DUPLICATE ALSO

NATIONAL METEOROLOGICAL LIBRARY

OCEAN APPLICATIONS TECHNICAL NOTE 6.

'DISTORTED PHYSICS', TIMESTEP SENSITIVITY AND THE SPINUP OF OCEAN GCMs.

by

R.A.WOOD.

Met Office
FitzRoy Road, Exeter, Devon. EX1 3PB

© Crown Copyright 1995

This document has not been published. Permission to quote from it must be obtained from the Head of Ocean Applications at the above address.

'DISTORTED PHYSICS', TIMESTEP SENSITIVITY AND THE SPINUP OF OCEAN GCMs

Richard A. Wood
Hadley Centre for Climate Prediction and Research,
Meteorological Office,
London Road,
Bracknell,
Berks. RG12 2SY,
U.K.

ABSTRACT

The distorted physics (DP) technique of Bryan and Lewis (1979) is applied to the Hadley Centre ocean GCM, to allow longer timesteps to be taken. Because of the complex nature of the mixing terms in the Hadley Centre model, a number of modifications are required to the model in order for DP to work correctly. These are:

i. increased background horizontal tracer diffusivity, in order to damp baroclinic waves in the Circumpolar Current in the DP system

ii. replacement of the standard maximum slope limitation in the model's isopycnal diffusion scheme with the diffusivity-limiting scheme of Gerdes et al (1991), in order to prevent a timestep sensitivity in the Greenland Sea

iii. removal of the Richardson number dependence of the vertical tracer diffusivity, in order to reduce a timestep sensitivity at the base of the tropical mixed layer.

When these three modifications are applied, the model can be integrated to equilibrium using DP, and shows an acceptably small drift when the integration is continued from this equilibrium state with DP switched off. DP is therefore a feasible technique to reduce the computational cost of ocean-only sensitivity studies. However, the use of DP in the spinup stage of transient climate change experiments is not feasible for the present model, because large drifts are observed on switching between ocean only and coupled modes of integration.

1. INTRODUCTION

Coupled atmosphere-ocean general circulation models (GCMs) are a powerful tool in the study of transient climate change in response to changes in climate forcing due to greenhouse gases, aerosols etc. Typical experiments to study such climate change involve at least two integrations of the GCM. The first integration, with unperturbed or 'present day' forcing, establishes the 'control' climatology of the model, while in the second integration the climate forcing is changed according to some scenario, and the gradual departure of the model climate from the control climate is considered to be the response to the change in forcing.

In order to establish the control climatology, a preliminary 'spinup' integration is usually needed. This allows the model to settle down from its prescribed initial conditions into some quasi-steady climatological state (perhaps including some seasonal and interannual variability). Note that the spinup stage is necessary even if the model is initialised with the present-day mean climate state, since all the present generation of climate models have systematic errors which mean that their climatology is somewhat different from the observed climate; this situation is likely to continue for some years.

An important practical question is, "How long a spinup integration is necessary?" This is important because typically a large fraction of the computing time cost of the experiment is tied up in the spinup phase. Two approaches have been taken in previous experiments:

1. Start the model from an initial state close to present-day climatology, and run it until the drift in the model's heat budget is much slower than the changes expected from the perturbation to be applied in the climate forcing (Murphy 1995).

2. Start the model from arbitrary initial conditions and integrate until there is no secular trend at all in the model's heat budget (Manabe et al 1991, Cubasch et al. 1992)

In each case the timescale to reach the chosen degree of equilibrium is controlled largely by the long timescales of the ocean circulation, since the atmosphere is thought to equilibrate fairly rapidly with any imposed sea surface temperature forcing . With method 1, a spinup time of a few centuries is typically required, whereas for method 2, integrations of several thousand years are necessary; the latter are only possible on the present generation of supercomputers if special techniques are employed to allow long timesteps to be taken - either a highly implicit timestepping scheme as in the 'Large Scale Geostrophic' model of Cubasch et al (see Maier-Reimer et al., 1993), or the 'distorted physics' technique used with the Geophysical Fluid Dynamics Laboratory (GFDL) ocean model, which is the subject of this paper.

Two further points should be made in passing. First, because the spinup time is essentially controlled by the slow physics of the ocean circulation, the spinup phase is usually run in 'ocean only' mode, in which the atmospheric component of the the GCM is replaced by a simple relaxation boundary condition on the surface temperature and salinity. For a model with the same horizontal resolution in atmosphere and ocean, the ocean component typically uses only 15-20% of the total CPU time, so running ocean-only reduces the CPU requirement by a factor of 5 or 6. However, the ocean circulation produced is not always stable to a switch in boundary conditions from ocean-only back to coupled, particularly in sea-ice regions, and this can mean that the control climate of the ocean-only integration is not the same as that of the coupled model (Murphy 1995). While this is not the main subject of the present paper, an example is shown in Appendix A.

Secondly there is the question of which of the strategies 1 and 2 above is preferable. Strategy 2 has the advantage that the control climate is at equilibrium, so that any changes when the climate forcing is perturbed can more confidently be ascribed to that perturbation. On the other hand, strategy 1 produces a climate state which has had less chance to drift away from the observed climatology, and is therefore preferable if the response to the pertubation is thought to be critically dependent on the basic state. Experiments with a simplified ocean GCM (Wood 1993) suggest that which strategy is adopted could make a difference to the results of a transient climate change experiment, but which method is preferable is an open question which is not the topic of this paper. In this paper we consider the so-called 'distorted physics' (DP) method first described by Bryan and Lewis (1979), which is a technique for allowing the long integrations that are necessary to implement strategy 2 with the widely used GFDL ocean model.

Section 2 introduces the DP technique, and discusses some potential pitfalls in its implementation. The Hadley Centre ocean GCM is described in section 3; it differs from other models which have used DP in that it includes a more sophisticated representation of vertical mixing processes near the surface. Associated with this the model also has rather fine vertical resolution (10m) in the upper layers. These features of the model lead to some potential difficulties in using the DP technique, particularly associated with timestep sensitivity in the solutions. Section 4 describes the timestep sensitivities and shows how the model can be formulated to avoid them. Section 5 is a brief summary.

2. DISTORTED PHYSICS

The GFDL ocean model (Bryan 1969a, Cox 1984) is a primitive equation ocean model which uses a largely explicit timestepping scheme. As a consequence of this, the maximum timestep possible in the model for numerical stability is limited by Courant-Friedrichs-Lewy (CFL)-type conditions for the various wave modes supported by the primitive equations (see, e.g., Smith 1985). The fastest of these is the external (barotropic)

gravity mode, which has a typical phase speed of 200ms^{-1}. This mode is filtered out of the model by means of a rigid lid boundary condition, but even so in the Hadley Centre OGCM running at a resolution of 2.5°x3.75°, the maximum timestep possible is currently about 1 hour. This is necessary to resolve fast processes such as barotropic planetary (Rossby) waves and geostrophic adjustment (internal gravity waves), as well as to avoid numerical instabilities associated with steep topography (Killworth 1987). These processes are present in the model but are not thought to have a direct bearing on the longer timescale adjustments (seasonal, interannual and longer) which we are interested in for climate purposes.

The idea of DP is to write down a set of equations which has the same steady (or quasi-steady) solutions as the primitive equations, but in which the above fast adjustment processes have been slowed down, thus allowing a longer timestep to be taken. Provided the fast processes are not slowed down so much that they become slower than the climate processes, the overall adjustment time of the DP model should be about the same as that of the undistorted model, and the amount of computing time saved is therefore in direct proportion to the lengthening of the timestep.

If we write the primitive equation momentum and tracer equations as

$$\frac{\partial u}{\partial t} = (RHS)_u \qquad (1)$$

$$\frac{\partial T}{\partial t} = (RHS)_T \qquad (2)$$

where u represents a horizontal momentum component and T represents a tracer such as temperature or salinity, and the RHS includes all the physical processes in the model, such as advection, diffusion, source and sink terms, then the distorted physics system is

$$\frac{\partial u}{\partial t} = (1/\alpha)\,(RHS)_u \qquad (3)$$

$$\frac{\partial T}{\partial t} = \gamma(z)\,(RHS)_T \qquad (4)$$

where $\alpha \geq 1$ is a constant and γ is a function of z, typically ≥ 1 for all z. Note that the

equilibrium solution ($\partial/\partial t \equiv 0$) of the DP system (3,4) is the same as the equilibrium solution of the original system (1,2), and setting $\gamma \equiv \alpha = 1$ in the DP system recovers the original system. The rationale behind the distorted equations is discussed by Bryan and Lewis (1979) and by Bryan (1984) (note however that $\gamma(z)$ in our notation is $1/\gamma(z)$ in theirs). Essentially the effect of α is to slow down inertia gravity and other fast waves by a factor α, allowing a longer timestep to be taken in the model (typically a factor α longer). γ is normally chosen to be 1 for depths less than 1000m or so, increasing with depth below that level. The effect of this, as can be seen from (4), is to accelerate the very slow advective/diffusive adjustment of the deep water masses, which is the main limiting factor which controls the timescale on which the model approaches its equilibrium state. With $\gamma=1$ at all depths, any choice of α can be made, but choosing $\gamma>1$ effectively speeds up internal waves, and in principle this could make it necessary to choose a larger value of α to maintain stability with a given timestep (Bryan 1984). However in practice it has proved possible to run GCMs with $\gamma=4$ or 8 at the bottom and $\alpha=1$, without shortening the timestep from its standard value (R.J. Stouffer, pers. comm., T.C. Johns, pers. comm.); this suggests that it is not the internal waves which are limiting the timestep in these models, but some other fast processes; in particular for the model used in the present study the 'Killworth instability' (Killworth 1987), an interaction of the barotropic flow with steep bottom topography, appears to be the limiting factor.

Potential Problems with the Distorted Physics Technique

DP is an attractive technique for reducing the computing time needed to spin up an ocean GCM to equilibrium. However there are several potential problems with the method.

1. Seasonal, Interannual and Synoptic Variability

For a model with steady forcing and a steady equilibrium solution it is clear that the equilibrium of (3,4) is the same as that of (1,2). However, most GCMs include a seasonal cycle and so the 'equilibrium' solution is more like a limit cycle, with a period of 1 year. In addition, models typically show some variability on interannual and longer timescales. In these cases, it is possible that the limit cycles of the distorted and undistorted systems may be different. However, this does not appear to be a problem in practice. Provided the value of α is not too large, seasonal 'equilibrium' solutions can be obtained which do not drift substantially when α is set back to 1 (undistorted) and the integration continued (Manabe et al 1991, Wood 1993).

How large can α be in a seasonally forced integration, without distorting the seasonal quasi-equilibrium? The wind stress and thermohaline forcing of the ocean surface vary on a seasonal timescale, and the leading order response of the ocean to this is largely through vertical mixing processes and geostrophic adjustment. Setting $\alpha>1$ in the momentum equation slows down the spinup time for Ekman currents, and for the gravity waves which accomplish geostrophic adjustment, by a factor of order α. These processes

are controlled by the inertial timescale, so at midlatitude the adjustment time is typically a day or so, much shorter than the seasonal timescale; it seems sensible to choose α so that the model can still respond to the forcing in this way. In this paper α is normally chosen to be 24; the geostrophic and Ekman adjustment timescales are thus still less than 1 month, so the model's response to the time-varying forcing should not be too grossly distorted. Other adjustments to the seasonal forcing, at least at low latitudes, may be through Rossby waves. Killworth et al. (1984) have shown that $\alpha>1$ distorts short Rossby waves more than long ones. For $\alpha=24$, the only baroclinic Rossby modes which are significantly distorted have wavelengths which are subgrid scale at the model resolution used here (3.75° long x 2.5° lat); however at higher resolution some resolved waves may be distorted, so care would be required if using DP for high resolution, seasonal integration. In coarse resolution models, because the resolved slow baroclininc adjustments are not much modified by setting $\alpha=24$, it may even be possible to use DP to study transient problems such as the variability of the thermohaline circulation on decadal and longer timescales, and this has been common practice (e.g. Huang 1994). In such cases, however, γ must be set to 1 since if $\gamma\neq1$ the model does not conserve heat or salt and its *transient* behaviour is likely to be incorrect.

It is common practice in some modelling centres to use the annual mean windstress to force the barotropic velocity in distorted physics models, even when seasonally varying heat fluxes are being used (R.J. Stouffer, pers. comm.). Seasonally varying stresses are retained for the baroclinic mode. The original thinking behind this was probably related to the idea that both the barotropic and the baroclinic response to the evolving stress can be thought of as taking place in part through planetary waves. As described above only the shorter waves (i.e. wavelengths of a few Rossby radii) are significantly distorted when $\alpha=24$; this means that the baroclinic modes that are distorted are typically subgrid scale in a coarse resolution model, whereas some resolved barotropic modes may be distorted (typical midlatitude Rossby radii are baroclinic: 50km, barotropic: 2000km). Topographically and coastally trapped waves are also thought to play some part in the ocean's response to seasonal variations in the wind stress (Anderson and Corry 1985), and it is possible that such processes may be distorted. However, test runs with the model used in the present study showed that the annual mean state of the model was very similar whether seasonally varying or annual mean stresses were applied (the seasonal cycle was not examined).

Vertical mixing is another process which controls the response of the upper ocean thermohaline structure to the seasonal cycle in surface forcing. This process is contained in the right hand sides of (2) and (4), and in order not to distort the seasonal cycle in the mixed layer, γ is normally set to 1 for depths down to 1000m. Time-varying mixing is of course possible at deeper levels than this due to deep convective events, but setting $\gamma>1$ below 1000m does not seem to cause problems in practice.

Finally, setting $\alpha>1$ distorts unstable baroclinic waves in such a way that growing modes are found at longer wavelengths than in the undistorted system (Bryan 1984). This

means that in a coarse resolution climate model such as the one used here, baroclinic instability can be resolved in the distorted system even though it is not resolved when $\alpha=1$. In the present model this causes some problems in the Southern Ocean, which are described in Section 4.

2. Multiple Equilibria

Although (3,4) have the same equilibrium solution(s) as (1,2), there is in principle no guarantee that the time evolution of the two systems from a given set of initial conditions will be comparable. This could cause problems in a case where there is more than one equilibrium solution, since a set of initial conditions may lie in the domain of attraction of, say, equilibrium 1 in the undistorted system, but lie in the domain of attraction of equilibrium 2 in the distorted system. DP is usually used to study systems with restoring boundary conditions on both temperature and salinity, and since multiple equilibria have not to date been observed in such systems, this has not been a problem (but see the comments in 1 above on the use of DP in studies of the transient behaviour of the thermohaline circulation).

3. Timestep Sensitivity of Solutions

It is possible (and we shall see that it is in fact the case) that certain aspects of the model solution are sensitive to the length of timestep taken, in the range allowed by the DP technique, so that the steady or quasi-steady solution obtained using DP and a long timestep is different to that obtained using the standard (short) timestep. Parts of the model physics which act in a quasi-instantantaneous way (i.e. which switch on or off within a timestep) are the ones most likely to cause this type of problem. For example, sea ice is formed 'instantaneously' within a timestep, and cuts off subsequent surface heat fluxes; convection and wind mixing mix the water column instantaneously. If there are strong feedbacks through changes in surface fluxes or stratification, a timestep sensitivity is possible.

3. THE MODEL

The model used here is similar to the ocean component of the coupled GCM described by Johns et al. (1995). It is based on the widely-used GFDL ocean model (Bryan 1969a, Cox 1984), and is configured on a regular spherical grid with a resolution of 3.75° longitude x 2.5° latitude. There are 20 levels in the vertical, at the depths given in Table 1. A Kraus-Turner mixed layer model (Kraus and Turner 1967) is embedded in the model levels in the manner described by Foreman (1990), and a Richardson number dependent vertical eddy diffusivity is applied to both momentum and tracers (Pacanowsi and Philander 1981). The isopycnal tracer diffusion scheme of Redi (1982) is also

included.

For the itegrations described here, forcing fields are as follows: wind stress and wind mixing power from Hellerman and Rosenstein (1983), heat fluxes from Esbensen and Kushnir (1981), precipitation from Jaeger (1976), evaporation from Esbensen and Kushnir (1981). In all cases annual mean values are used (no seasonal cycle). For the heat and freshwater fluxes, a relaxation term is also added. For the heat flux this has the form $\lambda(T_C-T)$, where T_C is the annual mean climatological SST from the Levitus (1982) atlas and T is the model SST; this has the effect of relaxing the model mixed layer temperature back towards the climatology, with a timescale determined by the constant λ. Here we take λ=80.18 $Wm^{-2}K^{-1}$, which gives a relaxation timescale of about 29 days for a 50m deep mixed layer; the corresponding coefficient for salinity is chosen to give the same timescale. The Levitus surface salinities are slightly modified by increasing the value to 35.0 psu around the edge of Antarctica; this value is thought to reflect more closely the wintertime values in the area, and has been shown to improve the representation of the deep water masses in a GCM similar to the present one (England 1993).

The methodology is as follows: the model is integrated to equilibrium using DP with α=24, γ=1 at all depths[1] and Δt=1 day. The model is deemed to have reached equilibrium when the global mean temperature at each model level is changing by less than 0.01°C per century, and the salinity by 0.004 psu per century (this typically takes between 400 and 1600 years, depending on how close the initial conditions are to the final solution). From the equilibrium state, the integration is continued for 10-20 years with DP switched off and Δt=1hour. Any drift from the DP equilibrium during the undistorted phase indicates a problem with the DP technique or a timestep sensitivity in the solution.

4. DEVELOPMENT OF MODEL PHYSICS TO OPERATE WITH DP

In order to use DP successfully it is found that a number of modifications are required to the model described in section 3. This section describes those modifications.

A series of preliminary integrations showed that a basic model could be produced which showed no drift on switching to undistorted physics. This model had the mixed

1Although γ=1 was chosen for simplicity, a number of integrations were also carried out with γ as in Table 1. Equilibrium solutions were always the same as the corresponding γ=1 solution, and no additional drifts were found in the continuation integrations.

layer and isopycnal diffusion schemes removed, and the standard convective adjustment scheme in the Cox (1984) code replaced by a scheme which ensures complete mixing of a convectively unstable water column (Marotzke 1991). Restoring these parts of the model one by one demonstrates the drift due to each part. In each case a modification to the model has been found which cures the drift, and these modifications are described below. The experiments are summarised in Table 2.

4.1 Isopycnal diffusion

Experiment A is identical to the basic model run described above, except that the isopycnal diffusion scheme of Redi (1982), as realised by Cox (1987) has been added. This scheme provides a component of diffusion for tracers which is oriented along the isopycnal surfaces. In the model used here the isopycnal diffusivity A_I depends on depth z and is given by $A_I=A_2+(A_1-A_2)\exp(-z/D)$, where $A_1=2000m^2s^{-1}$, $A_2=400m^2s^{-1}$ and D=500m. For reasons of numerical stability, a maximum slope of 10^{-2} is imposed on this component of the diffusive flux. A diapycnal diffusivity is calculated according to the Richardson number dependent scheme of Pacanowski and Philander (1981). In practice, this scheme reverts to a fixed background value κ_B below the mixed layer; we set $\kappa_B=1.0\times10^{-5}m^2s^{-1}$ at the surface, with a linear increase with depth to $1.5\times10^{-4}m^2s^{-1}$ at 5000m. In addition to the isopycnal diffusion, a background horizontal diffusion A_H must also be added to suppress gridscale noise. For the initial experiment this is set to $80m^2s^{-1}$. Note that both the background diffusivity and the maximum slope condition imply a significant diapycnal diffusive flux in regions where the isopycnals slope steeply. The isopycnal diffusion scheme can produce a very large ($O(20m^2s^{-1})$) vertical component of diffusivity where the isopycnal surfaces slope steeply. To maintain numerical stability with these diffusivities an implicit timestepping scheme is required for the vertical part of the diffusion equation; as we shall see in section 4.3 and in Appendix B, this scheme can result in a timestep sensitivity.

Fig. 4.1 shows the zonal mean temperature and salinity drifts after 8 years of undistorted integration (starting from the equilibrium state of the DP spinup). Strong drifts of up to 0.3 °C and 0.04 psu can be seen in two areas: the Antarctic Circumpolar Current (ACC) and the Greenland-Iceland-Norwegian (GIN) Sea. On further investigation the drift in the ACC region appears to be associated with some wavelike variability there in the DP spinup integration. Fig. 4.2 shows the temperature difference between the last year of the DP spinup and a previous 10-year mean, at a depth of 666m (near the depth of maximum drift in Fig. 4.1). A wavelike disturbance can be seen, with a wavelength of order 800km. The contours are distorted into a boomerang shape due to advection by the ACC mean flow. Further investigation suggests a period of around 20-30 years. Because of its wavelike character, this oscillation does not show up in the global mean temperature at this level, which is used to determine convergence of the spinup integration. On switching to the undistorted integration, the wave rectifies itself into a coherent warming trend around 45°S and a cooling around 60°S (Fig. 4.3).

The cause of the variability in the ACC appears to be a form of baroclinic instability. Bryan (1984) shows that in a simple two layer model, the effect of $\alpha>1$ on baroclinic waves is to increase the range of wavenumbers over which instability occurs, and to weaken the growth rate of the fastest growing waves. Taking a value of 50 km for the Rossby radius, Bryan's equation (6.9) suggests that wavelengths of order 1000 km are stable when $\alpha=1$, but unstable (with an e-folding time of 140 days) when $\alpha=24$. This is consistent with the behaviour of the model. Because the instability is weak, a moderate amount of diapycnal diffusion can damp it completely, and so in the basic run (in which the isopycnal diffusion scheme was replaced by a constant horizontal diffusivity of 2000 m^2s^{-1}) no variability or subsequent drift were observed.

With the above in mind, a cure for the drift was found by increasing the background horizontal diffusivity A_H from $80m^2s^{-1}$ to $400m^2s^{-1}$ (Experiment B). This appears to damp any variability in the spinup integration, and there is no drift in the ACC on switching to undistorted integration (Fig. 4.4). The higher value of A_H is more in line with values used by other workers (e.g. Manabe et al. 1991, England 1993), and it has the additional advantage of damping out some of the gridscale noise generated in the lower thermocline by the isopycnal diffusion scheme. On the other hand, it does imply a diapycnal diffusive flux which can be as large as the explicit diapycnal flux (due to κ_B) in regions of steep isopycnal slopes. This has a significant effect on the water mass properties in the model; Fig. 4.5 shows that the thermocline is thicker and deep ocean temperatures higher in Experiment B than in Experiment A, consistent with the well-known effects of increasing diapycnal diffusivity (Bryan 1987).

The only remaining problem in Experiment B is the drift in the GIN Sea (Fig. 4.4). The plus/minus signal in the vertical suggests that vertical mixing (perhaps by convection) may be involved. A preliminary investigation of this region in the model showed that although the water column is strongly stratified in terms of temperature and salinity, its potential density is almost uniform in the vertical. In other words, the isopycnal surfaces slope very steeply, steeper than the maximum permitted slope of 10^{-2}. In this case the slope of the diffusive tracer fluxes is reset to 10^{-2}; this preserves numerical stabilty, but at the cost of introducing a large diapycnal component to the diffusive flux. Detailed heat budget diagnostics suggested that this diapycnal flux was playing an important part in maintaining the equilibrium state, and the possible interaction of this flux with the convective adjustment was thought to be a potential source of timestep sensitivity.

In an attempt to reduce the diapycnal tracer fluxes, a modification to the isopycnal diffusion scheme, developed by Gerdes et al. (1991) (hereafter referred to as GKW), was introduced. In the GKW scheme the diffusive fluxes are always directed along the isopycnal surfaces, however steep these become (apart from the effects of the background diffusivity A_H, which is still required in this model). Numerical stability is retained by reducing the isopycnal diffusivity A_I in regions of steep slope. It was found necessary to include a 'complete mixing' convective adjustment scheme alongside the GKW scheme, since small static instabilities left by the standard Cox convection scheme led to numerical

overflows in the GKW code (see below for a discussion of convective adjustment schemes). The effect of these instabilities had previously been masked by the maximum slope criterion.

Experiment C is identical to Experiment B apart from the inclusion of the GKW and complete convective adjustment schemes described above. Comparison of Figs. 4.5b and 4.5c shows that the only significant difference between the DP equilibrium states of the two experiments is in the region just north of the Iceland-Faeroes-Shetland Ridge, where the drift occurs in Experiment B. In Experiment C, the cold (<0°C) water immediately adjacent to the ridge is absent, and the drift in the continuation is virtually eliminated, apart from some minor, gridscale features near the surface (Fig. 4.6). To check that this change is due to the inclusion of the GKW scheme, rather than the complete convective mixing scheme, Experiment B has been rerun with complete convective mixing (but without GKW). Results are almost identical to the original Experiment B, with the drift shown in Fig. 4.4b still present.

4.2 Convection Scheme

The standard algorithm for convective adjustment in the Cox (1984) model scans each water column for pairs of adjacent gridpoints which are statically unstable, and mixes each such pair. At the end of a double pass through the column (testing first levels 1 and 2, 3 and 4 etc., then levels 2 and 3, 4 and 5 etc.), stability of the whole column is not guaranteed, and the whole process is repeated NCON times. In general, all static instabilty is removed only in the limit NCON$\rightarrow\infty$. The process is expensive in terms of CPU time and typically a value of NCON around 5 is used. Marotzke (1991) has demonstrated a strong timestep sensitivity in a version of the Cox model configured in an idealised 'sector' basin, caused by interaction of the convection scheme (Marotzke used NCON=3) with the periodic forward timestep used in the model to prevent time splitting. However other workers running GCMs with more realistic geometry have reported no such problems when using the standard convection scheme with values of NCON in the range 3-7 (M.England and R.J.Stouffer, pers. comm.). The reason for this discrepancy is not clear, but there is evidently a possibility that the convective adjustment scheme could cause a timestep sensitivity in the present model.

Experiments A and B above were each run twice: once with the standard Cox scheme (with NCON=5), and once with an alternative scheme which ensures complete mixing of the water column at each timestep. This scheme removed the timestep sensitivity in Marotzke's (1991) study. The scheme is however computationally expensive, increasing the overall runtime of the model by 30% compared with the standard scheme with NCON=5. Both the equilibrium state and the drift on switching off DP were found to be almost identical with the two schemes, suggesting that the precise choice of convection scheme does not cause any timestep sensitivities in the present model. Heating rate diagnostics (not shown) suggest that much of the vertical mixing in the model is achieved

through the vertical component of the isopycnal diffusion rather than through convection. Danabasoglu et al. (1994) also find that isopycnal diffusion decreases the role of convection in a global GCM. This may be why the simplified model of Marotzke (1991) (which does not include isopycnal diffusion) appears to be more sensitive to the choice of convection scheme than do the GCMs (which normally do include it).

Since the above runs were completed, Rahmstorf (1993, unpublished manuscript) has developed a 'complete mixing' convection algorithm which is more computationally efficient than the original Marotzke algorithm. Experiment C was rerun using this algorithm, with results which were indistinguishable from those using the Marotzke scheme. The overall runtime of the model was slightly less than with the standard Cox scheme with NCON=5. Thus it is possible to implement the GKW scheme to eliminate the drifts associated with isopycnal diffusion, plus the complete convective adjustment scheme, at no extra computing cost.

4.3 Mixed Layer Model

Experiment D is identical to Experiment C except that the mixed layer model has now been added. This model configuration is thus very close to the standard one used at the Hadley Centre in climate change experiments (Johns et al. 1995), except for the removal of the sea ice model and seasonal cycle.

In the initial run of this experiment, the drift in the GIN Sea, as seen in Fig. 4.4, reappeared. This was because in the standard version of the isopycnal diffusion scheme currently in use at the Hadley Centre, isopycnal diffusion is switched off within the mixed layer. The diffusion reverts to a horizontal orientation with diffusivity A_I. This introduces a diapycnal component to the diffusion, in a similar way to the maximum slope criterion, down to the depth of the mixed layer (O(1500m) in the GIN Sea region). Therefore a modification to the standard code was introduced, in which the isopycnal diffusion is switched off in the top 50m only (the same code was in fact used in Experiments A-C, which did not include the mixed layer model. Leaving the isopycnal diffusion scheme switched on in the top 50m resulted in numerical instability). This modification removes the drift in the GIN Sea, leaving the drift shown in Fig. 4.7. The bases of the tropical mixed layer and polar halocline become shallower with the shorter timestep. The timestep sensitivity here appears to arise from a complex interaction betweeen the mixed layer scheme, the implicit vertical diffusion scheme, and the Richardson number dependent diffusivity described in section 4.1 above. This sensitivity is discussed in more detail in Appendix B.

The results of Appendix B suggest that the timestep sensitivity can be minimised by reducing the vertical variations in the diffusivity profile which arise from the low values of Richardson number (weak stratification) in the mixed layer and larger values (stronger stratification) below. Therefore in experiment E the vertical tracer diffusivity is

set to its background value κ_B everywhere, so that the Richardson number dependent diffusivity profile is applied only to momentum. This would be expected to have very little effect on the solution, since in practice the tracer diffusivity is found to take a value of approximately κ_B everywhere outside the mixed layer, whereas within the mixed layer vertical gradients, and hence diffusive fluxes, are zero. Appendix B shows that with an explicit diffusion scheme the effect is indeed negligible, but in the implicit case large diffusivity gradients result in a stronger timestep sensitivity than is seen with a uniform diffusivity.

Fig. 4.8 shows that turning off the Richardson number dependent part of the diffusivity in the GCM substantially reduces the drift in the undistorted continuation. The maximum drift over 8 years is now of order (0.05°C, 0.02psu) in the zonal mean.

6. SUMMARY AND CONCLUSIONS

In this paper the ocean part of the Hadley Centre coupled GCM (Johns et al. 1995) has been modified to accommodate use of the 'distorted physics' (DP) technique of Bryan and Lewis (1979), which allows timesteps of order 1 day to be used, rather than 1 hour which is the maximum posssible in the standard model. The Hadley Centre ocean model contains more complete physics than typical GCMs used in climate change studies, particularly in the area of vertical mixing, and it was found that a number of modifications are required to the standard model in order to prevent timestep sensitivities in the solutions. These modifications are in the schemes for isopycnal diffusion, convection and Richardson number dependent vertical mixing.

One of the main reasons for undertaking the work presented here was to allow a full spinup in ocean only mode in the transient climate change experiments reported by Mitchell et al. (1995). In view of the results of Appendix A, the ocean only spinup strategy does not appear to be appropriate for this model, and a coupled spinup must be used instead (Johns et al. 1995). However, it will still be of value to include the modifications presented here in the model, since DP can then be used to perform the sensitivity and process studies which are essential to the future development of the ocean part of the model.

ACKNOWLEDGEMENTS

I thank the many colleagues who made helpful suggestions on parts of this work and who shared their experience, particularly Mike Bell, Jenny Crossley, Matthew England, Steve Foreman, Chris Gordon, Jonathan Gregory, Tim Johns, John Mitchell, David Roberts and Ron Stouffer.

I also thank Claudia Köberle and Rüdiger Gerdes for supplying code to implement the GKW isopycnal diffusion scheme, Malcolm Roberts for implementing this code in the Hadley Centre model, Stefan Rahmstorf for code to implement his convective adjustment scheme, and Tim Johns for providing the forcing fields from the Hadley Centre coupled model used in Appendix A.

This work was carried out under DOE contract PECD 7/12/37.

APPENDIX A. COUPLED vs. OCEAN ONLY SPINUP

In the Introduction it was noted that the spinup phase of a transient climate change GCM experiment is usually run largely in 'ocean-only' mode in order to reduce the CPU time required to bring the model to equilibrium. In such an integration, the surface fluxes into the ocean are derived from a previous (relatively short) run of the atmosphere only or coupled model (see for example Manabe et al. 1991, Murphy 1995). Typically, monthly or 5-day mean fluxes are derived from the last 10 years or so of the atmospheric integration. In addition, a relaxation term is applied to the sea surface temperature and salinity, as described in section 3 of this paper; this ensures that the surface fields do not drift too far from climatology. The relaxation terms are usually diagnosed at the end of the spinup and used as 'flux correction' terms in the free coupled stage that follows.

It is clear that in the ocean-only phase the ocean model does not see the same surface fluxes as it would if it were being run in coupled mode, for two reasons:

i. Since multi-year monthly or 5-day means are used, the fluxes do not contain a realistic level of short timescale (synoptic) variability, and contain no interannual variability at all.

ii. The feedbacks between the changing SST and SSS and the surface fluxes (through changes in the atmospheric boundary layer and large scale circulation) are not accurately represented.

For these reasons, the spinup phase of the 2nd Hadley Centre transient experiment (Johns et al., 1995) was run in fully coupled mode, although still with the addition of relaxation terms. Since the atmospheric component of the model used in this study shows a higher degree of variability than many similar GCMs, and since the ocean component contains some highly nonlinear physics (for example, the sea ice and mixed layer models), it was anticipated that the ocean-only equilibrium state would be too far from the coupled equilibrium, leading to difficulty when the model was recoupled at the end of an ocean-only spinup. This was confirmed by the experiment described below.

Forcing fields for the ocean model were derived as 5-day means from the last 10 years of the (coupled) spinup phase of the model of Johns et al. 1995. At this stage of the spinup only a slow temporal drift remained in the ocean temperature and salinity fields. The ocean model was set up exactly as in the coupled model (the main differences from the model described in section 3 being the inclusion of the seasonal cycle and a simple dynamic/thermodynamic ice model - see Johns et al. 1995 - and use of a relaxation coefficient λ=163.76 $Wm^{-2}K^{-1}$ instead of 80.18$Wm^{-2}K^{-1}$). The model was run for 50 years (with Δt=1 hour and no DP) starting from the final state of the coupled spinup.

Fig. A.1 shows the annual mean sea ice thickness for year 50 of the ocean only integration, compared with the same field from the last year of the coupled spinup. In

general the ice is thicker in regions where it is formed thermodynamically and exported by advection (e.g. near the Antarctic continent and in the Canadian Basin) and thinner in regions where it is imported by advection and melts (e.g. the ACC and the East Greenland Current). These differences develop monotonically through the 50 year integration and are larger than the level of interannual variability in the coupled spinup.

The above results appear to confirm that an ocean only spinup phase is inappropriate for the Hadley Centre model as described by Johns et al. 1995, although the strategy has been used successfully in the past in other models (e.g. Manabe et al. 1991). The main differences are thought to be that the atmospheric part of the Hadley Centre model exhibits stronger variability over a range of timescales than most GCMs, and that the ocean part contains more strongly nonlinear physics (sea ice and mixed layer model) and has higher vertical resolution near the surface. These differences tend to increase the difference between the surface fluxes in ocean only and coupled modes (i. and ii. above). A detailed examination of this question is, however, beyond the scope of this paper. For further discussion see Johns et al. 1995.

APPENDIX B. TIMESTEP SENSITIVITY IN A COMPLEX VERTICAL MIXING SCHEME

Fig. 4.7 shows a cooling drift at the base of the tropical mixed layer in the undistorted continuation run of Experiment D. This drift is not present in Experiment C, which differs from Experiment D only in that the Kraus-Turner mixed layer model is absent, so the mixed layer model is implicated in the drift. In this appendix a simple one-dimensional model is used to show that the drift results from a timestep sensitivity involving the interaction between the mixed layer model, the Richardson number dependence of the vertical diffusivity (see section 5 'Isopycnal Diffusion') and the implicit scheme used to timestep the vertical diffusion equation.

The model consists of 3 active boxes of density ρ_i (i=1(top),2,3) above a 'reservoir' box whose heat capacity is assumed infinite and whose density therefore remains constant at ρ_0. For simplicity the boxes are taken to have equal volume and all variables are dimensionless. The model densities are driven by a surface buoyancy flux of the form $q_0 + \lambda(\rho_{ref} - \rho_1)$, i.e a basic flux plus a restoring term, as in the GCM, and adjust through diffusion and through the Kraus-Turner mixed layer physics. Thus

$$\frac{\partial \rho_i}{\partial t} = -\delta_{i1}[q_0 - \lambda(\rho_{ref} - \rho_i)] + [K_i(\rho_{i+1} - \rho_i) - K_{i-1}(\rho_i - \rho_{i-1})] + MLM \quad (5)$$

where K_i is the diffusivity at the interface between boxes i and i+1, and MLM represents the mixed layer physics. For all the runs presented here we take $q_0>0$ (positive buoyancy input), so that there is no convective adjustment, and the mixed layer is driven purely by an input of turbulent kinetic energy from wind stirring, of power w. Starting from the surface, two boxes are mixed if there is sufficient TKE available to provide the gravitational PE increase due to the mixing, and the TKE is then depleted by this amount. At the bottom of the mixed layer, there is in general a small amount of TKE remaining, and this is used to perform a partial mixing of the box below the mixed layer with the mixed layer (see Foreman 1990 for details). Parameter values for the runs presented here are chosen so that boxes 1 and 2 are mixed and there is partial mixing of box 3.

The diffusive part of the model is timestepped using an implicit scheme in which all the densities in the diffusive term on the right hand side of (5) are evaluated at time level n+1. This is the method used in the GCM. An explicit solution, in which the right hand side is evaluated at time level n, is also shown later for comparison.

We take q_0=0.1, λ=5.0, ρ_{ref}=9.8 w=0.4, ρ_0=10.0 as a typical example run (all dimensionless units). All three boxes are initialised with a density of 10.0 and the model is run to equilibrium. The diffusivities K_2 and K_3 are taken to be 1.0, but various values in the range 0.1-500 are used for K_1. This reflects the fact that in the GCM the Richardson number diffusion scheme produces diffusivities in the mixed layer which are O(500) times larger than the values outside the mixed layer.

Fig. B.1 shows the results from a number of runs of the model. The equilibrium mixed layer density is found to be relatively insensitive to the timestep Δt, because it is constrained by the strong surface relaxation term, so only the equilibrium value of ρ_3 (corresponding to the gridbox immediately below the mixed layer in the GCM) is shown, as a function of Δt and for various values of the mixed layer diffusivity K_1. The solution with an explicit timestepping scheme for the diffusion is also shown for comparison (the model can be solved analytically in this case, and the equilibrium solution is found to be independent of Δt and K_1; this was confirmed computationally). Considering first the curve for K_1=500 (the closest to the GCM in Experiment D), it can be seen that as Δt increases, ρ_3 decreases. This behaviour (decrease of density at the base of the mixed layer with increasing timestep) is consistent with the drift in the undistorted continuation run in Experiment D (Fig. 4.7a). Note that the timestep sensitivity (the slope of the curve) is stronger at *short* timesteps. As K_1 is decreased, the solution becomes closer to the explicit solution. For all values of K_1, the solution appears to tend to the explicit solution as $\Delta t \rightarrow 0$.

To see whether the timestep sensitivity in the simple model is strong enough to account for the temperature difference seen in Fig. 4.7a, the dimensionless timestep in the simple model must be associated with the dimensional timestep in the GCM. Taking a mixed layer depth of 30m and a diffusivity (outside the mixed layer) of $1.0\times10^{-5}m^2s^{-1}$ in the GCM gives a diffusion timescale of order 1000 days. Therefore a timestep of 1×10^{-3} in

the simple model (the shortest possible before roundoff error began to contaminate the solutions) corresponds approximately to 1 day in the GCM. The equilibrium density ρ_3 in this case is 9.8645, less dense than the explicit solution by about 0.0014 or about $0.007(\rho_0-\rho_{ref})$. Taking a temperature difference of 10°C across the mixed layer suggests a corresponding temperature difference due to the timestep sensitivity in the GCM of 0.07°C, compared with the 0.2°C seen in Fig. 4.7a (and of the correct sign, i.e. denser/cooler with shorter Δt). In view of the simplicity of the model (5), in particular the coarse resolution, fixed deep density, neglect of salinity effects and assumption of equal thickness gridboxes, this agreement is close enough to suggest that (5) captures the essential elements of the timestep sensitivity in GCM Experiment D.

To confirm that the timestep sensitivity is due to an interaction of the mixed layer model with the implicit diffusion scheme, the model (5) was rerun with the MLM switched off, so that the problem is a purely diffusive one. With implicit diffusion there is still a timestep sensitivity in the solution; in this case the sensitivity is strongest for $K_1=1$, and the difference between the implicit solution with $\Delta t=0.12$ and the explicit solution (= the implicit solution as $\Delta t\rightarrow 0$) is around 0.005 or a quarter of the corresponding difference in the case with the MLM included. Thus the timestep sensitivity appears to be an inherent property of the implicit diffusion scheme which is exacerbated by working on the type of vertical density profile produced by the MLM, especially when the (apparently irrelevant) diffusivity within the mixed layer is large.

In the GCM, and in the above simple model, the timestep is performed in the order diffusion, surface fluxes, MLM. To see if the order of calculations has any impact on the results, the simple model was rerun with the timestep ordered surface fluxes, MLM, diffusion; for an equilibrium problem this merely amounts to redefining where the end of the timestep is. For large timesteps the sensitivity is somewhat reduced, but as expected the solutions at short timesteps (<0.02) are very close to the original. Running the timestep in the order surface fluxes, diffusion, MLM results in a *stronger* sensitivity for large Δt, and again little difference from the original solutions for $\Delta t<0.02$. As shown above the timesteps relevant to the GCM are O(0.001) or less, so reordering the code in either of these ways is unlikely to reduce the sensitivity in the GCM.

If we take (5) as our model, we can use Fig. B.1 to suggest two possible cures for the drift in Experiment D:

(i) Use an explicit timestep for the vertical diffusion
(ii) Reduce K_1. Fig. B.1 shows that this makes the solution converge faster to the explicit solution as $\Delta t\rightarrow 0$, and the explicit solution is in any case independent of K_1.

Of these, (i) is thought to be impractical in the GCM because numerical stability constraints mean that an extremely short timestep (O(3s)) would be required, resulting in an unacceptable increase in computational cost. So the simplest approach appears to be

(ii). This was tried in Experiment E, and resulted in a considerable reduction in the drift (Fig. 4.8).

TABLE 1. Vertical levels used in the model and the corresponding values of γ.

Level	Depth of Gridpoint (m)	Thickness of Gridbox (m)	γ
1	5.0	10.0	1.00
2	15.0	10.0	1.00
3	25.0	10.0	1.00
4	35.1	10.2	1.00
5	47.9	15.3	1.00
6	67.0	23.0	1.00
7	95.8	34.5	1.00
8	138.9	51.8	1.00
9	203.7	77.8	1.00
10	301.0	116.8	1.00
11	447.1	175.3	1.00
12	666.3	263.2	1.00
13	995.6	395.3	1.00
14	1501.0	615.0	1.51
15	2116.0	615.0	2.13
16	2731.0	615.0	2.75
17	3347.0	615.0	3.38
18	3962.0	615.0	4.00
19	4577.0	616.0	4.00
20	5193.0	616.0	4.00

TABLE 2. The five experiments described in section 4.

Experiment	Isopycnal Diffusion Scheme	$A_H(m^2s^{-1})$	Mixed Layer Model	Tracer Diffusivity
A	Std	80	N	P&P
B	Std	400	N	P&P
C	GKW	400	N	P&P
D	GKW	400	Y	P&P
E	GKW	400	Y	κ_B

(Std = standard isopycnal diffusion scheme with maximum slope of diffusive flux set to 10^{-2}, GKW = Gerdes, Köberle and Willebrand (1991) scheme, P&P = Pacanowski and Philander Richardson number dependent scheme, κ_B = constant background value).

REFERENCES

Anderson, D.L.T. and R.A. Corry, 1985: Seasonal transport variations in the Florida Straits: a model study. J. Phys. Oceanogr., 15,773-786.

Bryan, F., 1987: Parameter sensitivity of primitive equation ocean general circulation models. J. Phys. Oceanogr., 17, 970-985.

Bryan, K., 1969a: A numerical method for the study of the circulation of the world ocean. J. Comput. Phys., 3, 347-376.

Bryan, K., 1984: Accelerating the convergence to equilibrium of ocean-climate models. J. Phys. Oceanogr., 14, 666-673.

Bryan, K. and L.J. Lewis, 1979: A water mass model of the world ocean. J. Geophys. Res., 85(C5), 2503-2517.

Cox, M.D., 1984: A primitive equation, three-dimensional model of the ocean. GFDL ocean group technical report no. 1, Princeton, NJ, USA, 143pp.

Cox, M.D., 1987: Isopycnal diffusion in a z-coordinate ocean model. Ocean Modelling, 74, 1-5 (Unpublished Manuscript).

Cubasch, U., K. Hasselmann, H. Hock, E. Maier-Reimer, U. Mikolajewicz, B. D. Santer and R. Sausen, 1992: Time-dependent greenhouse warming computations with a coupled ocean-atmosphere model. Climate Dynamics, 8, 55-69.

Danabasoglu, G., J.C. McWilliams and P.R. Gent, 1994: The role of mesoscale tracer transports in the global ocean circulation. Science, 264, 1123-1126.

England, M.H., 1993: Representing the global-scale water masses in ocean general circulation models. J. Phys. Oceanogr., 23, 1523-1552.

Esbensen, S.K. and Y. Kushnir, 1981: Heat budget of the global ocean: estimates from surface marine observations. Report no. 29, Climatic Research Institute, Oregon State University, Corvallis, Oregon, USA, 27pp.

Foreman, S.J., 1990: Ocean model: mixed layer formulation. UKMO Unified Model Documentation Paper no. 41.1, UK Meteorological Office, Bracknell, UK.

Gerdes, R., C. Köberle and J. Willebrand, 1991: The influence of numerical advection schemes on the results of ocean general circulation models. Climate Dynamics, 5, 211-226.

Hellerman, S. and M. Rosenstein, 1983: Normal monthly wind stress over the world's ocean with error estimates. J. Phys. Oceanogr., 13, 1093-1104.

Huang, X.H., 1994: Thermohaline circulation: Energetics and variability in a single-hemisphere basin model. J. Geophys. Res., 99, 12471-12485.

Jaeger, L., 1976: Monatskarten des Niederschlags für die ganze Erde. Bericht des Deutschen Wetterdienstes, 18, no. 139, Offenbach, Germany.

Johns, T.C., R. Carnell, J.F. Crossley, J.M. Gregory, J.F.B. Mitchell, C.A. Senior, S.F.B. Tett and R.A. Wood, 1995: The second Hadley Centre coupled ocean-atmosphere GCM: model description, spinup and validation. In preparation.

Killworth, P.D., 1987: Topographic instabilities in level model OGCMs. Ocean Modelling, 75, 9-12 (Unpublished Manuscript).

Killworth, P.D., J.M. Smith and A.E: Gill, 1984: Speeding up ocean circulation models. Ocean Modelling, 56, 1-4 (Unpublished Manuscript).

Kraus, E.B. and J.S. Turner, 1967: A one-dimensional model of the seasonal thermocline. Part II. Tellus, 19, 98-105.

Levitus, S., 1982:Climatological atlas of the world's ocean. NOAA Professional Paper no. 13, US Dept. of Commerce, National Oceanic and Atmospheric Administration, 173pp.

Maier-Reimer, E., U. Mikolajewicz and K. Hasselmann, 1993: Mean circulation of the Hamburg LSG OGCM and its sensitivity to thermohaline surface forcing. J. Phys. Oceanogr., 23, 731-757.

Manabe, S., R.J. Stouffer, M.J. Spelman and K. Bryan, 1991: Transient responses of a coupled ocean-atmosphere model to gradual changes of atmospheric CO2. Part 1: Annual mean response. J. Climate, 4, 785-818.

Marotzke, J., 1991: Influence of convective adjustment on the stability of the thermohaline circulation. J. Phys. Oceanogr., 21, 903-907.

Mitchell, J.F.B., T.C. Johns, J.M. Gregory and S.F.B Tett, 1995: Transient climate response to increasing levels of greenhouse gases and sulphate aerosols. Nature, 376, 501-504.

Murphy, J.M., 1995: Transient response of the Hadley Centre coupled ocean-atmosphere model to increasing carbon dioxide. Part I: control climate and flux adjustment. J. Climate, 8, 36-56.

Pacanowski, R.C. and S.G.H. Philander, 1981: Paramerterisation of vertical mixing in numerical models of tropical oceans. J. Phys. Oceanogr., 11, 1443-1451.

Rahmstorf, S., 1993: A fast and complete convection scheme for ocean models. Ocean Modelling, 101, 9-11 (Unpublished Manuscript).

Redi, M.H., 1982: Oceanic isopycnal mixing by coordinate rotation. J. Phys. Oceanogr., 12, 1154-1158.

Smith, G.D:, 1985: Numerical solution of partial differential equations: finite difference methods. Oxford University Press, 337pp.

Wood, R.A., 1993: A study of ocean spinup strategies for climate models. In Research Activities in Atmospheric and Oceanic Modelling, 18, ed. G.J. Boer, WMO/TD no. 533, WMO, Geneva.

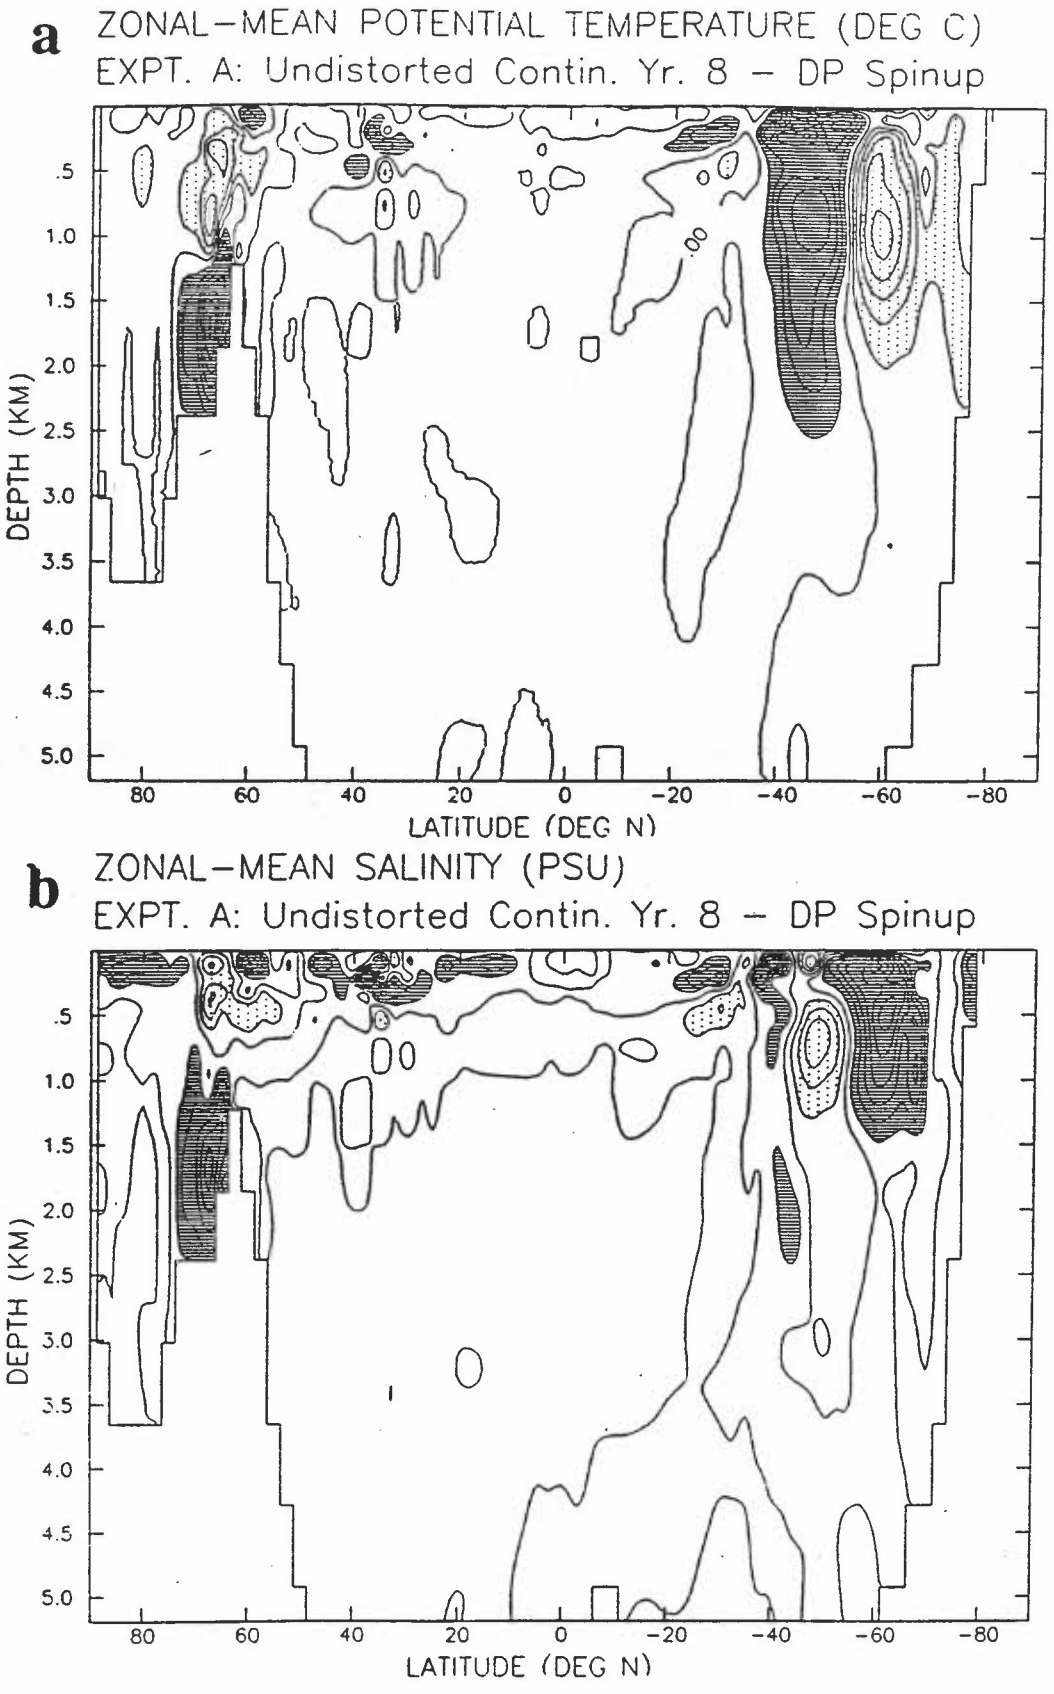

Fig. 4.1 Drift in the zonal mean potential temperature and salinity on switching from DP to undistorted integration, for Experiment A. The panels show the mean of year 8 of the undistorted run minus the mean of the last 10 years of the DP spinup from which it started.

a. Potential temperature. Contour interval 0.05 °C. Regions >0.05 °C shaded, regions <-0.05 °C stippled.

b. Salinity. Contour interval 0.005 psu. Regions >0.005 psu shaded, regions <-0.005 psu stippled.

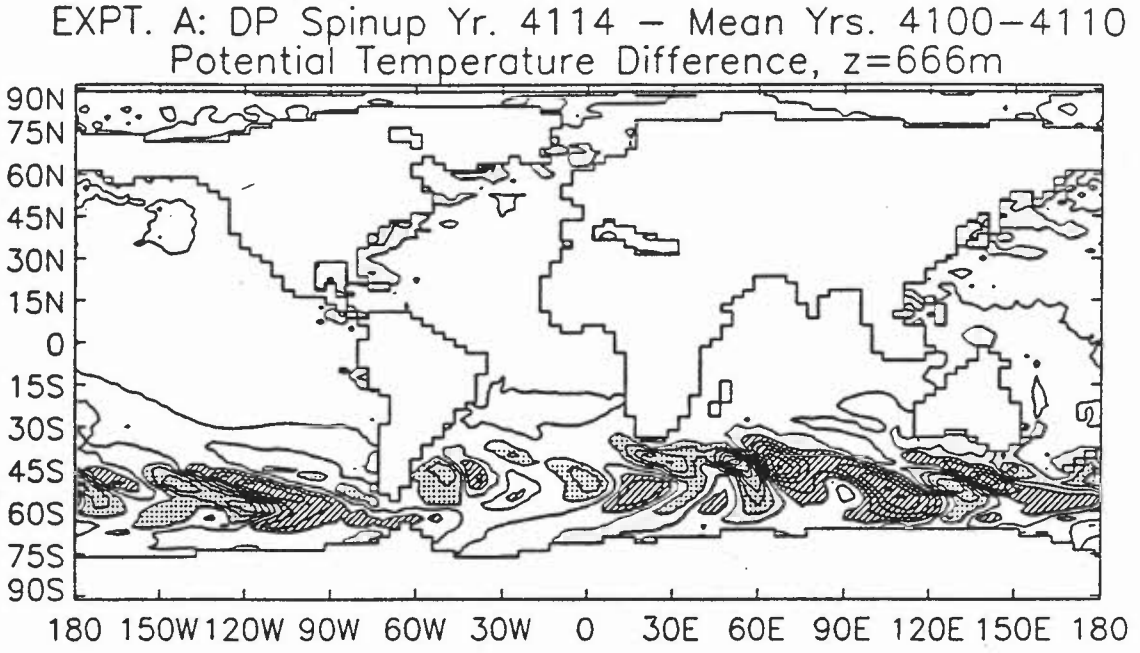

Fig. 4.2 Potential temperature difference at model level 12 (666.3m). Mean for model year 4114 (the final year of the DP spinup of Experiment A) minus the mean for years 4100-4110. Contour interval 0.3 °C. Regions >0.3 °C shaded, regions <-0.3 °C stippled.

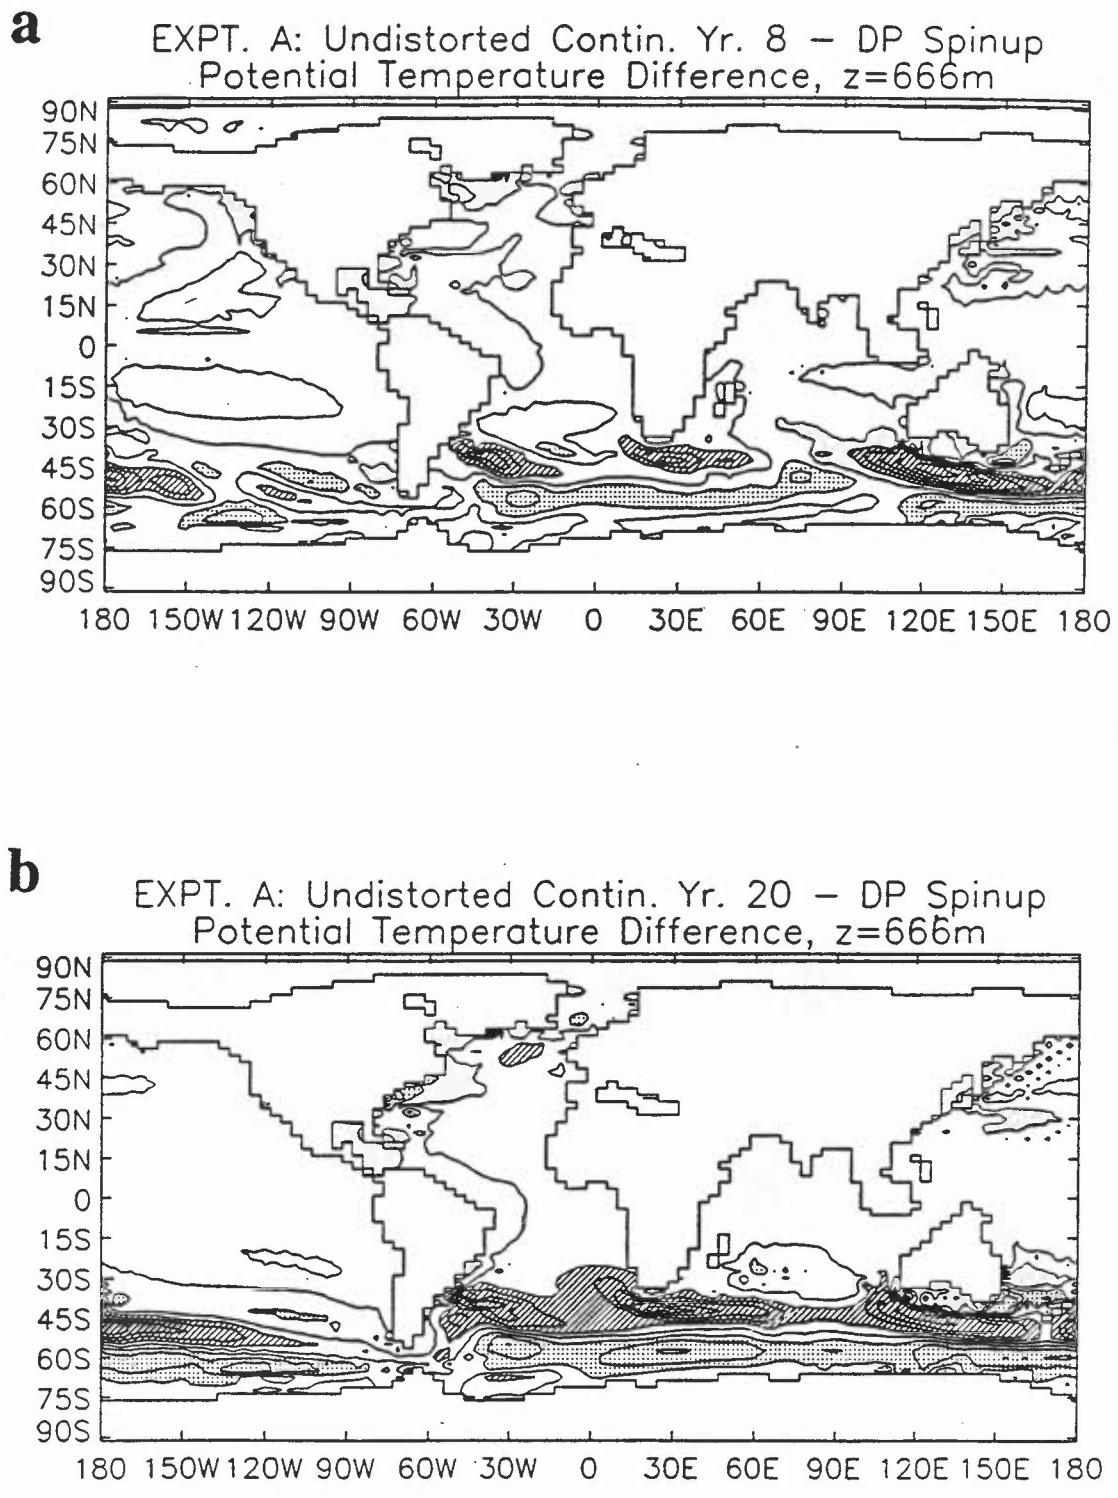

Fig. 4.3 As Fig. 4.2 but the difference is now for (a) year 8 and (b) year 20 of the undistorted continuation, again minus years 4100-4110 of the DP spinup.

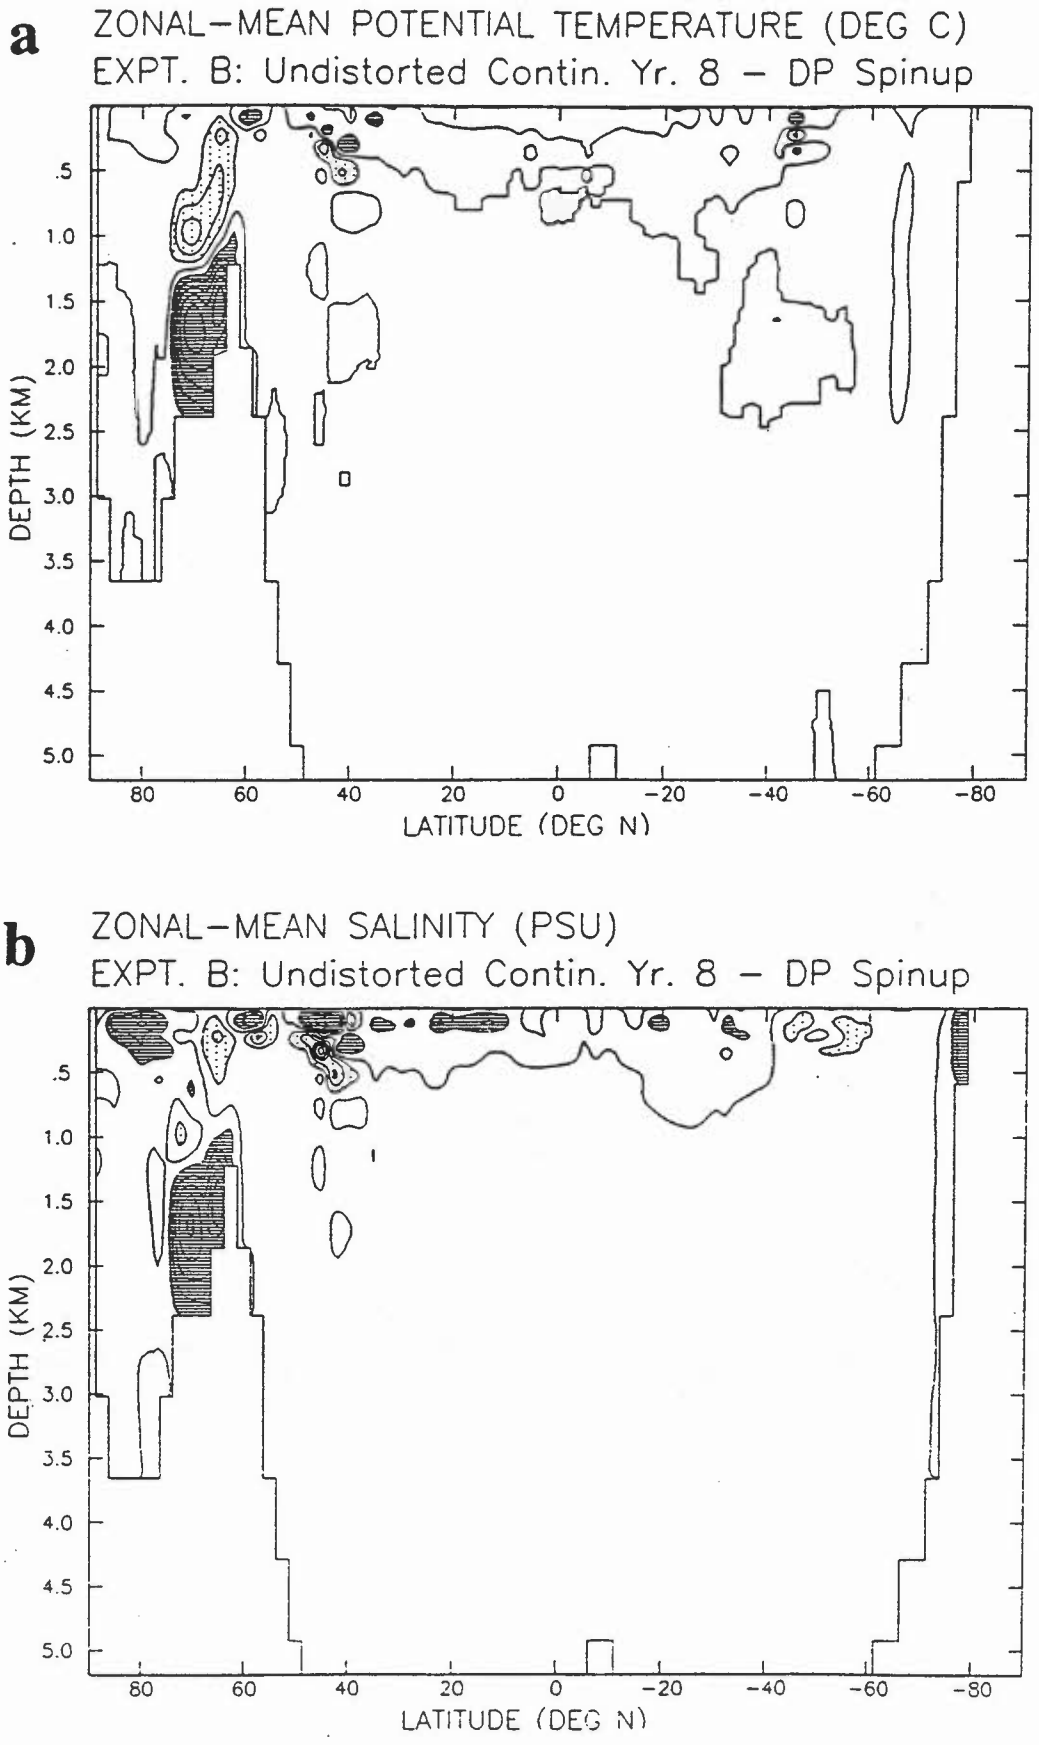

Fig. 4.4 As Fig. 4.1 but for Experiment B.

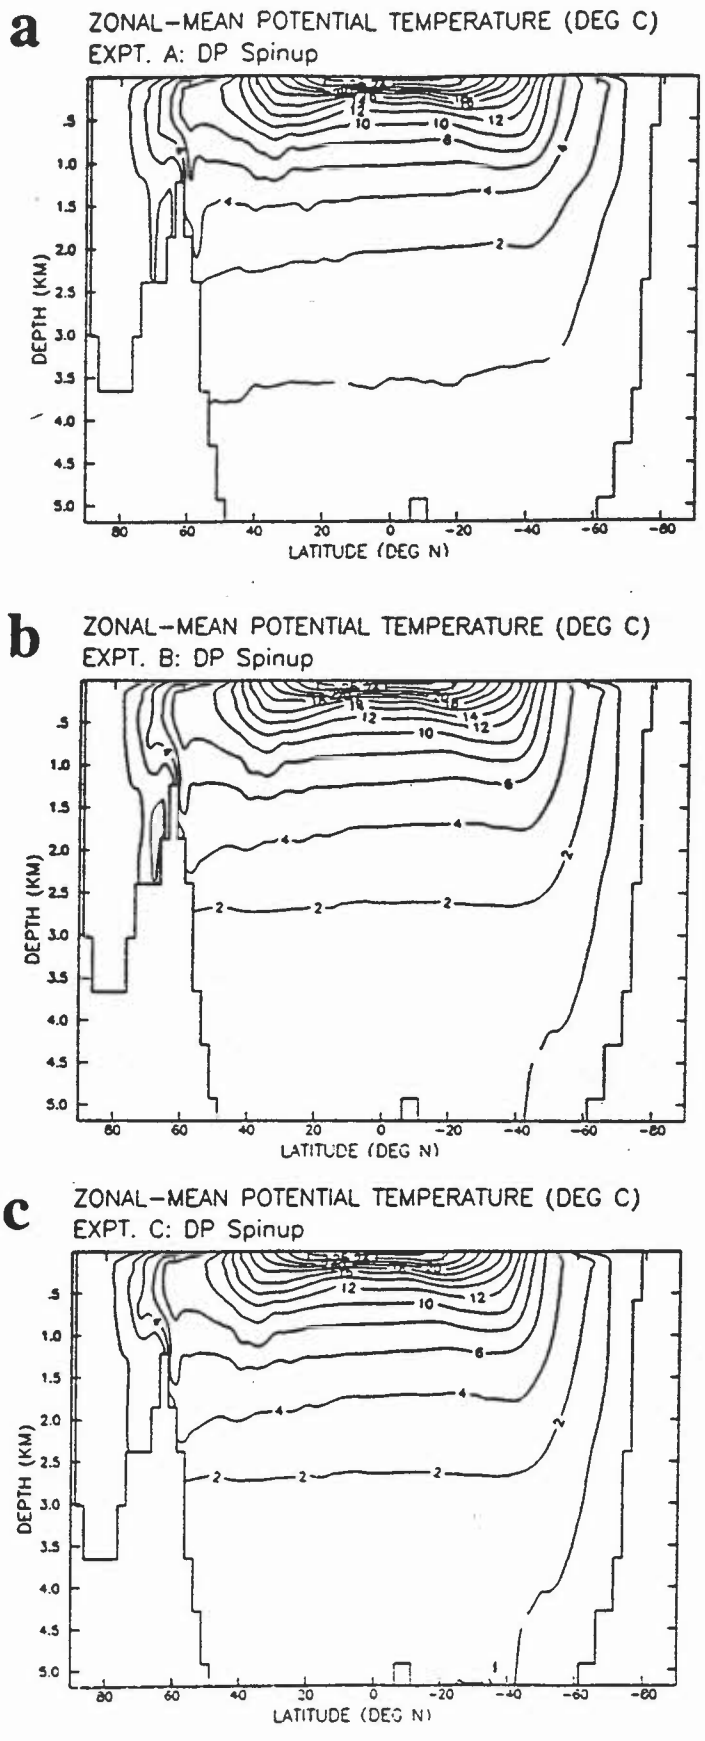

Fig. 4.5 Zonal mean potential temperature for the DP spinup states of (a) Experiment A, (b) Experiment B and (c) Experiment C. Contour interval 2 °C.

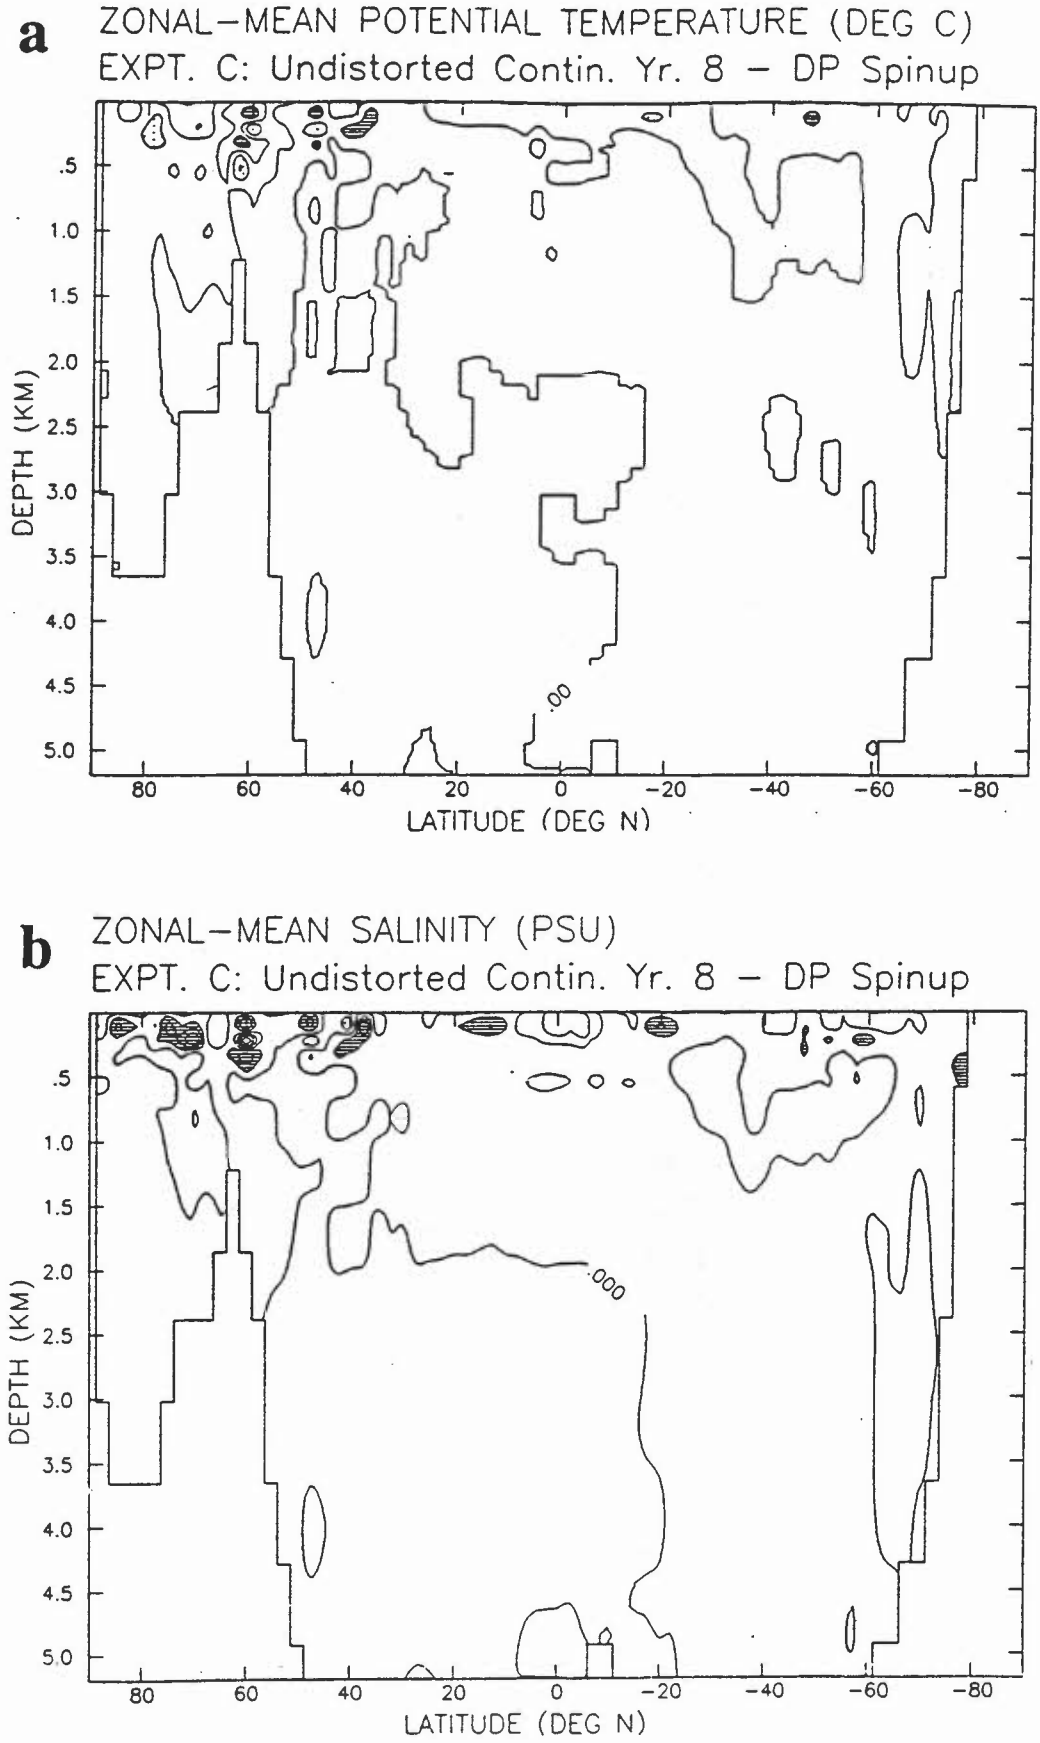

Fig. 4.6 As Fig. 4.1 but for Experiment C.

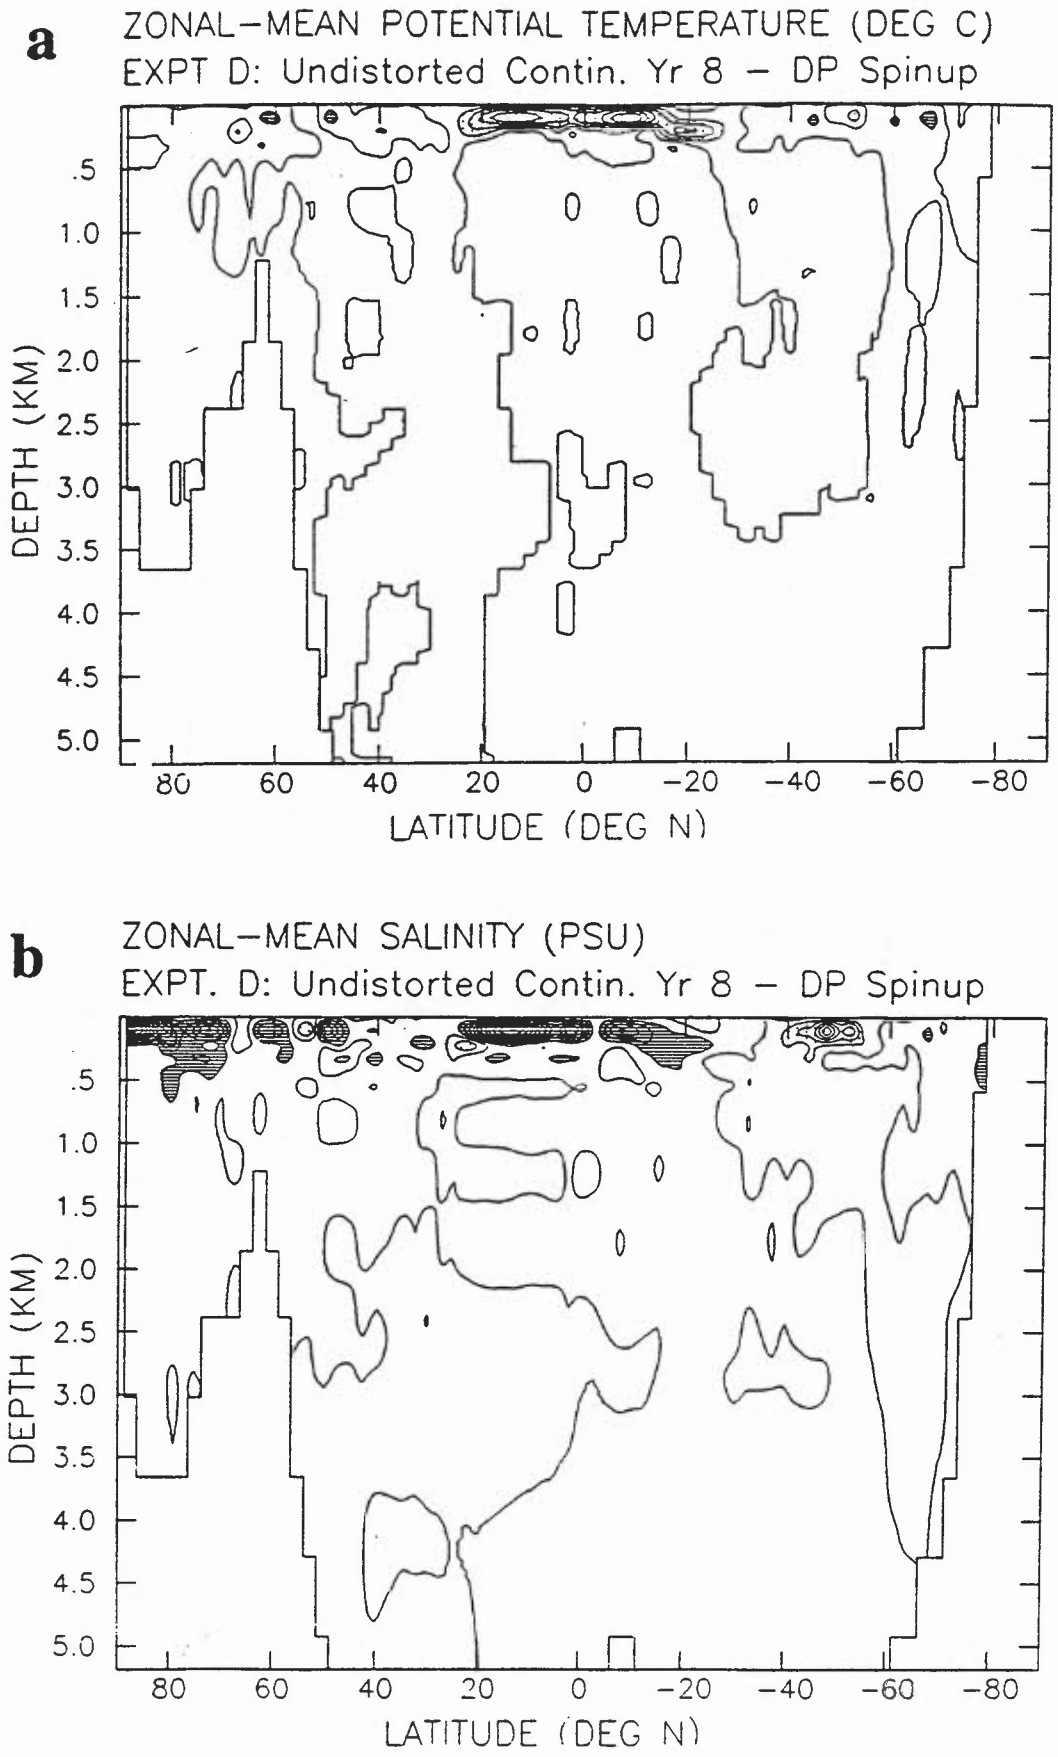

Fig. 4.7 As Fig. 4.1 but for Experiment D.

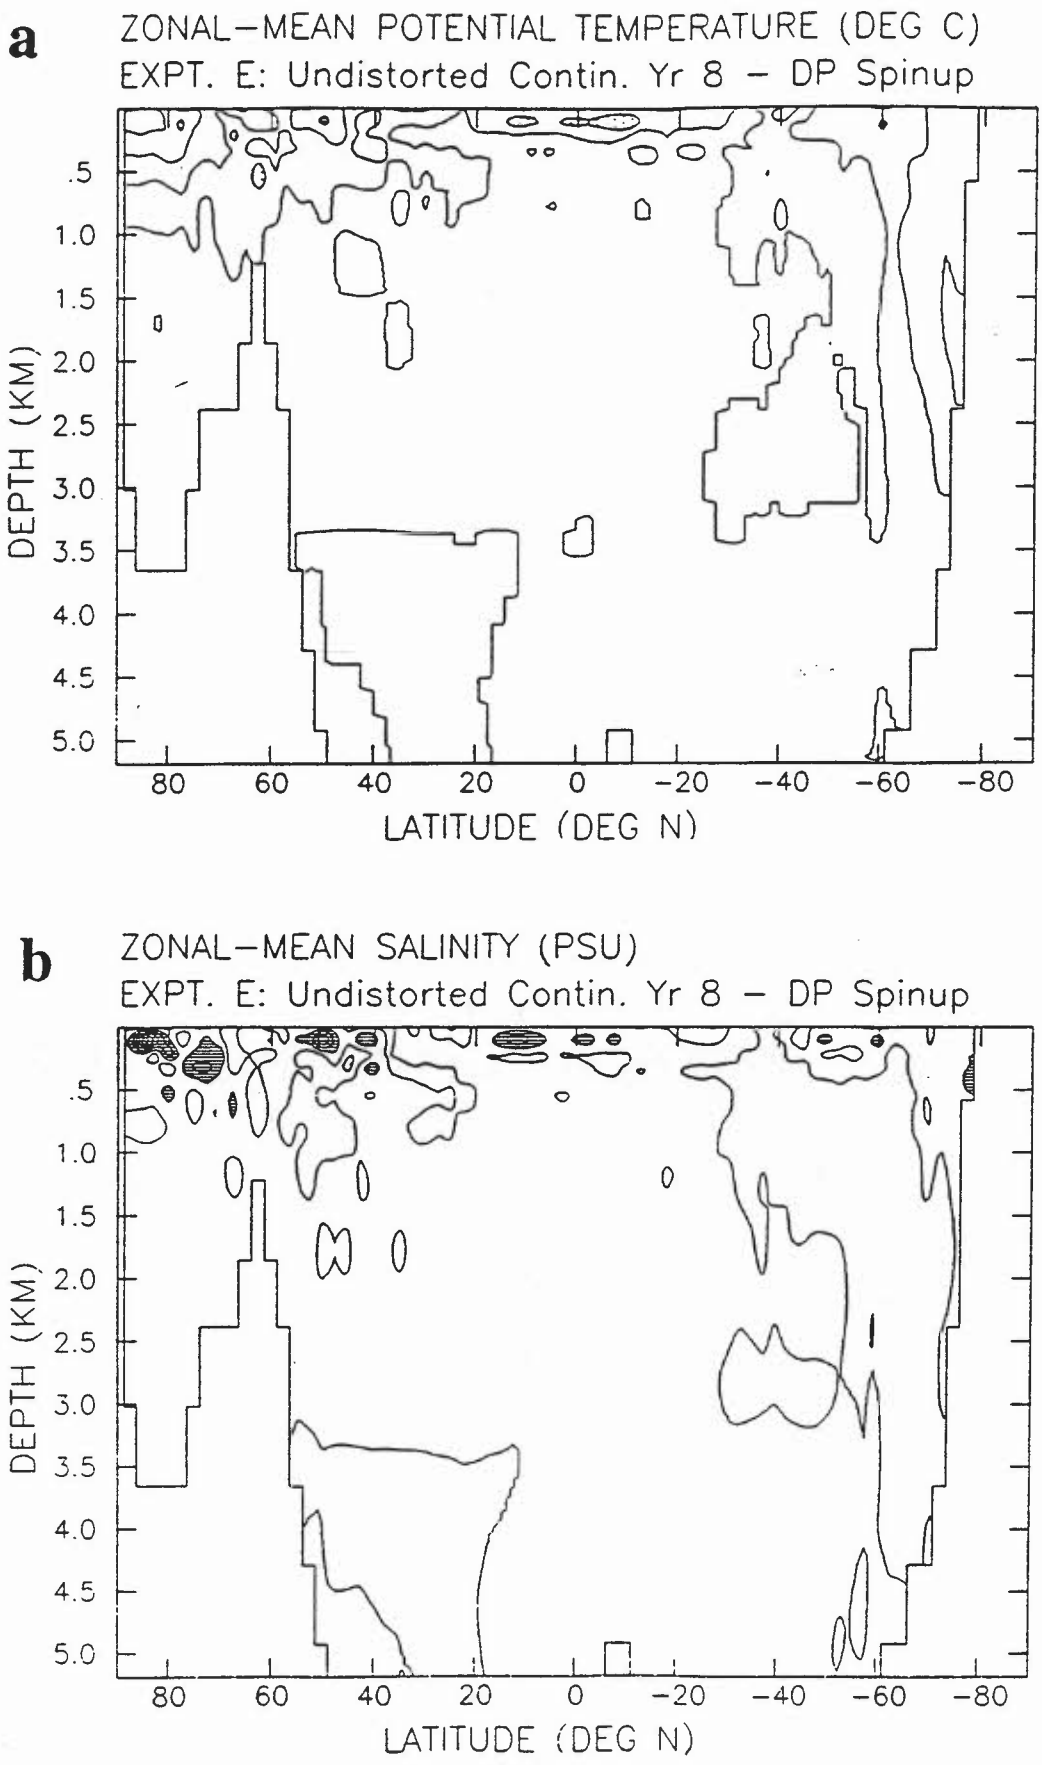

Fig. 4.8 As Fig. 4.1 but for Experiment E.

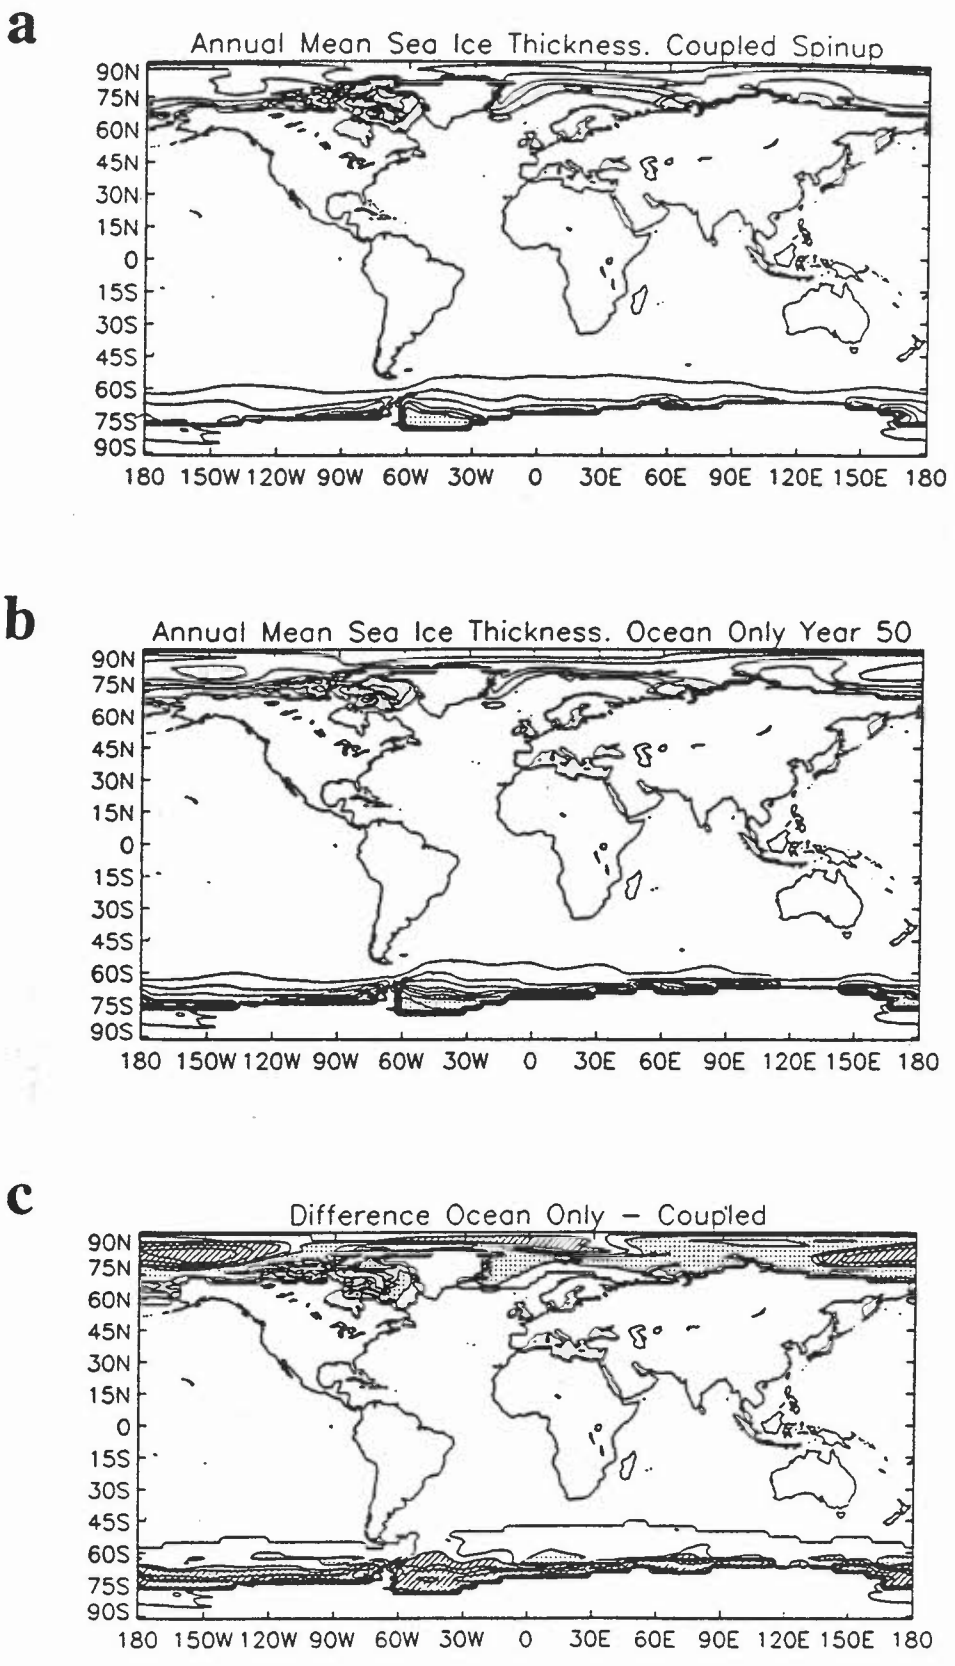

Fig. A.1 Annual mean sea ice thickness from (a) the final year of the coupled spinup of Johns et al. (1995), (b) year 50 of an ocean-only integration starting from (a), and (c) the difference (b)-(a). In (a) and (b) contours are at 0.0m, 0.1m, 0.5m, 1.0m and subsequently at intervals of 1m; regions >4m are stippled. In (c) contours are at 0m, ±0.1m, ±0.5m, ±1.0m and ±2.0m; regions >0.1m are shaded, regions <-0.1m are stippled.

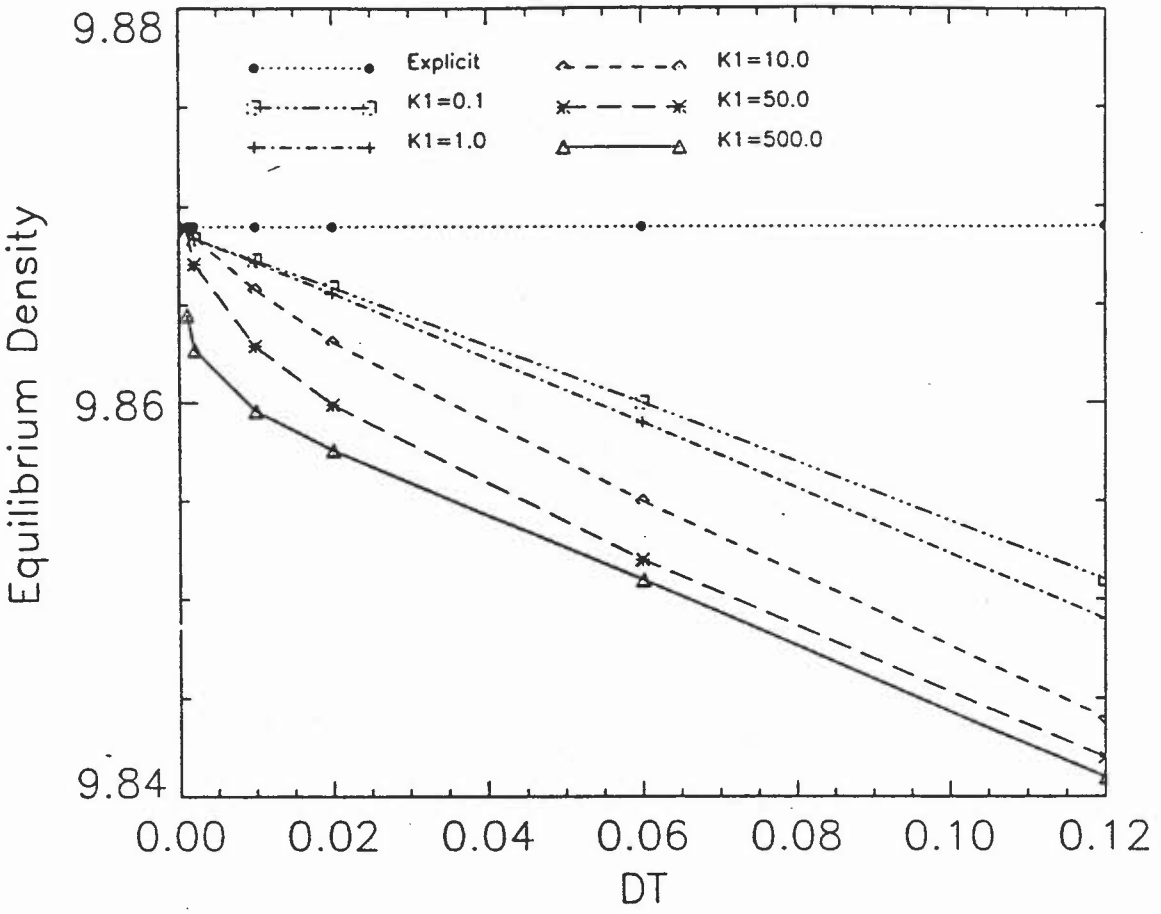

Fig. B.1 Equilibrium density ρ_3 as a function of timestep Δt, for the model of Appendix B. Solutions are shown for a number of values of the diffusivity K_1, and the solution with an explicit timestep (which is independent of K_1 and Δt) is also shown for comparison.