

Characterization of Millennial-Scale Climate Variability

GERARD H. ROE AND ERIC J. STEIG

Quaternary Research Center, and Department of Earth and Space Sciences, University of Washington, Seattle, Washington

(Manuscript received 5 May 2003, in final form 8 October 2003)

ABSTRACT

The oxygen isotope time series from ice cores in central Greenland [the Greenland Ice Sheet Project 2 (GISP2) and the Greenland Ice Core Project (GRIP)] and West Antarctica (Byrd) provide a basis for evaluating the behavior of the climate system on millennial time scales. These time series have been invoked as evidence for mechanisms such as an interhemispheric climate seesaw or a stochastic resonance process. Statistical analyses are used to evaluate the extent to which these mechanisms characterize the observed time series. Simple models in which the Antarctic record reflects the Greenland record or its integral are statistically superior to a model in which the two time series are unrelated. However, these statistics depend primarily on the large events in the earlier parts of the record (between 80 and 50 kyr BP). For the shorter, millennial-scale (Dansgaard-Oeschger) events between 50 and 20 kyr BP, a first-order autoregressive [AR(1)] stochastic climate model with a physical time scale of $\tau = 600 \pm 300$ yr is a self-consistent explanation for the Antarctic record. For Greenland, AR(1) with $\tau = 400 \pm 200$ yr, plus a simple threshold rule, provides a statistically comparable characterization to stochastic resonance (though it cannot account for the strong 1500-yr spectral peak). The similarity of the physical time scales underlying the millennial-scale variability provides sufficient explanation for the similar appearance of the Greenland and Antarctic records during the 50–20-kyr BP interval. However, it cannot be ruled out that improved cross dating for these records may strengthen the case for an interhemispheric linkage on these shorter time scales. Additionally, the characteristic time scales for the records are significantly shorter during the most recent 10 kyr. Overall, these results suggest that millennial-scale variability is determined largely by regional processes that change significantly between glacial and interglacial climate regimes, with little influence between the Southern and Northern Hemispheres except during those largest events that involve major reorganizations of the ocean thermohaline circulation.

1. Introduction

Earth's climate varies on all time scales. Most of our knowledge of its intrinsic variability is derived from the instrumental record, which is temporally very limited. For much of the globe, the instrumental record is considerably shorter than 100 yr. This is only marginally sufficient for understanding the statistics or dynamics of decadal or longer time scales (Gedalof et al. 2002; Percival et al. 2001). On these time scales we must include data from the paleoclimate record: proxy climate data from tree rings, corals, ice cores, and lake and ocean sediments (e.g., Bradley 1999). Several recent studies have used compilations of such data to examine climate variability over the last several hundred to one thousand years (e.g., Overpeck et al. 1997; Mann et al. 1998; Crowley 2000).

On time scales longer than centuries, climate variability is known only from the paleoclimate record. Perhaps as a consequence, millennial-scale variability has

received relatively little attention from the climate dynamics community. Yet evidence for millennial-scale variability is available in the form of hundreds of paleoclimate records that document the last 10–20 kyr in considerable detail, and dozens more that cover the last 100 kyr and longer. For example, pollen records from Europe invariably show a significant cooling event between about 11.6 and 12.8 ka (ka \equiv kyr before present) known as the Younger Dryas. Ocean sediment cores provide compelling evidence for significant changes in ocean circulation associated with the Younger Dryas and other similar events.

Despite the great number of available records, the use of paleoclimate data as a basis for a theoretical understanding of millennial variability has been problematic. The chief limitation is age control. For nearly all records, the time scale is known with insufficient precision to determine the relative timing of changes observed in different records. Questions such as whether or not the Younger Dryas event was global have been the subject of much controversy, still unresolved in the literature (e.g., Markgraf 1993; Steig et al. 1998; Denton et al. 1999; Mulvaney et al. 2000; Jouzel et al. 2001; Morgan et al. 2002). There is also unavoidable ambiguity in the

Corresponding author address: Dr. Gerard H. Roe, Dept. of Earth and Space Sciences, University of Washington, Box 351360, Seattle, WA 98195.
E-mail: Gerard@ess.washington.edu

interpretation of the climate proxy data in terms of conventionally understood climate variables (i.e., temperature, precipitation, pressure). Accordingly, significant weight has been placed on a relatively small group of datasets that are believed to suffer the least from these problems. In particular, current theories of the causes of millennial-scale climate changes draw substantially on evidence from three datasets: oxygen isotope profiles from ice cores at the summit of the Greenland ice sheet [the Greenland Ice Sheet Project 2 (GISP2): 72.6°N, 38.5°W and the Greenland Ice Core Project (GRIP): 73.1°N, 35.5°W], and from West Antarctica (Byrd Station: 80°S, 20°W). Oxygen isotope records are well understood measures of temperature variability (Jouzel et al. 1997).

The Greenland ice cores provided the first definitive evidence for rapid climate change events, which are apparent not only in the oxygen isotope ratios (Dansgaard et al. 1993; Grootes et al. 1993), but also in ice chemistry (Mayewski et al. 1993; Taylor et al. 1993), annual snow accumulation rates (Alley et al. 1993), and in gases preserved in trapped air bubbles (Severinghaus and Brook 1999; Lang et al. 1999). The GISP2 and GRIP cores are also remarkably well dated. Time scales were established by counting of visible annual layering, cross-validated with annual signals in ice chemistry and identification of fallout (sulfate ion, and in some cases ash) from volcanic events of known age. The age of the ice in the GISP2 core is known to within 2% during the most recent 40 kyr, and to better than 10% at 100 ka (Meese et al. 1997). The Antarctic Byrd record, while not dated by annual layer counting [except for the most recent 10 kyr; Hammer et al. (1994)], is important because it has been placed on a relative time scale tied directly to the GRIP and GISP2 records, using the identification of distinctive changes in the concentrations of trace gases extracted from air bubbles within the ice (Blunier et al. 1998). Uncertainty in the Byrd time scale, relative to GISP2 and GRIP, is estimated to be at most a few hundred years during the last 50 kyr (Blunier and Brook 2001). Due to their age control and high resolution, the Greenland and Antarctic ice core records provide a very powerful test of our understanding of the physics of climate variability on millennial time scales.

A prevailing approach has been to place emphasis on the rapid climate change events in the Greenland cores. These events are so large in magnitude and rapidity ($>10^{\circ}\text{C}$ changes in mean annual temperature and a doubling of net precipitation in a few decades or less) that understanding their cause must be considered one of the great challenges in climate research (Alley et al. 2002). They also occur with remarkably regularity, falling into two nominal groupings: the Dansgaard–Oeschger events with a spacing of about 1500 yr, and less regular, longer-lived events about 6000 yr apart. The latter are often referred to as “Heinrich events,” because of their apparent relationship to the Heinrich iceberg discharge

events inferred from deep sea sediment cores from the North Atlantic Ocean (Heinrich 1988). Proposed theories of millennial-scale climate have sought to account for the phase relationships between these events in Greenland and their apparent analogs in Antarctica. The observed frequency and phase relationships (see, e.g., Steig and Alley 2002; Hinnov et al. 2002) have been used extensively to support the attribution of millennial-scale variability to external (i.e., solar) forcing, oscillatory behavior in deep ocean circulation, ice sheet dynamics, nonlinear dynamics in the Tropics, or some combination of all of these (e.g., Broecker 1997; Stocker 1998; Clark et al. 1999; Seidov and Maslin 2001; Weaver et al. 2003).

Several specific models have been proposed as simple characterizations of the physical mechanisms responsible for the Dansgaard–Oeschger and Heinrich time scales identified in the Greenland and Antarctic records. Alley et al. (2001a,b) propose that these records can both be described by a global stochastic resonance mechanism, characterized by threshold crossings between two stable states, with the threshold crossing interval determined by a combination of noise and an underlying, periodic time scale of 1500 yr (the period taken from the highly significant spectral peak in the GISP2 time series). Schmittner et al. (2003) and Stocker and Johnsen (2003) have further proposed that Antarctic climate may be characterized as the integral of Greenland climate [or alternatively Greenland is taken as the derivative of Antarctic climate; see Schmittner et al. (2003) and Huybers (2003)]. This offers a simple mathematical description of proposed global climate physics, and in particular is an improved conception of the so-called interhemispheric seesaw mechanism (e.g., Stocker 1998), in which the thermal adjustment time scale of the thermohaline circulation acts to integrate the Northern Hemisphere climate anomalies. An alternative view is that of Wunsch (2003), who suggests that there is no significant relationship between the time series on millennial time scales and that, at least on average, regional processes dominate the character and pacing of millennial-scale climate variability.

In this paper, we examine the extent to which these different characterizations of millennial-scale variability can be reconciled. Our approach is the application of stochastic climate models to the Greenland and Antarctic isotope records. Because any geophysical time series will always reflect the presence of random forcing and inertia (Hasselmann 1976), even if more complex processes are also involved, stochastic climate models can offer considerable insight into the underlying physics driving the climate variability recorded by the time series. In particular, such models can be used to estimate the characteristic physical time scales associated with millennial-scale variability. Our results show that a considerable amount of the background variance in the Greenland records can be explained by the simplest sto-

chastic model, a first-order autoregressive model [AR(1)] with a characteristic time scale of 400 yr. Although such a model does not account for the inherent asymmetries in the record (rapid warming, slow cooling), nor the periodic nature of the rapid-warming events, it compares favorably against the stochastic resonance model in its explanatory power. We find that for Byrd, AR(1) is entirely adequate. For those time intervals that include the largest climate anomalies observed in the Greenland records, the integrated Greenland signal plus AR(1) provides a better explanation for the Antarctic record than AR(1) alone. For most of the time interval considered (approximately the last 80 000 yr), however, any interhemispheric communication contributes much less than the independent variability inherent to each record. Furthermore, in the Holocene (the most recent 10 kyr), the characteristic time scales of variability are quite different than during the preceding glacial climate period.

2. Data

The isotope records GISP2 and GRIP in Greenland, and Byrd in Antarctica (hereafter generally referred to simply as “GISP2,” “GRIP,” and “Byrd”), are probably the best paleoclimate datasets available for the study of millennial-scale climate change. This is not to suggest that other records are not of considerable importance. Ocean sediment records from the Cariaco basin (Hughen et al. 1998), from the Santa Barbara basin (Hendy and Kennett 1999), and from the Bermuda rise and elsewhere in the North Atlantic (Sachs and Lehman 1999; Bond et al. 1997) all show compelling evidence for rapid climate changes. In Antarctica, most other ice cores show very similar variations to those observed at Byrd (e.g., Petit et al. 1999). These records provide important corroborating evidence that GISP2–GRIP and Byrd are broadly representative of millennial-scale variability at least in the North Atlantic region and in a large portion of the Antarctic, respectively, and perhaps over much broader areas. They do not however, generally provide additional information that can be used in the quantitative analyses we pursue in this study, either because the records are too short, not highly resolved enough, or not well enough dated. In many cases they are also less straightforward to interpret than are the temperature proxy data from the ice core isotope profiles. We therefore focus in this paper entirely on the GISP2–GRIP and Byrd isotope records.

When referring to the oxygen isotope data, we use the conventional notation

$$\delta^{18}\text{O} = \left[\frac{(^{18}\text{O}/^{16}\text{O})_{\text{sample}}}{(^{18}\text{O}/^{16}\text{O})_{\text{SMOW}}} - 1 \right] 1000\text{‰}, \quad (1)$$

where the international isotope standard is Standard Mean Ocean Water (SMOW). We treat anomalies in $\delta^{18}\text{O}$ as anomalies in climate. The interpretation of iso-

tope ratios in terms of temperature depends on the empirical relationship between $\delta^{18}\text{O}$ and mean annual temperature, which arises because of the temperature-dependent fractionation of water during the evaporation, transport, and condensation processes (Dansgaard et al. 1973). This interpretation has been validated with inversion of borehole-temperature measurements (Cuffey et al. 1995; Johnsen et al. 1995; Dahl-Jensen et al. 1998) and with anomalies in gas-isotope concentrations due to diffusion under strong thermal gradients (Severinghaus and Brook 1999; Lang et al. 1999). Although it is well known that the isotope–temperature relationship is not strictly linear, and may have varied considerably over time, these independent calibrations provide a strong justification for treating it as linear for the purposes of our study here. The largest changes in the isotope–temperature relationship have probably occurred during transitions between glacial and interglacial climates (e.g., between 20 and 10 ka); we do not include such transitions in our analysis.

Isotope data from the GISP2 and GRIP cores, drilled only 30 km apart, may for the most part be considered identical between 0 and 100 ka (Grootes et al. 1993). Hinnov et al. (2002) have noted some important differences in the spectral properties of the time series that may result in part from the procedures used in relating time scales for GRIP, GISP2, and Byrd to one another. We base our analyses on GISP2, because it is the least problematic for the statistical methods we use and is most reliably matched with Byrd [for both cores we use the GISP2-based time scales in Blunier and Brook (2001)]. Repeated analyses show that use of GRIP or alternative time scales for GISP2 gives results that lie within the stated uncertainties of our estimates and so would not influence any of our fundamental conclusions.

3. Stochastic climate models

At its simplest, stochastic climate modeling assumes that climate variability is entirely made up of random, uncorrelated variations plus inertial terms that characterize the climate system’s tendency to “remember,” to some degree, its previous states (e.g., Hasselmann 1976; von Storch and Zwiers 1999). Autoregression modeling is a mathematical technique that provides a framework for this kind of analysis (e.g., Jenkins and Watts 1968; von Storch and Zwiers 1999).

In an autoregressive model, an observation at time t , u_t , is regarded as being due to a linear combination of its previous p states ($u_{t-i\Delta t}$, $i = 1, 2, \dots, p$), plus a random stochastic term v_t (i.e., Gaussian white noise), which represents the uncorrelated forcings acting on the system:

$$u_t = a_0 v_t + a_1 u_{t-\Delta t} + a_2 u_{t-2\Delta t} + \dots + a_p u_{t-p\Delta t}. \quad (2)$$

This is known as a discrete autoregressive process of order p , or AR(p), and it may be used as a statistical

model to try to explain a time series of observations (assumed discrete and equally spaced in time). The purpose of looking at data in this way is that an AR(p) process is the discretized equivalent of a p th-order ordinary differential equation (e.g., von Storch and Zwiers 1999). The underlying equation can reveal the dynamics of the system being observed. For example, an AR(1) process is equivalent to

$$\frac{du}{dt} + \frac{u}{\tau} = \nu(t), \quad (3)$$

where

$$\tau = -\frac{\Delta t}{\ln a_1}. \quad (4)$$

Equation (3) describes a system that is forced by random noise, and that has a tendency to relax back toward its equilibrium state with a time scale, τ , associated with the inertia of the system and the physical process that restores it to equilibrium. In simple systems τ might be the time scale of thermal damping by heat fluxes or momentum damping by frictional dissipation. In more complex systems it can be thought of as a “characteristic time scale” associated with the entire assemblage of processes that act to restore the system to equilibrium. For a (linear) combination of time scales, the effective time scale is equal to 1 over the sum of the reciprocals of the individual time scales (like resistors in parallel), and so is strongly weighted to the longest characteristic time scale. An AR(1), or Markov, process is also known as “red noise:” the system’s inertia integrates the white noise forcing and results in a time series with variance that is enhanced at long time scales relative to short time scales. Higher-order AR processes (i.e., $p \geq 2$) represent higher-order ordinary differential equations, and as such have solutions that are a combination of oscillations and decaying or growing exponentials.

4. Characterization of GISP2 and Byrd

We begin by identifying the AR(p) processes that best characterize each time series using the MATLAB algorithm *arfit* of Schneider and Neumaier (2001). This employs a least squares estimator that, for a given order p , finds the combination of a_p s that minimizes the magnitude of a_0 (i.e., the noise), and therefore is the best explanation of the data for that AR(p) process. The optimal order, p , of an AR model is chosen by using a selection criterion that penalizes higher orders because they have more degrees of freedom. The default selection criterion used by the *arfit* algorithm is Schwarz’s Bayesian criterion (Schwarz 1978). An important part of the model fitting process is to check a posteriori that the model is self-consistent. That is, after having subtracted from the observations what is explainable by the autocorrelations, is what is left over (the residuals) consistent with uncorrelated, normally distributed white




FIG. 1. The $\delta^{18}\text{O}$ records from Byrd and GISP2 over the last 90 kyr. Values at GISP2 have been displaced by 5 (‰).

noise? This can be evaluated using the modified Li–Mcleod portmanteau statistic (e.g., Schneider and Neumaier 2001) to test for correlations in the residuals, and the Kolmogorov–Smirnov test (e.g., von Storch and Zwiers 1999) to test for the normality of the distribution.

From Fig. 1 it is clear that both Byrd and GISP2 reflect the overall progression of the ice ages, and, as we expect, they are highly correlated on longer time scales (i.e., greater than 10 kyr; e.g., Steig and Alley 2002). It is important to minimize the influence of these long time scales on our analysis. One approach would be to use a high-pass frequency filter on the full 80-kyr records, but doing so introduces spurious autocorrelations between adjacent data points. Therefore, instead we concentrate initially on a shorter interval of time: 50–20 ka.

The data have an average spacing of 91 yr for Byrd and 118 yr for GISP2 over the 50–20-ka interval. To obtain uniformly spaced time series, we use nearest-neighbor interpolation. The choice of interpolation method is important because the analysis relies on the autocorrelation between adjacent data points. Also, each data point represents the isotope values averaged over some finite depth of the core, which risks spurious autocorrelations between neighboring data points. These effects were evaluated by generating arbitrarily high-resolution realizations of the optimal AR(p) process for the records using (2), and then averaging those realizations over the actual sampling intervals in the original data. The time series generated in this way was then refit as an AR(p) model. If averaging and interpolation in the data had no effect, we would expect to obtain the same AR(p) coefficients. Repeating this exercise many times suggests that the interpolation and inherent averaging in the data causes the characteristic time scale to be overestimated by about 20%. We repeated our analyses for a variety of different intervals during the glacial period of the record (taken to be 80–20 ka), and also for the whole 60-kyr glacial record and applying a high-pass Butterworth filter (using a standard MATLAB routine). As an additional test, we used the “TAUEST” routine of Mudelsee (2002), which pro-

TABLE 1. Best-fit AR(p) models for ice core $\delta^{18}\text{O}$ record, selected using Schwarz–Bains criterion. Here Δt is the interpolation interval used, τ is the time scale associated with the AR process, and Per. is the period of oscillation implied by the AR(2+) processes. A Y in the penultimate column means that the residuals are indistinguishable from uncorrelated to better than 95%, based on a modified Li–McLeod test (Schneider and Neumaier 2001). A Y in the last column means the residuals are indistinguishable from a normal distribution to better than 95%, based on a Kolmogorov–Smirnov test (e.g., von Storch and Zwiers 1999).

Core	Interval ka	Δt yr	AR(p)	$\tau \pm 95\%$ yr	Per. $\pm 95\%$ yr	Uncorrelated residuals?	Normally distributed?
Byrd	50–20	100	1	503 ± 204	—	Y	Y
		200	1	610 ± 299	—	Y	Y
		300	1	710 ± 402	—	Y	Y
Byrd	10.2–0.2	50	1	129 ± 50	—	—	—
		100	1	170 ± 86	—	Y	Y
		200	1	208 ± 149	—	Y	Y
GISP2	50–20	100	1	451 ± 174	—	Y	—
		200	1	432 ± 190	—	—	Y
		300	1	366 ± 187	—	—	Y
GISP2	50–20	100	2	210 ± 61	2257 ± 2149	Y	—
		200	2	303 ± 138	2026 ± 1025	Y	—
		300	2	411 ± 243	2016 ± 656	Y	Y
GISP2	10.2–0.2	50	0	—	—	Y	Y
		100	0	—	—	Y	Y
		200	0	—	—	Y	Y

vides an estimate of the best-fit AR(1) model for unevenly spaced datasets. These tests all give time scales consistent within our uncertainties, and so do not alter our basic conclusions.

a. Characterization of Byrd

Results are presented in Table 1 for a variety of interpolation intervals, Δt . In all cases, the glacial record in Byrd is best represented by an AR(1) process. The residuals are indistinguishable from uncorrelated white noise, and so we can conclude the AR(1) is a self-consistent explanation of the Byrd record. The best-fit value for τ varies somewhat depending on the interpolation interval, although all are in agreement to within their 95% confidence limits. The results suggest a value for τ of about 600 ± 300 yr.

A power spectrum estimate for the 50–20-ka interval of the Byrd record is shown in Fig. 2. Also shown is




FIG. 2. Spectral estimate of Byrd between 50 and 20 ka, using windowed periodogram (10-kyr Hanning window), and theoretical spectrum of best-fit AR(1) process and 95% confidence intervals (dashed lines). $\Delta t = 100$ yr.

the theoretical spectrum for the parameters of the best-fit AR(1) process, and its 95% confidence limits. That the Byrd data lie largely within the uncertainty limits is further support that the time series is well represented by an AR(1) process.

The basic impression from Fig. 2 is that the spectrum is red. Spectral power increases at longer periods, and it is readily calculated that half of the variance in Byrd occurs at periods of 4000 yr and longer. An important point is that most of the variance in an AR(1) process occurs at periods substantially longer than the characteristic time scale associated with the physics driving the system back to equilibrium. The theoretical spectrum shown in Fig. 2 is for a discrete AR(1) process, but it is closely related to the spectrum for the continuous first-order autoregressive process of (3) (e.g., Jenkins and Watts 1968), which we use here for clarity:

$$P(f) = \frac{a_0^2}{1 + (2\pi f\tau)^2}, \quad (5)$$

where $P(f)$ is the spectral power as a function of frequency, f . It is readily shown that half the variance occurs at periods greater than $2\pi\tau$. These results then are a strong suggestion that the ~ 3000 – 6000 -yr “millennial-scale” variability seen by eye and in the power spectra of the Byrd record arises from physical processes with characteristic time scales that are significantly less: about 600 yr. This is shorter than that commonly cited for the overturning of the global thermohaline circulation [one to three thousand years; e.g., Broecker (1991); Schiller et al. (1997)], or for the dynamical adjustments of continental-scale ice sheets [several thousand years; e.g., Paterson (1994); Steig et al. (2000)]. The time scale is, however, consistent with that estimated for adjustment of thermohaline circulation (e.g., Weaver and Sarachik 1991), or of basin-scale ocean circulation (e.g.,

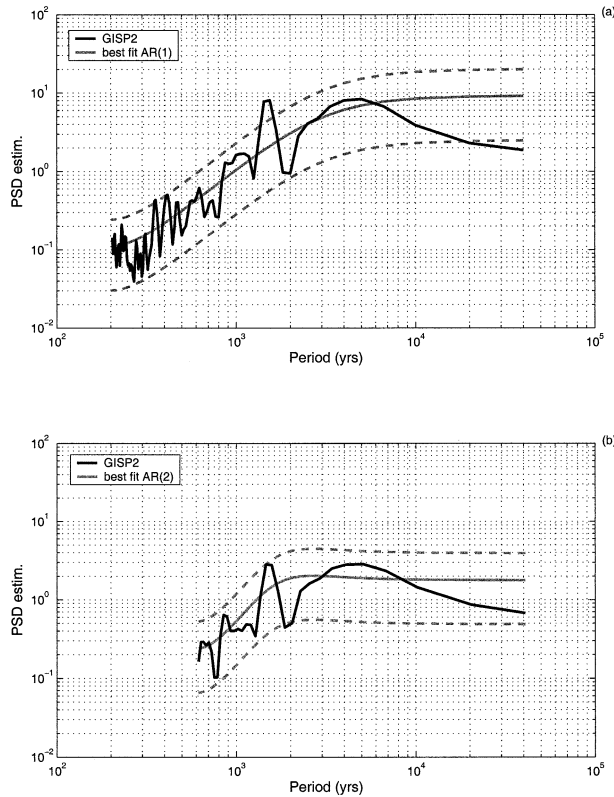


FIG. 3. (a) Same as Fig. 2, but for GISP2. (b) Same as (a) but best-fit AR(2) process, $\Delta t = 300$ yr, and linear interpolation.

Antarctic deep-water formation) to perturbations (e.g., Broecker 2000). It is also of the same order as suggested for switching of ice streams (e.g., Fahnestock et al. 2000; Conway et al. 2002), or the regional dynamics of smaller ice sheets and ice shelves (e.g., Jóhannesson et al. 1989). These comparisons are contingent on the inferred mechanisms being appropriately described as relaxation physics. Other physical mechanisms, such as the possible binge–purge cycles in ice sheets (MacAyeal 1993a,b; Bond and Lotti 1995), or delayed climate oscillators (e.g., Battisti 1988) have time scales that bear a different relationship to the power spectrum they generate.

b. Characterization of the GISP2 record

The autoregression analysis for GISP2 is less straightforward than for Byrd. For the interval between 50 and 20 ka, the analysis gives AR(1) as the optimal explanation of the data for all the interpolation intervals shown in Table 1, with a time scale of about 400 ± 200 yr. A spectral estimate for GISP2 is shown in Fig. 3, together with the theoretical spectrum for the best-fit AR(1) process for $\Delta t = 100$ yr. As noted by many previous authors (e.g., Stuiver et al. 1995), there is a prominent ~ 1500 -yr peak well above the 95% confidence limits for the theoretical spectrum, indicating that




FIG. 4. Histogram of the rate of change of GISP2 from 50 to 20 ka, $\Delta t = 100$ yr, and linear interpolation.

AR(1) is not a complete explanation of the data. However it is also clear that AR(1) does account for a large fraction of the variance in the spectrum as noted by Wunsch (2001) and Hyde and Crowley (2002). Importantly, for all of the interpolation intervals considered, the residuals resulting from an AR(1) fit are inconsistent with uncorrelated, normally distributed white noise; AR(1) is therefore not a self-consistent explanation of the data.

The results for this interval of GISP2 are quite sensitive to the interpolation method. If linear interpolation is used, the analysis procedure selects AR(2) as the best explanation of the data. This AR(2) process reflects relatively weak oscillatory behavior with a period of about 2000 yr and an exponential decay time scale of ~ 300 yr (Table 1). This decay time indicates how long phase information of the oscillation is remembered by the system. For $\Delta t = 300$ yr, the residuals are consistent with uncorrelated and normally distributed white noise and so, nominally, this AR(2) process could be a consistent explanation of the data. In this case the confidence limits for the theoretical AR(2) spectrum actually encompass the 1500-yr peak in the GISP2 spectral estimate (Fig. 3). It would therefore no longer be considered significant against this background spectrum (which, to be sure, contains a weak oscillation). However, by taking $\Delta t = 300$ yr, the strong asymmetry in the data between warmings and coolings (see below) has effectively been averaged out. Because the driving equations of AR(p) processes cannot themselves generate such asymmetries, we can reject a pure AR(2) process as the governing dynamics.

ASYMMETRY IN GISP2

The asymmetry clearly shows up in a histogram of the time derivative of the GISP2 records in Fig. 4. The shoulder on the positive flank of the distribution reflects the existence of rapid warming events. This kind of distribution can be imitated by incorporating a threshold behavior into an AR(1) process. Figure 5a shows a realization of the AR(1) process that is a best fit to GISP2

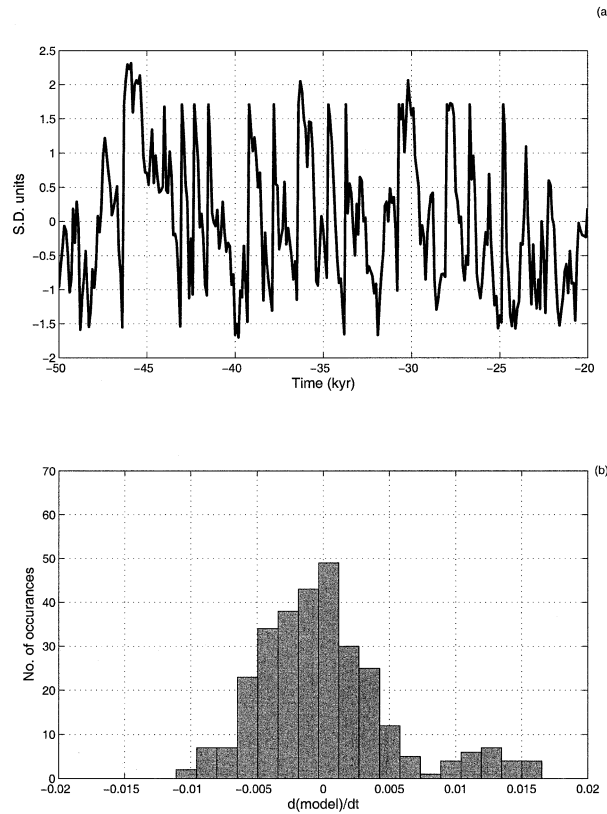


FIG. 5. Simple model of AR(1) process incorporating a threshold.

[i.e., $a_0 = 0.60$, $a_1 = 0.80$ in (2)] with a threshold condition that if u_t descends below a value of -1.0 it is immediately raised to a value of 2.0 . Figure 5a shows intervals of rapid changes and intervals of longer time scale variations. It does not however show the preferred 1500-yr spacing observed in the record. Other properties of the rapid warming events have been noted and tabulated such as an apparent clustering and a relationship to global ice volume (e.g., Schulz 2002a,b), but we are not aware of any statistical evaluation. A default (albeit incomplete) model of AR(1) plus threshold provides a null hypothesis against which to test these putative properties.

c. AR(1) versus stochastic resonance for Byrd and GISP2

As Rahmstorf (2003) has recently pointed out, the 1500-yr spacing in GISP2 is highly significant, not only because it produces a strong spectral peak, but because there is very long phase memory (more than 20 cycles). Obviously, a simple AR(1)-plus-threshold model will not capture this aspect of the record. Nevertheless the histogram of the rates of change for AR(1) shows a striking resemblance to that for the real data (i.e., cf. Fig. 5b with Fig. 4). As noted above, AR(1) is entirely self-consistent for Byrd. In this context it is illustrative




FIG. 6. Stochastic resonance histograms for (a) Byrd and (b) GISP2. The histograms shown are for the interval 80–20 ka, and include intervals between both warming and cooling events. Data was 20-kyr high-pass filtered using a Butterworth filter. The lines are estimates for the theoretical histogram based on a 10-Myr realization of that process (scaled to the 60-ka interval of the data), for the best-fit AR(1) process for that record (solid), and for a 1500-yr stochastic resonance process (dashed).

to consider how the stochastic resonance model proposed by Alley et al. (2001b) for both of these records compares with AR(1). The suggestion is that the climate acts as a stochastically resonant, bimodal system with transitions between the two states being paced by an underlying 1500-yr oscillation, perhaps due to solar variability. A characteristic property of stochastic resonance is the distribution of the time intervals between transitions into one of the states. Transitions into a given state are most likely to occur at a particular phase of the underlying cycle. A histogram of transition intervals is therefore peaked at 1500 yr, with subsidiary maxima at 3000, 4500 yr, etc.

Following Alley et al. (2001a) we determine a warming or cooling event as having occurred when the time series exceeds a particular fraction of its standard deviation, which we set to be 0.2 (although the conclusions do not depend on this value). The histogram of the intervals between such threshold crossings for the Byrd record is shown in Fig. 6a for the period 80–20 ka, and includes both warmings and cooling intervals. By generating 10-Myr realizations of the best-fit AR(1) and

TABLE 2. Comparison between the best-fit AR(1) and a 1500-yr stochastic resonance process as explanations of the data. The second column gives the interval examined. The third column denotes whether warming and cooling event intervals were included (W&C), or warming events only (W). The fourth column gives the number of events that were found in that interval. For each process [i.e., AR(1) or stochastic resonance], the χ^2 value is for the difference between the estimate for theoretical histogram of that process (based on a 10-Myr realization and the histogram of the data, using 30 histogram bins). The following column is the percentage of the record-length realizations of that process that had histograms *less* like the theoretical histogram than the data (based on their χ^2 values). A low percentage therefore suggests that the process is not a good description of the data. Finally the F test significance is the fraction of random realizations in which AR(1) is a better representation of the data than stochastic resonance (i.e., having a lower χ^2 value). A value of 50% would suggest that the two models were equally good explanations of the data. For the 60-kyr interval, the data is high-pass filtered using a Butterworth filter with a 20-kyr cutoff.

Core	Interval (ka)	Type	No. of events	AR(1)		Stochastic resonance		F test significance
				χ^2	(%)	χ^2	(%)	
Byrd	50–20	W&C	33	23.7	60%	171.6	1.2%	99%
Byrd	80–20	W&C	48	61.3	13%	276.0	0.2%	98%
Byrd	50–20	W	17	13.5	90%	55.0	16%	93%
Byrd	80–20	W	24	44.7	21%	172.9	0.5%	98%
GISP2	50–20	W&C	33	23.8	77%	51.6	22%	85%
GISP2	80–20	W&C	51	34.9	57%	72.8	11%	88%
GISP2	50–20	W	15	18.2	85%	19.9	65%	61%
GISP2	80–20	W	25	34.4	47%	39.3	29%	67%

stochastic resonance processes we can obtain a good approximation to the theoretical threshold crossing histograms. The histogram for an AR(1) process has a single peak that is directly related to the characteristic time scale, or memory, that AR(1) represents. As a measure of the closeness of the agreement between the data and each process, we can calculate χ^2 values for the difference between Byrd and each of these long realizations (e.g., Press et al. 1992). The results for the intervals 50–20 and 80–20 ka, as well as for histograms including intervals between warming events only, are shown in Table 2.

For the χ^2 values to be meaningful it is important to know what typical errors look like. That is, for each candidate process, what is the distribution of the errors between its theoretical histogram and random Byrd-length realizations of the same process? To evaluate this, we performed a Monte Carlo analysis with 10 000 random realizations of each process and calculated the distribution of χ^2 values for the deviations from the theoretical histogram. For both the best-fit AR(1) and stochastic resonance, the error distribution was close to χ^2 distributed with 28 degrees of freedom (i.e., 30 histogram bins with two model parameters). These calculated error distributions can then be used to attribute significance to the χ^2 values by asking what fraction of random realizations would be expected to have a ratio of χ^2 values exceeding that calculated for the data. This is equivalent to a standard F test (e.g., von Storch and Zwiers 1999), which also gives very similar results. As in Alley et al. (2001a), we can determine what fraction of the random realizations of each process have χ^2 values that exceed that between the data and the theoretical distribution. For example, if Byrd between 50 and 20 ka were equal to a typical 30-kyr realization of stochastic resonance, we would expect about 50% of random realizations of stochastic resonance to have higher χ^2 values (i.e., have less good agreement with the the-

oretical distribution). These fractions are also shown in Table 2.

Stochastic resonance does significantly worse than AR(1) in explaining Byrd. In all cases a much smaller fraction of the random stochastic resonance processes had χ^2 values showing worse agreement with the theoretical curve than the Byrd record, and thus Byrd is much more like a typical AR(1) process than a typical stochastic resonance process. Furthermore, the significance test based on the calculated χ^2 distributions shows that, with a high degree of confidence, it can be concluded that AR(1) is a better fit to the data than stochastic resonance. The reason is that, for stochastic resonance, transition intervals shorter than the oscillation period are very unlikely, and the χ^2 value is strongly penalized by the peak in the Byrd record at about 900 yr.

In the case of GISP2, the results are similar to, but slightly less conclusive than, those for the Byrd record. Figure 6b shows the results for AR(1) only, although very similar results were obtained for AR(2). As we have already concluded, neither AR(1) nor AR(2) are adequate explanations for GISP2. Nevertheless, Table 2 indicates that, like Byrd, the GISP2 record is more like a typical AR(1) process than a typical stochastic resonance process, although not quite at the 95% level of significance. As for Byrd, the stochastic resonance model is penalized by the peak in transitions shorter than the 1500-yr oscillation period. In other words, the GISP2 record exhibits large climate changes that are not accounted for by the stochastic resonance model. It is important to note, however, that many of these climate changes are cooling events. If only warming events are included, AR(1) is still preferred, but only by a small statistical margin. For Byrd, choosing warming intervals only still leaves AR(1) better than stochastic resonance with a high degree of confidence. For GISP2, the choice of warming events only may be reasonable, but this is

an implicit choice about the hypothesized physics that is being tested. For example, as Alley et al. (2001a) suggest, the physics of the transition may not be symmetric, meaning that one transition direction may be more easily identifiable against the background noise. We conclude that stochastic resonance cannot be ruled out as an explanation for the warming events at GISP2, but must be considered a rather incomplete description of the record. For Byrd, we can reject stochastic resonance as a parsimonious explanation.

d. Holocene variability

We now consider the character of both records during the Holocene, which we take as the last 10 kyr. For Byrd, it is again AR(1) that is the best explanation, and for $\Delta t \geq 100$ yr, the residuals are consistent with uncorrelated, normally distributed white noise. Importantly however, the time scale of the AR(1) process (175 ± 100 yr) is much shorter than for the glacial period. The changes at GISP2 between glacial and interglacial are even more dramatic. For the range of interpolation time scales shown in Table 1, the GISP2 Holocene record is consistent with uncorrelated, random white noise. For this data sampling, τ must be less than ~ 20 yr, meaning that GISP2 is essentially white during the Holocene. The same result was obtained for the GRIP and the North Greenland Ice core Project (NGRIP; Johnsen et al. 1995, 2001) records.

It is worth pointing out that other proxy indicators of Holocene variability are not necessarily inconsistent with our results, even though they do not all show the same response. For example, while the forcing climate might be white (or near white), proxy climate indicators that have multidecadal or longer response time scales, would have an enhanced reddened response to such forcing. Thus, our results do not necessarily contradict the well-known idea that glaciers have undergone substantial millennial-scale variability during the Holocene (e.g., Hormes et al. 2001; Bond et al. 1997).

5. The connection between Byrd and GISP2

There has been a considerable amount of discussion in the literature about the relationship between the GISP2 and Byrd records, and what it reveals about linkages in climate between the hemispheres at millennial time scales. Much of the analysis has focused on the direct connection between the records either by noting the visual similarities (e.g., Blunier et al. 1998; Blunier and Brook 2001), or by looking at the correlations between the two time series (e.g., Steig and Alley 2002). Figure 7a shows the two time series high-pass filtered (with a Butterworth filter and a 10-kyr cutoff). Treated in this way, there is a maximum correlation of -0.42 when GISP2 leads by 400 yr, and a close secondary maximum of 0.39 when Byrd leads by 1400 yr. Steig and Alley (2002) argued on this basis that the data did




FIG. 7. Relationship between GISP2 and Byrd. (a) GISP2 leading Byrd by 400 yr, (b) integral of GISP2 lagging by 200 yr. Both records have been high-pass filtered at 10 kyr using a Butterworth filter.

not discriminate between a “seesaw” climate mechanism or a “Southern Hemisphere lead.” While visually impressive, these correlations are not necessarily significant. Taking the unfiltered records between 50 and 20 ka, one finds that, allowing for a 6-kyr lead/lag window, 50% of random realizations of the AR(1) process that are the best fit for Byrd (for $\Delta t = 200$ yr) have a maximum lag correlation that exceeds that between Byrd and GISP2. That is, one would conclude that the apparent direct correlations between the records could well have arisen simply by chance. Wunsch (2003) arrived at the same conclusion using cross-spectral analysis.

Another way to examine this is to consider the relationship between the records on an event-by-event basis. There are 16 occasions between 80 and 20 ka where the rate of change of the GISP2 record exceeds two standard deviations (Fig. 8a), and it is clear that this criterion picks out most of what are generally considered the “rapid change” events in this interval. If large events in the GISP2 record are regularly either echoed or presaged in the Byrd record, we might expect to see an indication of this. Compiled in Fig. 8b are the amplitudes and rate of change of the Byrd record for the 1500 yr either side of each of the rapid change events identified in the GISP2 record (the same results hold




FIG. 8. (a) GISP2 record. The diamonds denote the 16 distinct events when the rate of change at GISP2 exceeds two standard deviations between 80 and 20 ka. Data was high-pass filtered at 10 kyr using a Butterworth filter. (b) Amplitude and rate of change in the Byrd record in a 1500-yr window either side of the rapid-change events in GISP2. There are thus 16 points plotted for each lead/lag. The lines denote 95% confidence intervals assuming a Student's t distribution.

true also if longer lags are allowed). There are thus 16 points plotted at each lead and lag. Confidence limits can be calculated assuming the points follow a Student's t distribution (e.g., von Storch and Zwiers 1999). At all leads and lags considered, the 95% limits (and the data) straddle zero for both the amplitude and rate of change at Byrd. Thus, there is no strong indication that rapid changes in GISP2 are reflected in Byrd in any consistent way.

In apparent contrast, several recent studies have argued that, in fact, Byrd and GISP2 are strongly related, but through an integral rather than a direct relationship. Schmittner et al. (2003) suggest an interhemispheric connection in which changes in the North Atlantic (as recorded at GISP2) have a cumulative effect on the Southern Ocean (as recorded at Byrd). Huybers (2004) further argues that the data support a Southern Hemisphere lead, although dating errors between the cores make the phasing somewhat ambiguous. Both of these studies base their analysis on simple models in which the GISP2 data (after bandpass filtering) is shown to correlate significantly with the (filtered) rate of change

at Byrd. In a similar vein, Stocker and Johnsen (2003) also show that after bandpass filtering there is a high correlation between the Byrd record and the time integral of the GISP2 record. This is mathematically similar to the Schmittner et al. (2003) and Huybers (2004) models, but has somewhat different implications for the physics as we discuss later.

The question thus arises whether we can reconcile the different results from direct comparison of the records on the one hand (showing no compelling evidence of significant shared variance, either in an average sense or an event-by-event basis), and on the other hand, the comparison of filtered Byrd and integrated GISP2 (showing evidence of a significant connection between the records). To address this we conducted analyses similar to those described above as follows. Figure 7b shows the Byrd record compared with the integral of the GISP2 record. Both records have been high-pass filtered at 10 kyr to remove the influence of Milankovich frequencies but not low-pass filtered. The integral is formed simply from the cumulative sum of the filtered record. Although there is a slight arbitrariness to this integral, the results do not depend strongly on filter cutoff frequency (between 20 and 5 kyr), nor on the starting point of the integration. The maximum lag correlation between the two time series is -0.6 . Consistent with Huybers' (2004) results, we find that this exceeds the maximum lag correlation obtained for random realizations of a red noise process that is the best fit to Byrd. These results strongly suggest that there is indeed some shared signal between the two records on millennial time scales. However, this connection is demonstrated only after heavy filtering of the data, which also removes much of the unshared variance. Furthermore, while the apparent relationship in Fig. 7b is striking, it appears that much of the shared signal is in the longer 6000-yr Heinrich events that dominate the earlier part of the record, prior to 40 ka. We therefore repeated the correlation analyses for sequential 30-kyr segments of the records (i.e., from 80–50 down to 50–20 ka), and find that the correlation decreases from 0.6 in the former to 0.3 in the latter. This suggests—in accordance with the qualitative interpretation of Blunier et al. (1998)—that while for the biggest events there is coupling between the hemispheres, for the majority of millennial-scale climate variations in GISP2 there is no demonstrable corresponding signal in Byrd. Thus, a substantial amount of the variance in Byrd must be thought of as unrelated to GISP2. In the next few sections, we further quantify these relationships by employing cross-spectral analysis, and a series of simple linear models to evaluate the character of this “unshared” variance and to estimate its magnitude.

a. Cross-spectral analysis

The changing nature of the relationship between the earlier and later halves of the records is borne out by

cross-spectral analysis. Figure 9a shows that the significant cross coherence between 1/2 and 1/10 kyr comes almost entirely from the 80–50-ky interval. While there is a slight hint of cross correlation between 50 and 20 ka at $\sim 1/1.5$ kyr, the different (and rapidly varying) relative phase relationship (Fig. 9b) would suggest caution in interpreting its significance (it is not at all clear, e.g., why ocean physics would integrate the signal only in a narrow band). For the whole record the relative phase of $\pi/2$ plus a linear trend in the band (1/7–1/2 kyr) is noted by Huybers (2004). This is consistent with the derivative of Byrd leading GISP2. However subdividing the data shows that this is not a robust relationship. Although spectral analysis is formally equivalent, analysis in the time domain may sometimes give a different interpretation, and to that end the next section presents results from linear regression modeling.

b. Linear regression modeling

From the results above, it appears that there is a significant connection between the Byrd and GISP2 records, but with some component of uncorrelated variance. In this section, we use linear regression modeling to estimate the magnitude of that variance. Specifically, we attempt to “explain” Byrd as the linear combination of AR(1), GISP2, and integral of GISP2. The four candidate models, following the discussion above, are

$$\begin{aligned} \text{Byrd}(t) &= \text{AR}(1), \\ \text{Byrd}(t) &= \phi_1 \text{GISP2}(t - t_{\text{lag}}) + \text{AR}(1), \\ \text{Byrd}(t) &= \int_{t_s}^{t-t_{\text{lag}}} \text{GISP2}(t') dt' + \text{AR}(1), \\ \frac{d\text{Byrd}}{dt} &= -\frac{\text{Byrd}(t)}{\tau} + \phi_2 \text{GISP2}(t - t_{\text{lag}}) \\ &\quad + \text{AR}(0). \end{aligned} \quad (6)$$

The first, default model is simply an AR(1) process. The second is a linear combination of GISP2 and AR(1). The third and fourth models would be favored by the results reviewed above that argue for a connection between the records. The third is a linear combination of the integral of the GISP2 record and AR(1). The fourth assumes that the rate of change of Byrd reflects its current state, the GISP2 record and white noise [AR(0)]. The third and fourth models are thus similar except that the third treats GISP2 and integrated memory separately, while the fourth integrates them together. In each model where it is used, we allow for a lag in the GISP2 record (or its integral). Equations (6) are discretized, model parameters are optimized to match the Byrd record, and the residuals are evaluated for consistency with white noise. The different models are compared using an F test for the ratio of the magnitudes of the unexplained variance, adjusting for the respective degrees of freedom (e.g., von Storch and Zwiers 1999). Table 3

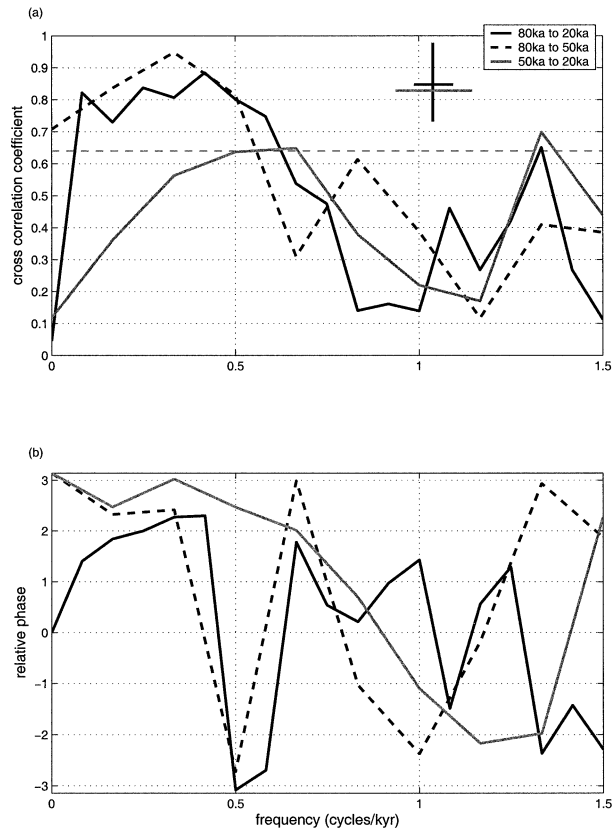


FIG. 9. Cross spectral analysis of the GISP2 and Byrd records for the full 80–20-ka interval, and subdivided into two halves: (a) cross-correlation coefficient, (b) relative phase. Here $\Delta t = 200$ yr, nearest-neighbor interpolation. Windowed cross periodogram was made using a Hanning window 1/5 the length of the time series. The dashed horizontal line in (a) indicates the level above which significant cross correlations can be inferred at the 95% confidence level. Also shown is the effective bandwidth for the two difference intervals and approximate 95% confidence limits on the cross-spectral estimates.

shows the results for each of the models for the 50–20- and 80–20-ka intervals. For the former, we present results for both the filtered and unfiltered time series (high-pass filtering is required for the full 80–20-ka interval to remove the influence of Milankovich forcing on the correlations).

None of the other models are better than pure AR(1) alone at the 95% significance level of the F test, if either unfiltered or high-pass-filtered data are used. However, this does not mean that there is no connection between the records; as expected from the correlation results discussed above, excluding the higher frequencies by redoing the analyses with a 1000-yr rather than 200-yr interpolation brings the F test statistics above 95%, though only for the full 80–20-ka record (the significance remains below 67% confidence if only 50–20 ka is considered). Consistent with the cross-spectral analysis, these results reflect the relatively small fraction of shared variance particularly in the more recent (50–20 ka) part of the record. The results for the Holocene also

TABLE 3. Results of linear regression modeling. Here B and G2 refer to Byrd and GISP2. When filtered, a 10-kyr high-pass Butterworth filter was used. Here amspe is the adjusted-mean-square error [the unexplained variance, adjusted for the number of model parameters (#P)]. The optimum lag and correlation coefficient are that between Byrd and GISP2 (or its integral), and for the best model fit. The last column gives the F test significances for whether the model is a better explanation than AR(1) alone. For all results shown nearest-neighbor interpolation and $\Delta t = 200$ yr were used (except for the Holocene when 100 yr was used). Other interpolation intervals and methods do not change the conclusions.

Model	Interval (ka)	Filt?	#P	amspe	Optimal lag (yr)	Correlation coef	$\tau \pm 95\%$	F test significance
$B = \text{AR}(1)$	80–20	Y	2	0.451	—	—	680 ± 239	—
$B = \text{G2} + \text{AR}(1)$	80–20	Y	4	0.441	–200	–0.38	621 ± 211	58%
$B = \int \text{G2} + \text{AR}(1)$	80–20	Y	4	0.403	200	–0.59	431 ± 132	83%
$\frac{dB}{dt} = \frac{-B}{\tau} + \text{G2} + \text{AR}(0)$	80–20	Y	4	0.412	0	–0.29	867 ± 294	78%
$B = \text{AR}(1)$	50–20	N	2	0.495	—	—	610 ± 299	—
$B = \text{G2} + \text{AR}(1)$	50–20	N	4	0.480	800	0.30	571 ± 273	57%
$B = \int \text{G2} + \text{AR}(1)$	50–20	N	4	0.483	0	–0.29	567 ± 276	56%
$\frac{dB}{dt} = \frac{-B}{\tau} + \text{G2} + \text{AR}(0)$	50–20	N	4	0.464	0	–0.16	938 ± 286	65%
$B = \text{AR}(1)$	50–20	Y	2	0.527	—	—	548 ± 256	—
$B = \text{G2} + \text{AR}(1)$	50–20	Y	4	0.510	800	0.29	520 ± 239	58%
$B = \int \text{G2} + \text{AR}(1)$	50–20	Y	4	0.479	0	–0.47	432 ± 190	72%
$\frac{dB}{dt} = \frac{-B}{\tau} + \text{G2} + \text{AR}(0)$	50–20	Y	4	0.480	0	–0.19	902 ± 274	71%
$B = \text{AR}(1)$	10.2–0.2	N	2	0.775	—	—	118 ± 59	—
$B = \text{G2} + \text{AR}(1)$	10.2–0.2	N	4	0.740	1400	–0.23	118 ± 59	59%
$B = \int \text{G2} + \text{AR}(1)$	10.2–0.2	N	4	0.749	0	–0.25	117 ± 59	57%
$\frac{dB}{dt} = \frac{-B}{\tau} + \text{G2} + \text{AR}(0)$	10.2–0.2	N	4	0.726	1400	–0.31	232 ± 74	63%

shown in Table 3 show low correlations and inconsistent optimum lags, which suggest that not only is the integral relationship weaker after 50 ka but that it also may not hold in the Holocene. Note that both models 3 and 4 are essentially equally favored, meaning that we cannot distinguish between a picture in which the Byrd records reflects the integrated GISP2 signal in an essentially direct way, from one in which GISP2 is part of the general climate forcing being integrated by regional processes. It is also important to note that these models do not allow us to simply rule out AR(1) as the complete explanation for Byrd; this would indicate however, that the GISP2 record is a response to Byrd, plus additional North Atlantic climate variability that is not expressed at Byrd. [As already noted, Huybers' (2004) results slightly favor this interpretation.]

The adjusted mean square prediction error (amspe) in Table 3 gives a measure of the minimum relative magnitude of the white noise forcing required to explain Byrd for each model. All tests show that more than 65% of the variance remains unaccounted for, even for the best model (number 3 for 80–20 ka, with filtering). This “unexplained” variance may be thought of equivalently as that climate variability that is independent (i.e., unshared) between the two records.

6. Discussion

While supporting the contention that there is a significant connection between the Northern and Southern Hemispheres on millennial time scales, our results also

demonstrate that there is a considerable amount of variance in both the GISP2 and Byrd ice core records that cannot be explained by a simple lag/lead or integral relationship between the two. Possibly, the fraction of shared variance might increase if the cross dating of the records could be further improved, particularly for the higher-frequency events between 50 and 20 ka. However, an equally plausible conclusion is that the relatively small magnitude of shared variance simply reflects the dominance of regional climate variability in both records. Indeed, because our results show that AR(1) is a self-consistent explanation for Byrd, a reasonable picture of millennial-scale variability is one in which Antarctic climate is essentially random (and does not depend on climate variability elsewhere on the globe). In this view, climate in Greenland would be affected by Antarctic climate variability—as indicated by the significance tests for the linear regression models discussed above—but regional, North Atlantic processes would nevertheless dominate the Greenland climate records. Notably, whatever mechanisms determine the pacing of the 1500-yr “cycle,” and of the threshold behavior in the Greenland climate, would not have to be present in the Southern Hemisphere to be consistent with these results. Although half the variance in the Byrd record occurs at periodicities of 4000 yr and longer, the characteristic time scale of the physics driving this variability is 600 ± 300 yr. This result (reflecting a basic property of red noise processes) serves to highlight the general point that long time scale climate var-

iability can reflect rather short time scale physics. For much of the record, this time scale is sufficient, by itself, to explain the magnitude of apparent correlations with GISP2. An alternative possibility, of course, is that the shared variance in the records reflects the integrated response of Antarctic climate to the Greenland climate, rather than the other way around. While in this case Byrd would consist of more than simply an AR(1) process, the fraction of variance accounted for by regional processes would still exceed about 70% for millennial time scales. Thus, at a minimum each of these records reflects a substantial amount of independent variability. We find that while there is strong evidence for a connection between the two records during intervals around the largest of the rapid events in GISP2 between 80 and 50 kyr, the bulk of the millennial-scale variability cannot be demonstrated to have a common origin in both records. From the standpoint of Antarctic climate variability therefore, any influence from Greenland may be best thought of as being a component of the noise forcing, which is integrated by regional processes in the southern high latitudes.

For Greenland, the asymmetries and the strong 1500-yr spacing in the GISP2 record preclude any simple autoregressive process being a complete explanation of the data. Indeed, it is important to note that while we rejected simple stochastic resonance, our results do not contradict the essential idea in Alley et al. (2001a): that the rapid climate change events in GISP2 reflect a combination of some threshold mechanism, plus an underlying process that sets the time scale of variability, plus noise. On the other hand, our results do show that noise plus relaxation physics alone [i.e., an AR(1) model] explains a substantial amount of the variance in GISP2. One conceptual picture therefore, is that North Atlantic climate variability, as reflected in GISP2, results from physics equivalent to an AR(1) process, plus a threshold mechanism. The characteristic physical time scale that provides the best fit of AR(1) to GISP2 is 400 ± 200 yr. Notably, various modeling studies have shown that ocean thermohaline dynamics might account for the characteristic rapid warming/slow cooling of the GISP2 events (e.g., Weaver and Sarachik 1991; Winton and Sarachik 1993; Ganopolski and Rahmstorf 2001): a rapid turn-on (or shift) of North Atlantic Deep Water circulation occurs after a gradual diffusive warming of the deep ocean. While characteristic time scales for this process in these models depends on model parameters and in particular on diffusive mixing rates, reasonable values for these parameters give rise to adjustment time scales that are not inconsistent with our results (i.e., 400 ± 200 yr). This does not, of course, account for the less clear-cut comparison between AR(1) and stochastic resonance if only warming events are included, nor the finding of Rahmstorf (2003) that the phase memory in GISP2 is remarkably precise. The Greenland climate records, therefore, remain enigmatic and deserving of further research.

These results have important implications for the interpretation of individual paleoclimate events. Our results strongly demonstrate that Northern and Southern Hemisphere climate (as inferred from Greenland and Antarctic ice core records, assumed to be broadly representative of the high latitudes in each hemisphere) are not linked in a direct way, as implied for example by the global antiphase relationship implied by the widely used term “seesaw.” As Fig. 8a illustrates, there is no consistent lag or lead relationship at Byrd with the instances of rapid warmings in GISP2. On the other hand, a more complete picture of millennial-scale climate variability in which linkage between the hemispheres is significant, particularly for large-magnitude or long-duration events, but in which regional climate processes nevertheless dominate, is consistent with the observation that different records in different parts of the Southern Hemisphere have given such apparently contradictory results. A case in point is the ongoing controversy over whether the Younger Dryas identified in European pollen records and in the Greenland ice cores should be thought of as a global event. As inferred solely by comparing Byrd with GISP2 or GRIP, one might conclude that the Younger Dryas was preceded some 1000 yr by the Antarctic Cold Reversal (ACR; see our Fig. 1). Yet the timing of the ACR in Byrd appears to be somewhat later at that location than in East Antarctica (Jouzel et al. 2001), while results from the East Antarctic Taylor Dome ice core, in particular, suggest an ACR roughly in phase with the Younger Dryas (Steig et al. 1998). While the latter conclusion may not withstand further examination of the Taylor Dome record [it has been suggested that the Taylor Dome time scale needs to be adjusted to better match the appearance of other East Antarctic records; e.g., Mulvaney et al. (2000)] there has been similar controversy over the presence or absence of the Younger Dryas elsewhere in the Southern Hemisphere (and also in the Northern Hemisphere outside the North Atlantic). Our results suggest that, because regional processes have a large influence on millennial time scales, climate anomalies around the time of the Younger Dryas may vary significantly in magnitude, sign, and timing, even if this is a global-scale event. The same argument would apply to other millennial-scale events during the last glacial period.

Our results also have implications for the underlying physics in millennial-scale climate variability. In particular, one interpretation of the results from our third linear regression model (section 5b) is that of Stocker and Johnsen (2003), who suggest that the correct physical interpretation of the “integral of GISP2” is the Southern Ocean acting as a heat reservoir that integrates the Greenland climate signals with some characteristic time scale (which would determine the phase relationship—with Byrd lagging behind GISP2 by several hundred years as shown in our analyses). Certainly, this is consistent with evidence for significant decreases in the Atlantic thermohaline circulation during millennial-

scale cold intervals in Greenland, and is also consistent with general circulation models that tend to show a global response to such decreases. Furthermore, the model experiments tend to show changes that are large in the North Atlantic region, but only modest elsewhere (e.g., Manabe and Stouffer 1988; Seager et al. 2002). Such experiments also result in patchy, spatially heterogeneous climate variations in the Southern Hemisphere, consistent with the conclusion that regional processes should dominate individual climate records.

Our results would appear, then, to support the general idea that linkage between the Northern and Southern Hemispheres—when it is large enough to overwhelm regional variability—occurs through ocean dynamics. In contrast, we find our results more difficult to reconcile with recent suggestions that atmospheric teleconnections with the Tropics, rather than variability in ocean thermohaline circulation, play the dominant role in determining millennial-scale climate variability at high latitudes. Yin and Battisti (2001) and Clement et al. (2001), for example, have argued that during glacial periods, extratropical climates may be acutely sensitive to small changes in tropical sea surface temperatures, in part because enhanced atmospheric potential vorticity gradients in a glacial climate allow for more effective atmospheric teleconnections. In the language of an autoregression model, the suggestion would be that the high-latitude climate variability recorded in Byrd and GISP2 reflects separate integrations (i.e., reddening) of the same realization of tropical forcing. Experimenting with the best-fit autoregressive processes for each record, we find this simple picture would give rise to much higher, direct correlations between the records than are actually observed. This does not of course rule out some role for the Tropics, but it does illustrate the point that, to be consistent with the observations, some reddening mechanism is required, and we find the simple argument that this mechanism is the ocean to be compelling.

Finally, we note that the characteristic time scales for the AR(1) processes that describe the unshared variance in the records, themselves have implications for the physics of the climate system. This aspect of the records—sometimes referred to as the “background variance”—can too easily be dismissed as uninteresting. Yet because it actually accounts for most of the variance in the records, it may reveal important information on the dynamics of the system. In this context, it is significant that, for both Byrd and GISP2, the characteristic time scales of variability are quite different between glacial and interglacial climates. This is strongly suggestive of a profound change in the processes giving rise to climate variability on millennial time scales. The absence of the great continental-scale ice sheets of the Northern Hemisphere is an obvious candidate (e.g., Schmittner et al. 2002). Also possible is a fundamental change in the dynamics and sensitivity of the thermohaline circulation, as suggested for instance by Ganopolski and Rahmstorf (2001). An important point is that our results

do not contradict evidence for significant variability in some components of the climate systems on relatively long time scales during the Holocene. They do, however, cast substantial doubt on the idea that there is an approximately constant pacesetter that persists through both glacial and interglacial times (Alley et al. 2001a; Bond et al. 1997). Understanding the processes that determine the magnitude and characteristic time scales of this “background variability” is an important and substantial challenge then, which we suggest is equal in importance to understanding the rapid climate change events in the Greenland records. Given the absence of evidence for significant rapid climate change events in the last 8000 yr, we further speculate that it may also prove more relevant to understanding how the longer time scale components of the climate system interact to create variability in the modern, interglacial climate.

Acknowledgments. The authors are grateful to Chris Bretherton, David Battisti, Peter Huybers, Carl Wunsch, Richard Alley, and Ed Sarachik for insightful conversations; to two reviewers, both of whom pointed out the integral connection between the records; and to Andrew Weaver, the editor. EJS was supported in part by NSF, and GHR acknowledges support from the Joint Institute for the Study of the Atmosphere and Ocean, and the Program on Climate Change at the University of Washington.

REFERENCES

- Alley, R. B., S. Anandakrishnan, and P. Jung, 2001a: Stochastic resonance in the North Atlantic. *Paleoceanography*, **16**, 190–198.
- , —, —, and A. Clough, 2001b: Stochastic resonance in the North Atlantic: Further insights. *The Oceans and Rapid Climate Change: Past, Present and Future*, *Geophys. Monogr.*, No. 126, Amer. Geophys. Union, 57–68.
- , and Coauthors, 1993: Abrupt accumulation increase at the Younger Dryas termination in the GISP2 ice core. *Nature*, **362**, 527–529.
- , and Coauthors, 2002: *Abrupt Climate Change: Inevitable Changes*. National Academy Press, 230 pp.
- Battisti, D. S., 1988: Dynamics and thermodynamics of a warming event in a coupled tropical atmosphere–ocean model. *J. Atmos. Sci.*, **45**, 2889–2919.
- Blunier, T., and E. J. Brook, 2001: Timing of millennial-scale climate change in Antarctica and Greenland during the last glacial period. *Science*, **291**, 109–112.
- , and Coauthors, 1998: Asynchrony of Antarctic and Greenland climate change during the last glacial period. *Nature*, **394**, 739–743.
- Bond, G. C., and R. Lotti, 1995: Iceberg discharges into the North Atlantic on millennial time scales during the last deglaciation. *Science*, **267**, 1005–1010.
- , and Coauthors, 1997: A pervasive millennial-scale cycle in North Atlantic Holocene and Glacial climates. *Science*, **278**, 1257–1266.
- Bradley, R. S., 1999: *Paleoclimatology: Reconstructing Climates of the Quaternary*. Academic Press, 613 pp.
- Broecker, W. S., 1991: The great ocean conveyor. *Oceanography*, **4**, 79–89.
- , 1997: Paleocene circulation during the last deglaciation; a bipolar seesaw? *Paleoceanography*, **13**, 119–121.

- , 2000: Was a change in thermohaline circulation responsible for the Little Ice Age? *Proc. Natl. Acad. Sci.*, **97**, 1339–1342.
- Clark, P. U., R. S. Webb, and L. D. Keigwin, Eds., 1999: *Mechanisms of Millennial Scale Global Climate Change*. *Geophys. Monogr.*, No. 112, Amer. Geophys. Union, 394 pp.
- Clement, A. C., M. A. Cane, and R. Seager, 2001: An orbitally driven tropical source for abrupt climate change. *J. Climate*, **14**, 2369–2375.
- Conway, H., G. Catania, C. F. Raymond, A. M. Gades, T. A. Scambos, and H. Englehardt, 2002: Switch of flow direction in an Antarctic ice stream. *Nature*, **419**, 465–467.
- Crowley, T. J., 2000: Causes of climate change over the past 1000 years. *Science*, **289**, 270–277.
- Cuffey, K., G. Clow, R. Alley, M. Stuiver, E. Waddington, and R. Saltus, 1995: Large Arctic temperature change at the Wisconsin-Holocene glacial transition. *Science*, **270**, 455–458.
- Dahl-Jensen, D., K. Mosegaard, N. Gundestrup, G. D. Clow, S. J. Johnsen, A. W. Hansen, and N. Balling, 1998: Past temperatures directly from the Greenland ice sheet. *Science*, **282**, 268–271.
- Dansgaard, W., S. J. Johnsen, H. B. Clausen, and N. Gundestrup, 1973: Stable isotope glaciology. *Medd. Groenl.*, **197**, 1–53.
- , and Coauthors, 1993: Evidence for general instability of past climate from a 250-kyr ice-core record. *Nature*, **364**, 218–220.
- Denton, G. H., C. J. Heusser, T. V. Lowell, P. I. Moreno, B. G. Andersen, L. A. Heusser, C. Schluchter, and D. R. Marchant, 1999: Interhemispheric linkage of paleoclimate during the last glaciation. *Geogr. Ann.*, **81A**, 107–153.
- Fahnestock, M. A., T. A. Scambos, R. A. Binschadler, and G. Kvaran, 2000: A millennium of variable ice-flow recorded by the Ross Ice Shelf. *J. Glaciol.*, **46**, 652–664.
- Ganopolski, A., and S. Rahmstorf, 2001: Rapid changes of glacial climate simulated in a coupled climate model. *Nature*, **409**, 153–158.
- Gedalof, Z., N. J. Mantua, and D. L. Peterson, 2002: A multi-century perspective of variability in the Pacific Decadal Oscillation: New insights from tree rings and corals. *Geophys. Res. Lett.*, **29**, 2204, doi:10.1029/2002GL015824.
- Grootes, P. M., M. Stuiver, J. W. C. White, S. Johnsen, and J. Jouzel, 1993: Comparison of oxygen isotope records from the GISP2 and GRIP Greenland ice cores. *Nature*, **366**, 522–554.
- Hammer, C. U., H. B. Clausen, and C. C. Langway, 1994: Electrical conductivity method (ECM) stratigraphic dating of the Byrd Station ice core, Antarctica. *Ann. Glaciol.*, **20**, 115–120.
- Hasselmann, K., 1976: Stochastic climate models. Pt. 1: Theory. *Tellus*, **28**, 473–485.
- Heinrich, H., 1988: Origin and consequences of cyclic ice rafting in the Northeast Atlantic Ocean during the past 130,000 years. *Quat. Res.*, **29**, 142–152.
- Hendy, I. L., and J. P. Kennett, 1999: Latest Quaternary North Pacific surface-water responses imply atmosphere-driven climate instability. *Geology*, **27**, 291–294.
- Hinnov, L. A., M. Schulzb, and P. Yiouc, 2002: Interhemispheric space-time attributes of the Dansgaard-Oeschger events between 100 and 0 ka. *Quat. Sci. Rev.*, **21**, 1213–1228.
- Hoerling, M. P., J. W. Hurrell, and T. Xu, 2001: Tropical origins for recent North Atlantic climate change. *Science*, **292**, 90–92.
- Hormes, A., B. U. Müller, and C. Schlüchter, 2001: The Alps with little ice: Evidence for eight Holocene phases of reduced glacier extent in the Central Swiss Alps. *Holocene*, **11**, 255–265.
- Hughen, K. A., J. T. Overpeck, S. J. Lehman, M. Kahsarian, J. Southon, L. C. Peterson, R. Alley, and D. M. Sigman, 1998: Deglacial changes in ocean circulation from an extended radiocarbon calibration. *Nature*, **391**, 65–68.
- Huybers, P., 2004: Comments on “coupling of the hemispheres in observations and simulations of glacial climate change” by A. Schmittner, O. A. Senko, and A. J. Weaver. *Quat. Sci. Rev.*, **22**, 659–671.
- Hyde, W. T., and T. J. Crowley, 2002: Stochastic forcing of Pleistocene ice sheets: Implications for the origin of millennial-scale climate oscillations. *Paleoceanography*, **17**, 19.1–19.8.
- Jenkins, G. M., and D. G. Watts, 1968: *Spectral Analysis and Its Applications*. Holden-Day, 523 pp.
- Jøhannesson, T., C. F. Raymond, and E. D. Waddington, 1989: Time scale for adjustment of glaciers to changes in mass balances. *J. Glaciol.*, **35**, 355–369.
- Johnsen, S. J., D. Dahl-Jensen, W. Dansgaard, and N. Gundestrup, 1995: Greenland paleotemperatures derived from the GRIP bore hole temperature and ice core isotope profiles. *Tellus*, **47B**, 624–629.
- , and Coauthors, 2001: Oxygen isotope and palaeotemperature records from six Greenland ice-core stations: Camp Century, Dye-3, GRIP, GISP2, Renland and North GRIP. *J. Quat. Sci.*, **16**, 299–307.
- Jouzel, J., and Coauthors, 1997: Validity of the temperature reconstruction from water isotopes in ice cores. *J. Geophys. Res.*, **102**, 26 471–26 487.
- , and Coauthors, 2001: A new 27 ky high resolution East Antarctic climate record. *Geophys. Res. Lett.*, **28**, 3199–3202.
- Lang, C., M. Leuenberger, J. Schwander, and S. Johnsen, 1999: 16°C rapid temperature variation in central Greenland 70,000 years ago. *Science*, **286**, 934–937.
- MacAyeal, D. R., 1993a: A low-order model of the Heinrich event cycle. *Paleoceanography*, **8**, 767–773.
- , 1993b: Binge/purge oscillations of the Laurentide ice sheet as a cause of the North Atlantic’s Heinrich events. *Paleoceanography*, **8**, 775–784.
- Manabe, S., and R. J. Stouffer, 1988: Two stable equilibria of a coupled ocean–atmosphere model. *J. Climate*, **1**, 841–866.
- Mann, M. E., R. S. Bradley, and M. K. Hughes, 1998: Global-scale temperature patterns and climate forcing over the past six centuries. *Nature*, **392**, 779–787.
- Markgraf, V., 1993: Younger Dryas in southernmost South America: An update. *Quat. Sci. Rev.*, **12**, 351–355.
- Mayewski, P. A., L. D. Meeker, S. Whitlow, M. S. Twickler, M. C. Morrison, R. B. Alley, P. Bloomfield, and K. Taylor, 1993: The atmosphere during the Younger Dryas. *Science*, **261**, 195–197.
- Meece, D. A., and Coauthors, 1997: The Greenland Ice Sheet Project 2 depth-age scale: Methods and results. *J. Geophys. Res.*, **102**, 26 411–26 423.
- Morgan, V., M. Delmotte, T. van Ommen, J. Jouzel, J. Chappellaz, S. Woon, V. Masson-Delmotte, and D. Raynaud, 2002: Relative timing of deglacial climate events in Antarctica and Greenland. *Science*, **297**, 1862–1864.
- Mudelsee, M., 2002: TAUEST: A computer program for estimating persistence in unevenly spaced weather/climate time series. *Comput. Geosci.*, **28**, 69–72.
- Mulvaney, R., R. Rothlisberger, E. W. Wolff, S. Sommer, J. Schwander, M. A. Hutterli, and J. Jouzel, 2000: The transition from the last glacial period in inland and near-coastal Antarctica. *Geophys. Res. Lett.*, **27**, 2673–2676.
- Overpeck, J. T., and Coauthors, 1997: Arctic environmental change of the last four centuries. *Science*, **278**, 1251–1256.
- Paterson, W. S. B., 1994: *The Physics of Glaciers*. Elsevier, 480 pp.
- Percival, D. B., J. E. Overland, and H. O. Mofjeld, 2001: Interpretation of North Pacific variability as a short- and long-memory process. *J. Climate*, **14**, 4545–4559.
- Petit, J. R., and Coauthors, 1999: Climate and atmospheric history of the past 420,000 years from the Vostok ice core, Antarctica. *Nature*, **399**, 429–436.
- Press, W. H., B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, 1992: *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, 933 pp.
- Rahmstorf, S., 2003: Timing of abrupt climate change: A precise clock. *Geophys. Res. Lett.*, **30**, 1510, doi:10.1029/2003GL017115.
- Sachs, J. P., and S. J. Lehman, 1999: Subtropical North Atlantic temperatures 60,000 to 30,000 years ago. *Science*, **286**, 756–759.
- Schiller, A., U. Mikolajewicz, and R. Voss, 1997: The stability of the North Atlantic thermohaline circulation in a coupled ocean–

- atmosphere general circulation model. *Climate Dyn.*, **13**, 325–347.
- Schmittner, A., M. Yoshimori, and A. J. Weaver, 2002: Instability of glacial climate in a model of the ocean–atmosphere–cryosphere system. *Science*, **295**, 1489–1493.
- , O. A. Saenko, and A. J. Weaver, 2003: Coupling of the hemispheres in observations and simulations of glacial climate change. *Quat. Sci. Rev.*, **22**, 659–671.
- Schneider, T., and A. Neumaier, 2001: Algorithm 808: ARFIT—A Matlab package for the estimation of parameters and eigenmodes of multivariate autoregressive models. *ACM Trans. Math. Software*, **27**, 58–65.
- Schulz, M., 2002a: On the 1470-year pacing of Dansgaard–Oeschger warm events. *Paleoceanography*, **17**, 4.1–4.9.
- , 2002b: The tempo of climate change during Dansgaard–Oeschger interstadials and its potential to affect the manifestation of the 1470-year climate cycle. *Geophys. Res. Lett.*, **29**, 1002, doi:10.1029/2001GL013277.
- Schwarz, G., 1978: Estimating the order of a model. *Ann. Stat.*, **6**, 461–464.
- Seager, R., D. S. Battisti, J. Yin, N. Naik, N. Gordon, A. C. Clement, and M. Cane, 2002: Is the Gulf Stream responsible for Europe’s mild winters? *Quart. J. Roy. Meteor. Soc.*, **128**, 2563–2584.
- Seidov, D., and M. Maslin, 2001: Atlantic Ocean heat piracy and the bipolar climate see-saw during Heinrich and Dansgaard–Oeschger events. *J. Quat. Sci.*, **16**, 321–328.
- Severinghaus, J. P., and E. J. Brook, 1999: Abrupt climate change at the end of the last glacial period inferred from trapped air in polar ice. *Science*, **286**, 930–934.
- Steig, E. J., and R. B. Alley, 2002: Phase relationships between Antarctic and Greenland climate records. *Ann. Glaciol.*, **35**, 451–456.
- , and Coauthors, 1998: Synchronous climate changes in Antarctic and the North Atlantic. *Science*, **282**, 92–95.
- , J. L. Fastook, C. Zweck, I. D. Goodwin, K. J. Licht, J. W. C. White, and R. P. Ackert Jr., 2000: West Antarctic Ice Sheet elevation changes. *Antarct. Res. Ser.*, **77**, 75–90.
- Stocker, T. F., 1998: The seesaw effect. *Science*, **282**, 61–62.
- , and S. J. Johnsen, 2003: A minimum model for the bipolar see-saw. *Paleoceanography*, **18**, 1087, doi:10.1029/2003PA000920.
- Stuiver, M., P. M. Grootes, and T. F. Braziunas, 1995: The GISP2 $\delta^{18}\text{O}$ climate record of the past 16,500 years and the role of the sun, ocean, and volcanoes. *Quat. Res.*, **44**, 341–354.
- Taylor, K. C., G. W. Lamorey, G. A. Doyle, R. B. Alley, P. M. Grootes, P. A. Mayewski, J. W. C. White, and L. K. Barlow, 1993: The “flickering switch” of late Pleistocene climate change. *Nature*, **361**, 432–436.
- von Storch, H., and F. W. Zwiers, 1999: *Statistical Analysis in Climate Research*. Cambridge University Press, 484 pp.
- Weaver, A. J., and E. S. Sarachik, 1991: The role of mixed boundary conditions in numerical models of the ocean’s climate. *J. Phys. Oceanogr.*, **21**, 1470–1493.
- , O. A. Saenko, P. U. Clark, and J. X. Mitrovica, 2003: Meltwater Pulse 1A from Antarctica as a trigger of the Bølling–Allerød warm interval. *Science*, **299**, 1709–1713.
- Winton, M., and E. S. Sarachik, 1993: Thermohaline Oscillations induced by strong steady salinity forcing of ocean general circulation models. *J. Phys. Oceanogr.*, **23**, 1389–1409.
- Wunsch, C., 2001: On sharp spectral lines in the climate record and the millennial peak. *Paleoceanography*, **15**, 417–424.
- , 2003: Greenland–Antarctic phase relationships and millennial time-scale climate fluctuations in the Greenland cores. *Quat. Sci. Rev.*, **22**, 1631–1646.
- Yin, J. H., and D. S. Battisti, 2001: The importance of tropical sea surface temperature patterns in simulations of Last Glacial Maximum climate. *J. Climate*, **14**, 565–581.