



# On box models of the North Atlantic thermohaline circulation: Intrinsic and extrinsic millennial timescale variability in response to deterministic and stochastic forcing

M. Stastna<sup>1</sup> and W. R. Peltier<sup>2</sup>

Received 19 September 2006; revised 22 May 2007; accepted 13 August 2007; published 25 October 2007.

[1] Millennial timescale variability in the North Atlantic Ocean circulation is often discussed in terms of concepts that are rooted in the dynamics of simple, low-dimensional box models. In this study we discuss possible explanatory mechanisms for the millennial timescale oscillations of the North Atlantic Thermohaline Circulation that have been revealed by the Greenland ice core record for the late glacial period. We subject three qualitatively different low-order models to stochastic and sinusoidal perturbations: (1) a bistable model, (2) a model with a single stable equilibrium point and a single stable periodic orbit (limit cycle) corresponding to a collapse-flush cycle in the North Atlantic circulation, and (3) a model with a single globally stable equilibrium point which nevertheless exhibits complex behavior when perturbed. We discuss both the physical nature of the model response and its parameter dependence. We conclude that the traditional definition of stochastic resonance should be expanded and that the temporal characteristics of the noise terms should be considered an integral part of model construction, as they profoundly affect the efficacy of a given explanatory mechanism.

**Citation:** Stastna, M., and W. R. Peltier (2007), On box models of the North Atlantic thermohaline circulation: Intrinsic and extrinsic millennial timescale variability in response to deterministic and stochastic forcing, *J. Geophys. Res.*, *112*, C10023, doi:10.1029/2006JC003938.

## 1. Introduction

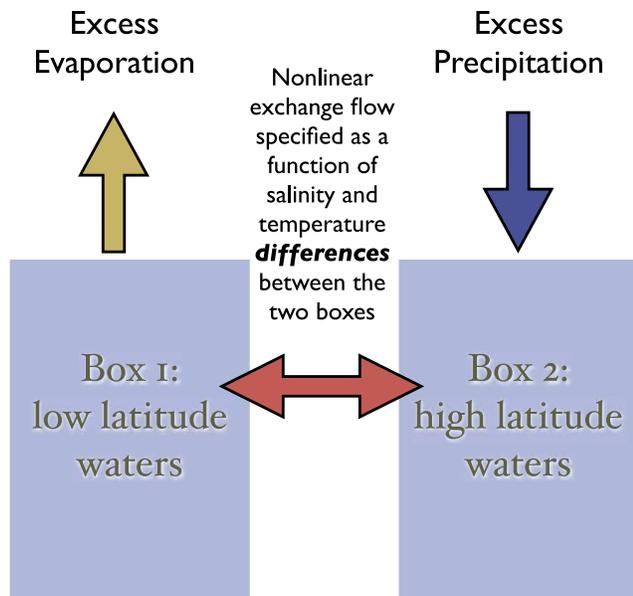
[2] It is well known that isotopic analyses of the Summit, Greenland ice cores [i.e., *GRIP members*, 1993; *GISP2 members*, 1993] reveal millennial timescale oscillations in atmospheric temperature during Oxygen Isotope Stage 3, a period of the most recent glacial period, 60–35 kya BP (1 kya = 1000 annums, BP is before present). These oscillations can exceed one half of the glacial-interglacial temperature change in amplitude. Inspection of these time series also reveals a saw-tooth modulation of the oscillations with a characteristic timescale of the order of 10 kya. In the following we will refer to the individual millennial timescale oscillations as D-O (Dansgaard-Oeschger) oscillations and their modulation as the Bond cycle [Bond *et al.*, 1993].

[3] It has been hypothesized that the D-O oscillations are intimately linked to changes in the strength of the North Atlantic thermohaline circulation (henceforth the NA THC), and that these changes are driven, in large part, by changes in the hydrological cycle, namely by variations in freshwater flux delivered onto the surface of the ocean at high

latitudes. The recent review by Broecker [2003] discusses both the evidence (based on ice and sediment cores) for this point of view, as well as the evidence for an alternative point of view that asserts the primacy of the tropical ocean-atmosphere system in driving millennial timescale climatic change, including changes in the NA THC. We will assume the former point of view throughout the following. The question then becomes how best to study the NA THC on relevant timescales. To date, a coupled ocean-atmosphere-ice sheet-land surface process general circulation model (GCM) study on timescales relevant to millennial timescale climate variations has not been possible. In the place of such a study a variety of reduced models, sometimes referred to as Earth System Models of Intermediate Complexity (EMICS), spanning a wide range of complexity and utilizing a variety of techniques, have been employed. At one end of the complexity spectrum are ocean circulation models coupled to simple atmosphere and sea ice models [Brix and Gerdes, 2003; Rahmstorf, 1995]. At the other end are low-order, conceptual models, generally involving a small number of nonlinear ordinary differential equations. Examples include the well-known Stommel two-box model (both in its original form [Stommel, 1961] and as developed by Cessi [1994]) and the three-box model of Sakai and Peltier [1999]. These two models are schematized in Figures 1 and 2, respectively. Other examples will be discussed at appropriate points in the text. The model discussed by Sakai and Peltier [1999] was developed with an eye to reproducing the

<sup>1</sup>Department of Applied Mathematics, University of Waterloo, Waterloo, Ontario, Canada.

<sup>2</sup>Department of Physics, University of Toronto, Toronto, Ontario, Canada.



**Figure 1.** Schematic of the Stommel two-box model. The model does not evolve individual salinities and temperatures, only their differences. In the majority of applications, including this study, the temperature difference is clamped to a constant value.

results of zonally averaged, multibasin models [Sakai and Peltier, 1997] in so far as their essential properties are concerned. The results of these multibasin model investigations suggested that the NA THC would be expected to be in a stable, time-independent state for high-latitude freshwater flux appropriate to modern conditions. However, for larger values of the high-latitude freshwater flux, this equilibrium point was shown to be replaced by a periodic orbit (or in other words the equilibrium point undergoes a Hopf bifurcation). Thus for high enough values of the freshwater flux the NA THC would undergo oscillations (corresponding to trips around a periodic orbit in phase space) that would continue even if the freshwater flux remained constant for subsequent times. This is in marked contrast to the two stable equilibrium points of the Stommel model in which variability requires external forcing and the NA THC tends to one of the two stable equilibrium points as soon as the forcing becomes time independent.

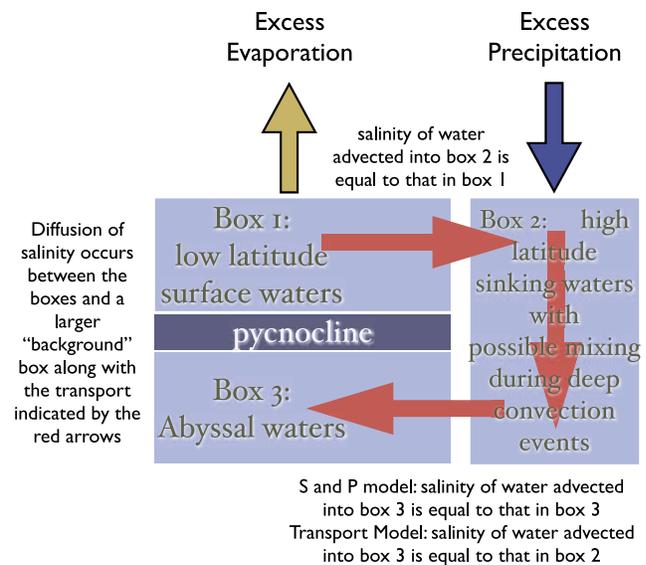
[4] The literature on the various spinoffs and extensions of the Stommel model is vast and it is not our intention to survey it here. However, one example worth mentioning is that which includes interhemispheric communication. Such a box model was first discussed by Rooth [1982]. Interhemispheric interactions are largely believed to behave according to the see-saw hypothesis [Broecker, 1998; Stocker, 2002] which asserts that a decrease in the intensity of the NA THC (i.e., due to a decrease of deep water production in the Greenland and Labrador Seas) will lead to an increase in the intensity of the SA THC (i.e., deep water production in Antarctic waters) and vice versa. The see-saw hypothesis is an active research area, see, for example, the box model work by Siddall *et al.* [2007] and the challenge to the see-saw hypothesis by Seidov *et al.* [2005] on the basis of numerical experiments with a fully coupled GCM.

As such, a detailed comparison of results discussed in the present with a corresponding study of a simple model of the see-saw type is a clear avenue for future research.

[5] In this paper we aim to address some aspects of low-dimensional dynamical system models of the NA THC. In particular we will focus on two competing explanations of millennial-scale variability of the NA THC, alternative explanations that find expression in models in this class. We will contrast systems for which time-dependent behavior is intrinsic, for example, through the existence of stable periodic orbits (limit cycles) produced by Hopf bifurcations of stable equilibrium points as parameters are varied, with systems for which time-dependent behavior is extrinsic, such as systems exhibiting stochastic resonance under a combination of sinusoidal and stochastic perturbations.

[6] We will focus the discussion herein on two models, namely the modification of the Stommel two box model as discussed by Cessi [1994] and a slight generalization of the three-box model due to Sakai and Peltier [1999]. We avoid discussion of models with a delay [i.e., Kurtze and Restrepo, 2001] since a delay greatly complicates the governing mathematics (indeed a single ordinary differential equation with delay is equivalent to an infinite system of ordinary differential equations without delay [Elsogols, 1966]).

[7] The layout of this paper is as follows. After a discussion of the modeling and simulation methodology, we begin by discussing stochastic perturbations of the



**Figure 2.** Schematic of the generalized Sakai and Peltier model. Salinity in each box diffuses between boxes and a larger “background” box or “bath.” The thermohaline circulation is modeled by a unidirectional (possibly nonlinear) transport of salinity from box 1 to box 2 to box 3 superimposed onto the diffusive exchange. The transport only occurs if the salinity in box 2 is larger than the background salinity. In box 2 varying amounts of mixing are assumed to occur during deep convection. In work by Sakai and Peltier [1999] the mixing with box 3 is complete, while in the “transport” model discussed in the present no mixing occurs.

version of the Stommel two-box model developed by *Cessi* [1994], henceforth referred to as C94. With an observation made in C94 as a starting point, we develop an elementary model for noise that is not white in time (i.e., has some memory) and outline several consequences of this memory. Further, we employ the Stommel model to develop an example of stochastic resonance. We next turn to a discussion of models with intrinsic variability (i.e., stable limit cycles). Using the three-box model developed by Sakai and Peltier we show how noise can destabilize an equilibrium point and lead to quasiperiodic behavior. Finally we utilize a variant of the Sakai and Peltier model to show that, contrary to one's intuition, stochastic resonance is possible in a system with a single, globally stable equilibrium point. However, the stochastically enhanced response in this case is richer than the traditional definition of stochastic resonance.

[8] A common thread throughout these discussions is the role of the temporal characteristics of the stochastic perturbations. While the vast majority of the mathematical literature on differential equations with stochastic terms concentrates on the case of white noise under the Ito interpretation [*Kloeden and Platen*, 1992], noise that has memory profoundly affects the regions of parameter space in which phenomena such as stochastic resonance can occur. For this reason we suggest that the noise memory should be treated as part of the model construction process as well as in evaluating the efficacy of any one particular explanatory mechanism.

[9] We utilize the relevant box models in question in dimensionless form, however the physical assumptions and limitations of each model are discussed at the beginning of section 2. In the conclusions section we discuss questions relating to the appropriateness of each of the models as descriptions of millennial scale variability of the North Atlantic. We offer criticisms of each of the various possible scenarios and integrate the present results into the literature on D-O oscillations.

## 2. Methodology

[10] Box models replace the complex dynamics of the ocean with a small number of communicating boxes, each with a well defined state (salinity and temperature, for example). We assume the reader is familiar with the idea behind box models (see, for example, the discussion by both C94 and *Sakai and Peltier* [1999]).

[11] The Stommel two-box model (see Figure 1 for a schematic), which we shall consider in the form presented in C94, consists of a low-latitude and a high-latitude box coupled by a nonlinear exchange function between the two boxes. The northern box is assumed to be forced by an excess of precipitation (negative salinity anomaly) while the southern box is forced by an excess of evaporation (positive salinity anomaly). There is no background pool of water and no assumed direction for the transport in the system, beyond that given by the choice of the mathematical form of the nonlinear exchange function. The three-box model due to *Sakai and Peltier* [1999] (see Figure 2 for a schematic), in contrast, has both a diffusive exchange with a background pool of water and an inherent transport direction, provided the THC is turned on: northward in the low-latitude surface box, sinking in the high-latitude box and

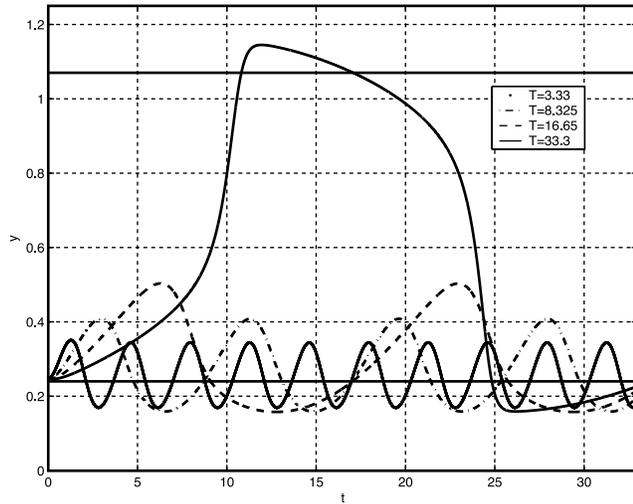
southward in the abyssal box. It is important to note that, as presented by *Sakai and Peltier* [1999], the model assumes two types of mixing. The first is given by diffusion between the boxes (as well as a larger 'bath'), while the second parameterizes the mixing that occurs during deep convection events. Indeed, as presented by *Sakai and Peltier* [1999], the model assumes that THC-transported water leaving the high-latitude sinking box has the same salinity as the water in the abyssal box. We shall weaken this assumption, and discuss the relevant mathematical development below. Physically, we find that a cessation of mixing with the abyssal waters during deep convection leads to a model with a single stable equilibrium. This is in better agreement with recent numerical experiments performed with a fully coupled GCM [*Peltier et al.* [2006] for anomalous freshwater forcing experiments of a modern climate. It is, however unclear whether the results for a modern climate should be extrapolated to the glacial climate, especially considering the suggestion made by *Sakai and Peltier* [1999] that the increase in steady freshwater flux onto the high-latitude North Atlantic maintains the glacial thermohaline circulation in a state perilously close to shutdown. As mentioned above, the Stommel model is schematized in Figure 1, while the generalized Sakai and Peltier model is schematized in Figure 2. The relevant mathematical equations are discussed in detail below.

[12] It is possible to systematically reduce the three-box model to the Stommel model via the following sequence of steps. First, eliminate the equation for the salinity of the abyssal box, by considering only the salinity of a high-latitude box and a low-latitude box, where the latter is now assumed to subsume the background pool and the original low-latitude surface box. Next reinterpret the transport and mixing functions as 'exchange' functions. Finally consider only the salinity difference between the two boxes. While the two models are thus formally similar, the Stommel model has the advantage of a smaller set of parameters, while the Sakai and Peltier model is more faithful in representing the unidirectional nature of the THC (at the cost of a larger number of parameters).

[13] The stochastic problems discussed in the following are formulated in terms of ordinary differential equations subjected to additive stochastic perturbations, as in

$$\frac{dy}{dt} = f(y, t) + \xi, \quad (1)$$

where  $y(t)$  is the variable one wishes to solve for and  $\xi$  represents the noise. There is a considerable literature on the mathematics, approximation and numerical solution of such equations [*Kloeden and Platen*, 1992; *Ottinger*, 1996; *Gardiner*, 1990]. Two divergent approaches are possible. The so-called Langevin approach, adopted in the following, considers individual integrations of (1). It requires an appropriate numerical method for the stochastic differential equation (applying standard techniques for ordinary differential equations such as the Runge-Kutta method gives erroneous results) and a reliable generator for the stochastic perturbation. In the following we apply the Milstein method [*Ottinger*, 1996] for the discretization of the stochastic differential equation and the Mersenne Twister algorithm [*Matsumoto and Nishimura*, 1998] to generate realizations



**Figure 3.** The  $y$  versus  $t$  curves for Stommel two-box model subjected to purely sinusoidal perturbations. The amplitude of the perturbations is fixed while the period,  $T$ , varies. Transitions between the two stable equilibria occur only for perturbations with a long enough period.

of the stochastic perturbation. The algorithm employed is order  $\Delta t$  in time in both the deterministic and stochastic parts of the governing equation and converges strongly [Ottinger, 1996]. The strong convergence implies we can make mathematically valid conclusions on individual paths, an essential property when discussing stochastic resonance. The Mersenne Twister algorithm allows for the rapid construction of large ensembles while maintaining confidence in the temporal characteristics of the stochastic perturbations. Signal to noise ratio values are computed from a projection onto the driving Fourier component. When necessary, full power spectra are computed via the FFT. A variety of spectral estimation techniques (windowing, multitaper method) were tried. The results reported in the following (being qualitative in nature) were found to be insensitive to the choice of spectral estimation method.

[14] The second approach to stochastic differential equations forgoes the individual paths of the Langevin approach in favor of a partial differential equation (the so-called Fokker-Planck equation) governing the probability density. Under certain mathematical conditions [Ottinger, 1996; Gardiner, 1990] the two approaches are equivalent. However, in many instances, the former approach is preferable from a computational viewpoint (see the discussion by Ottinger [1996]). The second approach is well described in the monograph by Gardiner [1990], and has been utilized to study a version of the Stommel box model by Timmerman and Lohmann [2000], in an article that we shall discuss below.

### 3. Stommel Two-Box Model

[15] We follow C94 and we can write the Stommel two-box model in dimensionless form as

$$\begin{aligned} \frac{dx}{dt} &= -\alpha(x-1) - x[1 + \mu^2(x-y)^2] \\ \frac{dy}{dt} &= p(t) - y[1 + \mu^2(x-y)^2], \end{aligned} \quad (2)$$

where  $x$  and  $y$  are temperature and salinity differences between the low-latitude and high-latitude boxes which are assumed to make up the North Atlantic,  $\alpha$  and  $\mu$  are physical parameters (the former sets the strength of temperature restoring, while the latter sets the strength of the exchange flow between the two boxes) and  $p(t)$  is a known forcing function. We adopt the parameter values of C94, namely  $\mu = 2.49$  (or  $\mu^2 = 6.2$ ) and  $\alpha = 3600$ . The large value adopted for  $\alpha$  allows one to essentially drop  $x$  from consideration as  $x = 1 + O(\alpha^{-1})$ . In the literature [i.e., Velez-Belchi *et al.*, 2001], there has been some discussion of this point, with the alternative value  $\alpha = 360$  being suggested. We have tested a range of  $\alpha$  values, including those mentioned above, and have found that the two equation model can be quite safely reduced to the one equation model for all discussions herein. As discussed in C94, the single equation model can be interpreted as governing the inertia-less motion of a particle in a double-well potential. Following C94 we write

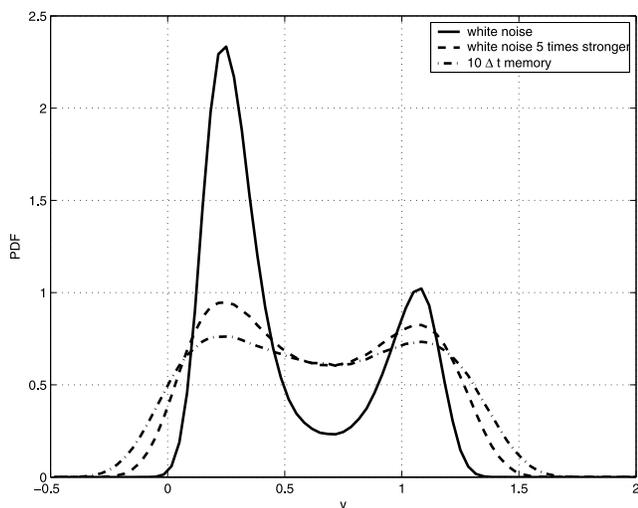
$$\frac{dy}{dt} = -y[1 + \mu^2(1-y)^2] + \bar{p} + p'(t) \quad (3)$$

and take  $\bar{p} = 1.1$ . The time-dependent perturbation  $p'(t)$  can consist of both deterministic and stochastic components.

[16] One of the most interesting findings reported in C94 is the response of the system to piecewise constant perturbations. The numerical experiments proceed as follows: start at the stable equilibrium point that corresponds to the global minimum of the potential (which we label as  $y_a \approx 0.24$ ), impulsively increase  $p'(t)$  to a fixed value  $\Delta$ , hold  $p'$  at this value for a time  $\tau$  and then impulsively decrease back to 0. The idea is to find the minimum value of  $\Delta$  that will lead to a transition to the shallower potential well (with local minimum at  $y_c \approx 1.1$ ) for a given value of  $\tau$ . Cessi reports (and we have confirmed) that as  $\tau$  increases from zero  $\Delta$  falls rapidly at first, but then asymptotes to a minimum value  $\Delta_0 \approx 0.2$  below which no transition occurs regardless of how large  $\tau$  is.

[17] The implications of this result are manifold. One particular example is given in Figure 3. In Figure 3 we show the response to perturbations of the form  $p'(t) = a \sin(\frac{2\pi t}{T})$ . We fix  $a$  and vary  $T$ . Note immediately that only the longest period perturbation shown leads to a transition between the two potential wells. A secondary point to note is the fact that when transitions between the potential wells do occur, they are phase locked to the sinusoidal perturbations (though the  $y$  versus  $t$  curve is not sinusoidal).

[18] Having established the above, we henceforth consider  $p'(t)$  to be stochastic. Following C94 we discretize the governing equation by the Millstein method (note that for the additive noise cases we discuss this is equivalent to the Euler method, but maintains the  $O(\Delta t)$  convergence [Ottinger, 1996]) and generate a Gaussian perturbation with variance  $\sigma^2$  at each time step (in other words white noise). In Figure 4 we show the probability distribution of  $y$ , computed by binning the time series. The two minima of the unperturbed potential occur near 0.25 and 1.1. The agreement of the solid curve with Figure 6 from C94 is excellent. The time series of  $y$  corresponding to Figure 4 consists of irregular oscillations around the two local minima with occasional transitions between the two poten-



**Figure 4.** Probability density function for  $y$  in the Stommel two-box model subjected to stochastic perturbations: solid line, white, Gaussian noise perturbations  $\sigma = 0.2\bar{\sigma}$ ; dashed line, white, Gaussian noise perturbations  $\sigma = 1.0\bar{\sigma}$ ; dash-dotted line, perfect memory for  $t = 0.001$  (or noise held fixed for 10 time steps) with no subsequent correlation, Gaussian perturbations  $\sigma = 0.2\bar{\sigma}$ .

tial wells. For the solid curve the standard deviation,  $\sigma = 0.2\bar{\sigma}$  while for the dashed curve  $\sigma = 1.0\bar{\sigma}$ . It is clear from the figure that the stronger noise case smears the probability distribution, in the sense that the distinction between the two potential wells is decreased and the region of nonzero probability density is extended.

[19] A natural extension of the results in C94 can be summarized in the following question: What happens if instead of generating a random perturbation at every time step we do so at only every 10th or 100th time step, holding the perturbation constant between switches. The idea is perhaps best illustrated with an example. Consider a series of 30 random integers between 0 and 9, 2383912078 4938576104 7492741006, with a space included after every ten digits. For the random perturbations at every 10th time step we would have the following sequence of numbers, corresponding to the above, 2222222222 4444444444 7777777777, again with a space included after every ten digits. In other words the noise used is no longer white in time, but has perfect memory for 10 (or 100) time steps, and then loses all memory. Indeed from a physical point of view it is not likely reasonable to expect freshwater forcing of the NA to be a white noise process and the above can be interpreted as a particular form of fading memory (such as that used in linear viscoelasticity [Pipkin, 1986], for example). In Figure 4 we show results corresponding to  $\sigma = 0.2\bar{\sigma}$  but with the stochastic perturbation held fixed for 10 time steps (corresponding to a perfect memory of duration 0.001). As an aside, we note that the same random vector is used to produce all curves in Figure 4. Note that this means that for the dash-dotted curve we use only every 10th entry of the random vector. The reader can readily confirm that increasing the temporal memory of the noise, while keeping  $\sigma$  fixed leads to a greater smearing of the probability distribution than keeping the temporal characteristics

white and increasing  $\sigma$  fivefold. The converse of this result is that with a nonwhite noise forcing applied to the system it is possible that a probability distribution like that of the solid curve in Figure 4 is obtained for much smaller values of the variance. The present results imply that when specifying a low-order model of the North Atlantic circulation that includes a stochastic component, a discussion of the memory of a stochastic perturbation is just as important as a discussion of the values of the physical parameters (the exchange functions between the two boxes, in the present context).

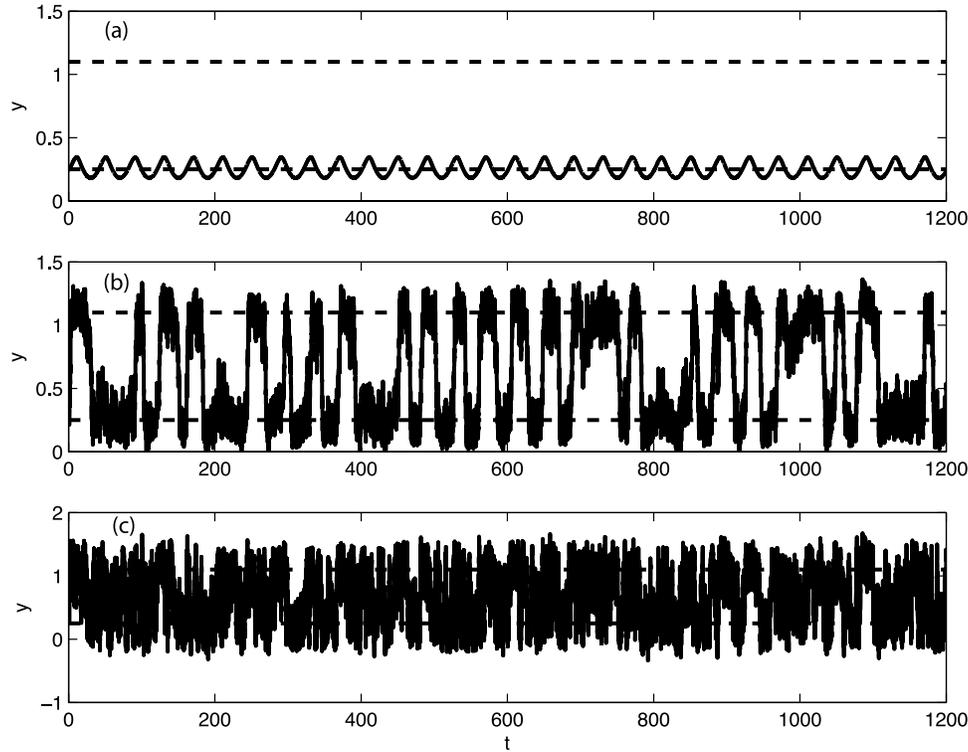
[20] The mathematically inclined reader will note that the type of memory outlined above renders the discretized differential equation non-Markovian. One important consequence of this fact is that there is no correspondence between the Langevin-type equation we integrate numerically and a Fokker-Planck equation for the probability density function. We have reproduced all relevant simulations using a Markovian red noise generated by an extension of the Box-Müller algorithm [Bartosch, 2001]. In the Markovian red noise case it is possible to reach similar conclusions to those discussed below using the appropriate Fokker-Planck equation. We have found no qualitative difference between the red noise results and the results employing the intuitive form of memory presented in the above and hence we will report results based on the latter only.

[21] For the Markovian red noise case the problem can be considered from the Fokker-Planck equation point of view. Indeed this has been done for a Stommel-type, two box model with multiplicative noise by Timmerman and Lohmann [2000] and interpreted from the point of view of catastrophe theory. The agreement with the present study is quite good; for example, compare our Figure 4 with Figures 6 and 7 of Timmerman and Lohmann [2000]. The advantage of the Fokker-Planck approach lies in its ability to find closed form expressions (albeit approximate) for the probability distribution [Timmerman and Lohmann, 2000, equations (33) and (34)]. In contrast, the Langevin equation approach employed in the present has the advantage of a transparent, highly parallelizable numerical implementation that, crucially, yields so-called strong approximations of the actual solution paths and applies both for cases with nonMarkovian noise and situations which combine deterministic and stochastic perturbations.

[22] We next turn to a discussion of stochastic resonance (SR) in the system (3). SR has received a great deal of attention in the literature [Rahmstorf and Alley, 2002; Ganopolski and Rahmstorf, 2001] as a possible explanatory mechanism for NA THC variability on millennial (and longer) timescales. The basic idea of SR (discussed in much greater detail by Gammaitoni et al. [1998]) is that a nonlinear system can exhibit strong response at a forcing frequency when noise is present even if the response to the same periodic forcing without noise is weak. We take the forcing  $p'(t)$  to be given by

$$p'(t) = a \sin\left(\frac{2\pi t}{T}\right) + p'', \quad (4)$$

where  $p''$  is stochastic. In Figure 5 we show three realizations of  $y$  versus  $t$  with the same initial conditions,



**Figure 5.** The  $y$  versus  $t$  curves for Stommel two-box model subjected to both stochastic and sinusoidal perturbations. The two stable equilibria in the unperturbed case are indicated by dashed lines. Both the amplitude,  $a$ , and period,  $T$ , of the sinusoidal forcing are fixed, and the standard deviation of the noise,  $\sigma$ , increases from top to bottom. In Figure 5a, no transitions occur; in contrast, Figure 5b exhibits transitions between the two stable regimes with the strongest phase locking to the driving frequency, while Figure 5c is almost completely dominated by noise.

sinusoidal forcing amplitude and period  $(a, T) = (0.15\bar{p}, 40.0)$  and varying strength of stochastic forcing. The value of  $\sigma$  is chosen as:  $0.004\bar{p}$  for Figure 5a,  $0.4\bar{p}$  for Figure 5b and  $2.0\bar{p}$  for the Figure 5c. The noise is white in time. From figure 5 it is clear (even to the naked eye) that the Figure 5b experiences the largest response at the forcing frequency. When the noise is weaker the transitions from one potential well to the other are rare (indeed do not occur over the time period shown). Of course, in the limit of no noise transitions do not occur at all. For strong noise (comparable to  $\bar{p}$  in the case of Figure 5c) the evolution of the system is completely dominated by the noise. Similar results to Figure 5 can be produced for other choices of memory for the stochastic noise.

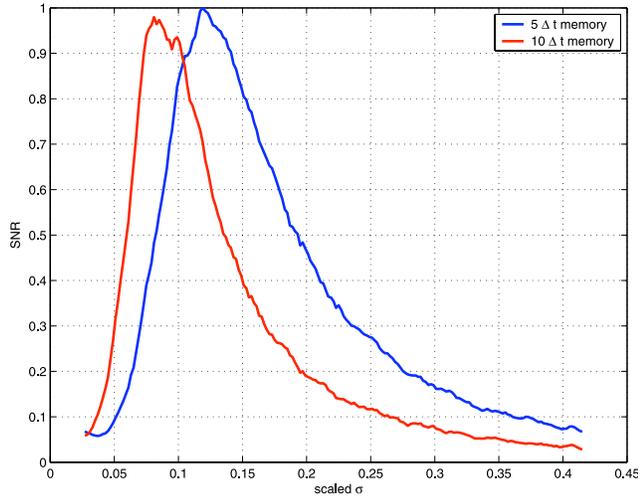
[23] To discuss SR quantitatively we must define a signal to noise ratio ( $SNR(\omega_0, a, \sigma)$ ) which measures the ratio between the strength of the output signal at the forcing frequency and the strength of the noise. There are a variety of ways to define signal to noise ratio [Gingl *et al.*, 2001]. In the following we employ

$$SNR = \frac{PSD(\omega_0)}{\sigma}, \quad (5)$$

where  $PSD(\omega)$  is the power spectral density at a given frequency,  $\omega$ , and  $\omega_0 = 2\pi/T$ . We have tried several other definitions of the signal to noise ratio and found no

qualitative change in the results presented in what follows. Figure 6 shows two signal to noise ratio versus  $\sigma/\bar{p}$  curves for  $(a, T) = (0.15\bar{p}, 40)$  and noise held 5 or 10 time steps, computed from a time series of  $10^7$  time steps. While both signal to noise ratio curves exhibit the sharp peak characteristics of SR, the location of this peak shifts from approximately  $\sigma = 0.07\bar{p}$  to  $\sigma = 0.12\bar{p}$ . This implies that, for a sinusoidal perturbation of fixed amplitude and period and stochastic perturbation with a fixed magnitude (or in other words, variance), the temporal characteristics of the stochastic perturbation exert a great deal of influence on whether stochastic resonance can take place in a physically meaningful parameter range. In particular, the usual interpretation of stochastic resonance in the literature asserts that even weak noise can lead to phase locking with the driving sinusoidal perturbation. Of course even if a reasonable model of the temporal characteristics of the noise was available a physical mechanism for the sinusoidal perturbation would need to be outlined.

[24] Finally, we note that SR in the context of the Stommel two-box model has been discussed by *Velez-Belchi et al.* [2001], however the numerical method employed is not clearly described in this paper. As most of the literature on stochastic differential equations makes the assumption of a stochastic perturbation that is white in time and Gaussian, we assume *Velez-Belchi et al.* [2001] make this choice. For such a choice, the above results suggest that the amplitude



**Figure 6.** Signal to noise ratio versus  $\sigma/\bar{p}$  curves for two types of memory. The peak that is characteristic of SR is clearly visible in both cases, however the location of the peak is a strong function of the memory of the noise, with longer memory leading to a peak at lower values of  $\sigma$ .

of the stochastic perturbation would need to be unphysically large.

#### 4. Generalized Sakai and Peltier Model

[25] The Sakai and Peltier model is a three-box model, with  $x$ ,  $y$  and  $z$  representing the nondimensional salinity anomalies, or in other words the differences of salinity between a large background box (or ‘bath’ in the language of thermal physics) and midlatitude surface, high-latitude and midlatitude abyssal boxes, respectively. The model involves five parameters:  $X$  and  $Y$  represent the midlatitude evaporation and high-latitude freshening, respectively,  $\mu$  is a diffusion parameter,  $\beta$  sets the degree to which the THC transport is nonlinear and  $m$  can be thought of as a parameter for mixing in the high-latitude box during deep convection. In nondimensional form the model reads

$$\begin{aligned} \frac{dx}{dt} &= X - x + \mu y - H(y)y^\beta x \\ \frac{dy}{dt} &= Y - 2y + \mu(x+z) + H(y)y^\beta[(x-y) + m(y-z)] \\ \frac{dz}{dt} &= -z + \mu y + H(y)y^\beta(y-z), \end{aligned} \quad (6)$$

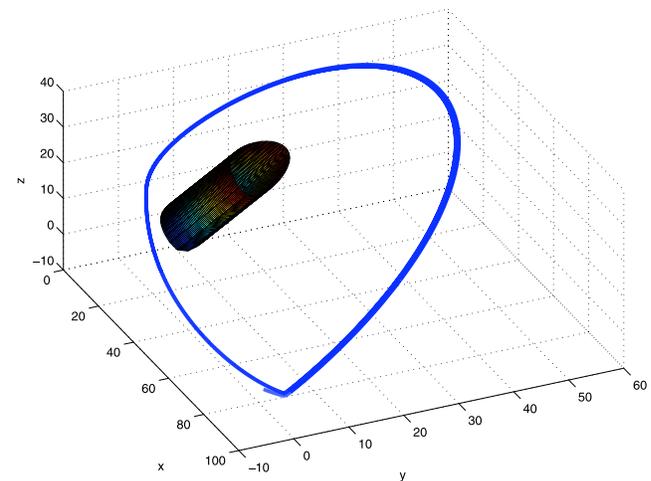
where  $H(p)$  is the Heaviside step function. The presence of  $H(y)$  in (6) implies the transport terms, due to the THC, turn off if  $y < 0$ . The strength of the transport is set by  $y^\beta$ , a function of the salinity anomaly in box 2, and varying  $\beta$  amounts to making a different choice for the exchange function (various choices are given by equation (2.4) of C94). There is no transport into box 1, but a transport of the salinity anomaly  $x$  out of box 1 and into box 2. There is transport into and out of boxes 2 and 3. Setting  $m = 1$  reduces the system (6) to the model extensively analyzed by Sakai and Peltier [1999] and referred to henceforth as the S and P model. In this case the water transported out of the

high-latitude sinking box is assumed to possess a salinity equal to that found in the abyssal box (presumably owing to mixing during deep convection events). The opposite limit  $m = 0$  corresponds to the case of no mixing during the deep convection events (the salinity of water entering the high-latitude box via THC transport equals that leaving it), and will be referred to in the following as the transport model. Note that the 2 in the  $-2y$  term in the evolution equation for  $y$  is carried over from Sakai and Peltier [1999], where it was used to simplify certain algebraic manipulations.

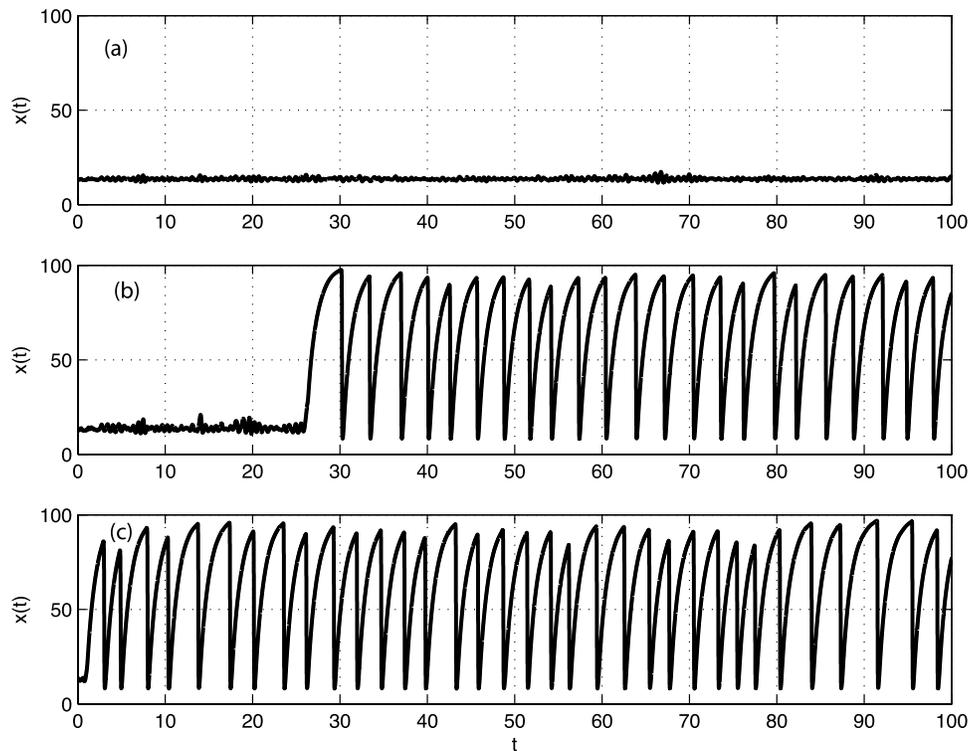
[26] As was suggested by an anonymous reviewer, one can view the differences between the S and P and transport models, as simply making different choices of transport functions with the aim of getting different types of ‘typical’ models. In this interpretation the S and P model is an example of a model in which a stable equilibrium point coexists with a stable periodic orbit, while the transport model is an example of a model in which only a single, globally stable equilibrium point exists, but the THC can be in either an ‘ON’ or ‘OFF’ state. While we prefer the interpretation of ‘transport’, we leave the choice of preferred interpretation to the reader.

#### 4.1. S and P Model Results

[27] In the following we illustrate the effect of stochastic perturbations by considering a subset of the cases studied by Sakai and Peltier [1999] and hence fix  $m = 1$ . We set  $(X, Y) = (-100, 45)$  and  $(\mu, \beta) = (0.5, 1.0)$  in order to elaborate on Figure 6 of Sakai and Peltier [1999]. For this set of parameters it has been shown [Sakai and Peltier, 1999] that the system has a single, stable equilibrium point (located at approximately  $(13.5, 6.5, 6.2)$ ) and a stable limit cycle. The stable periodic orbit (limit cycle) represents collapse-flush cycles in which the salinity at low latitudes builds up with the THC stopped, followed by a period of high THC transport. In Figure 7 we show both the domain of attraction of the stable equilibrium point (the shaded, approximately cylindrical region) and one orbit that tends to the periodic orbit (in blue). The domain of attraction of the equilibrium point is constructed by repeated integration with different



**Figure 7.** Stable periodic orbit (limit cycle), shown in blue, and the basin of attraction of the single, stable equilibrium point (shaded) for the S and P model.



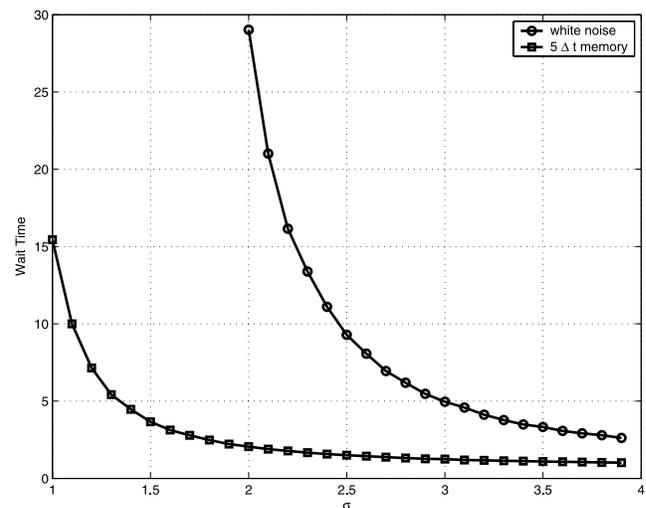
**Figure 8.** Noise induced destabilization of the stable equilibrium point in the S and P model as illustrated by  $x$  versus  $t$  curves: (a)  $\sigma = 1.0$  white noise, (b)  $\sigma = 2.0$  white noise, and (c)  $\sigma = 1.5$  noise with perfect memory for  $t = 0.0005$ , or in other words noise held fixed for five time steps.

initial conditions. It can be seen that the domain of attraction of the equilibrium point makes up a small region of the physically reasonable portion of phase space.

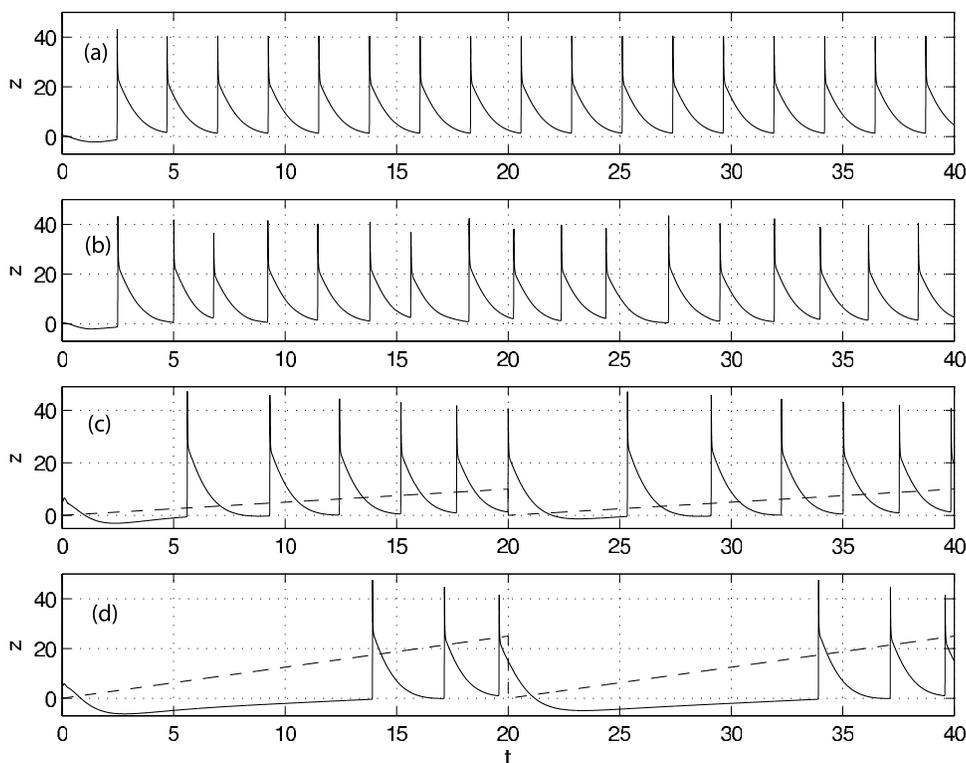
[28] We have subjected (6) to stochastic perturbations in the high-latitude fresh water forcing  $Y$  by employing the same methodology used above for the one equation model (3). In Figure 8 we show time series of  $x$  versus  $t$  with the initial conditions chosen to lie at the equilibrium point. In Figures 8a and 8b we use a white noise in time, with  $\sigma = 1.0$  in Figure 8a and 2.0 in Figure 8b. It can be seen that only the stronger noise forces the orbit away from the equilibrium point and toward the stable periodic orbit in the time interval shown. The transition occurs near  $t = 26$ . In Figure 8c we show a noise with a short memory (held for five time steps) with  $\sigma = 1.5$ . It is evident that the orbit is forced away from the equilibrium point and toward the periodic orbit much more quickly.

[29] Of course with a Gaussian noise of any nonzero variance all orbits will be pushed out of the domain of attraction of the stable equilibrium point provided one waits long enough (regardless of the type of memory the perturbations exhibit). This observation implies that a better characterization of phase space behavior as a function variance can be achieved by considering the waiting time to leave the domain of attraction of the stable equilibrium point for trajectories starting at the stable equilibrium point. In Figure 9 we show waiting times computed from an ensemble of 10000 trajectories. We show both white noise and noise with memory. The shape of both waiting time curves shown is qualitatively similar, with long waiting times for small  $\sigma$  decreasing rapidly with increasing  $\sigma$  and

eventually reaching a plateau near zero. However, it can again be seen that even a fairly weak memory (held for five time steps in this case) has a profound effect on the realizability of the destabilization for a given set of parameters. In particular, say it was known a priori that  $\sigma < 2$ , then



**Figure 9.** Waiting time versus  $\sigma$  curves for the destabilization of an orbit starting at the stable equilibrium point of the S and P model. The noise with memory (perfect memory for  $t = 0.0005$  in this case) shifts the curve considerably toward smaller values of  $\sigma$ .



**Figure 10.** Attempts to model the Bond cycle using the S and P model,  $z$  versus  $t$ : (a) unperturbed, (b) Gaussian, white noise perturbations, (c) deterministic saw tooth perturbation, and (d) large deterministic saw tooth perturbation.

the present results indicate that a white noise model would not predict destabilization on timescales less than  $t = 30$ .

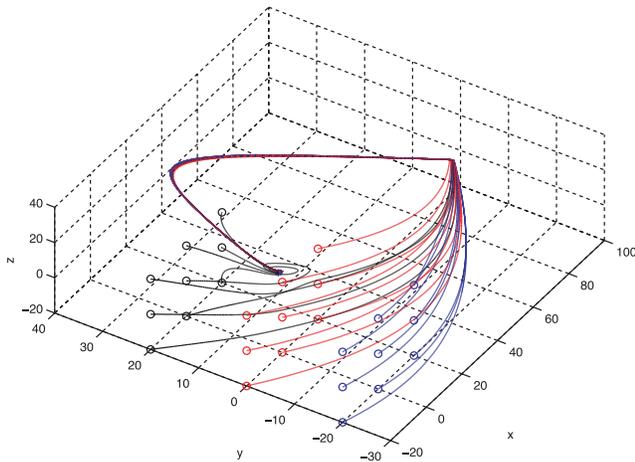
[30] A natural question to ask is whether noise can push an orbit out of the domain of attraction of the periodic orbit and to the fixed point. In all the simulations we have run this has not been found to occur, perhaps because of the large amplitude variation in  $x$  along the periodic orbit evident in Figures 7 and 8 and the small extent of the domain of attraction of the equilibrium point as seen in Figure 7.

[31] In work by *Sakai and Peltier* [1999] the S and P model was found to exhibit strong modulation of the periodic behavior when subjected to both sinusoidal and saw-tooth perturbations. In Figure 10 we reproduce the results reported by *Sakai and Peltier* [1999] for sawtooth perturbations and contrast them with modulation due to purely stochastic perturbations. In order to facilitate comparison with work by *Sakai and Peltier* [1999] we set  $(X, Y, \mu, \beta) = (100, -41, 0.5, 1.4)$ . Figure 10a shows an unperturbed  $z$  versus  $t$  time series. The trips around the periodic orbit (limit cycle) are clearly visible. Figure 10b shows a stochastically perturbed  $z$  versus  $t$  time series ( $\sigma = 1.5$  noise held 10 time steps). Small variations of the amplitude and period of the oscillations are visible. However, these variations are insignificant when compared with those in Figures 10c and 10d, both of which are induced by sawtooth shaped perturbations. In Figures 10c and 10d the saw-tooth perturbation is shown by a dashed line. The perturbations are not scaled in amplitude, but are shifted to fit onto the graph. The physical perturbation consists of a sudden large increase in freshwater flux ( $Y$  becomes more negative) followed by a slow relax-

ation to the undisturbed value ( $Y = -41$ ). The sudden increase in freshwater flux is intended to mimic the impact of the individual Heinrich event that causes the dramatic shutdown of the NA THC that characterizes each Bond cycle. As the amplitude of the sawtooth perturbation increases a modulation like the Bond cycle is clearly visible. Note also that Figures 10c and 10d indicate that no oscillations exist for very negative  $Y$ . These results suggest that it is the deterministic, as opposed to stochastic, perturbations that determine whether a modulation occurs. We note that the EMIC-based study of *Timmerman et al.* [2003], which finds a stochastically induced oscillation following a modeled Heinrich event in an ocean–atmosphere–sea ice model, does not yield a clear Bond cycle [see *Timmerman et al.*, 2003, Figures 5 and 7].

#### 4.2. Transport Model Results

[32] Recent experiments with the NCAR CCSM in its fully coupled configuration performed as part of the Paleoclimate Modeling Intercomparison Project (PMIP2) suggest that even for extremely strong anomalous freshwater fluxes into the North Atlantic region (1 Sv for 100 years) the climate system recovers, after several centuries, to a state that is very nearly its unperturbed state [*Peltier et al.*, 2006]. While results for modern climates should not be cavalierly extrapolated to the glacial climate, at the very least the modern results suggest that bistable box models may not be the most appropriate prism through which to view the millennial-scale oscillations in the North Atlantic climate. As an alternative, we consider the transport model ( $m = 0$ ) with the same set of parameters as the S and P



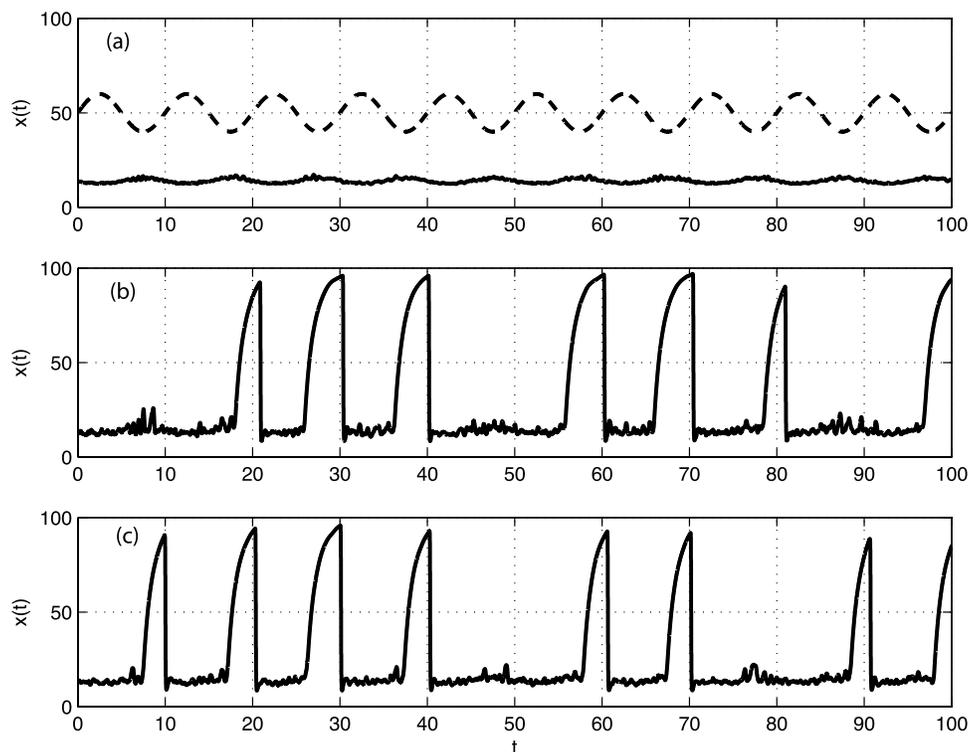
**Figure 11.** Twenty-seven sample orbits of the transport model showing the tendency to a single, globally stable equilibrium point. Colors are used only for presentation purposes.

model ( $m = 1$ ), namely  $(X, Y) = (-100, 45)$  and  $(\mu, \beta) = (0.5, 1.0)$ . It is a simple exercise in algebra to show that with this choice of parameters the system (6) has a single equilibrium point (found at approximately  $(13.88, 6.435, 6.0)$ ) that is globally stable. In Figure 11 we show the trajectories for 27 different initial conditions. The trajectories are color-coded according to the value of  $y(0)$  and the initial points in phase space are denoted by open circles. It can be noted that all the trajectories which pass through the  $y < 0$  region of

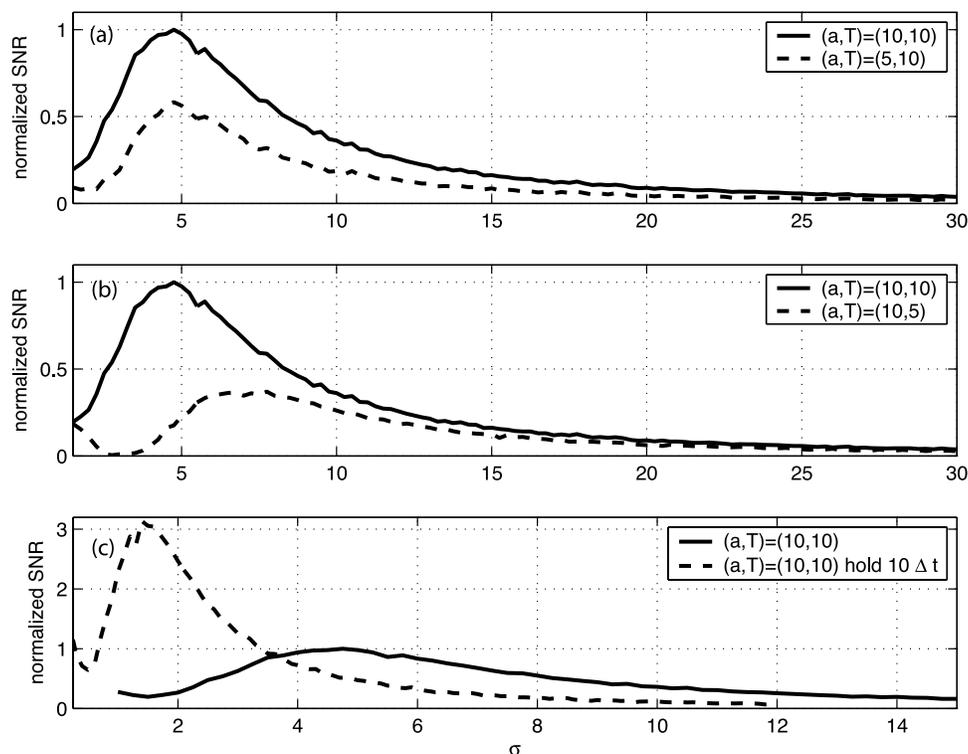
phase space tend to the stable equilibrium point along virtually the same path.

[33] We next consider perturbations to the high-latitude freshening parameter ( $Y$ ) of the form (4). In Figure 12 we show the evolution of  $x$  versus time. The sinusoidal portion of the perturbation is fixed to have  $(a, T) = (10, 10)$ . The top two panels have stochastic perturbations that are white in time. For Figure 12a, in which  $\sigma = 1$ , the system executes small amplitude oscillations around the equilibrium point. Figure 12a also shows a schematic of the sinusoidal forcing. In Figure 12b, for which  $\sigma = 4$  the increased noise leads to episodic large-amplitude excursions from the equilibrium point that appear quite similar to trips around the periodic orbit (limit cycle) in the S and P model (Figure 8). In Figure 12c we consider noise with  $\sigma = 1$ , but a longer memory (10 time steps). It is evident that the increase in memory allows even weak noise to induce phase locking with the driving signal.

[34] The apparent phase locking of the large excursions to the driving signal suggests that the transport model, a system with a single, globally stable equilibrium point can exhibit stochastic resonance. This rather surprising result is confirmed in Figure 13 where we plot the normalized signal to noise ratio versus  $\sigma$  curves for several combinations of  $a$  and  $T$ . In all cases the signal to noise ratio value is normalized by the maximum value for the  $(a, T) = (10, 10)$  case, and the reader is cautioned to note the change in horizontal scale between Figures 13a, 13b and 13c. It can be seen that for all cases shown the signal to noise ratio versus  $\sigma$  curve has a prominent peak (at  $\sigma^*$  say). A halving of the amplitude with the period fixed (shown as the dashed line in Figure 13a)



**Figure 12.** Phase locking onto a sinusoidal perturbation due to stochastic perturbations in the transport model,  $x$  versus  $t$  shown as solid, black lines, sinusoidal perturbation with  $(a, T) = (10, 10)$  schematized by dashed, black line: (a)  $\sigma = 1$  white noise, (b)  $\sigma = 4$  white noise, and (c)  $\sigma = 1$  noise held 10 time steps.

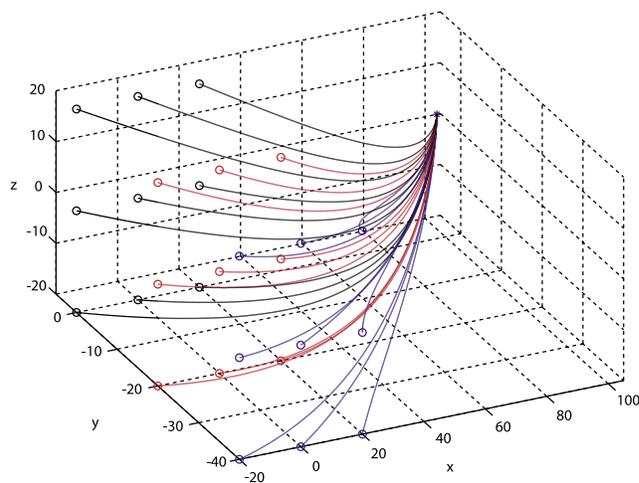


**Figure 13.** Signal to noise ratio versus  $\sigma$  curves showing the existence of SR in the transport model with a single, globally stable equilibrium point: (a) effect of changes in amplitude of sinusoidal forcing, (b) effect of changes in period of sinusoidal forcing, and (c) effect of increasing memory of the noise; note the change in horizontal scale for Figure 13c. The solid line represents the same case for all three plots.

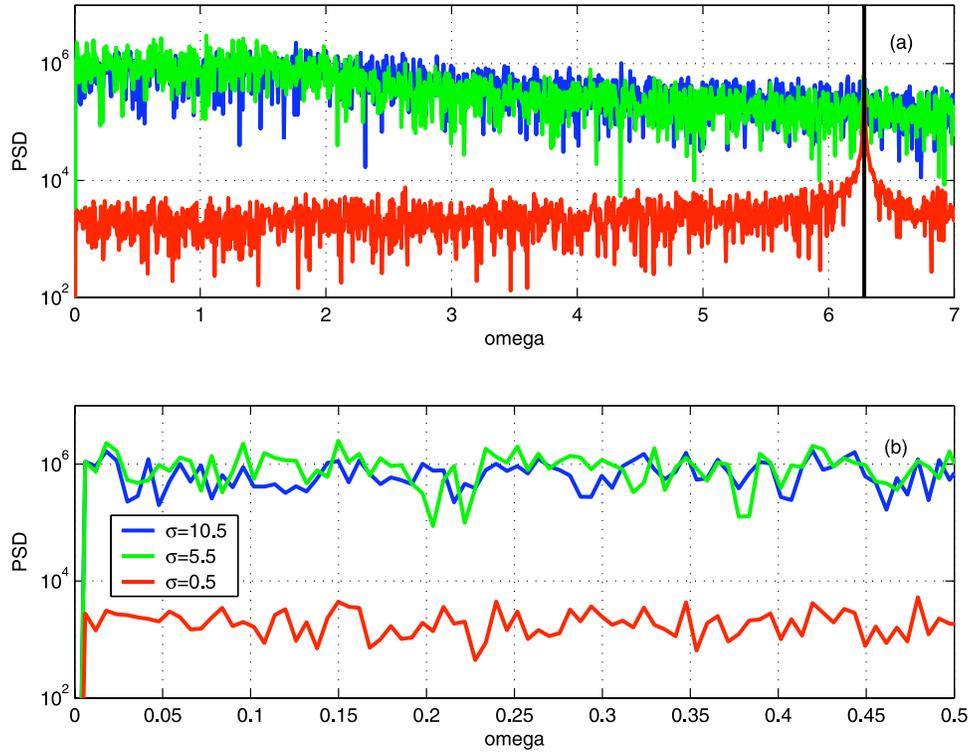
leads to a very slight increase in  $\sigma^*$  and a nearly 50% decrease in the height of the peak, while a halving of the period with the amplitude fixed (shown as a dashed line in Figure 13b) leads to both a substantial increase in  $\sigma^*$  and a decrease in the maximum signal to noise ratio value. Not unexpectedly, given the results shown in Figure 12, adding even a weak memory to the noise leads to a reduction of  $\sigma^*$  by more than 50% and a threefold increase in the maximum signal to noise ratio value (shown in Figure 13c). This is a striking example of the importance of having a reliable model of the noise memory when arguing for the physical applicability of a mechanism such as stochastic resonance.

[35] As SR is usually presented in terms of dynamics in bistable potentials (as in our section on the Stommel two-box model) it is important to identify the mechanism of SR in the present model. Toward this end, Figure 14 shows the evolution of 27 trajectories (color-coded as in Figure 11) for the system (6) with the nonlinear transport terms due to the THC turned off. In the integrations shown in figure 11 this corresponds to the governing equations when  $y < 0$ . From Figure 14 we can see (and indeed can confirm by a trivial analytical calculation) that the trajectories tend to a stable equilibrium point. Careful inspection of Figure 14, or a simple calculation of the equilibrium points of (6) with  $m = 0$  and the nonlinear terms switched off, indicates that this equilibrium point lies in the  $y > 0$  portion of phase space. This implies, for the full transport model, that any trajectories entering the  $y < 0$  portion of phase space tend toward a

‘ghost’ equilibrium point found in the  $y > 0$  portion of phase space. This ‘ghost’ equilibrium point corresponds to an ocean circulation that is purely diffusive in nature, and for which high salinities dominate the low-latitude box and



**Figure 14.** Twenty-seven sample orbits of the transport model with the THC turned off (i.e., purely diffusive model). A single, globally stable equilibrium point exists, however it occurs at a location with  $y > 0$ ; hence this equilibrium point is not an equilibrium point of the full transport model. Colors are used only for presentation purposes.



**Figure 15.** (a) Power spectral density for three values of  $\sigma$  (white noise); the sinusoidal perturbation has  $T = 1$ . It is clear that the largest noise induced response occurs at frequencies well below the frequency of the sinusoidal perturbation. (b) Detail of low frequencies.

both the sinking and deep boxes have values that are slightly larger than the background salinity (in other words slightly larger than zero). That the trajectories can never reach this “ghost” fixed point is thus clear, since the nonlinear terms immediately turn on when  $y > 0$ . This potential well without a equilibrium point thus explains both the clumping of the trajectories in Figure 11 and the SR shown in Figures 12 and 13.

[36] While we believe this to be the first demonstration of stochastic resonance in a simple climate model with a single, globally stable equilibrium point, we do note that *Timmerman et al.* [2003] used a considerably more complex model (involving ocean, sea ice and atmosphere components) to study a form of stochastic resonance for excitable systems. A detailed comparison between the present model and that employed by *Timmerman et al.* [2003] provides another clear avenue for future work.

[37] For the mathematically minded reader we note that the jump in the derivative found in (6) is not necessary to find SR. Indeed the above results can be reproduced with arbitrary accuracy by replacing  $H(y)$  with a smooth function such as

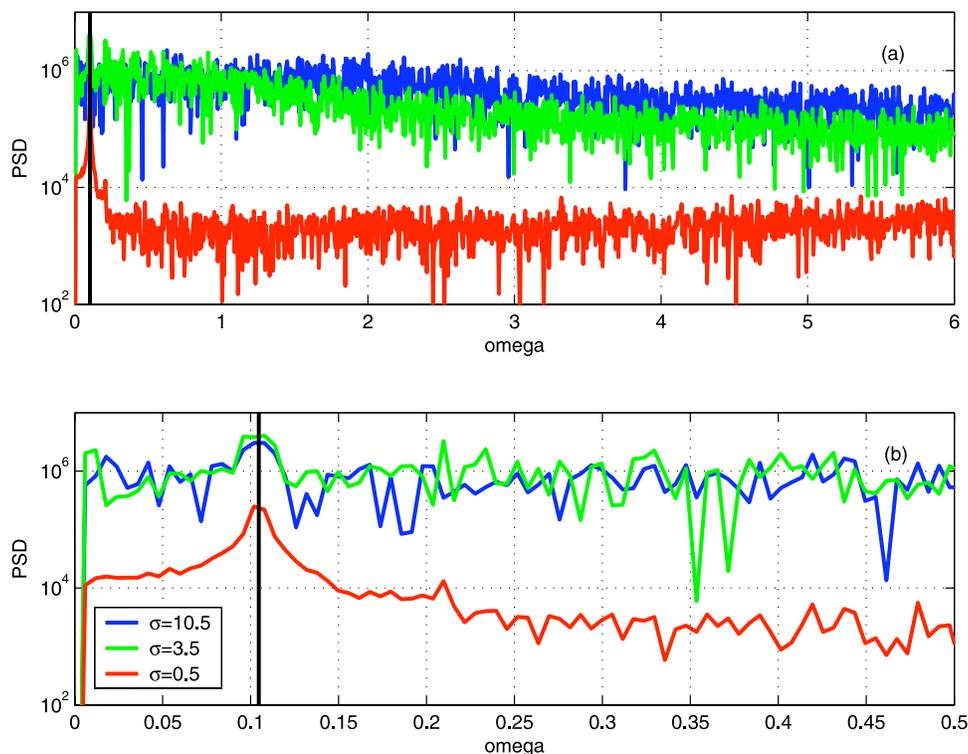
$$\frac{1 + \tanh\left(\frac{y}{d}\right)}{2},$$

where  $d$  specifies how rapid the transition is, and as such is chosen to be very small. *Stastna and Pogson* [2007] studied the mathematical properties of the S and P model from the point of view of dynamical systems with switching.

[38] The ‘ghost equilibrium point’ mechanism discussed above suggests that the individual collapse-flush cycles have an inherent timescale associated with them, much as individual trips around the periodic orbit (limit cycle) had an inherent timescale associated with them in the previous subsection. If this is true, however, then the definition of SR in terms of the power at the forcing frequency may not be entirely appropriate when there is a mismatch between the timescale inherent to the sinusoidal part of the forcing and the timescale inherent to the flush-collapse cycle. We demonstrate this by computing the full power spectral density (PSD) for  $T = 1$  (Figure 15) and  $T = 60$  (Figure 16). In each case we choose three values of  $\sigma$ , one which is well below the peak in the signal to noise ratio curve, one near the peak, and one well beyond it.

[39] In Figure 15 we show a semilog plot of the PSD for the case  $T = 1$  versus frequency. Figure 15a shows a segment of the frequency range that includes the driving frequency (indicated by a thick, vertical black line). It is clear the combination of stochastic and sinusoidal perturbations substantially increases the power in the low frequencies. However it is not true that the largest increase occurs at the driving frequency. The prominent increase in the power at low frequencies can be seen in Figure 15b, which considers only the lowest frequencies.

[40] In Figure 16 we show a semilog plot for the opposite case ( $T = 60$ ) for which the sinusoidal perturbations have a period that is much longer than individual flush-collapse events. As in Figure 15 a significant increase in spectral power at the low frequencies is evident. In Figure 16b only



**Figure 16.** (a) Power spectral density for three values of  $\sigma$  (white noise); the sinusoidal perturbation has  $T = 60$ . It is clear that the largest noise induced response occurs in a broad band of low frequencies, and not just at the frequency of the sinusoidal perturbation. (b) Detail of low frequencies confirming the broad peak of spectral power at low frequencies.

the lowest frequencies are shown, with the driving frequency indicated by a vertical, black line. In this case there is a clearly visible increase in PSD at the driving frequency, however this peak is not localized and similar values of the PSD occur for a range of higher frequencies.

[41] Finally, we wish to draw attention to the fact that despite their rather similar appearance (compare Figures 8 and 12) the oscillations in the S and P model are intrinsic and can thus occur with no noise (provided the initial condition does not lie in the domain of attraction of the stable equilibrium point), while the oscillations in the transport model are extrinsic (because the model has a single, globally stable equilibrium point) and as such crucially dependent on finite amplitude perturbations (be they deterministic or stochastic). It is only when the transport model is subjected to perturbations large enough to shut down the thermohaline circulation that its inherent timescale comes into play (as the salinity slowly builds up in the high-latitude box).

## 5. Conclusions

[42] At the most elementary level, the results discussed above indicate that millennial timescale variability can be reproduced in a variety of extremely simple models. However, this assertion must be accompanied by the caveat, that the temporal structure of the stochastic terms in the model should be determined as part of the modeling exercise, as even small changes in the memory of the noise can greatly alter the region of parameter space in which a phenomenon

of interest (e.g., stochastic resonance, destabilization to a periodic orbit) occurs. This caveat suggests that one should be particularly careful in extrapolating the results of simple models to more complex models. In this sense, it is encouraging that the intrinsic variability scenario, whether associated with an outright Hopf bifurcation due to changes in the high-latitude freshening parameter  $Y$  (as in work by *Sakai and Peltier* [1999]) or a noise-induced destabilization of an orbit from the neighborhood of a stable equilibrium point toward a stable periodic orbit that can occur before any Hopf bifurcation of the stable equilibrium point to a periodic orbit takes place (as in the above discussion), agrees well with numerical experiments performed on models of intermediate complexity (those discussed by *Sakai and Peltier* [1997] and references therein).

[43] A similar link of the stochastic resonance hypothesis to results produced by more complex models has been asserted by *Rahmstorf and Alley* [2002] on the basis of work employing a climate model of intermediate complexity [*Ganopolski and Rahmstorf*, 2001] of a rather different form from *Sakai and Peltier* [1997]. As our results have shown, some form of stochastically enhanced response is not unexpected in complex nonlinear models. However, it is our opinion that the appropriate formalism for ‘stochastic resonance’ in complex models remains to be worked out. This is due to the fact that complex models naturally exhibit complex phase space behavior with inherent timescales that may not match the timescales of the sinusoidal perturbation assumed to exist in the SR scenario. The mismatch of timescales leads to changes in the structure of the power

spectrum that are more complex than a simple increase in power at the driving frequency. Indeed, our results indicate that the enhanced response at lower frequencies (i.e., a red noise-like power spectrum) is a much more robust sign of stochastically enhanced response. It is worth noting that stochastic resonance has been discussed from a rather different point of view in the digital acoustics literature [Wannamaker *et al.*, 2000]. Our above results suggest that the issue of enhanced response at low frequencies should be reexamined, since in the case of the transport model SR does indeed depend on an approximate matching between the period of the driving frequency and the inherent time-scales in the model, as is the case for the driven sinusoidal oscillator.

[44] A second point to consider when assessing the efficacy of the SR mechanism is how sensitive it is to a departure from a purely sinusoidal forcing at a single frequency. In order to clarify this issue we have carried out simulations in which the deterministic perturbation has a slowly varying (as opposed to a constant) period. We have found that even for the bistable model, the standard definition of stochastic resonance in terms of a signal-to-noise ratio, such as (5), fails completely for variations in the period as small as 4%. This is a rather serious criticism, as it is unlikely that any realistic forcing is composed of a single sinusoid. Of course even if we do consider a forcing consisting of only a single sinusoid, we are left with the Herculean task of finding a physically meaningful forcing with the correct period. We will not attempt to address this matter here.

[45] This criticism does not in any way diminish the importance of the idea of stochastic resonance, or more generally, stochastically enhanced response. To put it quite simply, the issue of noise as a model for unresolved scales and variables in a reduced model cannot be ignored (a point made earlier by Timmerman and Lohmann [2000]). As such, the work by Ganopolski and Rahmstorf, [2001] indicating that stochastic resonance between two NA THC modes that vary in the physical location of the deep water production is encouraging. An important step in understanding stochastically enhanced response and stochastic resonance would involve the construction of a simple partial differential equation based model that would allow a thorough exploration of issues relating to the spatial (as well as temporal) distribution of the forcing. Toward this direction the EMIC study of Timmerman *et al.* [2003] is a useful step and provides an interesting counterpoint to the far simpler models presented above.

[46] A final open question is whether a stochastic model can yield a clear Bond cycle-type modulation. In the above we were able to induce some modulations of the modeled D-O oscillations in the S and P model, but such modulations were almost entirely due to large, deterministic (as opposed to stochastic) freshwater perturbations (contrast Figures 8b and 8d).

[47] In closing, we wish to reiterate that stochastic perturbations that are not white in time yield significantly different results from those that are white in time. We demonstrated this by employing a rather pathological (though convenient from a programming point of view) form of fading memory, and confirmed our results using a Markovian red noise generated by a standard algorithm

[Bartosch, 2001]. An interesting possibility for future work would involve the development of a physically reasonable form of fading memory, perhaps based on experiments performed with more complex models.

[48] **Acknowledgments.** We gratefully acknowledge useful discussions with Guido Vettoretti and Francis Poulin. This work was supported by the Natural Sciences and Engineering Research Council of Canada. The comments of three anonymous referees significantly improved our manuscript.

## References

- Bartosch, L. (2001), Generation of colored noise, *J. Mod. Phys.*, *6*, 851–855.
- Bond, G., W. Broecker, S. Johnsen, J. McManus, L. Labeyrie, J. Jouzel, and G. Nonani (1993), Correlations between climate records from North Atlantic sediments and Greenland ice, *Nature*, *365*, 143–147.
- Brix, H., and R. Gerdes (2003), North Atlantic Deep Water and Antarctic Bottom Water: Their interaction and influence on the variability of the global ocean circulation, *J. Geophys. Res.*, *108*(C2), 3022, doi:10.1029/2002JC001335.
- Broecker, W. S. (1998), Paleocan circulation during the last deglaciation: A bipolar seesaw?, *Paleoceanography*, *13*, 119–121.
- Broecker, W. S. (2003), Does the trigger for abrupt climate change reside in the ocean or in the atmosphere?, *Science*, *300*, 1519–1522.
- Cessi, P. (1994), A simple box model of stochastically forced thermohaline flow, *J. Phys. Oceanogr.*, *24*, 1911–1920.
- Elsgolts, L. É. (1966), *Introduction to the Theory of Differential Equations With Deviating Arguments*, Holden-Day, San Francisco, Calif.
- Gammaitoni, L., P. Hanggi, P. Jung, and F. Marchesoni (1998), Stochastic resonance, *Rev. Mod. Phys.*, *70*, 223–287.
- Ganopolski, A., and S. Rahmstorf (2001), Abrupt glacial climate change due to stochastic resonance, *Phys. Rev. Lett.*, *88*, 038,501.
- Gardiner, C. W. (1990), *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*, 2nd ed., Springer, Berlin.
- Gingl, Z., P. Makra, and R. Vajtai (2001), High signal-to-noise ratio gain by stochastic resonance in a double well, *Fluctuation Noise Lett.*, *1*, L181–L188.
- GISP2 members (1993), The ‘flicking switch’ of late Pleistocene climate change, *Nature*, *361*, 432–436.
- GRIP members (1993), Climate instability during the last interglacial period recorded in the GRIP ice core, *Nature*, *364*, 203–207.
- Kloeden, P. E., and E. Platen (1992), *Numerical Solution of Stochastic Differential Equations*, Springer, Berlin.
- Kurtze, D. A., and J. M. Restrepo (2001), Advective time lags in box models, *J. Phys. Oceanogr.*, *31*, 1828–1842.
- Matsumoto, M., and T. Nishimura (1998), A 623-dimensionally equidistributed uniform pseudorandom number generator, *ACM Trans. Comput. Model. Simul.*, *8*, 3–30.
- Ottinger, H. C. (1996), *Stochastic Processes in Polymeric Fluids*, Springer, Berlin.
- Peltier, W. R., G. Vettoretti, and M. Stastna (2006), Atlantic meridional overturning and climate response to Arctic Ocean freshening, *Geophys. Res. Lett.*, *33*, L06713, doi:10.1029/2005GL025251.
- Pipkin, A. C. (1986), *Lectures on Viscoelasticity Theory*, Springer, New York.
- Rahmstorf, S. (1995), Bifurcations of the Atlantic thermohaline circulation in response to changes in the hydrological cycle, *Nature*, *378*, 145–149.
- Rahmstorf, S., and R. Alley (2002), Stochastic resonance in glacial climate, *Eos Trans. AGU*, *83*(12), 129.
- Root, C. (1982), Hydrology and ocean circulation, *Prog. Oceanogr.*, *11*, 131–149.
- Sakai, K., and W. R. Peltier (1997), Dansgaard-Oeschger oscillations in a coupled atmosphere-ocean climate model, *J. Clim.*, *10*, 949–970.
- Sakai, K., and W. R. Peltier (1999), A dynamical systems model of the Dansgaard-Oeschger oscillation and the origin of the Bond cycle, *J. Clim.*, *12*, 2238–2255.
- Seidov, D., R. J. Stouffer, and B. Haupt (2005), Is there a simple bi-polar seesaw?, *Global Planet. Change*, *49*, 19–27.
- Siddall, M., T. F. Stocker, T. Blunier, R. Spahni, J. F. McManus, and E. Bard (2007), Using a maximum simplicity paleoclimate model to simulate millennial variability during the last four glacial periods, *Quat. Sci. Rev.*, *25*, 3185–3197.
- Stastna, M., and L. Pogson (2007), On simple models of the North Atlantic thermohaline circulation with switching, *Dyn. Continuous Discrete Impulsive Syst., Part B*, *13*, 305–320.
- Stocker, T. F. (2002), North-south connections, *Science*, *297*, 1814–1815.
- Stommel, H. (1961), Thermohaline convection with two stable regimes of flow, *Tellus*, *13*, 224–230.

- Timmerman, A., and G. Lohmann (2000), Noise-induced transitions in a simplified model of the thermohaline circulation, *J. Phys. Oceanogr.*, *30*, 1891–1900.
- Timmerman, A., H. Gildor, M. Schulz, and E. Tziperman (2003), Coherent resonant millennial-scale climate oscillations triggered by massive melt-water pulses, *J. Clim.*, *16*, 2569–2585.
- Velez-Belchi, P., A. Alvarez, P. Colet, J. Tintore, and R. L. Haney (2001), Stochastic resonance in the thermohaline circulation, *Geophys. Res. Lett.*, *28*, 2053–2056.
- Wannamaker, R. A., S. P. Lipshitz, and J. Vanderkooy (2000), Stochastic resonance as dithering, *Phys. Rev. E*, *61*, 233–236.
- 
- W. R. Peltier, Department of Physics, University of Toronto, 60 St. George St., Toronto, ON, Canada M5S 1A7. (peltier@atmos.physics.utoronto.ca)
- M. Stastna, Department of Applied Mathematics, University of Waterloo, 200 University Ave. West, Waterloo, ON, Canada N2L 3G1. (mmstastn@uwaterloo.ca)