Boundary impulse response functions in a century-long eddying global ocean simulation

Mathew Maltrud · Frank Bryan · Synte Peacock

Abstract   Results are presented from a century-long 1/10° global ocean simulation that included a suite of age-related passive tracers. In particular, an ensemble of five global Boundary Impulse Response functions (BIRs, which are statistically related to the more fundamental Transit Time Distributions, TTDs) was included to quantify the character of the TTD when mesoscale eddies are explicitly simulated rather than parameterized. We also seek to characterize the level of variability in water mass ventilation timescales arising from eddy motions. The statistics of the BIR timeseries are described, and it is shown that the greatest variability occurs at early times, followed by a remarkable conformity between ensemble members at longer timescales. The statistics of the first moment of the BIRs are presented, and the upper-ocean spatial distribution of the standard deviation of the first moment of the BIRs discussed. It is shown that variations in the BIR first moment with respect to the ensemble average are typically only a few percent, and that the variability slightly decreases with increasing ensemble size, implying that only a few ensemble members may be necessary for a reasonable estimate of the TTD. The completeness of the estimated TTD, i.e., the degree to which the century long BIRs capture the range of global ocean ventilation timescales is discussed, and the potential for extrapolation of the BIR to longer times is briefly explored. Several regional BIRs were also simulated in order to quantify the relative abundance of fluid parcels that originate in specific geographical locations.

Keywords   Modeling · Ocean circulation · Tracers · Timescales

M. Maltrud (✉)
Los Alamos National Laboratory, Los Alamos, NM 87545, USA
e-mail: maltrud@lanl.gov

F. Bryan · S. Peacock
National Center for Atmospheric Research, Boulder, CO 80307, USA

1 Introduction

Analysis of temperature and salinity distributions to characterize water mass properties, their formation rates, and pathways and timescales of transport in the ocean interior, has been at the foundation of observational oceanography for over a century. The availability of global scale observations of large numbers of natural and anthropogenic tracers during the last several decades has provided a more quantitative basis for characterizing the timescales associated with ocean transport and mixing. One particular timescale, the time since a water parcel was last in contact with the atmosphere (often referred to as "water-mass age") is of special interest in climate research. Tracer based estimates of water-mass age have been used to directly infer timescales of ocean ventilation, which gives bounds on the rate at which the ocean can sequester materials such as carbon dioxide.

In practice, however, various tracers yield different age estimates due to their distinct atmospheric time histories, their particular chemical properties, and oceanic sources and sinks. These differences can be understood through the conceptual framework of Transit Time Distributions (TTDs) (also called "age spectra" and "transit time probability density functions" [10,26]). The TTD is a probability distribution of age, which explicitly accounts for the fact that a water mass is comprised of a mixture of components with different ages, each with different advective and diffusive histories. Holzer and Hall [10] clarify the mathematical connection between this probability distribution function and the Greens function for the advection–diffusion equation governing tracer transport.

Idealized age tracers, and more recently TTDs, have been carried in simulations of the ocean ([4,11,12,14,25] and [24] (hereafter referred to as PM06)). Calculation of TTDs from ocean models with resolution typical of that used in climate research (horizontal grid spacing of order 1° latitude and longitude) is relatively straightforward, and achieved by exploiting the connection to the Greens function for the transport equation. The TTD in coarse resolution models can be obtained by a single forward integration of the model transport equation for a conservative tracer with an impulse boundary condition. This is possible because models in this resolution regime have essentially steady flow (excepting the seasonal cycle) and the transport by turbulent eddies is parameterized through some form of diffusive process. As we shall see, the simulation of a TTD-like tracer becomes much more complex when the model resolution becomes high enough to permit a fully turbulent, non-steady flow.

Both observational and modeling studies suggest that the ocean TTD typically has a long tail, an indication that mixing processes are important in determining the mean ventilation rate. Thus, the model estimates of the TTD from coarse-resolution models are critically dependent on the parameterization of turbulent mixing. However, it is not clear from coarse-resolution model simulations alone how to select an optimal mixing parameterization. Since the majority of the horizontal mixing in the ocean is due to turbulence from mesoscale eddies, directly assessing the effects of these motions requires a model that can resolve them. Here we describe the first ever simulation designed to estimate the ocean TTD using a global general circulation model in which mesoscale eddies are simulated directly rather than parameterized.

Computing the TTD from an unsteady turbulent circulation is complicated by the fact that the transit times themselves become time-dependent. The result of a single forward integration of the transport equation with an impulse boundary condition, formally referred to as the Boundary Impulse Response (BIR), can no longer be interpreted as a true probability distribution function of transit times, as is the case with steady flow. However, Haine et al. [8] (hereafter referred to as HZWH08) have shown that even for unsteady flow, the BIR and TTD share the same statistical moments under certain conditions. We exploit this property

to estimate the TTD as the average of a (modest sized) ensemble of BIRs, which are directly simulated in a global eddying simulation.

The high computational cost of running a global eddy resolving simulation precludes an integration of the several thousand years required to capture the full range of ocean ventilation time scales (PM06). We will present results from a century-scale simulation, which provides a useful basis for studying ocean ventilation processes in the more rapidly ventilated parts of the ocean. These are just the regions that are most relevant with respect to anthropogenic impacts on the ocean and the uptake of transient tracers such as CO_2 and Chlorofluorocarbons.

This paper begins with a brief overview of the TTD theory; the model configuration and experiment design is described, and the simulated BIRs are presented. There is a discussion of both the ensemble statistics of the BIR timeseries, the statistics of the first moment, and how these relate to the TTD. The issue of BIR convergence is also explored. Finally, preliminary results from a set of BIRs which are forced by boundary impulses in specific geographical regions are presented.

2 Boundary propagators and TTDs

The theory of TTDs has previously been developed and applied to oceanic flows [9,7], so only the most basic development will be presented here. Starting with the equation for the evolution of the concentration of a conservative passive tracer, C,

$$\left(\frac{\partial}{\partial t} + L(\mathbf{x}, t)\right) C(\mathbf{x}, t) = 0 \tag{1}$$

where $L(\mathbf{x}, t)$ represents advection and diffusion, an equation can be derived for the quantity $G'(\mathbf{x}, t)$,

$$\left(\frac{\partial}{\partial t} + L(\mathbf{x}, t)\right) G'(\mathbf{x}, t; \mathbf{x}_0, t') = 0 \tag{2}$$

which satisfies the convolution integral,

$$C(\mathbf{x}, t) = \int_0^t dt' \int_\Omega d^2\mathbf{x}_0 G'(\mathbf{x}, t; \mathbf{x}_0, t') C(\mathbf{x}_0, t') \tag{3}$$

where $\mathbf{x}_0$ is a two dimensional vector defined over an ocean surface region Ω. Thus, the boundary propagator, G', "propagates" the surface value of C from Ω into the interior of the ocean. If the convolution integral (Eq. 3) is rewritten as a function of $\tau = t - t'$ (the "transit time"), the kernel has a clear interpretation as a weighting based on the last time a water mass component was at the ocean surface.

Note that G' is seven dimensional, being a function of all locations in the ocean ($\mathbf{x}$), the "field time" (t), all locations at the surface ($\mathbf{x}_0$), and the "source time" (t'), so is typically intractable to calculate and analyze in simulated oceanic flows. A forward model simulation would need to carry a different tracer for each source time and surface location, which would need to be an extremely large number (order 1,000 even for a small number of surface locations) in order to accurately approximate the TTD. For the case of steady flow, the boundary propagator is invariant under time translation: $G'(\mathbf{x}, t; \mathbf{x}_0, t - \tau) = G'(\mathbf{x}, 0; \mathbf{x}_0, -\tau)$, and the TTD depends only on transit time, not the field time and source time individually, so a single forward integration from source time $t' = 0$ is sufficient to compute the TTD. For the case of

unsteady flow, the time translation invariance no longer applies. However, HZWH08 provide a proof that for statistically stationary flow, the distribution of $G'(\mathbf{x}, t_1; \mathbf{x}_0, t_1 - \tau)$ (the TTD) over t_1 is identical to the distribution of $G'(\mathbf{x}, t_2 + \tau; \mathbf{x}_0, t_2)$ (the BIR) over t_2. If we are primarily interested in the TTD as a statistical characterization of the turbulent flow rather than as a representation of a specific realization of the flow, as we are in this study, then we can estimate the mean TTD as an average of an ensemble of BIRs integrated forward from a set of source times t_2. This relationship between the TTD and BIR defines the computational approach used in our study: we simulate an ensemble of BIRs by forward integration in an eddying ocean model subject to a repeating annual cycle forcing, and estimate the TTD as their ensemble mean.

3 Model description

3.1 Model configuration

The fully global simulations described here were performed using the Parallel Ocean Program (POP, [3]) developed at Los Alamos National Laboratory. The configuration is based upon that described by Maltrud and McClean [18], hereafter referred to as MM05, but with several notable refinements. The horizontal grid retained the nominal longitudinal resolution of $0.1°$, but the grid design has been changed from a dipole, where the northern computational pole was smoothly displaced into Canada, to a tripole layout [23] with poles in Canada and Russia, resulting in a more isotropic distribution of resolution in high northern latitudes. The vertical resolution is the same as in MM05 with two additional deep levels that extend the maximum depth of the model from 5,500 to 6,000 m. The discretization of bottom topography has been changed from full-cell to partial bottom cells [1], creating a more accurate and smoother representation of topographic slopes.

Biharmonic viscosity and diffusivity of tracers are used with a cubic dependence on local grid size as described in Maltrud et al. [19]. In this simulation we have reduced the level of explicit dissipation by a factor of three relative to MM05, using equatorial values of $\nu_0 = -90 \, \text{m}^4/\text{s}$ for momentum and $\kappa_0 = -30 \, \text{m}^4/\text{s}$ for tracers. As in MM05, vertical mixing coefficients for momentum and tracers are calculated by an implementation of the K-Profile Parameterization [16] that includes the use of large diffusion coefficients ($0.1 \, \text{m}^2 \, \text{s}^{-1}$) to resolve gravitational instabilities. Background values for vertical tracer diffusion range from $10^{-5} \, \text{m}^2/\text{s}$ near the surface to $10^{-4} \, \text{m}^2/\text{s}$ at depth, with viscosity values an order of magnitude higher. Density is calculated from the equation of state for seawater derived by McDougall et al. [22]. Initial temperature and salinity fields were interpolated from the annual mean WOCE Global Hydrographic Climatology [6].

When simulating passive tracer concentrations, it is especially important that numerical dispersion errors arising from the discretization of advection don't result in spurious negative values. The advection scheme used in these experiments is similar to the third order direct space time with flux limiting scheme of Adcroft et al. [2]. The primary differences are that we replace the underlying third order scheme with the second order Lax–Wendroff scheme, to reduce numerical diffusion, and modify the flux limiters to take spatial variations in velocity and grid spacing into account. However, we found this scheme to be undesirably diffusive when applied to the active tracers, so standard centered advection (with no flux limiting) was used for temperature and salinity, as well as for momentum.

Since one of the fundamental aspects of this simulation was to quantify the eddy component of tracer transport, and the experimental design is predicated on a statistically

stationary flow field, we chose to restrict the timescales introduced by the surface forc-
ing. The atmospheric state was based on the repeat annual cycle (normal-year) Coordinated
Ocean Reference Experiment (CORE) forcing dataset [17], with the 6-hourly forcing aver-
aged to monthly. Wind stress was calculated offline using bulk formulae and the Hurrell Sea
Surface Temperature (SST) climatology [13,15]; evaporation and sensible heat flux were
calculated online using the same bulk formulae and the model predicted SST. Precipitation
was also taken from the CORE forcing dataset. Monthly river runoff from 46 major rivers [5]
was added to the fresh water flux at the locations of the actual outflow, with the remaining
ungauged runoff distributed evenly along the coasts of all of the continents. Ice cover was
prescribed based on the -1.8 C isoline of the SST climatology, with both temperature and
salinity restored on a timescale of 30 days under diagnosed/climatological ice. There is also
weak surface restoring of salinity in the open ocean with a timescale of 1 year to prevent
global salinity drift.

3.2 BIR implementation and simulation design

In theory, like the TTD, the BIR is the response to Delta function source. In practice, the
implementation of the BIR was achieved using a 1 year long top hat function in order to
average over the seasonal cycle (the same method was employed by Khatiwala et al. [14] and
PM06 in their TTD simulations). This is accomplished by fixing the top layer value of the
tracer to (1/year) for 1 year, then resetting to zero at all later times. As a result, the surface
becomes a sink for the BIR tracers after the first year.

Because TTDs and BIRs share the same statistics, an ensemble of BIR realizations is
necessary to estimate the TTD. Instead of running the model several times with a single BIR
tracer, each member of the ensemble was run within the same ocean simulation, but the start
times for the first four surface impulses were separated by 2 years. The fifth ensemble member
was added after an interval of 30 years. The decorrelation time for oceanic eddies is much
smaller than 2 years, so each BIR ensemble member can be considered to be transported by
an independent realization of the eddy flow field.

Figure 1 shows a schematic of the simulation design. The model was spun up from rest
with no passive tracers for 15 years, which is long enough for the upper ocean circulation to
reach a statistically steady state. At this point, tracers were added at the appropriate times
and the run continued for another 105 years. An ensemble of five realizations of "global"
BIRs (i.e., where the source region Ω is the entire global ocean surface) were simulated. The
majority of the presented results will focus on these global BIRs. A single realization of six
"regional" BIRs (where Ω is not the entire ocean surface, but a subset thereof) were included
to provide additional information on the relative contributions of different surface regions
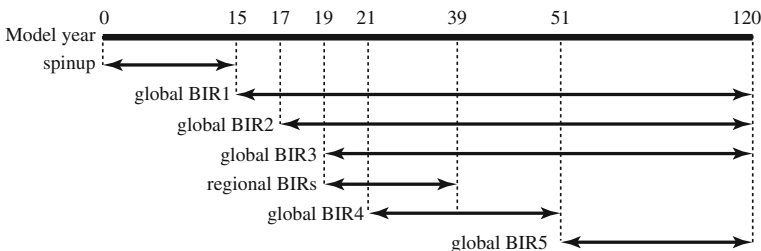

Fig. 1 Schematic of the simulation design showing the starting time and duration of each BIR

in ocean ventilation. Just as with the global BIRs, the surface layer in each defined region is reset to a value of (1/year) every timestep for a year. Outside of each region, the surface value is reset to zero for all times.

As can be seen in Fig. 1, each of the five global BIRs span a different amount of time, but there are two clear choices of record length when performing ensemble statistics. For long time scales, we will consider the first 100 years of BIR1, BIR2, and BIR3, where time is defined relative to January 1 of the year in which the particular BIR was introduced into the simulation. For shorter time scales, the ensemble consists of the first 30 years of all five BIRs.

4 Results

The results shown in this paper will focus on the characteristics of the simulated BIRs. Detailed analysis of the model circulation can be found in Maltrud et al. [20], where a wide array of metrics are applied to demonstrate a high degree of model fidelity when compared to data, theory, and previous high-resolution ocean simulations. As discussed by Maltrud et al. [20], the simulated geographical distribution and amplitude of eddy-activity shows a very close correspondence with observations, and there are many features of the simulated circulation that show a significant improvement over those seen in this model's predecessor (MM05, [21]).

This section begins with a general overview of the global BIR distribution; we then consider the statistics of the timeseries of the ensemble BIR members, and the statistics of the first moment of the BIRs. Finally, we briefly show some results from a set of regional BIR simulations.

4.1 BIR evolution

One way of looking at the BIRs is as a dye tracer with an impulse surface boundary condition. Figure 2 shows snapshots of two BIR realizations in the North and South Atlantic (left and right panels, respectively). In the North Atlantic snapshot, a tongue of high tracer values (red colors) can clearly be seen being advected around Flemish Cap, superimposed on a very turbulent background. The differences in tracer values between the two BIR realizations shown are clearly discernible by eye. In both realizations, very strong tracer gradients are evident across very small spatial scales. Much of this high-frequency variability is smoothed out in the lowermost panel (Fig. 2e), which shows the five-member ensemble average of the fourth year annual mean BIRs. The differences between the two realizations of the tracer in the South Atlantic are not quite so pronounced as in the North Atlantic; however, individual eddies can still clearly be seen as Agulhas rings and turbulence associated with Antarctic Circumpolar Current (ACC) fronts. Again, the ensemble-mean (Fig. 2f) shows a very smoothed version of the snapshots, and all semblance of individual eddies have been erased by the averaging.

The redistribution of the BIRs as the simulation progresses can be seen in the column integrals in Fig. 3. After only 10 years, the Atlantic Deep Western Boundary Current (DWBC) has transported the surface signal from the high-latitude North Atlantic into the Southern Hemisphere. By contrast, in the coarse resolution simulation of PM06 (their figure 1), the net mass transport of the DWBC is fairly realistic, but the model was not able to resolve the narrow width and high speeds of the core of the boundary current. As a result, several decades were required for the DWBC tracer signal to reach the Southern hemisphere in the coarse-resolution simulation.
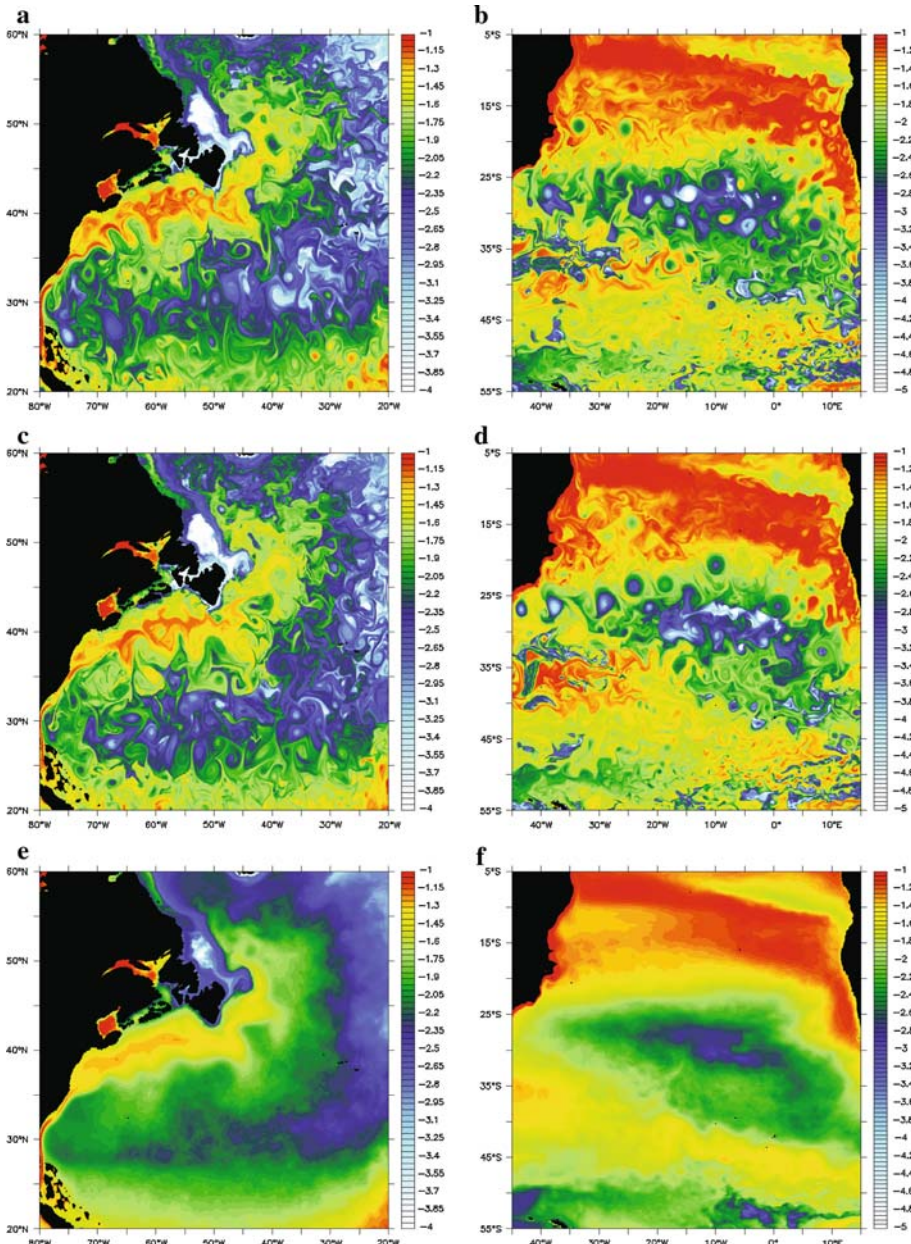
Fig. 2 Snapshots of log(BIR1) (**a**, **b**) and log(BIR2) (**c**, **d**) at 112 m depth on July 1 of year 4 in the Gulf Stream (*left*) and Southern Atlantic (*right*) regions. **e**, **f** Show the five member ensemble average of the year 4 annual mean BIRs for the same regions

After 50 years of model integration (Fig. 3b), some of the tracer transported southward by the Atlantic DWBC has entered the ACC as Circumpolar Deep Water (dark blue colors in the South Atlantic), and some Antarctic Bottom Water (AABW) has made its way into the North Pacific basin (dark blue colors in the Western Pacific). Notable signals of dye are
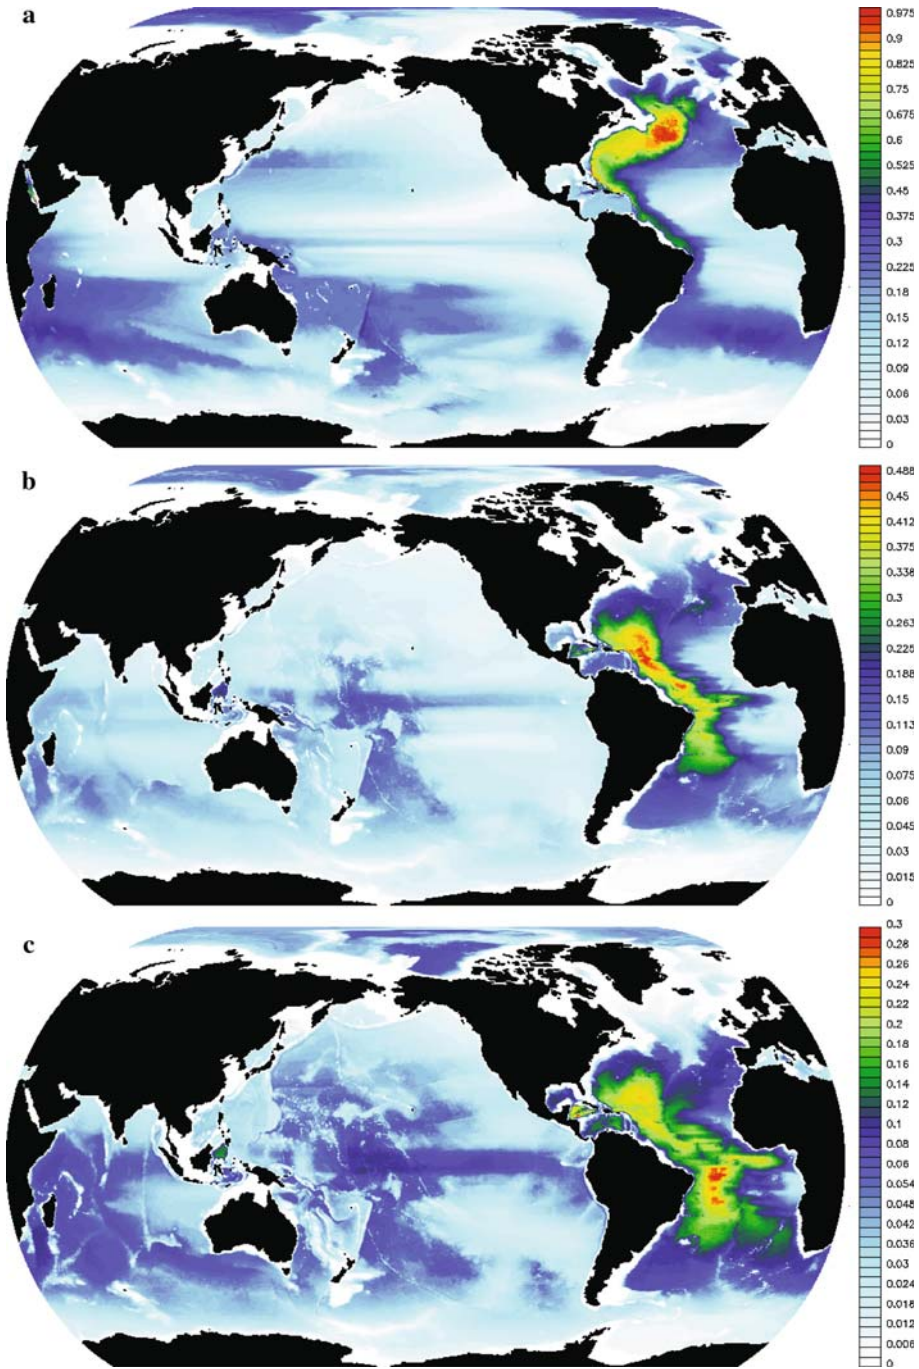
Fig. 3 Vertical integrals (weighted by the grid cell area) of the ensemble average annual mean global BIR for years 10 (five ensemble members), 50 (four ensemble members), and 100 (three ensemble members). Note the change in color scales between panels

also seen along the equator in the central and eastern Pacific, this being a reflection of rapid advection along shallow zonal currents in the equatorial Pacific basin.

At a time of 100 years (Fig. 3c), the highest column inventories are in the western half of the Atlantic, and AABW has spread through much of the abyss in both the Pacific and Indian basins. In PM06 (their figure 1), the DWBC signal has diffused across the entire North Atlantic basin, and the AABW has only begun to cross the equator by this time. It should be noted that it is not possible to distinguish from this figure alone between tracer added by thermocline ventilation and that added by deep ventilation; the inferences above were made by analyzing the distribution of the BIR on individual depth layers.

As described above, HZWH08 prove that TTDs and BIRs share the same statistics. An interesting question is thus how different the individual BIR timeseries are from their ensemble averages. Figure 4 shows the time dependent behavior of the simulated BIRs at several locations for each ensemble member (colored curves) and the ensemble mean (thick black curve). All show the greatest variability at early times, followed by remarkable conformity for times later than the peak in the distribution, (i.e. for the most probable transit time). At the Indian Ocean site, the response is dominated by the mixed layer processes related to the annual cycle (though there is significant eddy variability originating in the Leeuwin Current). In the northeast Pacific thermocline, eddies and differences in arrival time dominate the variability. The Southern Ocean location is exceptional in that it is bimodal, with an early signal from bottom water formed in the Ross Sea, followed by other Southern Ocean bottom water sources (a similar bimodal distribution was also reported in PM06 for most deep Southern Ocean locations). The most notable outlier in Fig. 4 is BIR5 (magenta curve) at the Arctic location. Recall that BIR5 was introduced into the simulation 30 years after BIR4, so this behaviour reflects a significant shift in the Arctic circulation during those three decades. This demonstrates that the flow cannot be considered stationary (at decadal time scales), thus the ensemble averaged BIRs are not a reliable estimate of the TTD in this particular region.

In addition to results from this simulation, a time series of the global BIR from PM06 can be seen in Fig. 4c. As noted above in reference to Fig. 3a, the signal arrives much later and has a broader leading edge in the low resolution (order 3°) case due to sluggish transport velocities and stronger diffusion. As one moves southward in the DWBC, the ratio of the coarse to fine grid arrival time remains fairly constant around a value of 5. For example, at 10°S, the BIRs from the current simulation arrive after about 10 years, compared to 50 years for the PM06 run (not shown).

4.2 Convergence to complete TTD

Because ocean ventilation times are thousands of years for much of the deep ocean, and many hundreds of years for the thermocline (PM06), it is important to determine over how much of the ocean the simulated BIR tracers have captured the full breadth of the TTD within the 100-year integration. To asses this, we define two criteria for locating areas of (near) convergence. Such a metric should prove useful in obtaining a global view of where the moments computed from the simulated TTD, including the mean age, are meaningful. We here show results using two different definitions of convergence:

(1) A *TTD normalization threshold*, which uses a criterion based on the TTD itself. By definition, TTDs must integrate to unity,

$$I(\mathbf{x}) = \lim_{\tau \to \infty} \int_0^\tau TTD(\mathbf{x}, t')dt' = 1, \tag{4}$$

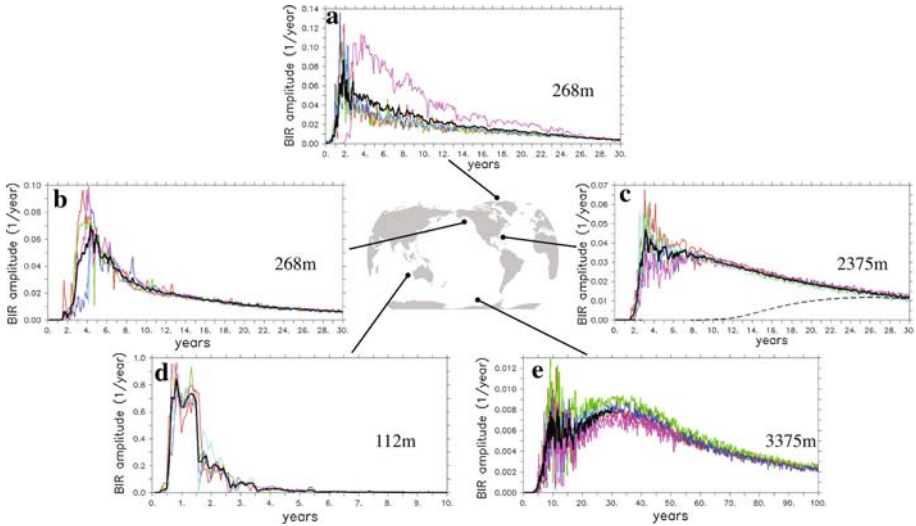**Fig. 4** Time series of monthly mean global BIRs at five locations. BIR1 through BIR5 are denoted by *red*, *green*, *blue*, *cyan*, and *magenta*, respectively, and the five-member ensemble average by the thick *black line*. Note that time axis varies in all of the plots, and that the ensemble mean is only defined for the first 30 years since that is maximum duration of the shortest time series (BIR4). The DWBC location (**c**) also includes a time series (*dashed curve*) from the 3° simulation of PM06. Note that the time axis varies in all of the plots

so a threshold of 0.95 for this integral has been chosen to define convergence. It is important to note that this relation only holds for BIRs in an ensemble average, that is, any given realization of the BIR may integrate to a value somewhat greater than or less than unity (HZWH08). The extent to which the average BIR integral exceeds unity may also be an indicator of how robust the estimate is for a given ensemble size, with highest values corresponding to lowest confidence. However, other factors such as numerical dispersion of the advection scheme or the finite length of the initial tracer impulse could contribute, and we are not able to distinguish between these possibilities with the available model output.

Figure 5 shows TTD integral, $I(\mathbf{x})$, after 100 years at 268 m depth. The colors in Fig. 5 change to yellow at a value of 0.95 (the normalization threshold), so areas that are not blue–green may be close to convergence. In general, the subtropical gyres are relatively close to convergence since they are regions with large amounts of rapidly ventilated mode waters. The North Atlantic subpolar gyre is also close to convergence due to annual wintertime convection reaching this depth. Areas that experience upwelling, such as eastern boundaries and equatorial regions, or that are strongly stably stratified, such as the northern Indian and Weddell Sea, are furthest from converging after 100 years of model integration.

(2) An *age-trend threshold* For direct comparison with ocean data, it is the *age* of a water parcel that is of primary interest, rather than the full BIR or TTD. As noted above, TTDs represent PDFs of transit times at a particular location, so by definition the first moment of the TTD is the mean age of the water mass:

$$Age(\mathbf{x}) = \lim_{\tau \to \infty} Age(\mathbf{x}, \tau) = \lim_{\tau \to \infty} \int_0^\tau TTD(\mathbf{x}, t')t'dt'. \tag{5}$$
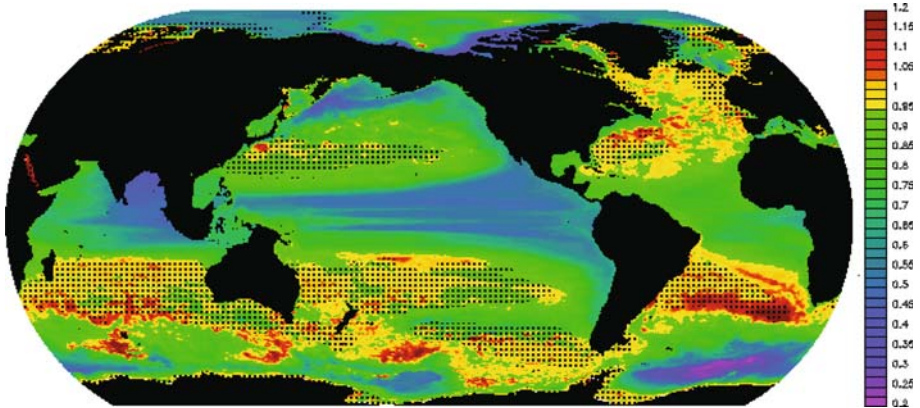
Fig. 5 Ensemble averaged BIR normalization integral (Eq. 4) at 268 m after 100 years of model integration (*colors*). Stipling shows regions of convergence using an age trend threshold (Eq. 5) of 1 year/decade

Because TTDs and BIRs share the same statistics, the first moment of the BIRs can be interpreted as an estimate of the true mean age of the water-mass. Figure 6 shows the running integral of the mean age, calculated using the BIRs at the same locations as in Fig. 4. Note that at each of the selected locations, the mean age of the water-mass changes rapidly with time, before asymptotically approaching the "true" mean age. Running integrals are thus useful as an indicator of the completeness of our century long TTDs, i.e, the degree to which they sample the complete range of transit times. Even after 100 years of model integration, only those locations with an age of less than 10 years (such as the Indian Ocean site in Fig. 4) appear to be have complete TTDs. The age-trend threshold is defined as a trend of 1 year of age increase per decade of simulation. The stipled regions in Fig. 5 show areas that meet this criterion. The same general patterns emerge using the normalization threshold and age-trend threshold metrics.

4.3 Spatial distribution of the first moment

Figure 7 shows the ensemble average, standard deviation, and relative standard deviation of mean age at 268 m for the three-member (i.e., each longer than a century) ensemble after 100 years of model integration. The standard deviation of mean age is typically only a few years at this depth, resulting in relative variation of only a few percent over most of the globe. In the regions where the age is close to its asymptotic value (Fig. 5), the standard deviation of age tends to be very small, indicating that the average BIR age is likely to be a good approximation to the TTD age in these regions.

There are several areas that stand out in Fig. 7 that indicate where the three member ensemble average BIR may not be a reliable estimate of the TTD. For example, in the south-eastern Pacific, the relatively high variability in the standard-deviation of mean age is due to BIR3 having a somewhat higher amplitude and wider distribution than either BIR1 or BIR2, and thus a greater age. This could be the result of an additional water mass contribution at this location, but without more flow pathway information (from regional BIRs or lagrangian particles, for example) it is difficult to be conclusive. The Weddell Sea, Sea of Japan, and Baffin Bay are all showing a sensitivity to the depth and timing of convection.
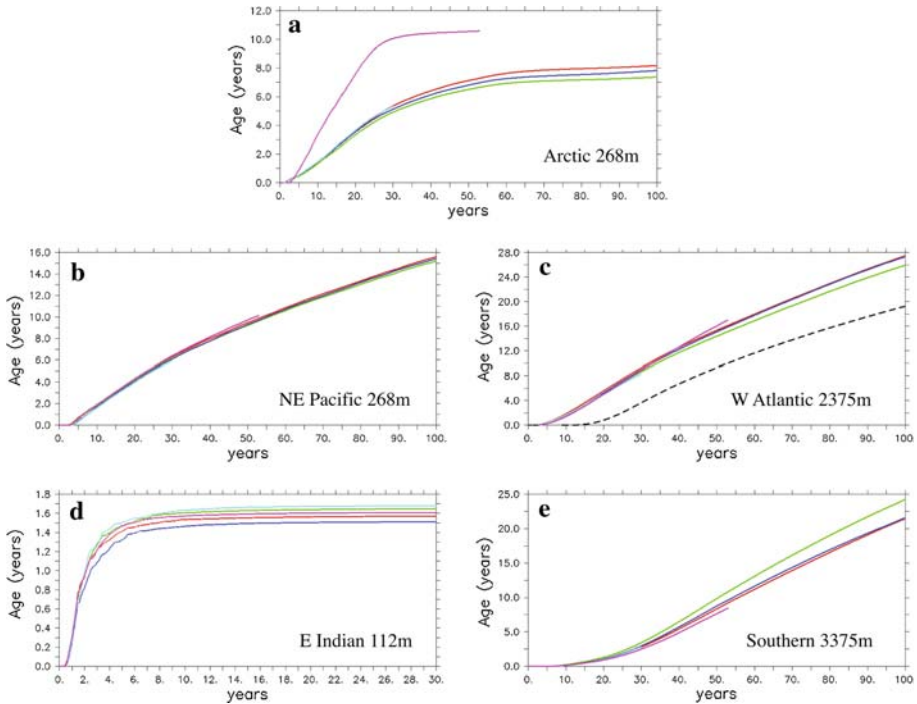
Fig. 6 Time series of the first moment ($Age(\mathbf{x}, \tau)$) of global BIR1 through BIR5, at the same five locations as in Fig. 4. Again, the individual realizations are denoted by *red*, *green*, *blue*, *cyan*, and *magenta*, respectively, and the five-member ensemble average by the *thick black line*. The *dashed curve* in **c** is from the 3° simulation of PM06

Figure 8a–c shows meridional sections of the mean age derived from the three-member ensemble average BIRs, again with the standard deviation and relative variation, at 55°W in the North Atlantic. The youngest water is found above the thermocline and in Labrador Current water making its way around the Grand Banks at the northern end of the section. The average BIR integral contours in Fig. 8 clearly show the more rapidly ventilated (younger) NADW in both the north and south centered around 2,500 m. Even though the deepest water (below 4,000 m) appears to be younger than at mid-depths, the BIR here is incomplete, so it is very likely that it would become the oldest water if the simulation was run for a substantially longer time. The variability along this section shows a relative minimum centered around 1,000 m associated with low salinity intermediate water from the south.

In the Pacific along 155°W (Fig. 8d), the youngest water is again found in the upper thermocline. The apparently low age water that occupies much of the range between 1,000 m and the bottom is actually a result of the fact that no signal has reached these depths after 100 years, consistent with the absence of large scale deep water formation in the north Pacific. The signature of AABW can clearly be seen originating from the south and flowing into the middle northern latitudes in the abyss.

One important question is that of whether increased ensemble size leads to an increase in our confidence in the statistics. Similar statistics to those in Fig. 7 may be obtained using the five-member ensemble, though due to the shorter duration it is necessary to move higher up in the water column to find BIRs that are complete. It was found that, at 112 m depth, after
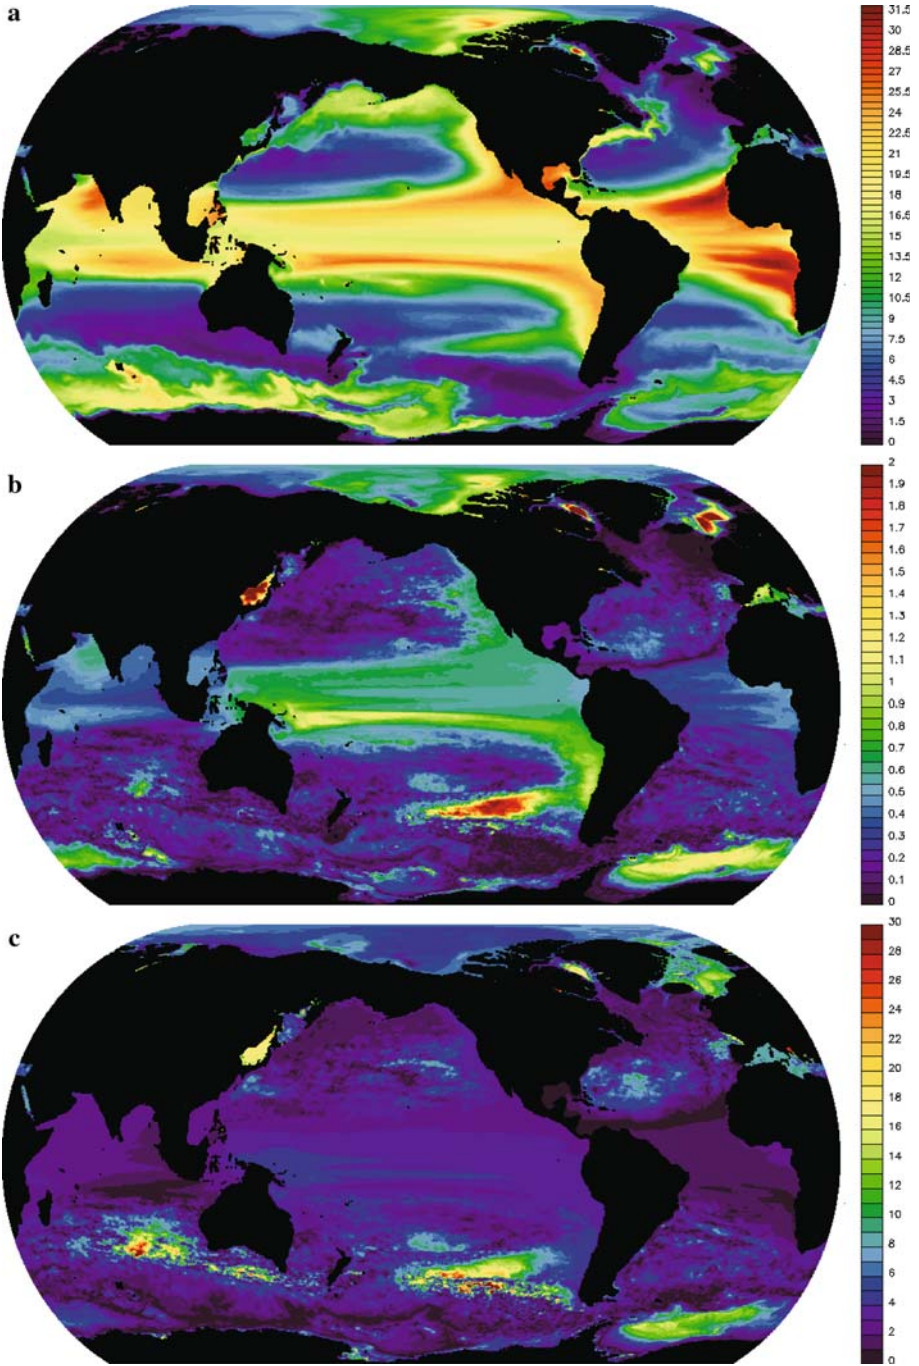
Fig. 7 **a** Ensemble average age (years), **b** standard deviation of age (years), and **c** relative standard deviation (standard deviation divided by the average, in percent) calculated from global BIR1, BIR2, and BIR3 at 268 m after 100 years
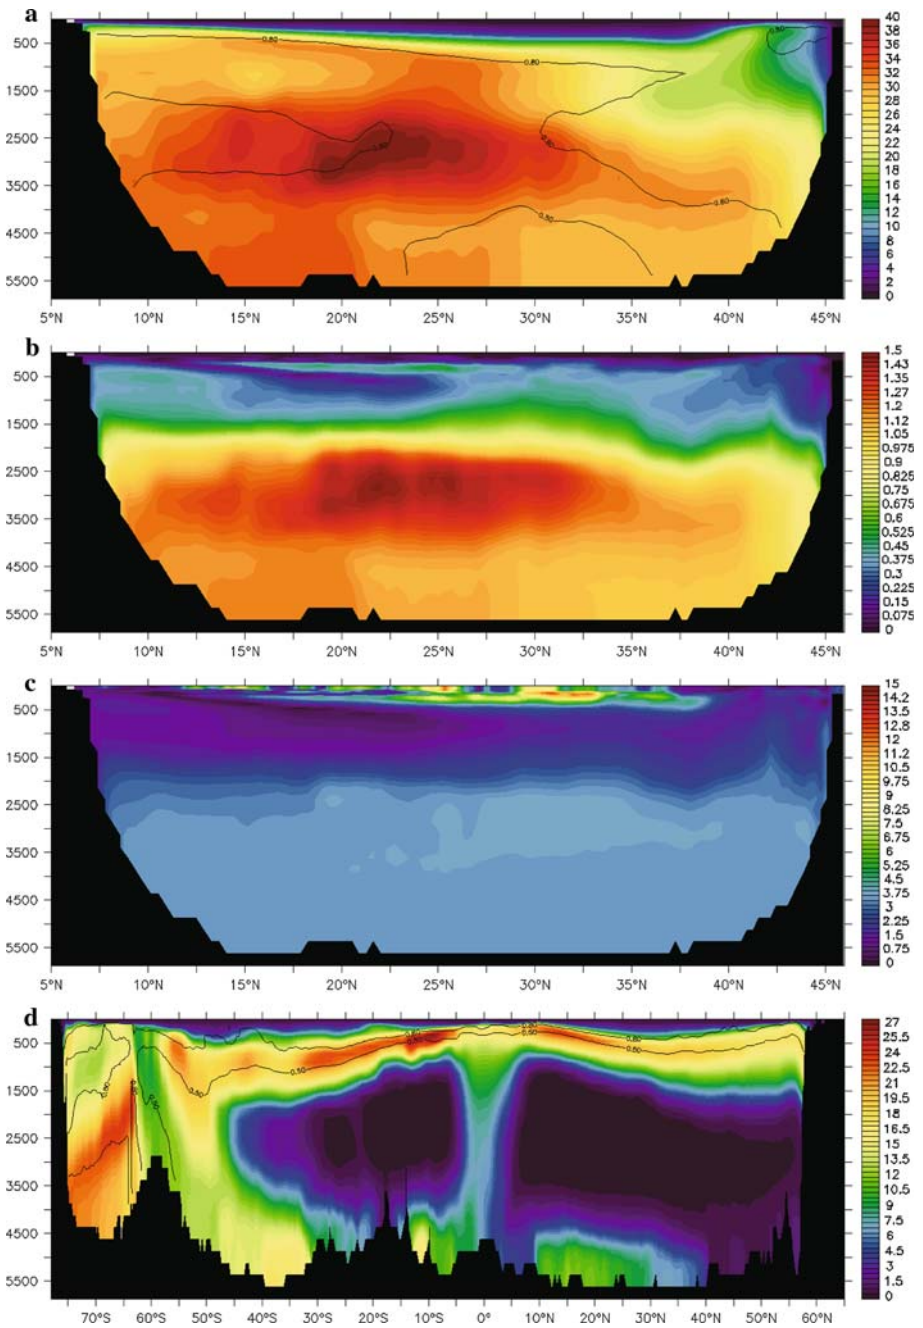
Fig. 8 **a** Ensemble average age (years), **b** standard deviation of age (years), and **c** relative standard deviation (percent) calculated from global BIR1, BIR2, and BIR3 at 55 W after 100 years. The average BIR integral is contoured with values of 0.5 and 0.8 in **a** and **d** to indicate which areas are relatively close to convergence. **d** Ensemble average age (years) at 150 W with same overlying contours as in **a**
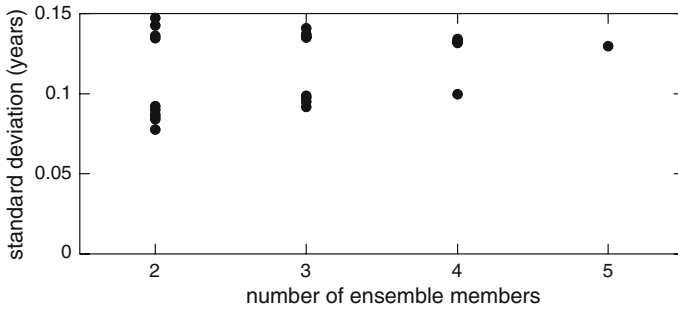
Fig. 9 Global area-averaged ensemble standard deviation of age (in years) at 112 m after 30 years using all possible unique combinations of between two and five ensemble members. The clusters of points with higher values each contain member BIR5 (see text)

30 years, the ensemble average, standard deviation, and relative standard deviation of age computed using five ensemble members was very similar to that seen in Fig. 7, with lowest mean age in the subtropical gyres and convection regions, and relative standard deviation of only a few percent in most areas.

Figure 9 shows the global area-averaged ensemble standard deviation of mean-age at 112 m as a function of ensemble size, using all possible combinations of ensemble members. For example, the ensemble size of three on the x-axis includes the ten possible unique combinations from (BIR1, BIR2, BIR3), (BIR1, BIR2, BIR4), etc. Since the spread in values decreases with the ensemble size, we conclude that a larger ensemble *does* improve our confidence, but that even an ensemble size of three can give a fairly good estimate of the statistical behavior, at least for this definition of a metric. It is worth noting that there are two distinct groupings of points in Fig. 9 for each ensemble size. The higher values all include contributions from BIR5 which has anomalously large departures from the ensemble mean in the Arctic (Fig. 4). If the Arctic is excluded from the calculation, the higher values drop by about 10%. This is an indication that the ideal of statistical stationarity of the flow field was not completely achieved over multi-decadal timescales.

4.4 Regional BIRs

As noted in Sect. 3.2, the simulation also included six regional BIRs for which the surface region Ω was not the entire ocean surface, but a specific area of interest (Fig. 10a). Regional BIRs can further elucidate the water mass properties at a given location by quantifying the relative importance of the source region on the parcel composition. These regional BIRs were simulated for 20 years, and only a single realization was performed for each source region.

Figure 10b–d shows the percentage of the global BIR3 integral (4) due to contributions from the regional BIRs. The global BIR3 integral is used as a reference since it was initiated at the same time as the regional BIRs (Fig. 1). By definition, the sum of all of the regional BIRs add up (within numerical error) to be identical to BIR3 in areas where there are no contributions from undefined regions, thus creating well-defined percentages. The colors map directly to the regions denoted in Fig. 10a (with yellow being the contribution from all regions *except* these), and the color saturation reflects the percentage. For example, Fig. 10b shows which source regions contribute the most to the total BIR3 integral after 20 years at 2,125 m depth. The NADW at this level is dominated (typically 70–90%) by contributions from the Labrador Sea (blue); only to the southeast of the Reyjknes Ridge is Irminger Sea
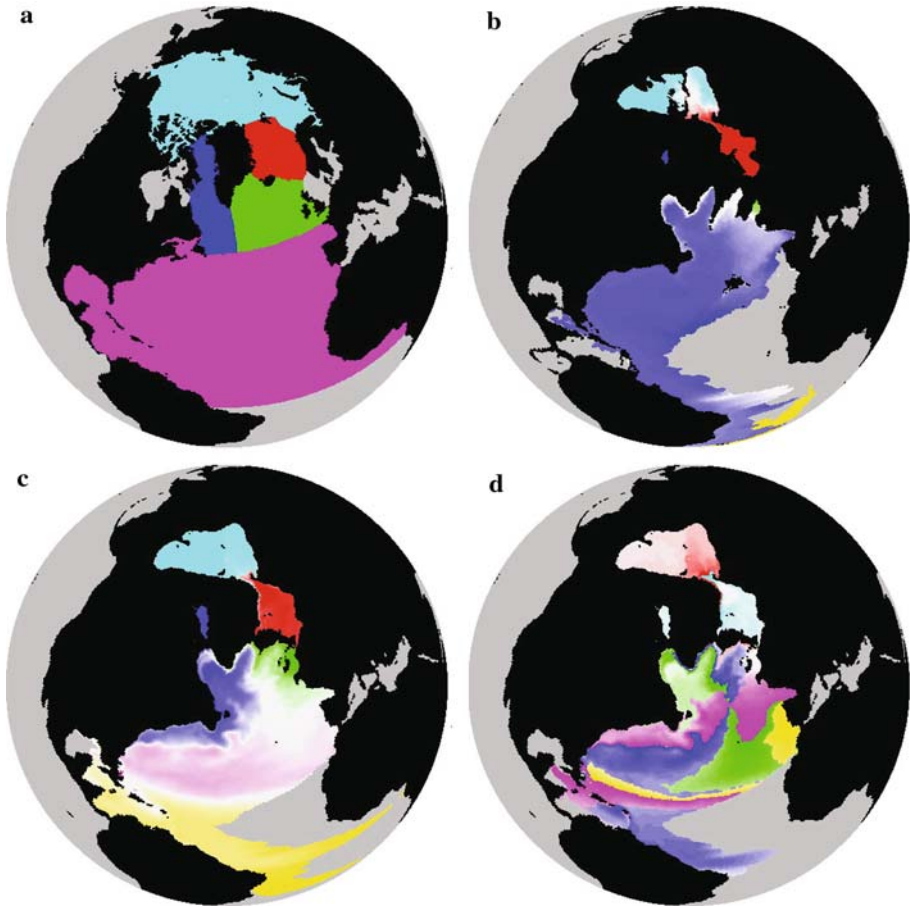
Fig. 10 **a** Locations of the surface ocean area which define the regional BIRs. Note that there is a sixth regional BIR that includes much of the Southern Ocean which is not considered here. **b** Primary contributions (color coded by the surface region seen in **a**) to the global BIR3 integral, $I(\mathbf{x})$, at 2,125 m depth after 20 years from each of the regional BIRs. The most saturated colors denote 100% of the total due to a particular regional BIR, and the least saturated (i.e., *white*) denote 50%. Areas that are grey have no BIR value, indicating that no ventilation has occurred yet. **c** Primary contributions to the global BIR3 integral at 918 m depth after 20 years from each of the regional BIRs. The most saturated colors denote 100% of the total, and the least saturated denote 40%. **d** Secondary contributions to the global BIR3 integral at 918 m depth after 20 years from each of the regional BIRs. The most saturated colors denote 40% of the total, and the least saturated denote 2%

water (green) the most plentiful component. If we consider which regional BIRs are the second most important in the total BIR3 integral, the patterns remain the same as in Fig. 10b, but the colors have switched (blue with green and cyan with red, not shown). The percentages have also dropped (since these are the secondary contributions), showing that Irminger Sea water accounts for 10–20% of the NADW at this depth. When considering the third most important contributors, we find a 5–15% contribution to the NADW from GIN Seas water (red, not shown).

At shallower depths (918 m), the primary contributions to BIR3 more directly correspond to the locations of the surface regions (Fig. 10c). In particular, the Subtropical water

(magenta) has been ventilated to the south of the Gulf Stream, with Labrador Sea water making its way south along the North American coast. Waters from the south (yellow) have penetrated through the equatorial current system into the Caribbean. The secondary contributions show a much more complex arrangement (Fig. 10d). Advection of Irminger Sea (green) and Subtropical (magenta) water through the gyre from the northeast can be seen. Labrador Sea water (blue) has been captured from the DWBC via the equatorial currents and is returning northward along the coast of South America. Mediterranean Sea outflow (yellow) can be seen clearly and is an important constituent (35%) in the eastern basin. The striped pattern to the east of Florida indicates the intersection of the different water masses, each associated with a range of densities, with this depth.

5 Discussion

In a highly turbulent flow, the use of ensemble averaging is fundamental. This is especially true here, where the statistics of the field we want to know (the TTD) are the same as for the field that is being simulated (the BIR). Due to the large computational burden imposed by the high horizontal resolution grid, a tradeoff between ensemble size and length of simulation was required a priori, with minor modifications being made during the run. Similarly, the length of the spinup run was chosen to be 15 years to allow for the maximum possible duration of the tracers. While this certainly is not enough time for the complete circulation to reach equilibrium, it is likely sufficient for the upper ocean processes, such as ventilation and formation of boundary undercurrents, that are of interest here (see, e.g., figure 1 of MM05).

Although performing statistics on an ensemble size of 5 (or 3) is far from ideal, it appears that we can obtain valuable information from them. Figure 9 suggests that any larger ensemble that could have practically been simulated (say, 10 or 15 members) would not have increased our confidence by a tremendous amount. Further, the 40-member simulations performed by Haine et al. [8] clearly show that a much smaller number of ensemble members are necessary at times sufficiently greater than the peak. On the other hand, Fig. 9 is only one metric at one model depth. The need to employ more statistical techniques, especially those specific to small ensemble sizes, is a path for future work.

As can be seen in Fig. 5, it appears that only a small percentage of the ocean has complete/converged BIRs. Figure 11 shows the depth dependence of convergence for both the TTD normalization and age-trend thresholds. These criteria are very consistent except at depth where the normalization threshold picks up the Atlantic DWBC, while the age-trend does not because the tails of the BIRs fall off relatively slowly with time (Fig. 4). Globally integrated, the TTD normalization threshold shows that 4.7% of the ocean volume has converged, while the age-trend threshold results in a somewhat lower value of 3.2% by the end of the 100 year simulation.

However, BIR convergence is likely to be too stringent a criterion for determining whether BIRs can be used as an estimate of the TTD. As seen in Fig. 4, the variability of the BIRs (both in time and among the ensemble members) decreases dramatically during their decay phase. It may then be possible to extrapolate the ensemble averaged BIR to much greater times, thus increasing the number of complete TTDs. For example, Fig. 12a shows the estimated TTD at a single location, showing that it exhibits a power law behavior for times following the peak of the distribution. Using this form in the first moment integral (Eq. 5), it is possible to estimate the asymptotic value of the age (Fig. 12b). This method likely provides an upper bound on age, since theory predicts that the BIRs decay exponentially at very large times [7], thus asymptoting faster than for any power law.
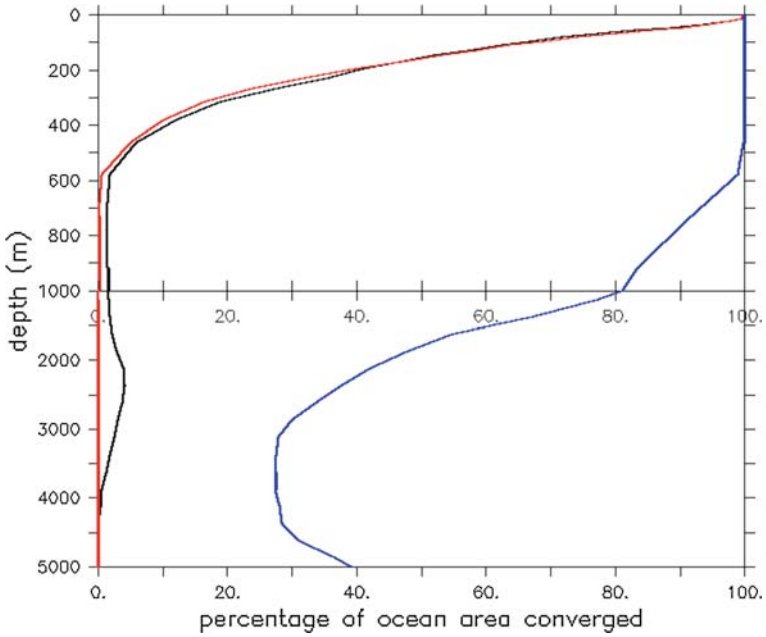
Fig. 11 Percentage of ocean area that has converged at each model depth using the normalization integral threshold (Eq. 4, *black*) and age-trend (Eq. 5, *red*). *Blue curve* shows the percentage of area at each depth that has passed the peak in the TTD by year 75
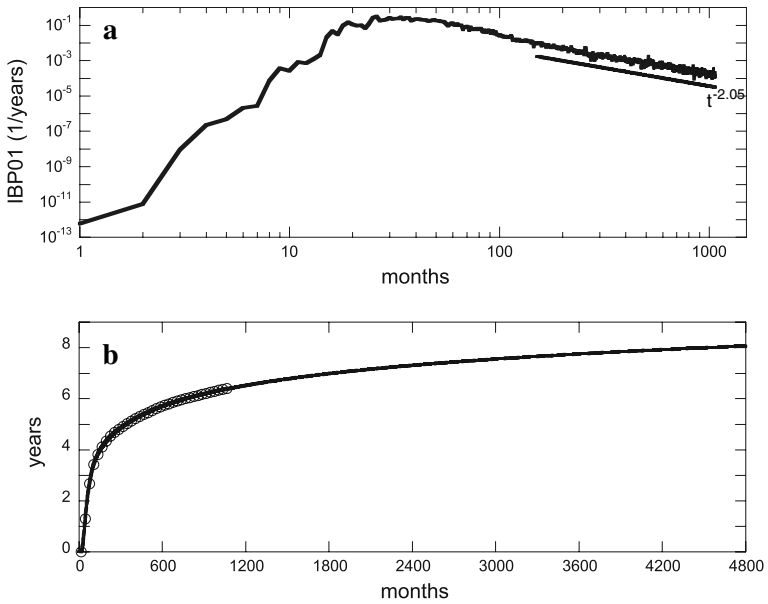

Fig. 12 **a** Time series of BIR1 at 21 N, 55 W, and 230 m depth in the subtropical Atlantic. The best least-squares fit to the slope for times greater than 12 years is −2.05. **b** First moment integral of the same BIR where *circles* denote the actual values and the *solid line* is the extrapolation based on the $t^{-2.05}$ power law fit

Although extrapolating all of the TTDs is beyond the scope of this paper, it is possible to estimate how many distributions may be amenable to extrapolation. The blue curve in Fig. 11 shows the percentage of TTDs whose peaks are reached by year 75, leaving 25 years of decay to guide the extrapolation. This results in a major increase in the volume of the ocean with potentially complete TTDs from 3.2 (using the age-trend threshold) to 50%. Note that essentially the entire ocean above 500 m has passed the peak of the TTD, and that 35–40% of even the deep ocean volume has achieved this threshold.

A second possible limitation of our experiments relates to the assumption of statistically stationary flow underlying the connection between the ensemble mean BIR and the TTD. While the simulation was forced with repeating annual cycle atmospheric conditions, and in many metrics (e.g. global mean kinetic energy) the flow is statistically stationary, there are signatures of low-frequency, multi-decaldal trends in some aspects of the circulation. These can be due to either model drift or intrinsic low frequency modes of variability arising in the turbulent solution. Evidence of this behavior was noted in the relatively large variability (as high as 40–50%) in the Arctic seen in Fig. 4. This may provide some insight into an upper bound, since it is due to a shift over the three decade period between the release of BIR4 and BIR5. Where the flow remains statistically steady, we would therefore expect the age variability to be much lower, perhaps on the order of 10%. We are planning a companion simulation forced with interannually varying atmospheric fields to quantify its effect.

6 Conclusions

We have presented results from a century-long simulation of the global ocean using a 1/10° grid that allows for spontaneous formation of mesoscale eddies. A suite of passive tracers, including an ensemble of five global BIRs were introduced in the simulation for the purpose of estimating distributions of water mass ventilation timescales. These provide the first estimate of the ocean TTD where the stirring and mixing effects of eddies have been accounted for directly rather than parameterized.

Consistent with the idealized basin study of HZWH08, one of our primary results is that the variance across the ensemble collapses dramatically for times following the peak in the TTD. At earlier times, significant variations due to chaotic processes can be seen in time series of the BIR for any given ensemble member (Fig. 4), but indications are that the modest size ensemble used here (five members for transit times less than 30 years) may be sufficient to obtain a robust estimate of the TTD. The TTD has important advantages over the individual BIRs, such as representing a true probability distribution function of transit times, and being able to be used in the convolution integral (Eq. 3) that propagates surface values of a tracer into the interior.

We also have found that the first moment (or age) of the BIRs is typically dominated by the long tail of the distribution where the variability is relatively small, so in those regions that have reached convergence the variations in the estimated age are typically only a few percent of the mean. Further, it appears that the long-time behavior of the BIRs may be amenable to extrapolation, resulting in even a more accurate assessment of water mass age.

While the variability across the ensemble arising from eddies is relatively small for intermediate and long transit times, there are fundamental differences in the character of the TTDs obtained here as compared to coarser resolution simulations. This is an indication of the deficiencies in the representation of sub-gridscale transport processes in the coarse resolution models. This result suggests that the TTD obtained from this simulation may provide a useful metric against which to judge eddy mixing parameterizations used in climate models.

References

1. Adcroft A, Hill C, Marshall J (1997) Representation of topography by shaved cells in a height coordinate ocean model. Mon Weather Rev 125:2293–2315
2. Adcroft A, Campin J-M, Heimbach P, Hill C, Marshall J (2005) MITgcm manual. http://mitgcm.org/sealion/online_documents/manual.html
3. Dukowicz JK, Smith RD (1994) Implicit free-surface method for the Bryan–Cox–Semtner ocean model. J Geophys Res 99:7991–8014
4. England M (1995) The age of water and ventilation timescales in a global ocean model. J Phys Oceanogr 25:2756–2777
5. Fekete BM, Vörösmarty CJ, Grabs W (2000) Global, composite runoff fields based on observed river discharge and simulated water balances. Technical Report, Global Runoff Data Centre.
6. Gouretski VV, Koltermann KP (2004) Woce global hydrographic climatology. A Technical Report. Tech. Rep. 35, Berichte des Bundesamtes für Seeschiffahrt un Hydrographi
7. Haine T, Hall T (2002) A generalized transport theory: water-mass composition and age. J Phys Oceanogr 32:1932–1946
8. Haine T, Zhang H, Waugh DW, Holzer M (2008) On transit-time distributions in unsteady circulation models. Ocean Model 21:35–45
9. Hall T, Haine T (2002) On ocean transport diagnostics: the idealized age tracer and the age spectrum. J Phys Oceanogr 32:1987–1991
10. Holzer M, Hall T (2000) Transit-time and tracer-age distributions in geophysical flows. J Atmos Sci 57:3539–3558
11. Holzer M, Primeau F (2006) The diffusive ocean conveyor. Geophys Rev Lett 33:L14618
12. Holzer M, Primeau F (2008) The path density distribution of ocean surface-to-surface transport. J Geophys Res 113:C01018
13. Hurrell JW, Hack JJ, Shea D, Caron JM, Rosinski J (2008) A new sea surface temperature and sea ice boundary data set for the community atmosphere model. J Clim 21:5145–5153
14. Khatiwala S, Visbeck M, Schlosser P (2001) Age tracers in an ocean GCM. Deep Sea Res 48:1423–1441
15. Large WG, Pond S (1982) Sensible and latent heat flux measurements over the oceans. J Phys Oceanogr 12:464–482
16. Large WG, McWilliams JC, Doney SC (1994) Oceanic vertical mixing–a review and a model with a nonlocal boundary layer parameterization. Rev Geophys 32:363–403
17. Large WG, Yeager SG (2004) Diurnal to decadal global forcing for ocean and sea-ice models: the datasets and flux climatologies. NCAR Technical Note TN-460+STR, National Center for Atmospheric Research
18. Maltrud ME, McClean JL (2005) An eddy-resolving global 1/10° ocean simulation. Ocean Model 8: 31–54
19. Maltrud ME, Smith RD, Semtner AJ, Malone RC (1998) Global eddy-resolving ocean simulations driven by 1985–1995 atmospheric fields. J Geophys Res 103:30825–30853
20. Maltrud ME, Peacock SL, Bryan FO (2009) A century-long global eddying ocean simulation (in preparation)
21. McClean JL, Maltrud ME, Bryan FO (2006) Measures of the fidelity of eddying ocean models. Oceanography 19:104–117
22. McDougall T, Jackett D, Wright D, Feistel R (2003) Accurate and computationally efficient algorithms for potential temperature and density of seawater. J Atmos Ocean Technol 20:730–741
23. Murray RJ (1996) Explicit generation of orthogonal grids for ocean models. J Comput Phys 126:251–273

24. Peacock S, Maltrud ME (2006) Transit-time distributions in a global ocean model. J Phys Oceanogr 36:474–495
25. Primeau F (2005) Characterizing transport between the surface mixed layer and the ocean interior with a forward and adjoint global ocean transport model. J Phys Oceanogr 35:545–564
26. Waugh D, Hall T, Haine T (2003) Relationships among tracer ages. J Geophys Res 108:1–16