Ocean Dynamics (2001) 52: 36–49

Till Kuhlbrodt · Sven Titz · Ulrike Feudel
Stefan Rahmstorf

A simple model of seasonal open ocean convection

Part II: Labrador Sea stability and stochastic forcing

Received: 25 April 2001 / Accepted: 10 July 2001

Abstract Aspects of open ocean deep convection variability are explored with a two-box model. In order to place the model in a region of parameter space relevant to the real ocean, it is fitted to observational data from the Labrador Sea. A systematic fit to OWS Bravo data allows us to determine the model parameters and to locate the position of the Labrador Sea on a stability diagram. The model suggests that the Labrador Sea is in a bistable regime where winter convection can be either "on" or "off", with both these possibilities being stable climate states. When shifting the surface buoyancy forcing slightly to warmer or fresher conditions, the only steady solution is one without winter convection.

We then introduce short-term variability by adding a noise term to the surface temperature forcing, turning the box model into a stochastic climate model. The surface forcing anomalies generated in this way induce jumps between the two model states. These state transitions occur on the interannual to decadal time scale. Changing the average surface forcing towards more buoyant conditions lowers the frequency of convection. However, convection becomes more frequent with stronger variability in the surface forcing. As part of the natural variability, there is a non-negligible probability for decadal interruptions of convection. The results highlight the role of surface forcing variability for the persistence of convection in the ocean.

Responsible Editor: Jörg-Olaf Wolff

T. Kuhlbrodt · S. Titz · S. Rahmstorf
Institute of Physics, University of Potsdam,
Postfach 60 15 53, 14415 Potsdam, Germany

U. Feudel
Institute for Chemistry and Biology of the Marine Environment,
University of Oldenburg, Postfach 2503,
26111 Oldenburg, Germany

T. Kuhlbrodt (✉) · S. Rahmstorf
Potsdam-Institut für Klimafolgenforschung,
Postfach 60 12 03, 14412 Potsdam, Germany
e-mail: kuhlbrodt@pik-potsdam.de

Keywords North Atlantic deep convection · Labrador Sea · Bistability · Stochastic climate model · Decadal variability

1 Introduction

Deep convection in the North Atlantic is a sensitive part of the world ocean's thermohaline circulation (THC). Its intensity, location, and variability influence the northward heat transport of the THC (Rahmstorf 1995a), the pathway of North Atlantic Deep Water (NADW) (Wood et al. 1999), and the total meridional overturning (Delworth et al. 1993; Rahmstorf 1995b) in models. Observational data from the past decades (Lazier 1980; Dickson et al. 1988; Belkin et al. 1998) display great interannual to decadal variability in the occurrence and depth of deep convection events. The role of deep convection in the THC is the removal of heat from the deep ocean, thus balancing the downward penetration of heat due to various diapycnal mixing processes (Munk and Wunsch 1998).

While freshwater fluxes are important for the background stratification, convection events are usually triggered by strong heat loss through the ocean surface, reducing the vertical density gradient until the water column becomes statically unstable. Vigorous vertical mixing down to 2000 m and more results (Send and Marshall 1995). Convection events have an extent of only a couple of days in time and some 10 km in space (Marshall and Schott 1999). As this is subgrid scale in most ocean general circulation models (OGCMs), convection is parameterized by convective adjustment schemes (Rahmstorf 1993; Klinger et al. 1996). These schemes remove static instability in the water column at each time step by mixing vertically adjacent grid cells. Although this may lead to a too intense mixing (Lilly et al. 1999) and grid-scale instabilities may occur (Cessi 1994; Molemaker and Dijkstra 2000), convective adjustment has proven to be a satisfactory parameterization for many applications.

The most simple conceptual model of deep convection consists of two boxes, one for the permanently mixed surface layer and one for the deep ocean (Welander 1982). It includes different boundary conditions for the surface fluxes of heat and salt, as well as a strongly nonlinear dependence of vertical mixing between the two boxes on the vertical density gradient. Under certain boundary conditions, this leads to a bistable regime of the water column where convection can be either permanently "on" or permanently "off", depending on the initial condition. This is due to a positive salinity feedback in the presence of surface freshwater input: once convection is interrupted (e.g., by a freshwater anomaly), the surface water will become less and less saline, making a restart of convection increasingly harder to achieve. Using Welander's box model, Lenderink and Haarsma (1994) showed that large regions of a model North Atlantic were bistable due to this feedback.

In Part I of this paper (Rahmstorf 2001), Welander's model is extended in two ways. First, temperature and salinity of the deep box are introduced as variables rather than as prescribed fixed values. Heat and salt can then accumulate in the deep box during nonconvecting phases and are released by convective mixing. In this way, the heat flow from the deep ocean through the convecting water column to the atmosphere is modeled. Second, a seasonal cycle is introduced in the boundary conditions for the upper box. This reduces convection to a short period in winter, rather than occurring permanently. The model demonstrates why the convection events are so short: slow processes like advection and diapycnal mixing replenish the heat store of the deep layer throughout the year, while a few days' time is enough to release the accumulated heat to the atmosphere via the much faster processes of convective mixing and surface exchange.

Welander's model has been used by some authors (Lenderink and Haarsma 1994; Pierce et al. 1995; Lenderink and Haarsma 1996; Hirschi et al. 1999) to analyze output from OGCMs. Here, we use the modified version to analyze observational data from the Labrador Sea. The data from Ocean Weather Ship (OWS) Bravo (Lazier 1980) show how convection was switched off and on again in the course of the Great Salinity Anomaly (GSA) of the years 1968–1972 (Dickson et al. 1988). By fitting the model to the OWS Bravo data we can locate the Labrador Sea in a stability diagram. The results suggest that the convecting state is only marginally stable; anomalies in the surface forcing can trigger state transitions very easily.

The effect of short-term variability on the ocean surface layer can be represented by adding a stochastic term to the surface forcing. The most simple form of oceanic response to stochastic forcing was analyzed by Hasselmann (1976). In his stochastic climate model, an ocean surface layer of fixed thickness acts as a reservoir and integrates white noise temperature forcing from the atmosphere, which leads to a red noise variance spectrum of sea surface temperature similar to that seen in observed time series (Frankignoul and Hasselmann 1977). Hasselmann's mechanism was found to trigger decadal variability of the THC in an OGCM (Weisse et al. 1994), where the water column in the Labrador Sea integrated noisy freshwater forcing. In other models, decadal variability is generated by stochastic excitation of internal ocean modes (Delworth et al. 1993) or by a local convective oscillator (Pierce et al. 1995; Rahmstorf 1999).

Here, we effectively combine the basic ideas of Hasselmann's stochastic climate model and Welander's two-box model to a stochastic climate model of deep convection. Adding stochastic perturbations to the surface forcing of the extended two-box model (developed in Part I) mimics the essential role of buoyancy forcing variability in triggering deep convection (Marshall and Schott 1999; Sathiyamoorthy and Moore 2001). The frequency of jumps and the residence times in both model states (i.e., with and without winter convection) are studied, together with their dependence on model parameters. The interplay of the different timescales involved is shown to result in decadal variability.

The OWS Bravo data are analyzed and interpreted in the following section. The third section gives a short summarizing description of the box model and explains how it was fitted to the OWS Bravo time series. The stability diagram in the non-stochastic case is discussed in section 4. The fifth section is devoted to the influence of stochastic forcing on the model dynamics, and the paper ends with some conclusions.

2 Observational data

Long time series of hydrographic data that show clear signs of deep convection events are rare. The data from Ocean Weather Ship (OWS) Bravo are exceptional due to their location and their sampling rate. OWS Bravo was located in the central Labrador Sea close to the area of the deepest convection events. In the time from January 1964 through September 1974 the sampling rate of the data varied between 6 h and 2 months. This enables the derivation of a time series of monthly means that clearly reflects the winter open-ocean deep convection events.

The original data (Lazier 1980) were interpolated to standard depth levels. Potential temperature (θ) and potential density (σ_0) were computed with the standard formulas (Fofonoff and Millard 1984). To obtain monthly mean values, the data of each month were binned and averaged at each depth level. Missing monthly means were interpolated linearly. Subsequently, the data were averaged for an upper layer (0–50 m) and a deep layer (200–2000 m). The intermediate level (50–200 m) was left out because, on the one hand, this layer still shows substantial seasonal variations, but on the other hand, it is not part of the surface mixed layer throughout the year.

The resulting time series of monthly means are given in Fig. 1. The winters 1969–1971 show the impact of the Great Salinity Anomaly (GSA, described by Dickson et al. 1988) that suppressed deep convection in the Labrador Sea by the advection of a large freshwater anomaly. In consequence, temperatures and salinities followed different trends in both layers: cooling and freshening in the upper layer leading to less dense waters, while the deep ocean is becoming slightly warmer and saltier. The upper layer values show a strong seasonal cycle. When the potential density difference between the two layers is small, this indicates deep convection. For a number of reasons, this difference is not exactly zero: the mixing of the layers occurs only during a few days, but the data are averaged monthly; the mixing does not occur necessarily exactly at the ship site and throughout the whole water column. In some winters the upper layer overshoots in temperature (though not very visible in the monthly means): to compensate for the deep ocean being saltier than the surface waters, the upper layer must become colder than the deep ocean before the vertical density gradient vanishes and deep convection starts (see Part I, section 4).

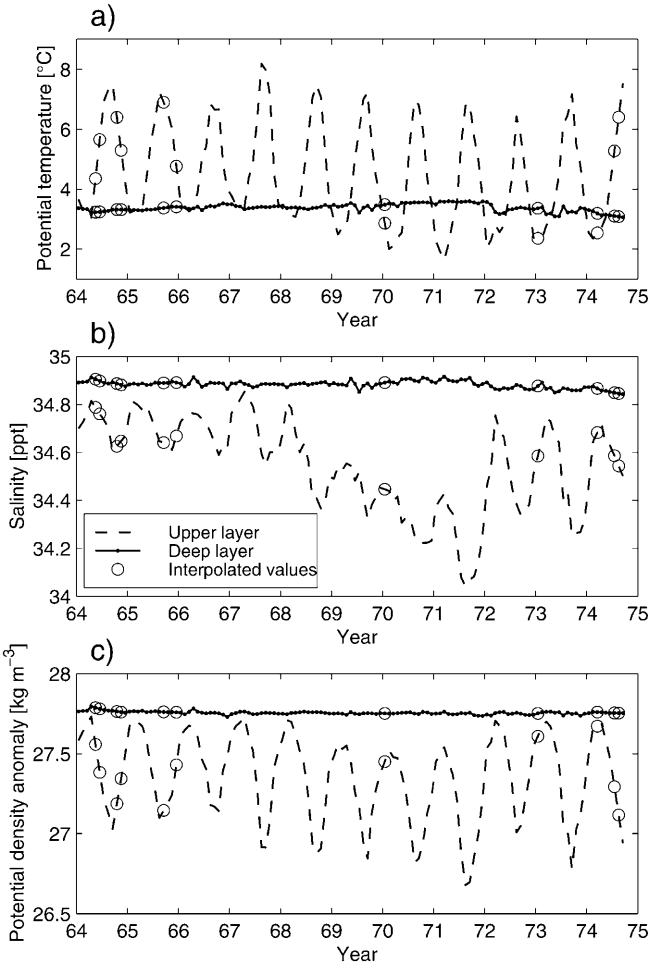

Fig. 1a–c Time series of monthly means obtained from the OWS Bravo data set: potential temperature (**a**), salinity (**b**), and potential density (**c**) of the upper layer (*dashed*) and the deep layer (*dots*). Interpolated values are indicated by *circles*. The large minimum density difference in the winters from 1969 to 1971 is an indication for the absence of deep convection, which led to the cooling and freshening of the upper layer

In order to consider seasonal and interannual variability separately, we first computed the mean seasonal cycle of the time series (Fig. 2). In the upper layer, the temperature cycle has its minimum in February and its maximum in September, with an amplitude of 2.2 °C. The salinity cycle lags by about 1 month: the cycle with an amplitude of 0.13 psu peaks in March and reaches its minimum in October. The σ_0 cycle lies in between with an amplitude of 0.29 kg m^{-3}. The deep layer seasonal cycles (not shown) are almost 2 orders of magnitude smaller. They show an annual warming of 0.12 °C and a density decrease of 0.015 kg m^{-3}, both of which start in April after the convection season, and reach their extremum in November/December. Deep-layer salinity variations are very small. The temperature cycle is mostly forced by fluxes of latent and sensible heat in winter and short-wave radiation in summer (Smith and Dobson 1984). Freshwater sources of poorly known strength (Canadian runoff, meltwater, local precipitation, and low-salinity inflow from the Arctic Sea) account for the salinity cycle (Lilly et al. 1999).

In a second step, we subtracted the seasonal cycle from each time series in Fig. 1. The resulting time series (Fig. 3) show the variability excluding the seasonal cycle. Three distinct phases stand out, marked by either the occurrence or the absence of convection. In the first convective phase (phase 1) from January 1964 to March 1968 the values in the upper box fluctuate with hardly any interannual trend. However, there are trends in the deep layer. Time series of single depth levels reveal that these trends are more pronounced in deeper layers. With maximum convection depth varying from year to year, the deeper layers sometimes remain untouched and accumulate heat over more than 1 year.

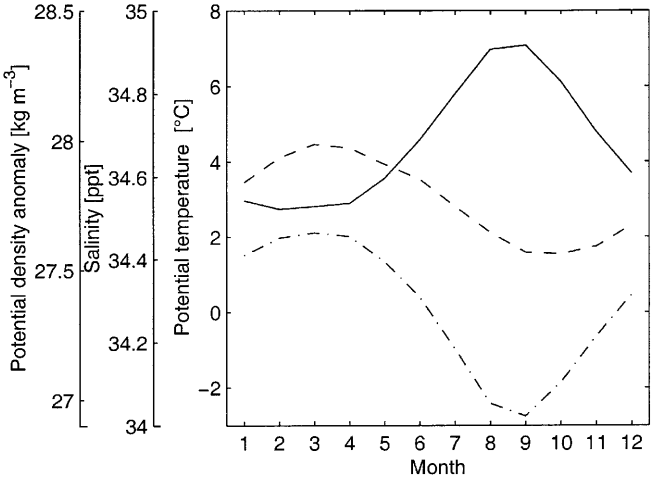

Fig. 2 Mean seasonal cycle of the upper layer (0–50 m) from OWS Bravo data for potential temperature (*solid*), salinity (*dashed*), and potential density (*dash–dotted*)

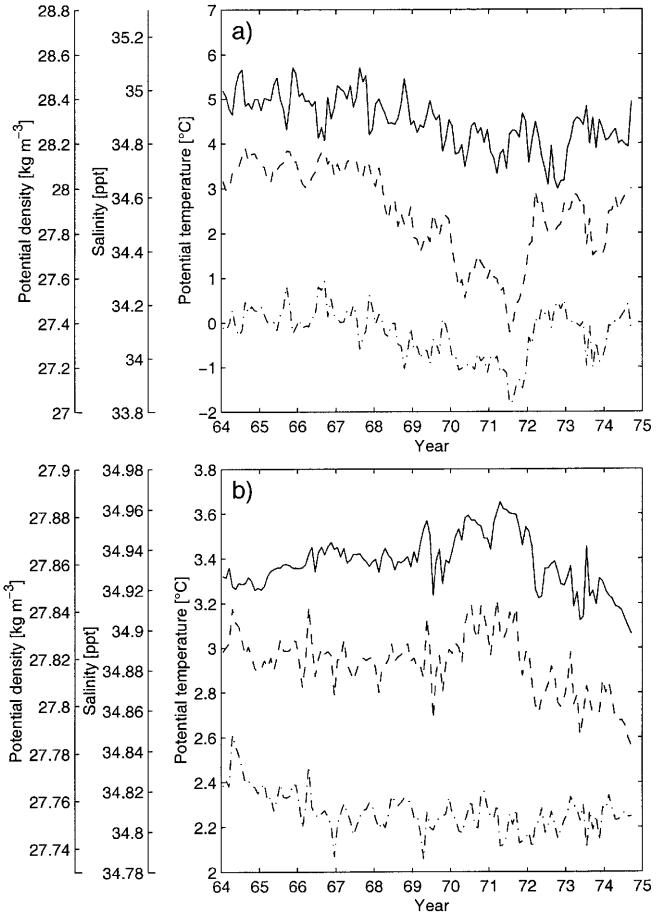

Fig. 3a, b Time series of monthly means with subtracted seasonal cycle for the upper layer (**a**) and the deep layer (**b**) of potential temperature (*solid*), salinity (*dashed*), and potential density (*dash–dotted*)

The phase from April 1968 to September 1971 (phase 2) is characterized by the GSA passing the Labrador Sea and suppressing convection there. Phase 2 begins after the last convection event and ends with the upper-layer salinity starting to rise again. Annual trends of all quantities during phase 2 are given in Table 1. The decoupling of the two layers induced by advective freshening (Dickson et al. 1988) in the upper layer leads there to cooling, further freshening, and a density decrease. While the cooling comes to a halt already in early 1970, the strong freshening continues until late 1971. In the deep layer, the waters are becoming slightly warmer and more saline, mostly by lateral mixing from adjacent water masses. Potential density in the deep layer shows a weak decreasing trend that is clearly smaller than in phase 1 as the warming and salinification partly compensate in their effect on density.

Phase 3 is again characterized by annual convection events. Starting in October 1971, strong wind mixing of the surface mixed layer caused it to deepen and entrain salt from below (Dickson et al. 1996). Additional strong cooling then achieved a vigorous deep convection event in early 1972. Afterwards, convection occurs again every

Table 1 Trends during the GSA. Trends of potential temperature, salinity and potential density in the upper layer (*index* 1) and the deep layer (*index* 2) during the GSA (04/68 to 09/71) from the time series without seasonal cycle depicted in Fig. 3

Quantity	θ_1	S_1	$\sigma_{0,1}$
Trend (yr^{-1})	−0.28 °C	−0.11 psu	−0.060 kg m^{-3}
Quantity	θ_2	S_2	$\sigma_{0,2}$
Trend (yr^{-1})	0.070 °C	0.0072 psu	−0.0013 kg m^{-3}

year until the end of the time series. The upper layer returns to more saline and dense conditions but remains cool, in a state clearly different from phase 1. Possibly this is a consequence of the deep convection chimney being farther away from the ship now; the larger winter gap between upper and lower layer salinity (compared to phase 1, see Fig. 1b) suggests this. The deep layer jumps back to a colder and less saline state, and a further cooling and freshening trend sets in, albeit with little effect on density.

In summary, the OWS Bravo data show a transition from a state of annual convection to stable stratification and back to convection. We will now discuss to what extent the simple box model can help to understand these transitions.

3 The simple box model

3.1 Model equations

Before discussing the fitting of the convection box model to the observational data, we recall the equations of the model and extend them slightly. The reader is referred to Part I of this paper for a detailed description of the model and the derivation of its equations. The model consists of two boxes, a shallow upper box (index 1) representing the annual mixed layer and a large deep box (index 2) representing the waters below the seasonal thermocline. The ratio of the box depths is termed h^*. For the sake of simplicity, the box depths are fixed, so effects like variable convection depth or mixed layer deepening are not included. The model variables are the temperatures T_1, T_2 and the salinities S_1, S_2 in both boxes. These four variables are relaxed towards prescribed relaxation temperatures and salinities T_1^*, S_1^*, T_2^*, S_2^*. Through the use of three different time scales τ_{1T}, τ_{1S}, and τ_2, the different coupling strengths of the heat and salt fluxes into box 1 and 2 are accounted for. The two time scales τ_{1T} and τ_{1S} of the upper box forcing represent the different feedbacks affecting temperature and salinity forcing (similar to mixed boundary conditions); included in this forcing are surface fluxes as well as lateral fluxes due to advection and mixing. The use of a single time scale τ_2 for the deep box is motivated by the eddy transfer fluxes at depth. In Part I, the basic equation system (Eq. 6) was refined by adding a seasonal cycle with amplitude A_T to the upper box relaxation temperature T_1^*, in order to include seasonality and achieve a

short winter convection period instead of year-round convection, in accordance with the observations. To facilitate comparison with observations and given the pronounced seasonal cycle in upper layer salinity in the data, we add a seasonal cycle (with amplitude A_S and zero mean) to the upper box salinity forcing as well. A phase shift ψ between temperature and salinity cycles is introduced as an adjustable parameter. Finally, keeping the dimensions of the variables clarifies the physical meaning of the parameters. Time is in units of 1 year. The full equation set now reads:

$$\frac{dT_1}{dt} = \frac{1}{h^*\tau_c(\Delta\rho)}(T_2 - T_1) + \frac{1}{\tau_{1T}}\left[T_1^* - A_T\cos(2\pi t) - T_1\right] \tag{1}$$

$$\frac{dS_1}{dt} = \frac{1}{h^*\tau_c(\Delta\rho)}(S_2 - S_1) + \frac{1}{\tau_{1S}}\left[S_1^* - A_S\cos(2\pi t + \psi) - S_1\right] \tag{2}$$

$$\frac{dT_2}{dt} = \frac{1}{\tau_c(\Delta\rho)}(T_1 - T_2) + \frac{1}{\tau_2}\left(T_2^* - T_2\right) \tag{3}$$

$$\frac{dS_2}{dt} = \frac{1}{\tau_c(\Delta\rho)}(S_1 - S_2) + \frac{1}{\tau_2}\left(S_2^* - S_2\right) \ . \tag{4}$$

In each of the equations, the first term on the right-hand side represents convective vertical mixing, and the second term the horizontal and surface heat/salt flux. We recall that the four relaxation temperatures and salinities, the three relaxation timescales, the two amplitudes, and the phase shift are model parameters (see Table 2 for a detailed list), but that the time scale of vertical convective mixing τ_c is a strongly nonlinear function of the model variables through the density difference

$$\Delta\rho = \rho_1 - \rho_2 = -\alpha(T_1 - T_2) + \beta(S_1 - S_2) \ , \tag{5}$$

where α and β are the thermal and haline expansion coefficients of the linearized equation of state of seawater, see Eq. (4) in Part I. For stable stratification ($\Delta\rho \le 0$) the vertical mixing is very weak, so τ_c has a large value. In the case of unstable stratification ($\Delta\rho > 0$) convection starts, i.e., vigorous vertical mixing with a timescale τ_c of the order of days. Since this value of τ_c is much smaller than the other time scales involved, we assume $\tau_c \to 0$ and use the common parameterization for deep convection, known as convective adjustment (Rahmstorf 1993; Klinger et al. 1996). The water column is checked at each time step for hydrostatic stability (we used $dt = 2$ days). Nothing is done in the case of stable stratification, but any occurring instability is instantaneously removed by complete mixing. Thus, the numerical integration scheme has two parts.

1. Integrate forward Eq. (1) to Eq. (4) one time step without the vertical mixing terms. If we start at time i, this gives preliminary values $\hat{T}_1^{i+1}$, $\hat{S}_1^{i+1}$, $\hat{T}_2^{i+1}$, $\hat{S}_2^{i+1}$ for the variables at time $i+1$.
2. Apply the convective adjustment scheme to obtain the final values T_1^{i+1}, S_1^{i+1}, T_2^{i+1}, S_2^{i+1} of the variables. If $\Delta\rho \le 0$, the final values are identical to the preliminary ones; if $\Delta\rho > 0$, the two columns are mixed:

$$T_1^{i+1} = T_2^{i+1} = h^*\hat{T}_1^{i+1} + (1-h^*)\hat{T}_2^{i+1} \quad \text{and}$$
$$S_1^{i+1} = S_2^{i+1} = h^*\hat{S}_1^{i+1} + (1-h^*)\hat{S}_2^{i+1} \ . \tag{6}$$

3.2 Fitting the box model to the OWS Bravo data

We adjust the model parameters (see Table 2 for a complete list) to find the best fit of the model to the OWS Bravo data, and use a least-squares fit procedure for this purpose. We define a cost function K as the sum of the quadratic distances of each monthly averaged model variable time series T_1, S_1, T_2, S_2 (weighted by the thermal and haline expansion coefficients to have a common density unit) to the observed values. The optimal parameter set minimizes K. As discussed above the OWS Bravo data show two different states with convection (phases 1 and 3, possibly due to different surface forcing and/or convection locations), but the model can have only one convecting solution with the same forcing. Hence we restrict our analysis to phases 1 and 2 of the OWS Bravo data. (An analysis based on phases 2 and 3 gives similar results). The 10-dimensional parameter space is spanned by a 10-dimensional matrix. For every

Table 2 Model parameters. The model parameters with the values determined from fitting the model to OWS Bravo data. The set of these parameter values is called the 'optimal parameter set'. The uncertainty range is spanned by all parameter sets whose cost function value exceeds the minimum by less than 10%. The parameters marked with # were not determined through the cost function, but directly from observational data (see section 2 and 3)

Parameter	Value	Uncertainty	Description
T_1^*	4.4 °C	4.0–4.6 °C	Upper box restoring temperature
S_1^*	33.5 psu	#	Upper box restoring salinity
T_2^*	4.1 °C	3.9–4.3 °C	Deep box restoring temperature
S_2^*	34.97 psu	#	Deep box restoring salinity
τ_{1T}	5 months	3–9 months	Restoring timescale of upper box temperature
τ_{1S}	8 years	6–11 years	Restoring timescale of upper box salinity
τ_2	20 years	14–28 years	Restoring timescale of deep box
A_T	6.4 °C	5.0–7.8 °C	Amplitude of seasonal cycle added to T_1^*
A_S	4.5 psu	3–6 psu	Amplitude of seasonal cycle added to S_1^*
ψ	0.6 months	−0.5–1.5 months	Phase shift of the seasonal cycles
h^*	1/36	#	Ratio of box depths

possible parameter combination out of this matrix the cost function K is computed from a model run with convecting state initial conditions. The onset of the GSA in the model is achieved by adding an anomalous salt flux of -0.8 psu a^{-1} to the upper box for a period of 3 months in spring 1968 to mimic the arrival of an advected freshwater anomaly. The crucial idea here is that we aim to find one single parameter set that yields a realistic model behavior in both states (convecting and nonconvecting) with the same forcing; the prescribed salt flux anomaly provides a brief "kick" which induces a state transition in the model.

It turns out that K and its derivatives with respect to the parameters are smooth functions and behave in a physically understandable manner. However, both physical intuition and the objective analysis show that the ten free model parameters are underdetermined by the fit, i.e., the problem is ill-posed. This is because the OWS Bravo data contain a steady convecting state (phase 1), but not a steady nonconvecting state. Phase 2 of the OWS Bravo data displays the initial trends after cessation of convection, but does not reveal which equilibrium values the three variables S_1, T_2, and S_2 would eventually reach in the nonconvecting state. Only the trend in T_1 stops in 1970, so that T_1^* can be determined. The impact of the missing nonconvecting state on the parameter determination can be clarified through the model equations. Take Eq. (2) for the upper box salinity. We neglect the seasonal cycle here because it has no impact on the long-term trend. If no convection occurs, Eq. (2) then reduces to:

$$\frac{\mathrm{d}S_1}{\mathrm{d}t} = \frac{1}{\tau_{1S}}\left(S_1^* - S_1\right) \quad . \tag{7}$$

When $\mathrm{d}S_1/\mathrm{d}t$ and an initial S_1 are known from the data, then on the right-hand side of Eq. (7) for any arbitrary choice of S_1^* a corresponding value of τ_{1S} can be found to fulfil the equation. Thus, one of the two parameters in Eq. (7) is free. The situation is similar for T_2^* and S_2^*; and since Eqs. (3) and (4) are coupled through the common time scale τ_2, one second degree of freedom arises here. In short, the least-squares fit procedure constrains the ten-dimensional parameters space to a two-dimensional subspace. This null space is clearly seen when attempting to minimize the cost function.

Two further constraints are thus needed to close the problem, i.e., to obtain a global minimum in K. One could assume arbitrary values of τ_{1S} and τ_2; we opt for making assumptions about S_1^* and either T_2^* or S_2^*. This option is equivalent to assuming the equilibrium mean values of the three variables S_1, T_2, S_2 in the nonconvecting state. The upper box salinity is expected to lie between the 34.7 psu of the convecting Labrador Sea and the approximately 32 psu of the North Pacific at the same latitude. From the mean salinity distribution in the Labrador Sea (Levitus 1982) a value of $S_1^* = 33.5$ psu seems plausible. For the deep box parameters T_2^* and S_2^*, the average values for the waters in the North Atlantic in the latitude band of the Labrador Sea (but excluding the Labrador Sea itself) and between 200 and 2000 m depth are $T_2^* = 4.3\,^\circ\mathrm{C}$ and $S_2^* = 34.97$ psu. These values are almost equivalent to assuming Irminger Sea conditions for the nonconvecting Labrador Sea. They yield a deep box time scale of $\tau_2 = 20$ years. In comparison with other studies on the exchange rates in the deep Labrador Sea (Khatiwala and Visbeck 2000), this value is rather large; yet it is in agreement with the small trends in the deep layer during the GSA (see Table 1). Finally, we applied one further assumption, namely to restrict the length of the winter convection event t_{wce} to less than 20 days, as low cost function values were in some cases reached also with excessively long convection periods. All the major conclusions of this paper are insensitive to the somewhat arbitrary assumptions described in this paragraph and hold for a wide range of S_1^*, T_2^*, S_2^*, and maximum t_{wce}. The constraints from the OWS Bravo data are sufficient to determine the stability properties discussed below.

The fitting procedure now arrives at a global minimum of the cost function K and is repeated with parameter matrices of higher resolution (in parameter space) to localize the global minimum more exactly. The optimal parameter set thus determined is shown in Table 2. The value for h^* is a result of our analysis in the previous section and not of the fitting procedure. Changes in h^* (by assuming an upper box depth of 100 m, say) lead to quantitatively slightly different results. To measure the parameter uncertainty we determined all parameter combinations for which the cost function remains within 10% of its minimum value. This defines an uncertainty range for each parameter (except the fixed ones).

A comparison of a model run using the optimal parameter set with the OWS Bravo data (Fig. 4) shows that many relevant features of the data are captured by the model. This includes the two upper box seasonal cycles during phase 1, leading to a winter convection event each year. The minimum difference between S_1 and S_2 is very small in the model time series because the model mixes the two boxes completely, whereas for a number of reasons the complete mixing is not visible in the observational data. In this cyclostationary state of the model, without stochastic forcing, there is no interannual variability in any variable. Phase 2 starts with a negative salt flux anomaly in the model run. Convection ceases, and the model reproduces the observed trends of all four variables: the upper box cools and freshens strongly, while the deep box warms and becomes more saline. (The trends in the deep box are hardly visible on the scale of Fig. 4). In accordance with the observational data, the model deep box is not an infinite reservoir of heat and salt (as in former versions of the box model, e.g., Lenderink and Haarsma 1994), but receives diffusive fluxes from the neighbouring waters. The end of phase 2, marked by the beginning salinification of the upper layer in the OWS Bravo data, is achieved in the model by a cold and saline anomaly in the upper box

forcing. Only a strong anomaly in surface forcing is capable of turning convection on again. It is clearly seen from the trends in Fig. 4 that the longer convection is off, the lighter the surface layer becomes and the stronger an anomaly must be to restart convection. This point is studied in greater detail in the next section. After convection is started again, the model returns to its previous convecting state. The different phase 3 state of the OWS Bravo data cannot be captured by the model.

4 Stability of the Labrador Sea

The stable states of the model under varying parameters are now explored. Of particular interest are changes in the upper box buoyancy forcing, that is in T_1^* and S_1^*. Changing other parameters will lead to similar pictures. (However, reducing the model time scales by 1 order of magnitude leads to decadal oscillations which will be described elsewhere.) The parameter space section along these two axes (Fig. 5) shows the stable model states: the convecting state, the nonconvecting state, and a bistable domain. The states are cyclostationary due to the presence of the seasonal cycles. The domain boundaries are given in a good approximation by the analytical expressions for the necessary conditions for the nonconvecting state:

$$\alpha\left(T_1^* - A_T - T_2^*\right) > \beta\left(S_1^* - A_S - S_2^*\right) \ , \tag{8}$$

and for the convecting state:

$$\alpha\left(T_1^* - A_T - T_2^*\right) < \beta\left(S_1^* - A_S - S_2^*\right)\frac{\tau_{1T} + \tau_2 h^*}{\tau_{1S} + \tau_2 h^*} \ , \tag{9}$$

compare Eqs. (9) and (12) in Part I. Since we introduced seasonal cycles such that convection now occurs around the temperature *minimum* and the salinity *maximum* in the upper box, the respective amplitudes are *subtracted* and *added* in Eqs. (9) and (12) in Part I. The shape of Fig. 5 is in agreement with earlier studies (Lenderink and Haarsma 1994). We can thus conclude that the

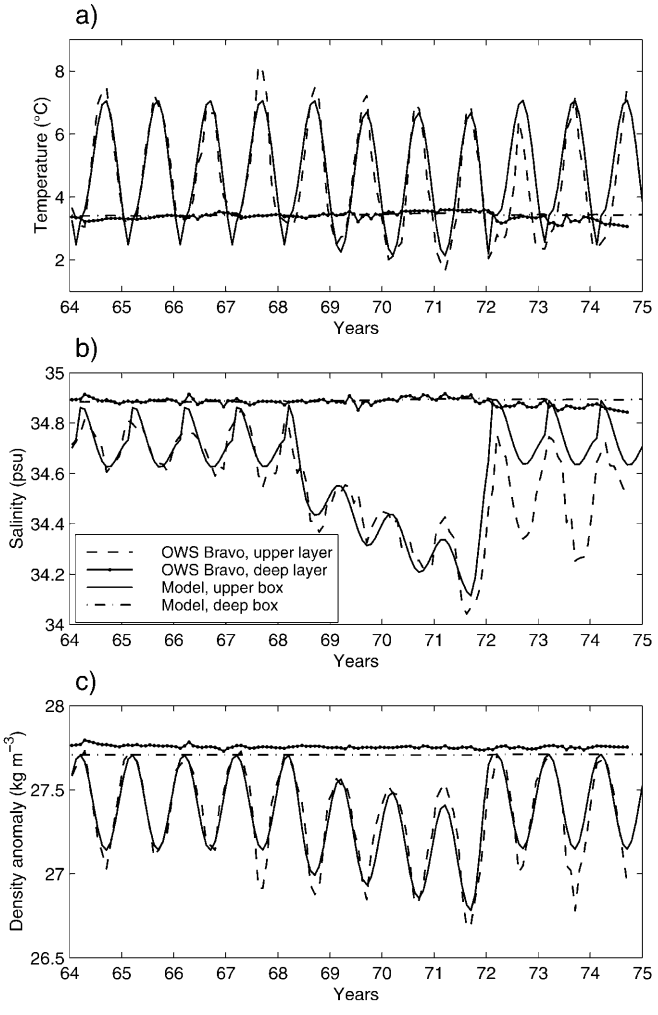

Fig. 4a–c Comparison of model output (*solid* upper box; *dash–dotted* deep box) and observed time series (*dashed* upper layer; *dotted* deep layer) for temperatures (**a**), salinities (**b**), and densities (**c**). The model was run with the optimal parameter set. All graphs are in monthly means. The GSA, also called phase 2 in the text, was started by adding a negative salt flux anomaly in the upper box during April to June 1968. The GSA was stopped by a positive salt flux anomaly during October to December 1971 accompanied by a cold anomaly in the upper box temperature forcing. See text for further explanation

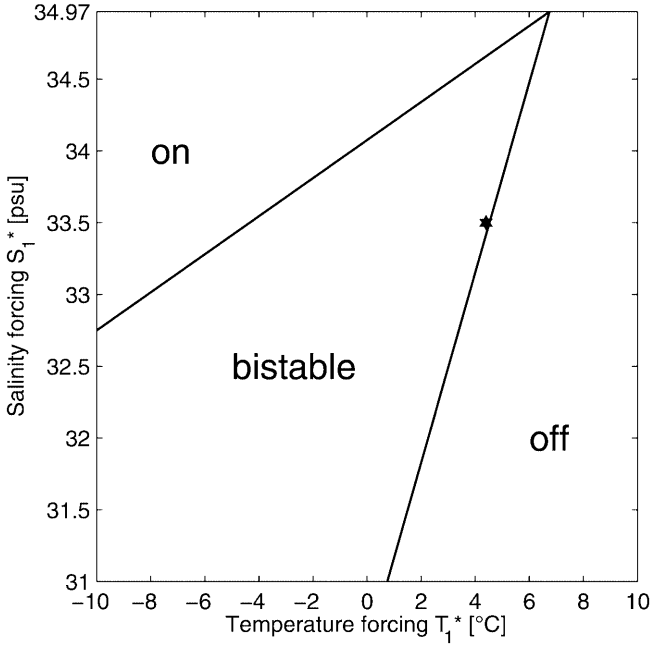

Fig. 5 Stable model states depending on the upper box buoyancy forcing parameters. Convection is either "on" or "off". These two stable states overlap in a bistable domain. The *asterisk* denotes the position of the optimal parameter set in the parameter space. The stable states were determined for varying only T_1^* and/or S_1^*, keeping the other parameters from the optimal set constant. For $S_1^* > 34.97$ psu the upper box would be forced to a higher salinity than the deep box; this case does not occur in high-latitude deep convection as in the North Atlantic. For other parameter sets within the uncertainty range or with different values of S_1^* or S_2^*, the slopes of the domain boundaries change, but the marginal position of the parameter set itself is a robust feature. The inclusion of sea-ice effects would affect the diagram for very low T_1^* values

presence of the seasonal cycles does not change the basic stability properties of Welander's (1982) box model.

The position of the model Labrador Sea – represented by the optimal parameter set – in the parameter space section is marked by the asterisk in Fig. 5. It is in the bistable domain, i.e., both states, convecting and nonconvecting, are steady states of the model under the given conditions. A sufficiently large anomaly can switch convection on or off. Moreover, the model is located very close to the domain where *only* the nonconvecting state is stable. Changing the buoyancy forcing by a few tenths of a degree or a few tenths of psu will lead to the convecting state becoming unconditionally unstable. In other words, there are two possible ways for suppressing convection. In the first way, convection is temporarily switched off by an anomaly, but can be restarted later by an opposite anomaly, while the average properties of the buoyancy forcing (i.e., the model parameters) do not change. This is the GSA case depicted in Fig. 4, and this may be the case with a 50 year-long spell of convection in the Labrador Sea being switched "on" among centuries without convection that Tett et al. (1997) found in a coupled GCM. In the second way, under a slowly changing buoyancy forcing (e.g., T_1^* or S_1^*) the convecting state eventually becomes unstable and convection stops. This scenario could apply to the global warming GCM run of Wood et al. (1999), in which Labrador Sea convection stops early in the 21st century. The role of the anomalies triggering state transitions in both cases is examined in detail in the next section.

We checked systematically how the stability diagram changes for different choices of S_1^*, S_2^* (or T_2^*) and maximum t_{wce}, as well as for parameter sets in the uncertainty range defined by a 10% change in the cost function. Equations (8) and (9) show that those parameter changes affect the width of the bistable domain: for instance, for a larger difference $(S_1^* - S_2^*)$ the bistable domain (in terms of T_1^*) is wider. However, the distance of the model solution from the border of the bistability domain varies only by some tenths of a degree on the T_1^* axis and similarly small amounts on the S_1^* axis. This is a consequence from the OWS Bravo data which explicitly show that a freshwater anomaly equivalent to 0.2 psu induced a transition from the convecting to the nonconvecting state. The precarious position of the Labrador Sea, in the bistable domain but close to the nonconvecting domain, is therefore a robust feature of the fitted model.

5 A stochastic climate model of deep convection

5.1 Motivation

In the previous section, the essential role of surface buoyancy forcing anomalies switching convection on and off became clear. We look for a meaningful way to introduce variability that generates these anomalies in the context of the simple box model used here. Observations (Lilly et al. 1999; Marshall and Schott 1999) suggest that heat flux anomalies (due to weather activity) and freshwater flux anomalies (by advection) act together to trigger or suppress convection. Yet the heat fluxes clearly prevail in their contribution to the overall buoyancy forcing (Sathiyamoorthy and Moore 2001). Hence, we focus on heat flux variability as the primary variability component.

The characteristic time scale of synoptic cyclones is a few days. Following the concept of Hasselmann's (1976) stochastic climate model, we parameterize synoptic-scale variability by a noise term added to the surface heat flux forcing. To estimate this noise term quantitatively, a 52-year-long time series of daily net surface heat flux from NCEP reanalysis data (Kalnay et al. 1996) was analyzed. The analysis – given in detail in the Appendix – motivates a noise term consisting of red noise ξ_t with a decorrelation time of about 6 days times a standard deviation σ, so that Eq. (1) is extended to:

$$\frac{\mathrm{d}T_1}{\mathrm{d}t} = \frac{1}{h^* \tau_c (\Delta\rho)} (T_2 - T_1) + \frac{1}{\tau_{1T}} \left[T_1^* - A_T \cos(2\pi t) + \sigma \xi_t - T_1 \right] . \qquad (10)$$

With this step, two classical conceptual models – for deep convection (Welander 1982) and for high-frequency atmospheric forcing of the ocean (Hasselmann 1976) – are combined to give a simple model that could be called a *stochastic climate model of deep convection*. Extended by the seasonal cycles and fitted to the OWS Bravo data, the box model is now suitable to study the variability of deep convection in the Labrador Sea.

5.2 Stochastic forcing and state transitions

Thirty years from a model run with stochastic forcing are displayed in Fig. 6. Several times, convection is interrupted for a few years, indicated by the small minimum density difference $\Delta\rho$ between the upper and the deep box (Fig. 6c). The upper box temperature T_1 (Fig. 6a) is the only variable directly influenced by the noise, so it shows the strongest variability. Apart from the convective mixing induced by that variability, the other three variables evolve in an unperturbed way. Similarly to the GSA in the OWS Bravo data (Fig. 1), in the non-convecting years the upper box tends to freshen (Fig. 6b) and cool (Fig. 6a), until a cold anomaly restarts convection again. In contrast to the observed GSA, the nonconvecting phase in this model run is kicked off by a warm anomaly in the upper box, not a freshwater anomaly. This is due to the fact that the stochastic variability appears in T_1 only, not in S_1, so by construction only temperature anomalies can appear. However, through the Hasselmann mechamism the

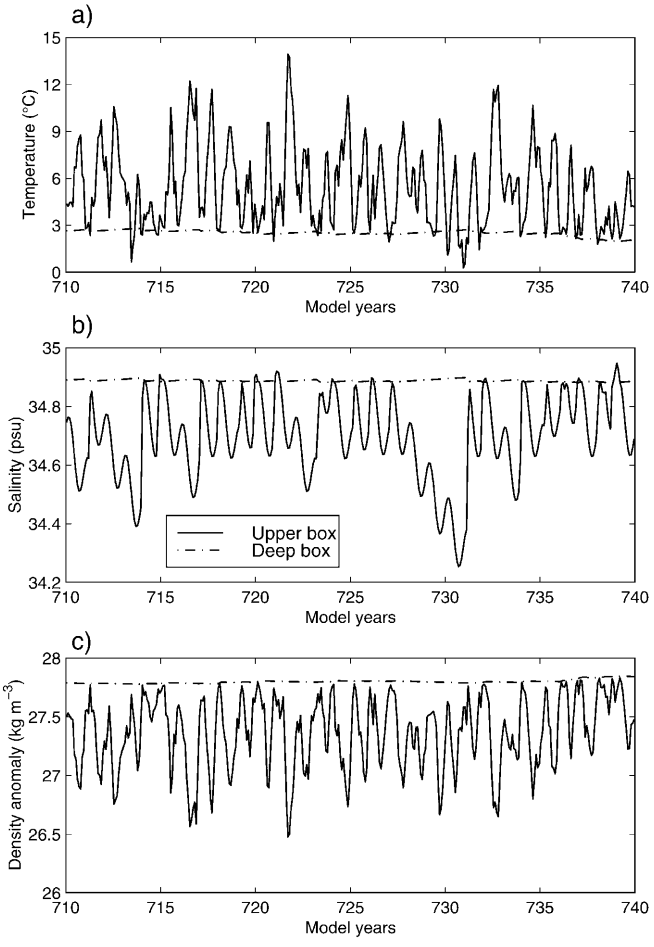

Fig. 6a–c Time series of monthly means from the stochastically forced model: temperature (**a**), salinity (**b**), and density (**c**) of the upper box (*solid*) and the deep box (*dash–dotted*). The model was run with the optimal parameter set and a standard deviation $\sigma = 18\,^{\circ}\mathrm{C}$ of the stochastic forcing. The 30 model years shown include several interruptions of the convecting state. The difference to the observed GSA (cf. Fig. 1) is that the positive buoyancy anomaly needed to stop convection is achieved here by local heat fluxes rather than advective freshwater fluxes. See text for details

upper box integrates the weather noise to (intra-)seasonal temperature anomalies (because $\tau_{1T} = 5$ months). Thus, integrated synoptic heat flux anomalies and advected intraseasonal freshwater anomalies have the same impact on $\Delta\rho$.

The key point here is that a comparatively short anomaly (lasting for a couple of months) triggers the state transition, and that the long-term trends evolving afterwards are due to the internal dynamics until the next anomaly triggers the next state transition. We hypothesize that this picture of a bistable water column holds for the Labrador Sea. This implies that the falling surface salinity from 1968 to late 1971 does not result from anomalous freshwater input during this whole time, nor does the return to convecting conditions in 1972 result from an end to the anomalous forcing. Rather, the falling salinity requires only a brief trigger anomaly (which could even have been thermal rather than freshwater) that prevents convection in 1968. Convection then cannot recover by itself but requires another substantial trigger event. The longer convection has been off, the larger the trigger needs to be. Had the winter of 1972 not been such a harsh one, subsequent winters would have needed to be even colder to restart Labrador Sea convection.

This hypothesis is consistent with the conclusion of Dickson et al. (1996), who analyzed the 1972 convection onset in the Labrador Sea in detail in the observed data. They conclude that the jump-like rise of the upper layer salinity is explicable only by anomalously strong wind forcing that mixes saline intermediate waters into the mixed layer; advective processes cannot lead to such strong changes. In other words, observational data show that the termination of the GSA in the Labrador Sea was achieved by anomalous weather conditions at the ocean surface, not by its internal dynamics.

The presence of noise leads to a qualitative change in the model's stability behavior: the sharp domain boundaries depicted in Fig. 5 are replaced by more gradual changes in the frequency of the occurrence of convection. As a measure we use n_c, the fraction of years with convection out of all years in a model run. Figure 7 shows how n_c depends on the noise strength σ. Using the optimal parameter set with the convecting state as initial condition and increasing σ (Fig. 7a), n_c drops close to zero for weak noise. This reflects the marginal position of the optimal parameter set in parameter space (Fig. 5). Any small perturbation shifts the model into the nonconvecting state, but the small perturbations are not able to induce a jump back to convection. For $\sigma > 12\,^{\circ}\mathrm{C}$, the convecting state is reached in some cases. The fraction of convective years rises quickly and asymptotically reaches a value of $n_c \approx 0.75$. In this regime of strong noise, the noise tends to override the deterministic stability properties. This feature becomes clearer when considering changes in the surface buoyancy forcing T_1^* in addition. The contour plot in Fig. 7b displays how n_c depends on σ and T_1^*. For $\sigma = 0$, Fig. 7b corresponds to the deterministic (not stochastically forced) parameter space section (Fig. 5) with only the nonconvecting state being stable for $T_1^* > 4.5\,^{\circ}\mathrm{C}$. For low noise and large T_1^* there is a large, wedge-shaped domain where almost no convection events occur. The shape of this domain can be understood when thinking of the convecting state becoming less and less stable for larger T_1^*. Then, for larger T_1^* a smaller amount of noise is needed to trigger a jump into the nonconvecting state. Since the nonconvecting state becomes more stable for larger T_1^*, a larger amount of noise is necessary to trigger jumps back into the convecting state. Figure 7b also shows that beyond this wedge-shaped domain the noise is capable of keeping the model in the convecting state for more than half of the time even when this state is unstable in the deterministic case. The exact extent of this domain depends on the respective parameter set: for instance, a lower S_1^* leads to a less dense upper box

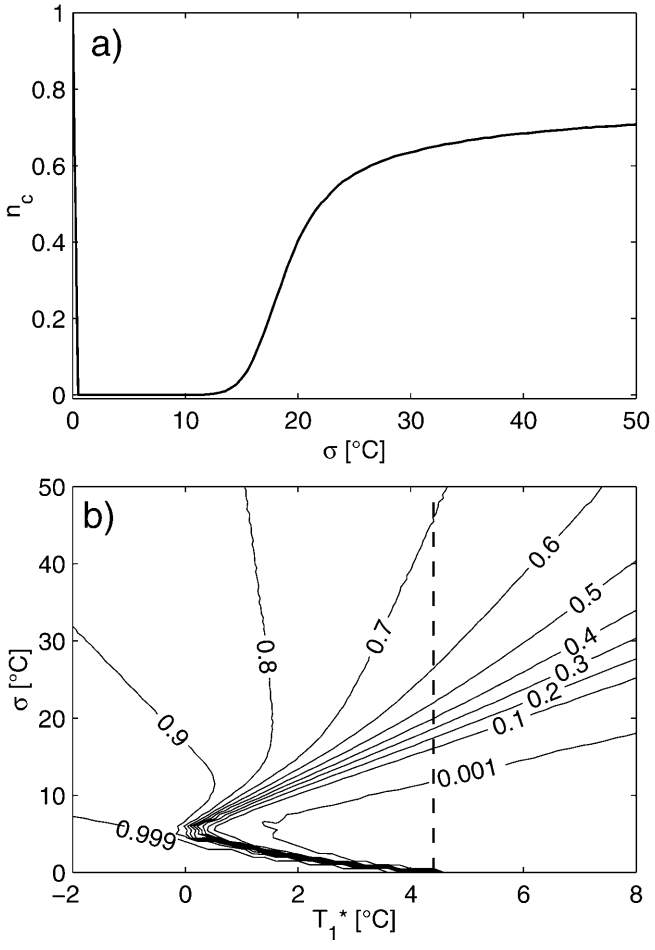

Fig. 7 **a** Dependence of the fraction n_c of convection years in a long model run (10^6 years) on the noise level σ. The model was run with the unchanged optimal parameter set and convecting state initial conditions. The deterministic case, without stochastic forcing, is at $\sigma = 0$. **b** Contour plot of n_c as a function of σ and additionally of the upper box temperature forcing T_1^*. The *dashed line* denotes the position of the graph shown in **a**. With a nonconvecting initial condition, $n_c = 0$ for $\sigma = 0$ for all values of $T_1^* > -4\,^\circ\mathrm{C}$ (i.e. in the bistable regime, see Fig. 5). The lower part of **b** would change accordingly

in the nonconvecting state, so convection is harder to trigger, and the domain becomes wider.

In the Appendix it is shown that for the Labrador Sea conditions in the box model σ is likely to be near or larger than 15 °C. This means that the model is located in a domain where n_c is sensitive to changes in the surface forcing. There are two ways of making convection occur less often: either by decreasing the variability σ or by increasing the surface temperature T_1^* (or equivalently decreasing S_1^*); but convection can still occur even when the convecting state is unstable in the deterministic case. If an increase in T_1^* is taken as a crude representation of a global warming scenario, then these results suggest that the frequency of Labrador Sea convection could decrease substantially due to a future warming (and/or freshening) unless variability increases strongly at the same time.

5.3 Stability and residence times

It is straightforward to estimate how the stability of the states (with respect to perturbations) changes quantitatively with varying model parameters. Consider a particle in an ideal double-well potential. If the particle is initially in one well, then added noise will rattle it. To leave the initial potential well and jump to the other one, the particle has to overcome a potential difference ΔU. If the noise is Gaussian distributed, eventually one perturbation, occurring after time t_r, is large enough for the particle to hop into the other well. In this ideal case there is a simple relation between the particle's mean residence time $\langle t_r \rangle$ in one well, the potential difference ΔU and the standard deviation σ of the added noise:

$$\Delta U = \sigma^2 \log\langle t_r \rangle + C \ . \tag{11}$$

This equation directly follows from Arrhenius' formula (Gardiner 1994). The constant C is a function of the potential's curvature at the well bottom and at the potential hill that separates the two wells.

The two stable states of the box model (with four variables) cannot be expressed as minima of an (one-dimensional) potential. However, there is a way to use Eq. (11) for our purposes. The potential U can be interpreted as a *quasipotential* (see Freidlin and Wentzell, 1998, for a rigorous definition). Then, the potential difference ΔU is the necessary perturbation strength for a state transition. The larger ΔU is, the more stable is the state. The mean residence times $\langle t_c \rangle$ and $\langle t_n \rangle$ in the convecting and the nonconvecting state can be estimated from long model runs. We do this for leaving σ constant and varying T_1^* only. With the help of Eq. (11) we are then able to give a quantitative estimate of the relative stability of the two states as a function of the surface forcing. Figure 8 shows how the logarithms of the residence times change with varying T_1^*. For low values of T_1^*, the convecting state is clearly the more stable one. Conversely, for warm surface forcing the stability of the nonconvecting state increases strongly. In the case of the optimal parameter set (dashed line in Fig. 8), the nonconvecting state is about twice as stable as the convecting one. This gives a quantitative understanding of Fig. 7a: weak noise can provide the anomalies to jump into the nonconvecting state, but anomalies twice as large, necessary for the jump back, occur only extremely rarely. Note the contrast between the sharp stability domain borders of the deterministic model (Fig. 5) and the smooth shape of the stability curves for the stochastic case.

From Fig. 8 we see that the mean residence times for the optimal parameter set and a standard deviation of $\sigma = 18\,^\circ\mathrm{C}$ in the stochastic forcing are $\langle t_c \rangle = 3.5$ years and $\langle t_n \rangle = 11$ years. Thus, the average time for the model to jump from one state to the other and back is about 15 years. In other words, the typical time scale for the variability is in the decadal range. This is clearly different from the synoptic time scale of the stochastic surface forcing. The effect of the weather noise is here to

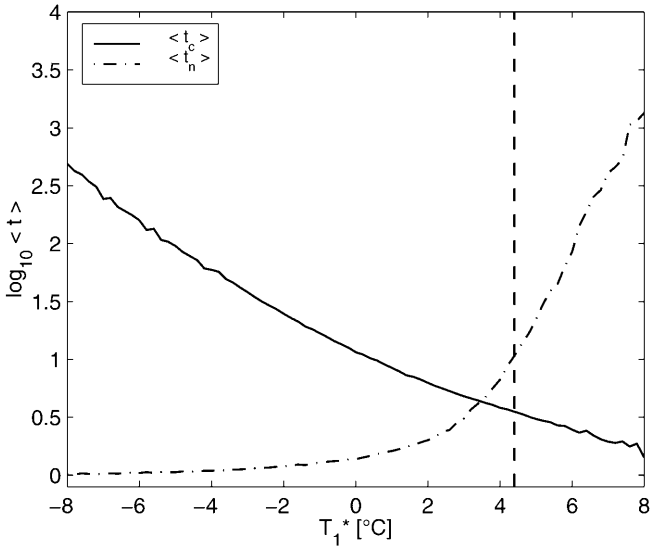

Fig. 8 Mean residence times in the convecting state ($\langle t_c \rangle$, *solid*) and in the nonconvecting state ($\langle t_n \rangle$, *dash–dotted*) in dependence on the mean surface temperature forcing T_1^* for constant standard deviation ($\sigma = 18\,^\circ\mathrm{C}$) of the stochastic surface temperature forcing. The *dashed line* indicates the position of the optimal parameter set: the convecting state lasts 3.5 years and the nonconvecting state 11.2 years, on average

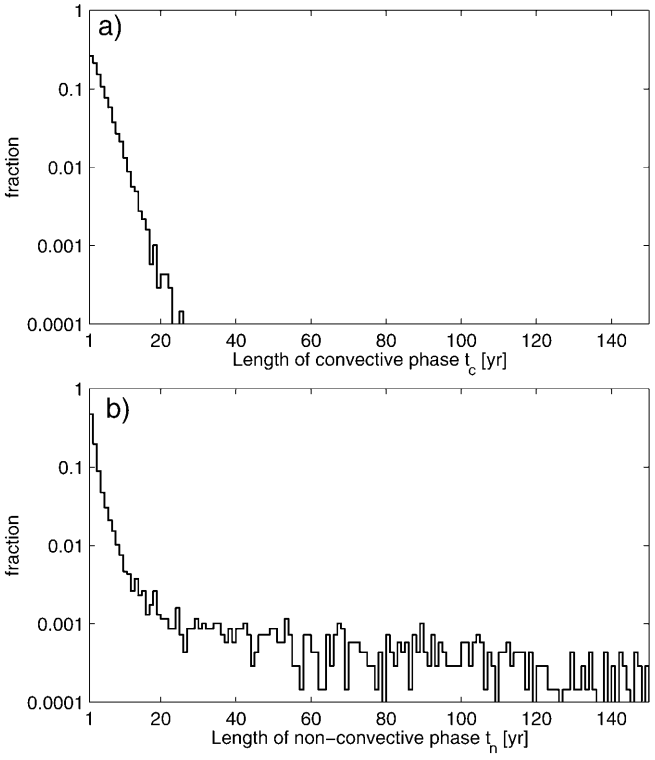

Fig. 9a, b Distribution of residence times in the convecting state (**a**) and the nonconvecting state (**b**). The frequency of every single residence time is given as a fraction of the total number of residences during a 10^5 year run. Hence, the distributions are approximate probability density functions. The fraction axis is logarithmically scaled; the small bar in panel **a** at $t_c = 25$ years corresponds to one single occurrence. The model was run with the optimal parameter set and a standard deviation of $\sigma = 18\,^\circ\mathrm{C}$, yielding a fraction of convective years $n_c = 0.26$. The tail of the distribution in (**b**) was cut arbitrarily; the maximum t_n is 526 years

excite intraseasonal variability in the mixed layer, which in turn triggers interannual to decadal variability. The deep ocean, being isolated from the atmosphere nearly at all times, ''sees'' the synoptic variability through the ''window'' of deep convection events – but responds to this forcing with its own typical decadal time scale. In this way, deep convection is a prominent example of time-scale interactions in the climate system.

The analysis can be carried one step further by extending our view from the *mean* residence times $\langle t_c \rangle$ and $\langle t_n \rangle$ to their *distributions* $p_s(t_c)$ and $p_s(t_n)$ that are equivalent to the stationary *probability density functions* (pdfs). This draws a more complete and accurate picture of the variability. The different shapes of the two pdfs in Fig. 9 stand out. The residence times of the convecting state (Fig. 9a) are distributed following a straight line, with $t_c > 20$ years occurring only rarely during a 100 000-year model run. In comparison, $p_s(t_n)$ has a bent shape, with high probability density for very short residence times and some occurrences of t_n exceeding 20 years. These features are obscured when considering the mean value only: while the *mean* residence time is $\langle t_n \rangle = 11$ years, we learn from $p_s(t_n)$ that the time series will contain many cases of only a few years without convection, but also some occasions where convection is interrupted for more than 100 years. With the *probability distributions* of the residence times (Fig. 10) this can be quantified. For instance, there is a 10% probability for the nonconvecting state to last longer than 13 years, but the convecting state will do so only with a probability of 1.5%. Hence, the observation of two decades without deep convection, as in the years 1982–2001 in the Greenland Sea (Rhein 1996; Visbeck and Rhein 2000; J. Holfort personal communication), is not necessarily a sign of a global climatic trend, but could be within the natural variability properties of a convective water column.

The difference between the two distributions in Fig. 9 can be understood qualitatively in the framework of ''runs'' introduced by von Storch and Zwiers (1999). A ''run'' is defined as the time that a stochastic process spends uninterruptedly on one side of its mean value. Von Storch and Zwiers analyzed AR(1) processes with varying autocorrelation coefficient α_1, and found that for $\alpha_1 = 0$ (white noise) the run length pdf decreases exponentially. In a logarithmic plot this pdf of the run lengths, or residence times, appears as a straight line, as in Fig. 9a. For red noise ($\alpha_1 > 0$), long residence times are more probable at the expense of intermediate times, which gives the bent shape of the graph in Fig. 9b. In other words, the linear shape of $p(t_c)$ in Fig. 9a means that the probability for a convection stop is equal in all years, whereas the exponential shape of $p(t_n)$ in Fig. 9b stems from the diminishing probability of leaving the nonconvecting state with increasing residence time t_n. This feature again reflects the positive salinity feedback.

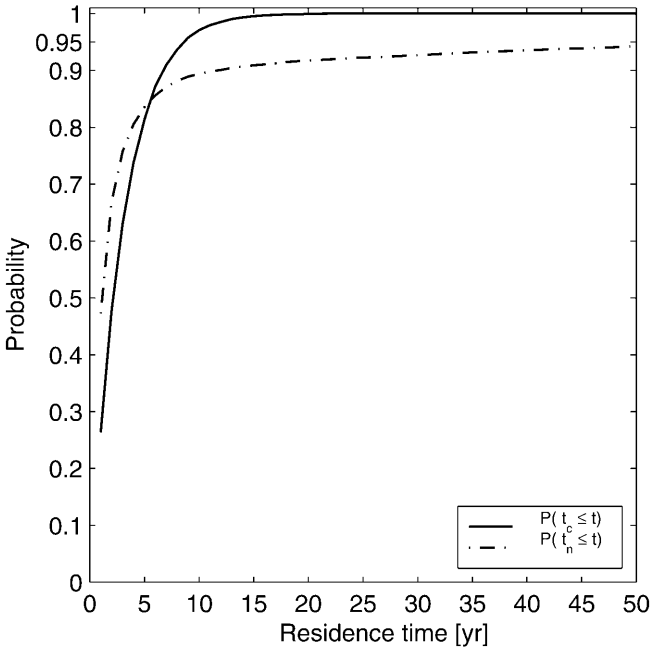

Fig. 10 Probability distribution of residence times in the convecting state (t_c, *solid*) and the nonconvecting state (t_n, *dash–dotted*) from the same model run as in Fig. 9. For any time t, the probability distribution gives the probability of the residence time being smaller than or equal to t

6 Conclusions

The aim of this paper is to understand some basic stability and variability properties of open-ocean deep convection. The simplest possible model for this purpose is a two-box model of a water column in a potentially convective part of the ocean. The simplicity of the model allows the exploration of large parts of the parameter space. Observational data from the Labrador Sea show phases with and without convection and are used for the parameter determination of the model. Not all of the model parameters are well constrained by the dataset. Some properties of a steady nonconvecting state in the Labrador Sea have to be assumed in order to close the problem of parameter determination.

With the parameter set obtained, the position of the Labrador Sea in a stability diagram can be determined. For a certain region in the parameter space the model has two stable states, with convection being either ''on'' or ''off'' each winter. With the Labrador Sea parameters the model is located in this bistable domain, so that anomalies in the forcing are capable of triggering jumps from the convecting to the nonconvecting state and back. The model shows that lasting anomalies similar to the Great Salinity Anomaly (1968–1972) can be triggered by short-term anomalies in the surface conditions suppressing convection in one winter. This mechanism of transitions between two stable states can explain the basic properties of the time series from OWS Bravo. The longer a nonconvecting phase lasts, the harder it is to interrupt it, suggesting that convection might have ceased for much longer if 1972 had not been an anomalously harsh winter over the Labrador Sea. Note the contrast with the hypothesis of Dickson et al. (1996), and Lilly et al. (1999), who conclude that anomalous surface conditions (be it in freshwater advection or local heat fluxes) lasted throughout the 4 years of the GSA and were needed to suppress deep convection. Our results suggest that a different mechanism might have been working here: a short-term perturbation switched convection off, and the subsequent evolution was governed by internal, local dynamics until another perturbation switched convection on again.

We found that the position of the model Labrador Sea in the bistable domain is very close to the border to the monostable domain without stable convecting state, irrespective of the parameter assumptions about the stable nonconvecting state. This position is precarious: changing the ocean's surface forcing by about 1 °C towards warmer conditions leads to the convecting state becoming unconditionally unstable. Such a shutdown of Labrador Sea convection occurred in a global warming scenario computed with a coupled AOGCM (Wood et al. 1999). It is not clear yet whether the additional feedbacks in complex models lead to a less precarious stability of deep convection.

With the model being in the bistable domain, anomalies in the forcing are essential to excite state transitions. Therefore, we included stochastic variability in the model as in the stochastic climate model of Hasselmann (1976), to parameterize weather variability over the Labrador Sea. In this way, two seminal conceptual models, for deep convection (Welander 1982) and for high-frequency atmospheric forcing of the ocean (Hasselmann 1976), are combined to give a ''stochastic climate model of deep convection''. The parameters of the noise term are estimated from a daily time series of surface heat flux.

We conclude that realistic noise amplitudes are large enough to blur the clear picture of the stability diagram. Even when the convecting state is unstable in the deterministic case, the variability excites frequent convection events. Conversely, the observation of intermittent convection in the real ocean or in a model does not allow direct inferences about the stability of the underlying deterministic state.

When stochastic variability is taken into account, a warming (and/or freshening) at the surface will not lead to a complete stop of convection at a certain threshold, but rather to a decline in the frequency of convection events. The frequency of state transitions depends on the noise level in a highly nonlinear way. There are two plateaus for the frequency value: weak noise triggers hardly any jumps, and with strong noise the jumps become very frequent and occur every few years. In between is a small range of noise levels where the jump frequency steeply rises. For the plausible parameter range, the position of the Labrador Sea is in the region

of this steep rise, which reflects its sensitivity to changes in surface forcing.

The interannual to decadal jumps are triggered by the intraseasonal density anomalies in the upper box, which in turn are excited by weather noise in the surface heat flux. Advective freshwater anomalies have not been modeled here, but would have the same impact on the model dynamics. The probability distribution of the residence times in the nonconvecting state shows that there is a small, but not negligible probability for the nonconvecting state lasting a decade or longer. In contrast to a deterministic understanding of the system, this means that convection may start again after a long break due to natural variability. The recent 19-year-long cessation of deep convection in the Greenland Sea is thus not necessarily due to a long-term climatic trend but could be part of the normal stochastic variability properties of convection. On the other hand, a surface warming or freshening trend as may result from anthropogenic global warming can be expected to substantially reduce the frequency of convection in the Labrador Sea and possibly elsewhere.

Acknowledgements The work on this paper was supported by the Deutsche Forschungsgemeinschaft (Sfb 555 and Heisenberg-Programm). The paper is a contribution to a cooperation of the universities of Potsdam, Germany, and Lisbon, Portugal, promoted by the Deutscher Akademischer Auslandsdienst. We are grateful to Eva Bauer, Miguel Morales Maqueda, Adam Monahan, Axel Timmermann, and Udo Schwarz for inspiring discussions. Comments of two anonymous reviewers helped to clarify the paper.

Appendix

The analysis of a 52-year-long time series of daily net surface heat flux at the OWS Bravo site is presented here. This time series was extracted from the NCEP database (Kalnay et al. 1996). The aim is an estimate for the standard deviation and the decorrelation length of the noise term in Eq. (10). Assuming that the heat flux time series Q can be decomposed into an average flux Q_0, a seasonal cycle with amplitude A_Q, and a noise term ξ_t with standard deviation σ_Q,

$$Q = Q_0 + A_Q \cos(2\pi t) + \sigma_Q \xi_t \ , \tag{A1}$$

it turns out that, for averaging intervals of 1 to a few days, the autocorrelation function of the noise process ξ_t falls off to zero only after the first few lags. In other words, ξ_t can be modeled by an AR(1) process

$$\xi_t = \alpha_1 \xi_{t-1} + \zeta_t \ , \tag{A2}$$

where the value ξ_t at time t is determined by the value ξ_{t-1} at the previous time step times the autocorrelation at lag 1, α_1, plus a random value ζ_t from a Gaussian white noise process. With values for α_1 estimated from the NCEP time series, the decorrelation time as defined in von Storch and Zwiers (1999):

$$\tau_D = \frac{1+\alpha_1}{1-\alpha_1} \ , \tag{A3}$$

lies between 5 and 7 days, depending on the averaging interval of the time series. This decorrelation time is just the typical time scale for synpotic activity. Using this red noise forcing in the model (instead of pure white noise) is a more realistic parameterization of the high-frequency heat flux variability and renders the model results more robust to changes in the time step of the numerical integration scheme.

Apart from the decorrelation length, the second parameter we need to estimate for the noise term in Eq. (10) is the standard deviation σ. Using the heat flux time series in Eq. (A1) to force the temperature of an ocean surface layer we write:

$$c_p \rho_0 h_1 \frac{\mathrm{d}T}{\mathrm{d}t} = Q_0 + A_Q \cos(2\pi t) + \sigma_Q \xi_t - \lambda T \ , \tag{A4}$$

with the specific heat capacity $c_p = 3990\,\mathrm{J\,kg^{-1}\,K^{-1}}$, a reference density $\rho_0 = 1028\,\mathrm{kg\,m^{-3}}$, the surface layer depth $h_1 = 50\,\mathrm{m}$, and a restoring constant λ. Setting $c_0 = c_p \rho_0 h_1$, this reads:

$$\frac{\mathrm{d}T}{\mathrm{d}t} = \frac{Q_0}{c_0} + \frac{A_Q}{c_0}\cos(2\pi t) + \frac{\sigma_Q}{c_0}\xi_t - \frac{\lambda}{c_0} T \ . \tag{A5}$$

Comparison with Eq. (10) shows how the standard deviation σ of the stochastic forcing in Eq. (10) is related with the standard deviation σ_Q of the heat flux time series:

$$\sigma = \sigma_Q \frac{\tau_{1T}}{c_0} \ . \tag{A6}$$

With a dependence on the averaging interval again, we find $\sigma_Q = 120$–$140\,\mathrm{W\,m^{-2}}$, which translates into $\sigma = 8$–$9\,^\circ\mathrm{C}$. This is a rather high value, larger than the average flux Q_0 – and still the seasonal cycle of the standard deviation itself, reaching its maximum in winter, has not been accounted for here. Sathiyamoorthy and Moore (2001), in their analysis of the buoyancy flux at OWS Bravo derived from weather ship data, obtain a similar result. The dominant role of the synoptic-scale variability of the heat flux in the Labrador Sea is highlighted here again. For simplicity we have added noise only to the surface temperature forcing. To obtain a realistic variability in the whole surface buoyancy flux, pronouncedly higher values of σ have to be assumed since the variability of surface freshwater fluxes and of advective transports of heat and salt have not been considered here. A range of $\sigma = 15$–$20\,^\circ\mathrm{C}$ seems therefore plausible. For the numerical integration we applied a semi-implicit Milstein scheme following Kloeden and Platen (1999, Chap.10). Its theoretical convergence is twice as good as with an ordinary Euler scheme, which is relevant regarding the discontinuities arising from the convective adjustment.

References

Belkin IM, Levitus S, Antonov J, Malmberg S-A (1998) Great Salinity Anomalies in the North Atlantic. Prog Oceanogr 41: 1–68

Cessi P (1994) A simple box model of stochastically forced thermohaline flow. J Phys Oceanogr 24: 1911–1920

Delworth T, Manabe S, Stouffer RJ (1993) Interdecadal variations of the thermohaline circulation in a coupled ocean–atmosphere model. J Clim 6: 1993–2011

Dickson RR, Lazier J, Meincke J, Rhines P, Swift J (1996) Long-term coordinated changes in the convective activity of the North Atlantic. Prog Oceanog 38: 241–295

Dickson RR, Meincke J, Malmberg S-A, Lee AJ (1988) The "Great Salinity Anomaly" in the northern North Atlantic 1968–1972. Prog Oceanog 20: 103–151

Fofonoff NP, Millard RC Jr (1984) Algorithms for computation of fundamental properties of sea water. Unesco Tech Pap Mar Sci 44, UNESCO, Paris

Frankignoul C, Hasselmann K (1977) Stochastic climate models, Part II: Application to sea surface temperature anomalies and thermocline variability. Tellus: 29: 289–305

Freidlin MI, Wentzell AD (1998) Random perturbations of dynamical systems. 2nd edn. Springer, Berlin, Heidelberg, New York

Gardiner CW (1994) Handbook of stochastic methods for physics, chemistry and the natural sciences. Springer, Berlin, Heidelberg, New York

Hasselmann K (1976) Stochastic climate models, Part I: Theory. Tellus 28: 473–485

Hirschi J, Sander J, Stocker TF (1999) Intermittent convection, mixed boundary conditions and the stability of the thermohaline circulation. Clim Dyn 15: 277–291

Kalnay E, Kanamitsu M, Kitler R, Collins W, Deaven D, Gandin L, Iredell M, Saha S, White G, Woollen J, Zhu Y, Chelliah M, Ebiszusaki W, Higgins W, Janowiak J, Mo K, Ropelewski C, Wang J, Leetmaa A, Reynolds R, Jenne R, Joseph D (1996) The NCEP/NCAR 40-years reanalysis project. Bull Am Meteorol Soc 77: 437–471

Khatiwala S, Visbeck M (2000) An estimate of the eddy-induced circulation of the Labrador Sea. Geophys Res Lett 27: 2277–2780

Klinger BA, Marshall J, Send U (1996) Representation of convective plumes by vertical adjustment. J Geophys Res 101(C8): 18175–18182

Kloeden PE, Platen E (1999) Numerical solution of stochastic differential equations. Applications of Mathematics, Vol. 23, Karatzas I, Yor M (eds). Third printing. Springer Berlin, Heidelberg, New York

Lazier JRN (1980) Oceanographic conditions at Ocean Weather Ship Bravo, 1964–1974. Atmosphere–Ocean 18: 227–238

Lenderink G, Haarsma RJ (1994) Variability and multiple equilibria of the thermohaline circulation associated with deep-water formation. J Phys Oceanogr 24: 1480–1493

Lenderink G, Haarsma RJ (1996) Modeling convective transitions in the presence of sea ice. J Phys Oceanogr 36(8): 1448–1467

Levitus S (1982) Climatological atlas of the world ocean. Technical report NTIS PB83–184093, NOAA/ERL GFDL, Princeton, NJ

Lilly JM, Rhines PB, Visbeck M, Davis R, Lazier JRN, Schott F, Farmer D (1999) Observing deep convection in the Labrador Sea during winter 1994–1995. J Phys Oceanogr 29: 2065–2098

Marshall J, Schott F (1999) Open-ocean convection: observations, theory, and models. Rev Geophys 37: 1–64

Molemaker MJ, Dijkstra HA (2000) Stability of a cold core eddy in the presence of convection: hydrostatic versus nonhydrostatic modeling. J Phys Oceanogr 30: 475–494

Munk W, Wunsch C (1998) Abyssal recipes II: energetics of tidal and wind mixing. Deep-Sea Res I 45: 1977–2010

Pierce DW, Barnett TP, Mikolajewicz U (1995) Competing roles of heat and freshwater flux in forcing thermohaline oscillations. J Phys Oceanogr 25: 2046–2064

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

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

Rahmstorf S (1995b) Multiple convection patterns and thermohaline flow in an idealized OGCM. J Clim 8: 3028–3039

Rahmstorf S (1999) Decadal variability of the thermohaline ocean circulation. In: Navarra A (ed) Beyond El Niño: decadal and interdecadal climate variability. Springer, Berlin, Heidelberg, New York, pp 309–332

Rahmstorf S (2001) A simple model of seasonal open ocean convection. Part I: Theory. Ocean Dyn 52: 26–35

Rhein M (1996) Convection in the Greenland Sea 1982–1993. J Geophys Res 101(C8): 18183–18192

Sathiyamoorthy S, Moore GWK (2001) Buoyancy flux at Ocean Weather Station Bravo. J Phys Oceanogr (in press)

Send U, Marshall J (1995) Integral effects of deep convection. J Phys Oceanogr 25: 855–872

Smith SD, Dobson FW (1984) The heat budget at Ocean Weather Station Bravo. Atmosphere–Ocean 22: 1–22

Tett SFB, Johns TC, Mitchell JFB (1997) Global and regional variability in a coupled AOGCM. Clim Dyn 13: 303–323

Visbeck M, Rhein M (2000) Is bottom boundary layer mixing slowly ventilating Greenland Sea deep water? J Phys Oceanogr 30(1): 215–224

von Storch H, Zwiers F (1999) Statistical analysis in climate research. Cambridge University Press, Cambridge

Weisse R, Mikolajewicz U, Maier-Reimer E (1994) Decadal variability of the North Atlantic in an ocean general circulation model. J Geophys Res 99: 12411–12421

Welander P (1982) A simple heat-salt oscillator. Dyn Atmos Oceans 6: 233–242

Wood RA, Keen AB, Mitchell JFB, Gregory JM (1999) Changing spatial structure of the thermohaline circulation in response to atmospheric CO_2 forcing in a climate model. Nature 399: 572–575