning circulation and the gyre circulation are largely compensating (open symbols, Fig. 7). The slight domination of the latter results in the anomalous southward transport, linked to the recovery of high-latitude salinity.

We currently do not understand the near-invariance of the freshwater transport at 25^oN. However, a linear regression analysis between freshwater transport and salinity (Fig. 8), and velocity (Fig. 9) fields at 25^oN reveals interesting features of changes that are occur in all experiments. Not unexpectedly, one finds a strong correlation between salinity anomalies and freshwater transport anomalies: fresh anomalies occur when transport anomalies by the meridional overturning are positive (i.e. when the MOC is weak). Salinity anomalies tend to be restricted to the upper $1000\,m$ of the ocean. They tend to be strongest in the eastern part of the basin. Anomalies in the velocity field have a more complex structure: when the MOC is weak, the mean current system in the western part of the basin (i.e. Gulf Stream and Deep Western Boundary Current, Fig. 9, top panel) is reduced (Fig. 9, middle and bottom panels). However, the (southward) Sverdup return flow, east of about 65^oW increases in strength. Together with the spatial distribution of the salinity anomalies it is most likely this extra southward flow that is responsible for the compensating southward gyre freshwater transport at 25^oN, Fig. 7.

It is clear that in experiments F02-04 (Fig. 7, red triangles) the excess freshwater transport across 25^oN is less than the amount of freshwater applied via the anomalous surface flux. E.g.

¶ changes in net surface flux ('SF') between $25 - 75^oN$ are very small in all runs, and $\Phi(y = 25^oN)$ is dominated by $FWT(y = 25^oN)$ in Eqn. 2

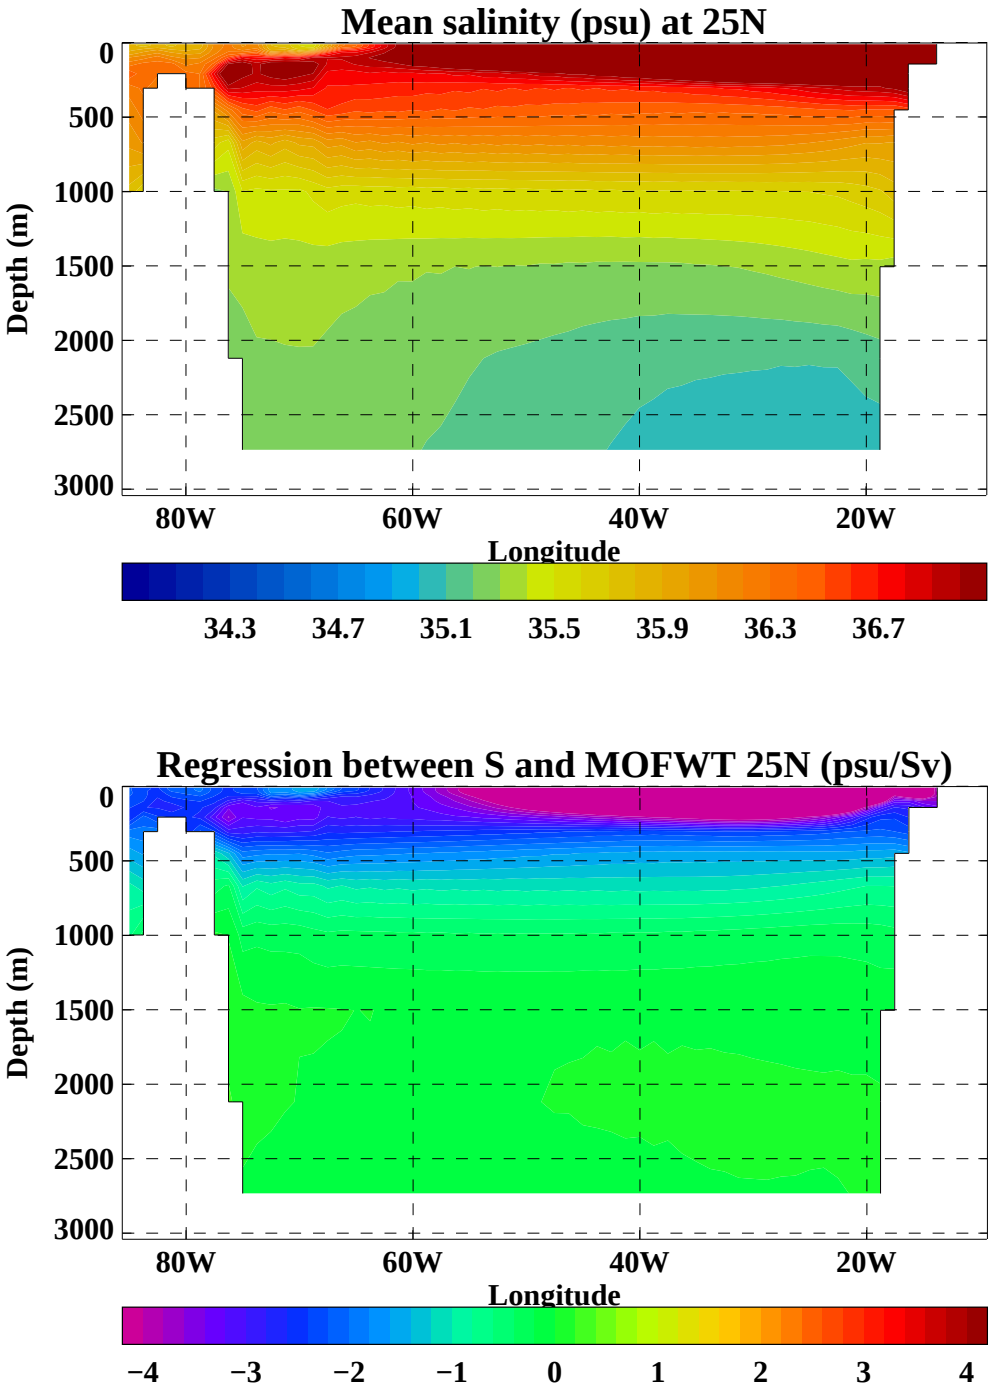

Figure 8: Cross section at 25^oN in the Atlantic of: **top**: time-mean salinity of control integration; **bottom**: linear regression between salinity anomalies and anomalies in meridional freshwater transport by the meridional overturning. Data from exps. P01-10 and F01-04 were converted into anomalies relative to their respective control simulations (about 190 decades of data in total) to calculate the regression.

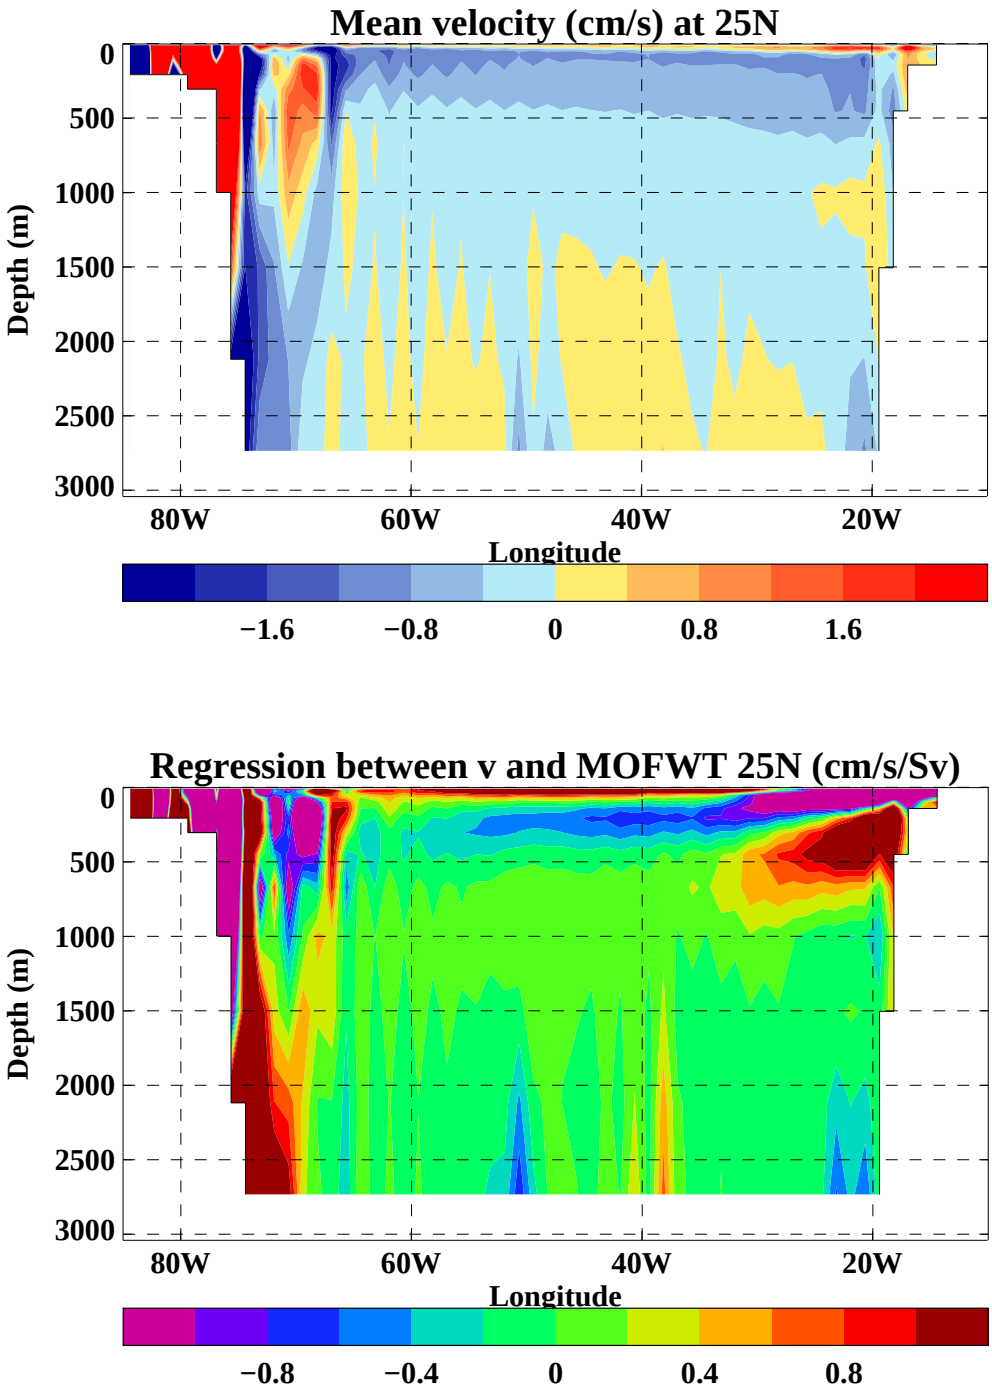

Figure 9: As Fig. 8, but for meridional velocity.

in F04 the net export converges to a value of about $0.6\,Sv$, whereas the applied flux is $1.1\,Sv$. It indicates that the ocean circulation is unable to remove the extra freshwater applied through the anomalous surface flux forcing, and the accumulation of freshwater north of 25^oN causing the gradual slow-down of the MOC.

c Sea surface temperature response

Here we will assess the sea surface temperature ('SST') response in the various experiments, and its dependence on the degree of MOC weakening. As mode 1 we have chosen the normalised SST anomaly of P03 in years 20-30 (Fig. 10a). It has a dipole structure around the equator, with a strong Northern Hemisphere signal (warming in this representation), and a weak Southern Hemisphere signal of opposite sign. The signal is particularly strong over the zonal extension of the North Atlantic Current. These areas correspond to places where under unperturbed conditions there is strong heat loss from the ocean to the atmosphere.

We have projected the SST anomalies from all experiments onto this anomaly (Fig. 10b). The anomalies project negatively onto pattern $\hat{e}_1$, i.e. a weak MOC corresponds to a cold (warm) Northern (Southern) Hemisphere. There exists a near-linear relationship between the degree of MOC weakening, and the projection coeffi cient $\alpha_1(t)$ for those MOC perturbations that exceed the amplitude of natural variability typical of the control simulation. Furthermore, for MOC disruptions that are stronger than about 4 Sv, the squared fractional norm ($\|\Pi_1\|^2/\|\mathcal{A}\|^2$) exceeds 70%, which means that mode 1 strongly dominates the total SST anomaly.

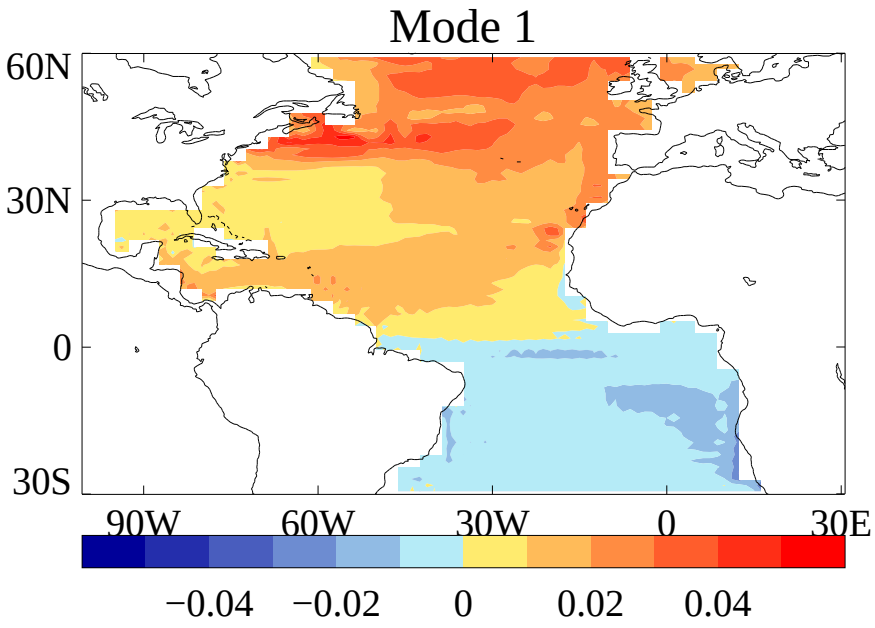

(a) $\hat{e}_1$

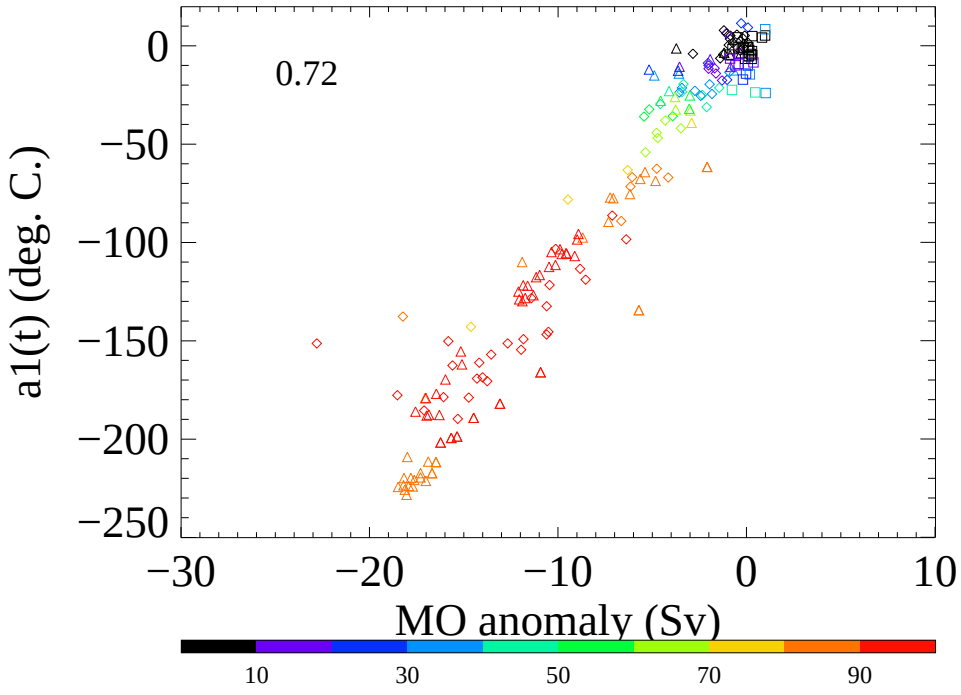

(b)

Figure 10: Sea surface temperature anomalies. (a): Normalised pattern (b) Scatterplot of expansion coeffi cient $\alpha_1(t)$ (vertical) vs. MOC anomaly (horizontal). Diamonds indicate data from P01-10, triangles data from F01-04, and squares data from control integration. Colours stratify data according to the ratio $\|\Pi_1\|^2/\|\mathcal{A}\|^2$, expressed as a percentage. This quantity is a measure of the prominence of mode 1 in the total anomaly $\mathcal{A}$. The linear correlation coeffcient is shown in the left.

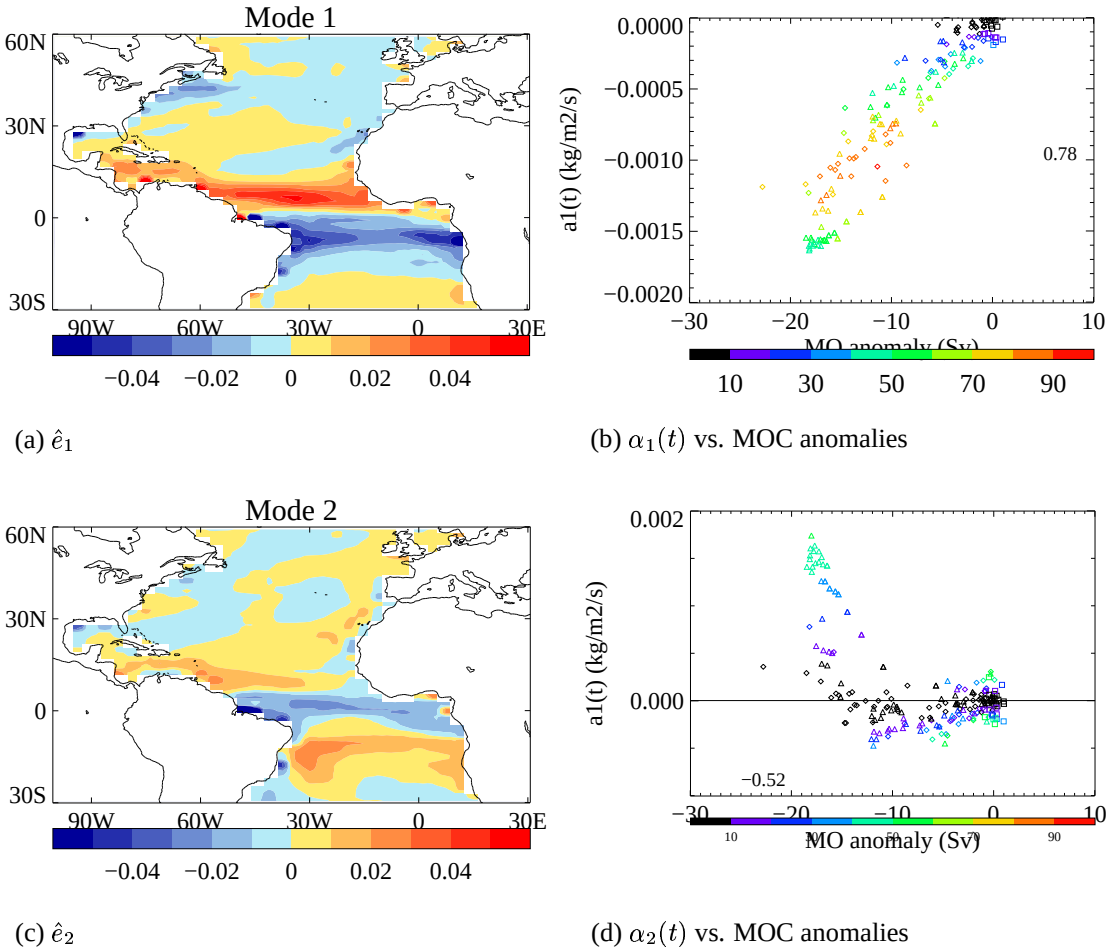

(a) $\hat{e}_1$

(b) $\alpha_1(t)$ vs. MOC anomalies

(c) $\hat{e}_2$

(d) $\alpha_2(t)$ vs. MOC anomalies

Figure 11: Surface freshwater flux anomalies (P-E+R-I). (a): Mode 1 (b) Scatterplot of expansion coeffi cient $\alpha_1(t)$ (vertical) vs. MOC anomaly (horizontal). Diamonds indicate data from P01-10, triangles data from F01-04, and squares data from control integration. Colours stratify data according to the ratio $\|\Pi_1\|^2/\|\mathcal{A}\|^2$, expressed as a percentage. This quantity is a measure of the prominence of mode 1 in the total anomaly $\mathcal{A}$. The linear correlation coeffcient is shown in the left. (c) Mode 2 (d) As (b), but for α_2, and colouring shows values of $\|\Pi_2\|^2/\|\mathcal{A}\|^2$.

d Precipitation minus evaporation response

The way in which the net surface freshwater flux responds to disruption of the MOC is important, because it can support MOC recovery, as happens in P03 in (Vellinga *et al* (2002)). Anomalies in net surface flux (abbreviated as 'PERI' [‖]) are projected onto the anomaly taken from P03 in years 20-30 (Fig. 11a). The projection coeffi cients $\alpha_1(t)$ have a strong (negative) linear relationship with the amount of MOC weakening (Fig. 11b): for a strong MOC weakening, there is strong net input (removal) of freshwater from the tropical south (north) Atlantic, and a net input over the western part of the North Atlantic Current. By and large the PERI anomaly of mode 1 (experiment P03) is typical of the response in the other experiments, and its squared fractional norm exceeds 70%. However, in some cases of strong MOC anomalies the squared fractional norm of PERI mode 1 is about 50% at most. Therefore, we have identifi ed a second mode, orthogonal to mode 1, onto which most of the residuals project (Fig. 11c-d). The effect of adding the tri-banded pattern $\hat{e}_2$ to the bi-polar pattern $\hat{e}_1$

[‖] i.e. the sum of Precipitation minus Evaporation plus River-runoff minus brine-rejection due to sea Ice formation.

(with $\alpha_1(t) < 0$ and $\alpha_2(t) > 0$), is to shift the bi-polar pattern of mode 1 further to the south. Maximum and minimum PERI anomalies now occur further south. Positive projection onto $\hat{e}_2$ is relatively important for the data in the cluster of triangles, which are from F02 and F04 (these are states with essentially a dead MOC). For some of the more moderate MOC disruptions, there is a signifi cant negative projection onto $\hat{e}_2$ (i.e. $\alpha_2(t) < 0$), which, when added to mode $\hat{e}_1$, shifts the PERI anomaly slightly further north.

We conclude that the basic PERI response to MOC disruptions is described by pattern $\hat{e}_1$, its amplitude depends linearly on the amount of MOC weakening. Most of the residual PERI response is a slight northward or southward displacement of the bi-polar pattern of $\hat{e}_1$.

4 Response to MOC collapse under greenhouse gas forcing

It is worthwhile to determine if climate anomalies due to a sudden MOC collapse are similar in a pre-industrial and in a GHG climate. To that end we carried out experiment P10, in which the perturbation of P03 was applied again, but to a climate state taken from year 2049 of an IS92a scenario (the original experiment from which this state was taken will be referred to as 'GHG'). If climate response is linear one expects the anomaly "P10-GHG" to be similar as the anomaly "P03-control". The analysis here will be limited to surface air temperature and precipitation.

Time series of area-averaged surface air temperature show that the amount of global cooling after an MOC collapse in GHG in year 2049 is stronger than what it would be in a pre-industrial climate (-1.4^oC vs. -0.93^oC, Fig. 12a). This difference in global cooling is inconsistent with internal climate variability under pre-industrial conditions. The enhanced cooling under GHG conditions can only be explained by different relative strengths of climate feedbacks that are caused by anthropogenic climate change. Exploring the details of these differences are beyond the scope of the present report, but may be analysed at a later time. Presently we restrict the analysis to the conclusion that over much of the globe the temperature anomaly pattern "P10-GHG" is generally colder than the cooling pattern "P03-control", perhaps reflecting different cloud properties in GHG and control (Fig. 12d.) Red areas indicate where there is less cooling in GHG than in control are in the northern subpolar regions and the Arctic. Presumably, the warm SSTs in GHG prevent the strong increase in sea ice cover that follows the MOC collapse in a pre-industrial climate (Vellinga *et al* (2002)). This lack of increased sea ice cover would prevent the strong cooling of the overlying air observed in P03.

The rate of increase of global temperature in the course of P10 is comparable to that of the standard IS92a 'climate change' of GHG. It is not in any way augmented by the rate of temperature increase associated with the recovery of the MOC, as apparent in the temperature curve for P03.

Average cooling over the Northern Hemisphere and over Europe after MOC collapse in 2049 is also stronger under GHG conditions than it is under pre-industrial conditions (Fig. 12b-c). However, the effect of enhanced cooling over Europe is weak, and it is consistent with internal variability of pre-industrial climate. Effects of enhanced warming and cooling over Europe largely cancel (Fig. 12d).

To assess the combined effect of global warming under an IS92a scenario and an MOC collapse hypothetically occuring in 2049, we have calculated the difference between the average temperature in years 0-10 of P10, and the control simulation (Fig. 13). Over much of the North Atlantic, the greenhouse gas-induced warming is annulled by the cooling induced by the MOC collapse. This effect is still noticeable over the British Isles, and Scandinavia.

In contrast to surface air temperature, the reduction in global mean and Northern Hemisphere precipitation due to a collapse of the MOC is similar in GHG and in pre-industrial conditions (Fig. 12a-b); differences are small and consistent with internal variability. Only for the average over Europe is there a significant difference in the response between GHG and pre-industrial conditions (Fig. 12c): under GHG conditions the reduction is relatively weak, about half of what is is under pre-industrial conditions. It is mainly in Southwest Europe that the anomaly "P10-GHG" is significantly smaller than the anomaly "P03-control" (i.e. there is less 'drying' of SW Europe under GHG than under pre-industrial conditions).

Although in this section we have attempted to discriminate between significant and non-significant differences of anomalies under GHG and pre-industrial conditions, a proper assessment of the effect of internal variability would require an ensemble of model experiments. Currently, though, this is beyond our computational resources, and some care must therefore be taken when interpreting the results in this particular section.

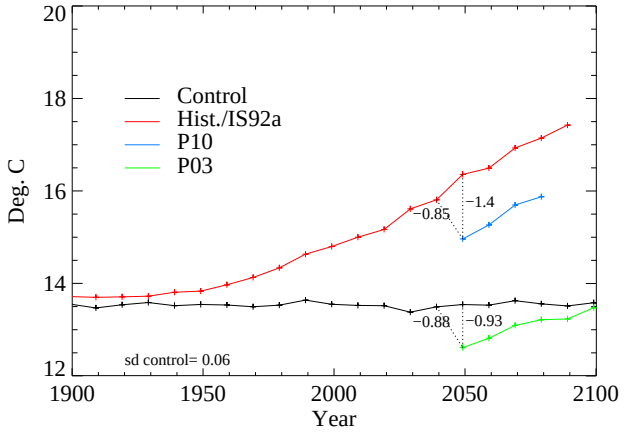

(a) Global

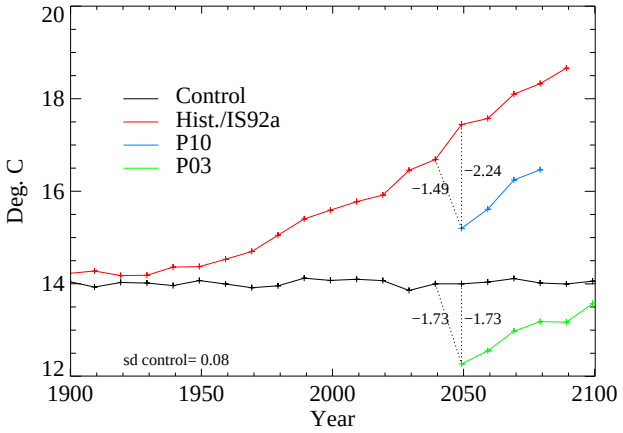

(b) Northern Hemisphere

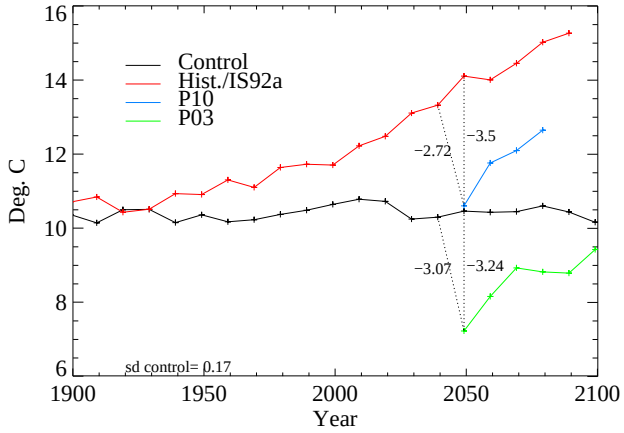

(c) Europe

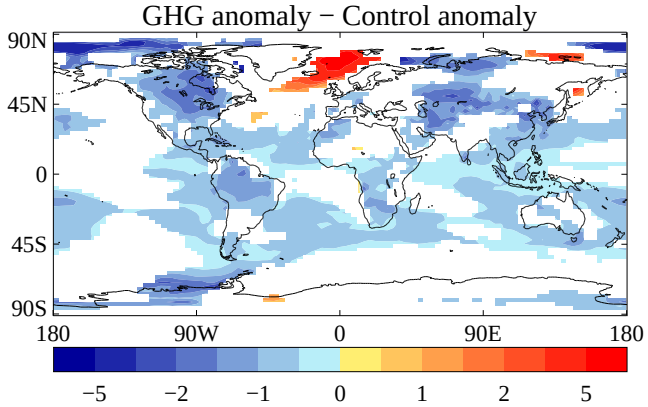

(d)

Figure 12: Time series of decadal mean surface air temperature for (a) the global domain (b) Northern Hemisphere (c) Europe for the HadCM3 control simulation (black), historical/IS92a scenario (red), P03 (green) and P10 (blue). Shown near the dotted lines are: drop in temperature after MOC collapse in 2049, as experienced in consecutive decades relative to respective controls (left number); temperature change in 2049-59 (right number). St dev of the control timeseries is shown in lower left (d): Spatial pattern of difference in air temperature anomaly: "P10-GHG" - "P03-control". Values not exceeding two standard deviations of the control run are consistent with internal variability and have been masked.

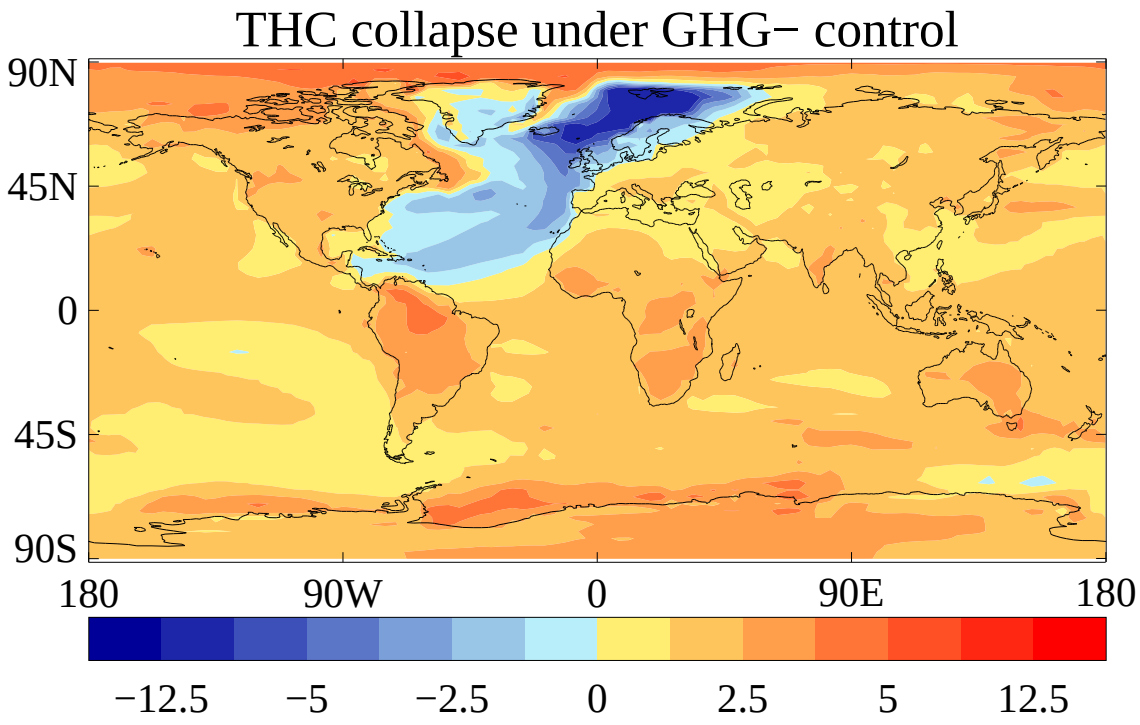

Figure 13: Difference in surface air temperature between year 10 of P10 (i.e. artifi cally induced MOC collapse in IS92a scenario run in 2049) minus long-term average air temperature under pre-industrial conditions.

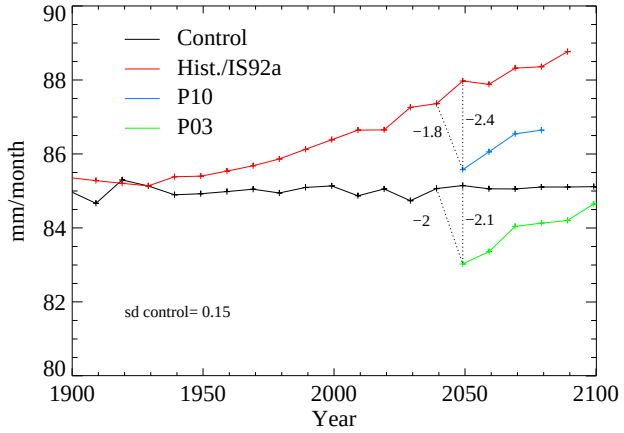

(a) Global

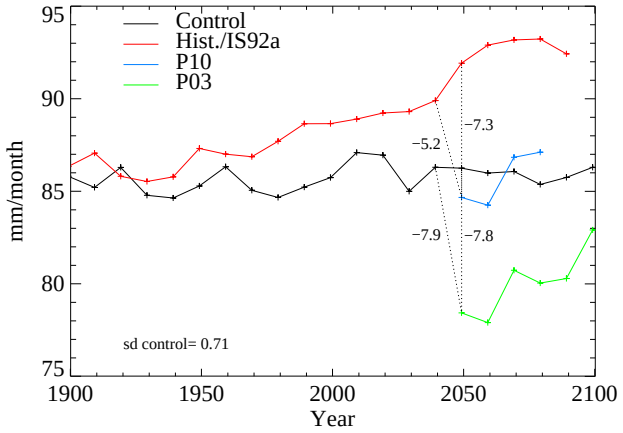

(b) Northern Hemisphere

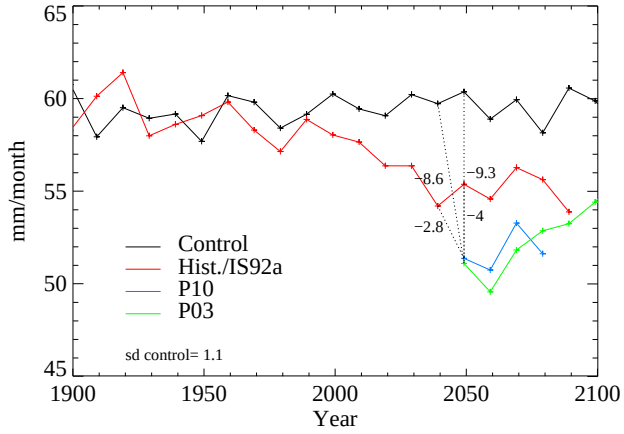

(c) Europe

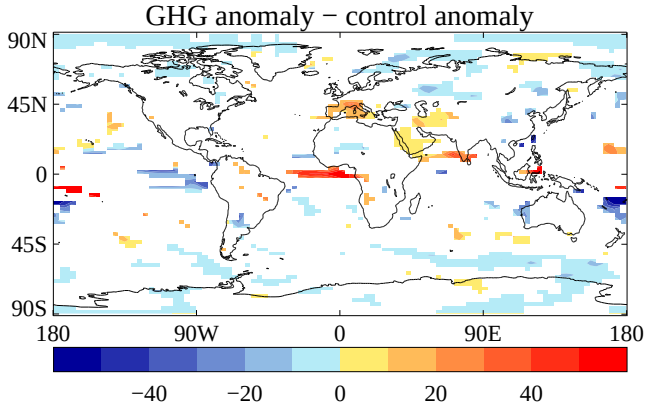

(d)

Figure 14: As in Fig. 12, but for total precipitation. All units are in $mm/month$.

5 Conclusions and discussion

We have analysed climate response in HadCM3 in a number of experiments in which the meridional overturning circulation ('MOC') was disrupted. The main conclusions that can be drawn from this analysis are:

1. In experiments where the perturbation is applied as instantaneous freshening of the North Atlantic there is generally a rapid weakening of the MOC (provided the freshening is sufficiently strong). This weakening is generally followed by a gradual recovery of the MOC (Fig. 1). If the perturbation is applied as continuous forcing by anomalous surface freshwater flux the weakening is gradual, and typically lasts for as long as the forcing is applied. If the forcing is removed, the MOC recovers, albeit slowly (Fig. 2).

2. The degree of MOC weakening depends near-linearly on the magnitude of the perturbation. Perturbations are most effective if applied as instantaneous freshening of northern high latitudes (Fig. 3a). Rate of MOC recovery is approximately constant (Fig. 1), independent on the magnitude of the initial perturbation (Figs. 3b).

3. The dominant process by which high-latitude salinity is restored to original values after the perturbation is a re-distribution of freshwater within the Northern Hemisphere through ocean transports (Figs. 4 and 5). Changes in ocean currents and salinity in the subtropical Atlantic when the MOC is weak are such, that there is little net change in the southward freshwater transport, typical of the unperturbed system (Figs. 7– 9). The invariance of ocean freshwater transport in the subtropical Atlantic is not understood, and is an important model property to be considered more carefully.

4. Climate response in surface variables (e.g. sea surface temperature, precipitation etc.) after disruption of the MOC can be characterised by a low number of generic spatial patterns, that account for much of the total climate anomaly. Typically, the amplitude of these patterns depends linearly on the degree of MOC weakening (Figs. 10-11).

5. Climate response after a forced MOC disruption by 2049 under an IS92a greenhouse gas ('GHG') scenario was compared with a similar disruption under pre-industrial conditions. In the case of surface air temperature the difference of the anomaly patterns in these two cases has a global signature. The amount of global cooling due to MOC collapse under GHG is about 0.5^oC stronger than under pre-industrial conditions (Fig. 12). Over the North Atlantic and on the western fringes of continental Europe, cooling caused by an MOC collapse can be stronger than the warming caused by the increase in greenhouse gas forcing. (Fig. 13). For precipitation, differences are predominantly regional, and the global reduction of precipitation after MOC disruption (corresponding to the colder temperatures) is similar under GHG and pre-industrial conditions (Fig. 14).

We have identified a near-invariance in the freshwater transport in the subtropical Atlantic as an important aspect of the recovery of the MOC. For, if the freshwater transport were strongly reduced in conjunction with a weakened MOC, there would be no obvious process by which high-latitude salinity, and hence the MOC itself, could recover. To what extent this near-invariance reflects a real property of the coupled ocean-atmosphere system, or is a model artefact is at present unclear. This is an important issue, though, that needs to be investigated further. It would be useful if a comparison could be made between modelled and observed variability of the meridional freshwater transport. Lack of observational data on

the one hand, and modelling issues ** on the other hand will make such a comparison a non-trivial task. Possibly, in the medium-to-long term, sufficient observational data will become available from the NERC RAPID thematic programme (e.g. through a planned observational array at 24^oN).

Another possible strategy is to analyse northward freshwater transport in a set of ocean and climate models that are known to have an equilibrium state with a collapsed MOC. At least one other model has a similar invariance as HadCM3, with the meridional freshwater transport apparently independent of the strength (or even presence or absence) of the MOC (UVic. model, J. M. Gregory, *pers. comm*).

The magnitude of the climate response caused by a disruption of the Atlantic MOC does not outweigh the global response to anthropogenic greenhouse gas forcing. But at regional scales, response to MOC collapse can be comparable in size when measured in a number of variables. This means that research attempting to understand to the physics and implications of MOC collapse is relevant, given that a collapse would have implications for climate of the same order of magnitude as 'global warming.'

** Interpretation of freshwater transports in a rigid-lid ocean model is not straightforward.

Appendix: Orthonormal expansion of anomalies

When analysing the temporal and spatial evolution of anomalies, expansion methods are often useful to reduce the number of spatial dimensions ('M'). E.g. for a HadCM3 sea level pressure field $M \approx 7000$, yet a large fraction of sea level pressure variability may be described by the time evolution of only one or two fixed spatial patterns. Expansion into empirical orthogonal functions ('EOFs') consists of constructing a basis of patterns that has the property that the amount of variance obtained when projecting the total anomaly on only a subset of the basis is maximised (e.g. von Storch and Zwiers (1999)). The disadvantage of EOF analysis is that the patterns comprising the orthogonal basis are not necessarily related to physical processes.

Here we will use the knowledge about typical anomaly patterns in baseline experiment P03 (described by Vellinga *et al* (2002)) to reduce the dimensionality of the anomalies. We extend the baseline (or 'canonical') anomaly pattern $\mathcal{C}$ with $M-1$ orthonormal patterns $\hat{e}_i$ $(i = 2, \ldots, M; \|\hat{e}_i = 1\|)$ to form an orthonormal basis that spans the full M-dimensional space. Then, at any time t, a given anomaly $\vec{\mathcal{A}}(t)$ can be expanded as in Eqn. 1

$$\vec{\mathcal{A}}(t) = \sum_{i=1}^{M} \alpha_i(t)\hat{e}_i$$

$\hat{e}_1$ is the normalised canonical pattern $\mathcal{C}$. The $\hat{e}_i$ $(i = 2 \ldots, M)$ can be calculated using Singular Vector Decomposition ('SVD') analyis (Press *et al* (1992)). The $\hat{e}_i$ are orthonormal, i.e. $\langle \hat{e}_i \cdot \hat{e}_j \rangle = \delta_{ij}$. The expansion coefficients $\alpha_i(t) \equiv \langle \vec{\mathcal{A}}(t) \cdot \hat{e}_i \rangle$, i.e. they are formed by the dot product of anomaly and basis function $\hat{e}_i$.

It is possible to supplement $\mathcal{C}$ with n other anomaly patterns, and again use SVD analysis to supplement these with $M - n$ basis elements to form a complete M-dimensional orthonormal basis. The SVD analyis will rotate the n 'input' patterns slightly, so that they are truly orthonormal (within the numerical accuracy used in the SVD algorithm). Because of the large dimensionality M of the problem these rotations are small, and lead only to small modifications of the 'input' patterns.

References

Cubasch, U., G. A. Meehl, G. J. Boer, R. J. Stouffer, M. Dix, A. Noda, C. A. Senior, S. C. B. Raper, and K. S. Yap, 2001: Projections of future climate change. *Climate change 2001: The scientific basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, Houghton, J. T., Y. Ding, D. J. Griggs, M. Noguer, P. van der Linden, X. Dai, K. Maskell, and C. I. Johnson, Eds., Cambridge University Press, 525–582.

Manabe, S., and R. J. Stouffer, 1997: Coupled ocean-atmosphere model response to freshwater input: comparison with Younger Dryas event. *Paleoceanography*, **12**, 321–336.

Pardaens, A. K., H. T. Banks, J. M. Gregory, and P. R. Rowntree, 2003: Freshwater transports in HadCM3. *Clim. Dyn.*, **21**, 177–195.

Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, 1992: *Numerical Recipes in Fortran: the Art of Scientific Computing*. Cambridge University Press, 2nd edition.

Rahmstorf, S., 1995: Bifurcations of the Atlantic thermohaline circulation in response to changes in the hydrological cycle. *Nature*, **378**, 145–149.

Thorpe, R. B., J. M. Gregory, T. C. Johns, R. A. Wood, and J. F. B. Mitchell, 2001: Mechanisms determining the Atlantic thermohaline circulation response to greenhouse gas forcing in a non-flux-adjusted coupled climate model. *J. Climate*, **14**, 3102–3116.

Vellinga, M., and R. A. Wood, 2002: Global climatic impacts of a collapse of the Atlantic thermohaline circulation. *Climatic Change*, **54**(3), 251–267.

——, ——, and J. M. Gregory, 2002: Processes governing the recovery of a perturbed thermohaline circulation in HadCM3. *J. Climate*, **15**(7), 764–780.

von Storch, H., and W. Zwiers, 1999: *Statistical analysis in climate research*. Cambridge University Press.

Wood, R. A., A. B. Keen, J. F. B. Mitchell, and J. M. Gregory, 1999: Changing spatial structure of the thermohaline circulation in response to atmospheric CO_2 forcing in a climate model. *Nature*, **399**, 572–575.

List of Figures