Advection-Diffusion Process Inference via Statistical Oceanographic Methods in the North Atlantic and Southern Oceans

David S. Trossman

A dissertation submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Washington

2011

Program Authorized to Offer Degree: Oceanography

University of Washington
Graduate School

This is to certify that I have examined this copy of a doctoral dissertation by

David S. Trossman

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

Chair of the Supervisory Committee:

_____

LuAnne Thompson

Reading Committee:

_____

LuAnne Thompson

_____

Sabine Mecking

_____

Mark J. Warner

Date: _____

In presenting this dissertation in partial fulfillment of the requirements for the doctoral degree at the University of Washington, I agree that the Library shall make its copies freely available for inspection. I further agree that extensive copying of this dissertation is allowable only for scholarly purposes, consistent with "fair use" as prescribed in the U.S. Copyright Law. Requests for copying or reproduction of this dissertation may be referred to Proquest Information and Learning, 300 North Zeeb Road, Ann Arbor, MI 48106-1346, 1-800-521-0600, to whom the author has granted "the right to reproduce and sell (a) copies of the manuscript in microform and/or (b) printed copies of the manuscript made from microform."

Signature_____

Date_____

University of Washington

Abstract

Advection-Diffusion Process Inference via Statistical Oceanographic Methods in the North Atlantic and Southern Oceans

David S. Trossman

Chair of the Supervisory Committee:
Professor LuAnne Thompson
Oceanography

Quantities pertaining to ventilation, the process by which water is exchanged between the surface mixed layer and the interior of the ocean, are estimated here: To do this, the utility of the thin-plate spline (TPS) as an interpolation procedure is described as a way to estimate mapping errors. Second, using the concept of transit-time distribution (TTDs), or mass fractions of water getting from the surface to an interior location over some time interval, the volume transports of Labrador Sea Water (LSW), Antarctic Intermediate Water (AAIW), and the LSW component of North Atlantic Deep Water (NADW) are estimated through three different repeated cross-sections of the ocean. Evidence is found that both the single and mixture inverse Gaussian representations of a TTD are insufficient, that at least some of the AAIW seen south of Tasmania forms from Subantarctic Mode Water (SAMW) that is entrained and detrained in the southeast Indian Ocean, and that at least some of the LSW seen south of Iceland is formed in the Irminger Basin. Third, estimates for rates at which the North Atlantic thermocline is ventilated advectively are given using profiler, scatterometer satellite, and altimetric satellite observations from several consecutive years. Ekman pumping is a secondary effect relative to the horizontal and temporal structure of the mixed layer depths (MLDs) for controlling ventilation. Fourth, the trends of mode water subduction rates and the extent to which along-isopycnal diffusion contributes to subduction in the North Atlantic and Southern Oceans are estimated by utilizing TTDs. Using a

statistical technique called Bayesian model averaging (BMA), observations of CFC-11 are incorporated into a model simulation-derived TTD-based estimates of ventilation rates, which are compared with diapycnal velocity-based estimates, and uncertainties are quantified. It is found that along-isopycnal diffusion has increasingly contributed to SAMW ventilation over time and that both ventilated North Atlantic Subpolar Mode Water (SPMW) and, to a lesser extent, ventilated North Atlantic Subtropical Mode Water (STMW) erode by processes with at least two distinguishable time scales. Last, using BMA at each grid point, a density dependence of the model parameter for the along-isopycnal diffusivity, which has yet to be determined to be the path-averaged along-isopycnal diffusivity of the real ocean, is inferred. The model parameter for the along-isopycnal diffusivity is suggested to be smaller for water masses that form by deep convective processes than for water masses that form by turbulent processes that mix distinguishable water masses at intermediate depths.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

The author wishes to express sincere appreciation to Jonathan Wakefield, Susan Hautala, and his committee members: LuAnne Thompson, Sabine Mecking, Mark J. Warner, Eric D'Asaro, Vladimir Minin, and Peter Guttorp for their helpful comments and advice and the National Science Foundation for their generous support. Emphasis is placed on gratitude for having LuAnne as a mentor and adviser.

Chapter 1

INTRODUCTION

The communication between the atmosphere and interior of the ocean (i.e., the ocean below the deepest depth affected directly by atmospheric forcing over a year) is essential for understanding the temporary sequestration of tracers such as carbon and atmospheric signals such as heat. Without this communication between the atmosphere and interior of the ocean, the sea surface would have one less repository for the heat it transports, the buffering effect of the ocean due to its heat capacity which is about 850 times as large as that of the atmosphere would be dampened, and the photosynthetic biota of the ocean would deplete their environment of the nutrients they need to live more quickly. The ocean would not be able to transport as much heat poleward because the heat would be concentrated in a relatively thin, highly stratified surface layer without being able to communicate with the oceanic interior. Instead, large-scale transfer of fluid between the surface layer of the ocean and the interior of the ocean makes communication between the atmosphere and the oceanic interior possible.

Fluid parcels of distinct characteristics (water masses) can form by two different physical processes: one is to emerge from the surface layer that is being mixed and the other is to mix already existing water masses together. Similarly, water masses can be destroyed or eroded by two different physical processes: one is to enter the surface mixed layer and the other is to be mixed with another existing water mass. The process of emerging from the surface mixed layer, typically thought to be via movement along constant density (isopycnal) surfaces until the water mass is beneath the surface mixed layer where water properties are approximately conserved, is referred to as subduction. The process of entering the surface mixed layer is also typically thought to be along isopycnal surfaces until the water mass is above the depth to which the surface mixed layer extends (i.e., the mixed layer depth) and is referred to as obduction. However, in addition to the convergence of wind-driven waters

(Ekman pumping) and passage of waters through deepened mixed layers (lateral induction), diffusive processes along isopycnals contribute to the ventilation process to varying extents for particular water masses [67].

The extent to which processes are diffusive has been difficult to ascertain for accurate numerical simulation, but a combination of isopycnal coordinate models and observations of transient tracers (i.e., tracers that have concentrations that are not steady, even on multi-year time scales due to non-constant sources) provides an ideal testbed for testing the hypothesis that the along-isopycnal component to diffusion is a density-dependent process. Previous studies [64, 1] have suggested that the along-isopycnal component to diffusion should be different depending upon location. The average of the along-isopycnal component of the diffusivity of a water parcel over its path is referred to as the path-averaged diffusivity here. Isopycnal coordinate models are ideal because circulation fields can be held fixed while background diffusivity is varied and background diffusivity can be held fixed while circulation fields are varied. Transient tracers provide sufficient data to constrain the uncertainty in diffusion across isopycnals [56, 27], but have not been used to constrain the uncertainty in diffusion along isopycnals. Since water masses attain their characteristics through different physical processes in which diffusion plays a role, it is hypothesized that a water mass formed one way will influence downstream eddy (i.e., departure from time mean flow) activity differently from a water mass formed another way. This, in turn, influences the large-scale circulation.

The ultimate oceanographic goal of this thesis is to infer properties of the large-scale circulation of the ocean such as the degree to which a water mass is advected as opposed to diffused along its path and where the water that composes a water mass comes from. This requires recognizing the boundary between the surface mixed layer and the oceanic interior below as a communication medium in which both advection and diffusion play a role [67, 60, 45]. The tools for inference here are rooted in the advection-diffusion equation with a source function at the surface mixed layer. The one-dimensional version of this equation is unique in that it has an exact analytical solution that is a distribution characterizing the fraction of water that gets from the surface to any given point in the interior over a given time interval (called an 'age') [38]. However, the ocean is not a perfectly non-leaky

pipe. Because of this, we can only estimate an approximation to the actual distribution of water parcel ages by using observations. A better approximation to the distribution of water parcel ages was made with a method that did not characterize all contributing factors of uncertainty [39] and is made by utilizing a combination of numerical models and observations via Bayesian analysis that can characterize all known contributing factors of uncertainty here. We ultimately relate the distribution of water parcel ages to a rate at which mass is transferred from one domain to another [31] to quantify the ventilation process.

This thesis mainly focuses on applying statistical estimation methods to oceanographic contexts. These contexts include ventilation rates, along-isopycnal diffusivities, and water mass pathways. The distribution of ages since a fluid parcel last was at the ocean surface serves as a link between all of these contexts. The methods of estimation include numerically solving differential equations, two types of Bayesian estimation techniques (Markov Chain Monte Carlo and Bayesian model averaging), and one machine learning estimation technique (thin-plate spline). For each of these methods, a physical ocean model is necessary to input physical information into what would otherwise be strictly mathematical. These physical models are almost always a combination of numerical and analytical approaches.

Chapter 2

INTERPOLATION AND QUANTIFYING MAPPING ERROR

2.1 Introduction

Putting irregularly spaced data on a regularly spaced grid is not only a problem of finding optimal interpolation estimates in certain situations, but also of inferring how surprising it would be to measure a value outside of an estimated range, or confidence interval, at a particular point. There are occasions that arise where a quantitative estimate of mapping error is very important. For example, oceanographers are often interested in spatial patterns of specific water masses, which are characterized by minima or maxima in certain tracer quantities, but we need to determine whether these signals are above the noise.

Here, like McIntosh (1990) [63], we investigate how an interpolation method that aims to minimize the error of the fitted estimates (objective mapping, or OM) compares with a spline-based method that does not necessarily minimize the error of the fitted estimates. We use a spline-based method (one that makes use of the thin-plate spline, or TPS) that has not, to our knowledge, been applied in an oceanographic context. Spline-based methods [108, 40] generally aim to find an optimal balance between bias and variance of the fitted estimate by penalizing the wiggliness (as measured by the square of the second, or more, derivatives of a function) of the fit as opposed to minimizing error. Two examples of spline-based based methods are spline interpolation, which integrates the biharmonic differential equation (requiring that the interpolated field is smooth and knowledge of the intrinsic length scale of the field), and the multiquadric-biharmonic method, which solves for weighting factors on different parametric basis functions (requiring a potentially large number of weights to be estimated from the data) [98]. Sokolov and Rintoul (1999) [98] use the term TPS to refer to a technique similar to the multiquadric-biharmonic method, but with a different parametric basis function.

Here, to be consistent with Wood (2003) [116], we use the term TPS to refer to a

type of spline-based method [108, 40] that, in contrast with previous spline-based methods, trains a statistical model that can be written as a non-parametric regression equation (a Generalized Additive Model, or GAM) and generalizes the penalty on wiggliness of the fit to the data found in other non-parametric regression techniques. By non-parametric, we are referring to placing a functional relationship between the covariates (also sometimes referred to as predictors or independent variables), and this relationship is found algorithmically as opposed to specifying an analytical form. Two examples of non-parametric regression techniques include ridge regression, which uses an L_2 norm penalty (sum of squares of coefficient estimates), and the lasso, which uses an L_1 norm penalty (sum of absolute values of coefficient estimates) on wiggliness [34]. A GAM can be used to 'predict' data, which can be applied in an oceanographic context as an interpolation estimate. One advantage of the more generalized penalty of spline-based methods, such as the TPS, over other penalties is that the generalized penalty is invariant to rotation or translation in a two-dimensional domain where the coordinate directions and origin position can be arbitrary. This penalty is always non-negative and is zero if and only if the fitted estimate is linear in all of its covariates. Another advantage of the TPS in our context is that knots (fixed endpoints connected by numerically-derived, typically polynomial-based, fits) are placed at the data locations. This avoids the typical problem with regression splines of choosing where to place knots, a decision that can greatly influence the fitted estimate.

OM [2, 8], also referred to as objective interpolation [98] and more generally referred to as objective analysis [63], interpolates irregularly spaced data onto a grid by weighting the data with parameterized covariance functions. What we refer to here as OM is relatively easy to implement on observational data [85] with covariance function parameters that are estimable [7] but variable in boundary current regions [12]. OM is regularly used in the atmospheric sciences and oceanography and has been shown to be more or less similar to previous spline-based methods in other contexts depending upon factors such as spatial density of the data distribution and how well the boundaries are sampled [98]. However, OM has been shown by McIntosh (1990) [63] to produce uninterpretable estimates with root-mean-squared error as large as the standard deviation of the field when data spacing becomes as large as the correlation length scale of the interpolated field. Furthermore, OM

only gives the error associated with its mapping procedure in terms of a percent of a specified variance which is an ad hoc estimate of the true variance. Even if one has confidence in the form of the specified covariance function, OM error estimates do not account for the error associated with the parameters that must be specified in the covariance function. The TPS technique, on the other hand, can be used as a surrogate mapping procedure, requires much less *a priori* knowledge of the field statistics, and is able to provide errors in the form of a standard error.

While there are other general field reconstruction approaches (e.g., based on spectral representations by Chu et al. (2003) [10] and based on variational inverse models by Rixen et al. (2000) [84]), we limit this study to a comparison of OM and TPS methods. We focus, more specifically, on the merits and nuances of using the TPS-based error estimates with either TPS- or OM-based field estimates, particularly when data are sparsely sampled spatially. Here, we compare TPS estimates with OM estimates along several World Ocean Circulation Experiment (WOCE) transects by using bottle data from the CLImate VARiability and Predictability (CLIVAR) Repeat Hydrography Program that is a follow-up on WOCE (Table 2.1). We use bottle data rather than CTD (conductivity-temperature-depth) data for comparison for two reasons: First, computational constraints on both mapping methods arise when the high-density CTD data are mapped. Second, we want to evaluate how using relationships between the data samples (as is the case with using the TPS) rather than the spatial statistics of the data samples (as is the case with using OM) affect estimates and associated errors of a particular field.

We focus here on maps and associated errors of CFC-11 (chlorofluorocarbon- or freon-11) concentrations, and associated partial pressure-based (pCFC-11-)ages. CFC-11 was used in refrigerators and air conditioners until outlawed, along with CFC-12, in the 1980s due to their catalytic disruption of stratospheric ozone. The measurement errors of the tracers of interest (Table 2.2) are held to strict standards and are typically swamped by small-scale natural variability as well as by errors associated with mapping the fields onto a regular grid. pCFC-11-ages are determined from the measured bottle data by taking the measured CFC-11 concentrations in [pmol/kg], converting them into partial pressures by dividing by their solubility (temperature and salinity dependent via Warner and Weiss, 1985 [111]), and

finding the year in the atmospheric time series (via Walker et al., 2000 [110]) that matches this partial pressure. The pCFC-11-ages are not true ages due to mixing and non-linearity in the atmospheric time series of CFC-11 [64].

The sections that we focus on here are SR03, A16(N), and A05. The SR03 section (from 67^oS to 43^oS, between 135^oE and 147^oE) is a transect that goes from Tasmania to Adélie Land of Antarctica. The A16 section (from 63^oN to the southern hemisphere along 20^oW) is a transect that goes southward from the southern side of Iceland. The A05 section runs along 24.5^oN from about 14^oW to about 79^oW in the North Atlantic (Table 2.1).

2.2 Methods

We review the methods of OM and the TPS for the reader's convenience. The variable y refers to the observations of the quantity either method is estimating on a grid and the variable f is reserved as the field that is fitted to estimate y. Subscripts will be added in the OM section, whereas both functional parentheses and a vector notation will be added in the TPS section. We want to distinguish how OM and TPS use a large amount of correlated data differently from one another, while we would also like to preserve much of the conventional notation from the literature.

2.2.1 Objective Mapping

OM uses the Gauss-Markov theorem, subject to its assumptions, to estimate values of a scalar variable, $y(\vec{x}) = y_x = f_x + \epsilon_x$, (e.g., temperature, salinity, CFC-11 concentrations, or pCFC-ages) at points on a grid, $\vec{x}$, between (rather than at) observations. The Gauss-Markov theorem says that the optimum linear unbiased estimator, $\hat{f}_x$, of f_x has a particular form, given that the errors, $\overline{\epsilon_i} = 0$ and $\overline{\epsilon_i \epsilon_j} = \delta_{i=j} E$ where $E = \epsilon_E \overline{y_i^2}$ is a constant, ϵ_E are the signal-to-noise amplitudes, $\delta_{i=j}$ is one for all points $\vec{x_i} = \vec{x_j}$, and $\delta_{i=j}$ is zero for all points $\vec{x_i} \neq \vec{x_j}$ where we have observations. In other words, the average of all of the errors is zero and the errors are uncorrelated at points where we have observations [8]. The form estimator, $\hat{f}_x$, takes is given by

$$\hat{f}_x = \frac{\sum_{k,l} A_{k,l}^{-1} y_l}{\sum_{k,l} A_{k,l}^{-1}} + \sum_{i=1}^{N} C_{i,x} [\sum_{j=1}^{N} A_{i,j}^{-1} (y_j - \frac{\sum_{k,l} A_{k,l}^{-1} y_l}{\sum_{k,l} A_{k,l}^{-1}})]. \tag{2.1}$$

Here, each of the summations over i, j, k and l refer to the N points at which there are observations. The covariance between observations, $A_{i,j}$, is assumed to be the sum of the error covariance $\overline{\epsilon_i \epsilon_j}$ and a function $F(q)$ of the distance $q = ||\vec{x_i} - \vec{x_j}||$ between the observations at $\vec{x_i}$ and $\vec{x_j}$. To be explicit,

$$A_{i,j} = \overline{y_i y_j} = F(q) + \overline{\epsilon_i \epsilon_j}.$$

Likewise, the covariance $C_{i,x}$ between the estimated f_x at a point on the grid, $\vec{x}$, and the ith measurement, y_i, at point $\vec{x_i}$ is given by,

$$C_{i,x} = \overline{f_x y_i} = F(||\vec{x} - \vec{x_i}||).$$

The first term in (2.1) represents a weighted mean of observations depending on the choice of covariance function between pairs of observations. When using OM, the true covariance should always be non-negative (preferably positive definite). If the true covariance is negative definite, the 'mean square error' that is given by [85],

$$\overline{(\hat{f}_x - f_x)^2} = C_{x,x} - \sum_{i,j} C_{x,i} A^{-1}_{i,j} C_{x,j} + \frac{(1 - \sum_{i,j} C_{x,i} A^{-1}_{i,j})^2}{\sum_{i,j} A^{-1}_{i,j}} \tag{2.2}$$

can be negative and may have an absolute value that is large, which indicates a bad estimator. The last term in (2.2) represents a weighted mean contribution of the departure of the covariance between observations and the estimates.

In practice, we follow Roemmich (1983) [85] to construct the OM in which both large-scale and small-scale parametrized representations of the field are included. He assumed that there is a maximum length scale over which the mapped fields are correlated; the estimated field covariance, $\overline{y_i y_j}$, is

$$\overline{y_i y_j} = \overline{y_i^2}^{1/2} e^{-q^\eta / d_C^\eta} \tag{2.3}$$

with

$$q^2 = [\frac{(\Delta x)^2}{L^2} + (\Delta z)^2]. \tag{2.4}$$

The OM procedure occurs in two steps with separate parameters chosen for the large- (small-) scale fields: the decay scale, $d_C = 1000$ m ($d_C = 200$ m); the horizontal length

scale, $L = 10^5$ m ($L = 2 \times 10^4$ m); the signal-to-noise ratio amplitude, $\epsilon_E = 0.1$ ($\epsilon_E = 0.02$); and $\eta = 2$ ($\eta = 1$). During the first step, the large-scale field is estimated and subtracted from the raw data to leave the small-scale field plus residual noise. During the second step, the small-scale field is estimated, leaving just residual noise. In practice, to make the amount of data tractable by Matlab, we typically process only twelve stations of bottle data at a time and repeat.

2.2.2 Thin-Plate Splines

As an alternative approach to OM, we fit a field, $\mathbf{f}$, to the observations, $\mathbf{y} = \mathbf{f} + \epsilon$, with a sum of linear and non-linear functions of covariates, called marginal smooths, that are penalized for being too wiggly. A simple example is a smoothing spline. Smoothing splines essentially minimize the penalized sum of squares,

$$||\mathbf{y} - \mathbf{f}||^2 + \lambda \times penalty, \tag{2.5}$$

in order to prevent the fit, $\mathbf{f}$, from being too wiggly [116]. Here, λ is a smoothing parameter that must be estimated (e.g., by leaving one datum out, determining how well the missing datum is estimated from the rest of the data, and repeating for each datum, which is referred to as a cross-validation method [118]), $||\cdot||$ is the Euclidean norm, $\mathbf{f}$ is a vector approximating $\mathbf{y}$, and each component of $\mathbf{f}$ can be an additive sum of different functions of combinations of covariates, $\mathbf{c}_i$, which define the GAM.

TPS example

As an example, we can define a GAM that closely approximates a single, continuous, one-dimensional profile of salinity from discrete bottle samples. Using the above notation, we let $\mathbf{c}_0$ = pressure (p), and estimate $\mathbf{y}$ = salinity (S), with the functional form,

$$S = f_*(p) \tag{2.6}$$

where f_* is called a 'smoother'. Smoothers can be thought of as additive effects of single (nonlinear functions of) covariates. The smoother (s), in this case, places a form,

$f_*(p) = \sum_k \beta_k h_k(p)$, on the salinity, S, with basis functions, $h_k(\cdot)$ and regression coefficients, β_k. As described below, we have to find appropriate basis functions by finding a matrix representation with rank, κ, one higher than the worst-case change from going to a rank, $\kappa-1$, and β_k are estimated with a penalized least squares estimate, optionally taking a correlation structure in a particular covariate (due to unexplained variability in unobserved quantities) into account (see below). For all practical purposes here, the smoother functions always contain only one covariate as an argument. In the case of estimating a functional relationship between salinity and pressure using a smoother, the fit will look like a spline fit using knots at each of the observations.

Another example of a GAM is to estimate the actual equation of state for the density of a water parcel. Using the above notation, we let $\mathbf{c}_1$ = temperature (T), $\mathbf{c}_2$ = salinity (S), and $\mathbf{c}_3$ = pressure (p), and estimate $\mathbf{y}$ = density, both by neglecting (σ_t) and considering (ρ_θ) compressibility, with the functional forms,

$$\sigma_t = f_0(T, S) \tag{2.7}$$

$$\rho_\theta = f_1(T, S) + f_2(p, T) + f_3(p, S), \tag{2.8}$$

where each f_i ($i = 0, 1, 2, 3$) is a 'tensor product' of two covariates. Tensor products can be thought of as sums of multiplicative effects between two (nonlinear functions of) covariates. The tensor products (te), in this case, take on the forms, $f_0(T, S) = \sum_i \sum_j \xi_{ij} h_{1,i}(T) h_{2,j}(S)$, $f_1(T, S) = \sum_i \sum_j \xi'_{ij} h_{1',i}(T) h_{2',j}(S)$, etc. Here, $h_{1,i}$ and $h_{2,j}$, $h_{1',i}$ and $h_{2',j}$, etc. are basis functions to be determined and ξ_{ij}, ξ'_{ij}, etc. are regression coefficients that are estimated taking a wiggliness penalty and, optionally, a correlation structure into account (see below). The tensor product functions always contain more than one covariate as an argument. Here, the functions, f_0 and f_1 may not be the same, even though $\sigma_t = \sigma_t(T, S, p = 0)$ can be thought of as a special case of the equation of state, $\rho_\theta = \rho_\theta(T, S, p)$. This is because some of the variability that explains ρ_θ by f_0 is also explained by f_2 and/or f_3, which leaves f_1 to explain a subset of the variability in ρ_θ that f_0 would explain. The use of a single non-additive function of three covariates, $f_4(T, S, p)$, is computationally demanding and is not handled by the `mgcv` package in R, as described in the Appendix. A more mathematical description of these types of functions will be explained in more detail below.

Estimation with the TPS

The TPS can be used to approximate $\mathbf{f}$ with a function, $\mathbf{g}$, that minimizes the penalized sum of squares,

$$||\mathbf{y} - \mathbf{g}||^2 + \lambda J_{m,v}(g), \tag{2.9}$$

where $\mathbf{g} = (g(\mathbf{c}_1), g(\mathbf{c}_2), ..., g(\mathbf{c}_n))^T$, $\mathbf{c}_j$ $(j = 1, ..., n)$ are vectors of v observed covariates, m is the order of differentiation (typically $m = 2$), and the wiggliness penalty is

$$J_{m,v}(g) = \int dc_1 ... \int_{R^v} dc_v \sum_{\nu_1,...,\nu_v:\nu_1+...+\nu_v=m} \frac{m!}{\nu_1!...\nu_v!}(\frac{\partial^m g}{\partial c_1^{\nu_1}...c_v^{\nu_v}})^2. \tag{2.10}$$

This penalty was developed by Wahba and Wendelberger (1980) [108], where m controls the rate at which the power spectrum of g tapers off and λ then controls the half-power point of the resultant effective filter. Both m and λ need to be estimated, which is accomplished with generalized cross-validation, or GCV [70]. GCV finds the global minimum of the function,

$$V(\lambda) = n_{obs} \frac{||(\mathbf{I} - \mathbf{A}_{proj}(\lambda)\mathbf{y})||^2}{tr(\mathbf{I} - \mathbf{A}_{proj}(\lambda))^2} \tag{2.11}$$

where n_{obs} is the number of observations, $\hat{\mathbf{f}} = \mathbf{A}_{proj}(\lambda)\mathbf{y}$, $\mathbf{I}$ is the $(n_{obs} \times n_{obs})$ identity matrix, $\mathrm{tr}(\cdot)$ is the trace function, and the estimated λ (with chosen m) will minimize the average squared error of the fit. For example, using two covariates like potential temperature, θ, and salinity, S, we have $J_{2,2} = \int \int d\theta dS\{(\partial^2 g/\partial\theta^2)^2 + 2(\partial^2 g/\partial\theta\partial S)^2 + (\partial^2 g/\partial S^2)^2\}$ over all possible values of θ and S. The wiggliness penalties can expressed in a quadratic form [116], $\gamma^T \mathbf{D}\gamma$, where γ is a vector of coefficients of the marginal smooths of $\mathbf{f}$. For example, in a generalized linear model, $\lambda = 0$ and the marginal smooths would be the linear covariates themselves.

Imposing the restriction that $2m > v$, where m is the order of differentiation and v is the number of covariates, we can transform this problem, following Wood (2003) [116], into one in which we minimize

$$||\mathbf{y} - \mathbf{D}\gamma - \mathbf{T}\alpha||^2 + \lambda\gamma^T\mathbf{D}\gamma, \tag{2.12}$$

subject to $\mathbf{T}^T\gamma = \mathbf{0}$ by finding that

$$g(\mathbf{c}) = \sum_{i=1}^{n} \gamma_i \eta_{m,v}(||\mathbf{c} - \mathbf{c}_i||) + \sum_{j=1}^{M} \alpha_j \phi_j(\mathbf{c}). \tag{2.13}$$

The general solution to (2.12) is (2.13) with γ and α coefficients to be determined. Here, $M = (m+v-1)!/((m-1)!v!)$ is the number of functions that are linearly independent polynomials, $\phi_j(\mathbf{c}_i) = T_{ij}$, spanning the space of polynomials for which $J_{m,v} = 0$ and

$$\eta_{m,v}(r) = \frac{(-1)^{m+v/2+1}}{2^{2m-1}\pi^{v/2}\Gamma(m)\Gamma(m-v/2+1)} r^{2m-v} log(r) \quad \text{v is even} \tag{2.14}$$

$$\eta_{m,v}(r) = \frac{\Gamma(v/2-m)}{2^{2m}\pi^{v/2}\Gamma(m)} r^{2m-v} \quad \text{v is odd}$$

where $\Gamma(\cdot)$ is the Gamma function with $\Gamma(a) = (a-1)!$ for any non-negative integer a, $r = ||\mathbf{c}_i - \mathbf{c}_j||$ and $D_{ij} = \eta_{m,v}(||\mathbf{c}_i - \mathbf{c}_j||)$ [118]. The reason why there are two terms in (2.13) is that there is a condition on the wiggliness of the fitted function which makes the solution more complicated than least-squares. The solution can be broken up into two terms because one of them, $\mathbf{D}$, can be written analytically while the other, $\mathbf{T}$, places a condition on γ and α.

In practice, (2.12) is solved by transforming to a family of low-rank smoothers in which we find γ such that the worst possible simultaneous changes in basis (to a lower rank, κ),

$$max_{\gamma \neq \mathbf{0}}\{||(\mathbf{D} - \tilde{\mathbf{D}}_\kappa)\gamma||/||\gamma||\}$$
$$max_{\gamma \neq \mathbf{0}}\{\gamma^T(\mathbf{D} - \hat{\mathbf{D}}_\kappa)\gamma/||\gamma||^2\},$$

are minimized. Here, $\tilde{\mathbf{D}}_\kappa = \mathbf{D}\mathbf{\Gamma}_\kappa\mathbf{\Gamma}_\kappa^T$, $\hat{\mathbf{D}}_\kappa = \mathbf{\Gamma}_\kappa\mathbf{\Gamma}_\kappa^T\mathbf{D}\mathbf{\Gamma}_\kappa\mathbf{\Gamma}_\kappa^T$, and $\mathbf{\Gamma}_\kappa$ is defined by $\gamma = \mathbf{\Gamma}_\kappa\gamma_\kappa$ [116]. Finally, estimation of α follows. The effective degrees of freedom as determined by estimation (or fixing) of the smoothing parameter, λ, via cross-validation are constrained above by κ. If λ is too close to zero, κ is increased. One reason 'spline' appears in the method's name is that when $\lambda = 0$, this can be thought of as an extension of a regression spline. In summary, the parameters estimated, in order, are: m is selected (subject to $2m > v$), a suitable basis is found through unit decrements of κ (selecting the κ one greater than the rank at which there is the greatest change in least-squares error), γ (through its relationship with $\mathbf{\Gamma}_\kappa$ and ultimately minimizing the scaled difference between $\mathbf{D}$ and $\hat{\mathbf{D}}_\kappa$), α (by minimizing (2.12)), and λ (with GCV). This is repeated many times in order to find joint estimates of each of these quantities. In this study, we use an R implementation (see Appendix), where the estimation is performed within the `gamm` function when there is not

too much data and `bam` when there are thousands of data samples. The default standard error estimation procedure is described below.

Confidence intervals on the TPS estimates

To form confidence intervals on the estimates of $\mathbf{f}$, following Wood (2006a) [117], we choose bases for the smooth functions in our GAM such that the penalized least squares objective, (2.5), can be written as

$$||\mathbf{W}^{1/2}(\mathbf{g}_0(\mathbf{E}[\mathbf{y}]) - \mathbf{X}\beta)||^2 + \sum_{i=1}^{n_{obs}} \lambda_i \beta^T \mathbf{S}_i \beta. \tag{2.15}$$

Here, $\mathbf{W}$ is a positive definite weighting matrix, $\mathbf{E}[\cdot]$ is the expectation operator, $\mathbf{S}_i$ is a positive semi-definite matrix that is used to quantify departure from smoothness (i.e., (2.10) after choosing bases such that the penalty term is of the form given in (2.15)), $\mathbf{g}_0(\cdot)$ is a known link function (i.e., specifying a relationship between the covariates and the mean of the distribution function), $\mathbf{y}$ is a realization of $\mathbf{Y}$ ($Y_i \sim$exponential family such as normal or Gamma), $\mathbf{X}$ is a design matrix (e.g., first column is all ones for an intercept), and β are our unknown parameters. Then using a Bayesian estimation technique with many samples,

$$\beta|\mathbf{y} \sim N(\hat{\beta}, (\mathbf{X}^T\mathbf{W}\mathbf{X} + \sum_{i=1}^{n} \lambda_i \mathbf{S}_i)^{-1}\sigma^2) \tag{2.16}$$

for σ^2 being the scale parameter (of the variance), $\hat{\beta}$ maximizing $-l(\beta) + 1/2 \sum_{i=1}^{n_{obs}} \lambda_i \beta^T \mathbf{S}_i \beta$ with respect to β for $l(\cdot)$ being the logarithm of the likelihood function for the exponential family of the Y_i. An estimator for σ^2 is known for exponential family distributions:

$$\hat{\sigma}^2 = \frac{1}{tr(\mathbf{I} - \mathbf{B})} ||\mathbf{W}^{1/2}(\mathbf{y} - \hat{\mathbf{E}}[\mathbf{Y}])||^2 \tag{2.17}$$

$$\mathbf{B} = \mathbf{X}^T(\mathbf{X}^T\mathbf{W}\mathbf{X} + \sum_{i=1}^{n_{obs}} \lambda_i \mathbf{S}_i)^{-1}\mathbf{X}^T\mathbf{W}$$

where $\mathbf{I}$ is the $(n_{obs} \times n_{obs})$ identity matrix, $\mathrm{tr}(\cdot)$ is the trace function, and $\hat{\mathbf{E}}[\cdot]$ is the expectation vector of the Y_i. Finally, to find the cumulative distribution function, $F_{cdf}(\cdot)$ of a function, $G(\beta)$, determined by the bases of our GAM, because the distribution of $\beta|\mathbf{y}$ is known, a set of random vectors, $\{\beta_i^* : i = 1, ..., n^*\}$ is simulated, from the distribution of

$\beta|\mathbf{y}$ so that the estimate of $F(\cdot)$ is

$$\hat{F}_{cdf}(g) = \frac{1}{n^*} \sum_{i=1}^{n^*} H(g - G(\beta_i^*)) \tag{2.18}$$

where

$$H(x) = \begin{cases} 1, & \text{if } x \geq 0 \\ 0, & \text{if } x < 0. \end{cases}$$

It is from this estimated cumulative distribution function that the TPS standard error estimates (and ultimately confidence intervals) are derived. The lack of dependence upon length-scales or signal-to-noise ratios is an advantage of the TPS over OM when data are sparse. The non-requisite use of restricted parametric forms for covariance functions associated with using the TPS is another advantage over OM when estimating mapping errors.

The way to introduce a correlation structure for estimation of standard errors is similar to that of linear mixed model theory, as $\mathbf{y} = \mathbf{g} + \mathbf{R}\mathbf{b} + \epsilon$. Here, $\mathbf{g}$ is our TPS estimate from a GAM as above, and $\mathbf{R}$ are random effects with $\mathbf{b}$ coefficients that determine the correlation structure, which can be left unstructured at the expense of losing degrees of freedom for each coefficient estimate [118]. For example, we could regard the effects of potential density, ρ_θ, on [CFC-11] in our data as a random sample of the effects of all potential densities across an entire ocean basin on [CFC-11] by including potential density as a random effect. Our expression for $\mathbf{B}$ in (2.17) would then place a particular form on $\mathbf{W}$ other than the identity or inverse variance from an ordinary least squares estimator.

TPS implementation

In implementing TPS-estimation, we use the public-domain R package, `mgcv`, as described in the Appendix. First, a GAM is specified, reasoned *a priori* instead of *ad hoc*, and trained. Raw data is used to train this model. Next, we estimate the quantity of interest according to the trained model at the points where the values of the covariates are known. Because the tracer quantities of interest also tend to be spatially distributed along isopycnals, when we use tracers to estimate other tracers in a GAM, we expect that the unexplained variability will be more correlated within a constant potential density layer than between different potential density layers. Therefore, when using tracers to estimate other tracers, we use

an AR(1) correlation structure in potential density. We first define a vector $X_{1:N-1}$ that is composed of the raw tracer values, ordered such that potential density is non-decreasing where N is the number of data points. We then find the correlation, c_*, between this vector and its one-index lagged counterpart $X_{2:N}$. For an AR(1) correlation in ordered potential density, $\rho_{\theta,1:N}$, the tracer field along the $\rho_\theta = \rho_{\theta,i}$ isopycnal will have correlation, $c_*^{|i-j|}$, with the tracer field along the $\rho_\theta = \rho_{\theta,j}$ isopycnal.

The first fields to be estimated are potential density, (potential/*in situ*) temperature, and salinity by assuming the functional form,

$$G = f_{te}(l,p) + f_s(d) \tag{2.19}$$

where G is the potential density/temperature/salinity, the scaled distance d is $d = \sqrt{(lcos(l\pi/180) \times 6.3781 \times 10^6/s_0)^2 + p^2}$ where $s_0 = max_l\{|lcos(l\pi/180) \times 6.3781 \times 10^6|\}/max_p\{|p|\}$ for A16 and SR03 (and is multiplied by an additional $l/\sqrt{24.5^2 + l^2}$ for A05), p is the pressure in decibars, and l is the latitude (longitude for A05) in degrees. The scaled distance appears quadratic in pressure, which approximately mimics the vertical structure of isopycnals. Isopycnal, isothermal, and isohaline surfaces generally appear to slope along constant d surfaces for our sections, and we pick up the rest of the variance in G by including the tensor product between latitude and pressure. The tensor product (te), in (2.19), in this case, places a form on the covariates, $\sum_i \sum_j \xi_{ij} h_{1,i}(l) h_{2,j}(p)$, for the $f_{te}(l,p)$ term, where $h_{1,i}$ and $h_{2,j}$ are basis functions determined in (2.13) and ξ_{ij} are regression coefficients that are estimated taking the AR(1) in ordered potential density correlation structure into account. The smoother (s), in this case, places a form, $\sum_k \beta_k h_k(d)$, on $f_s(d)$ in (2.19), likewise with h_k basis functions and β_k regression coefficients.

When the potential temperature, θ, and salinity, S, fields have been estimated on a grid via either OM or (2.19) for a GAM in using the TPS, we can estimate the rest of the fields of interest on that grid as well. First, we estimate the CFC-11 concentrations by assuming the functional form:

$$CFC = f_{1,te}(d,\theta) + f_{2,te}(d,S). \tag{2.20}$$

The physical reasoning for this form is two-fold: First, the CFC-11 concentrations are approximately distributed along isopycnals where the density, $\rho \approx \rho_0(1 - a\theta(d) + bS(d))$, for

reference density ρ_0, thermal expansion coefficient a, and haline contraction coefficient b. Second, the solubility function for CFCs, $F_{sol}(\theta, S)$, according to Warner and Weiss (1985) [111] is dependent upon both potential temperature and salinity.

Next, we estimate the pCFC ages by assuming the functional form:

$$\tau_{CFC} = f_{1',te}(CFC, \theta) + f_{2',te}(CFC, S). \tag{2.21}$$

The physical reasoning for this form is that the pCFC ages have been determined through conversion of concentrations to partial pressures through $pCFC = CFC/(F_{sol}(\theta, S)f_{sat})$ for saturation level, f_{sat} (here, assumed to be 100%), which in turn are compared to a non-linear atmospheric time series.

2.3 Results and Discussion

When we refer to OM estimates, we are referring to the estimates from (2.1), and when we refer to TPS-derived estimates, we are referring to the solution to the variational problem given in (2.12). In this section, we first examine the most sensitive specifications of the GAM, as implemented in R (see Appendix) using the choices described in "TPS implementation" of Section 2.2.2, and how the TPS method compares with a classic multilinear regression technique (ordinary least squares) as a statistical model. We next consider how well the TPS and OM methods perform when only a subset of the bottle data is used, leaving large spatial data gaps. Finally, we demonstrate what can happen if a slightly biased mapped quantity (interpolated with TPS rather than with OM) is used as a covariate with TPS in the context of pCFC-ages. This is intended to highlight how the errors do not become inflated by bias, but rather the errors reflect what we expect from uncertainty associated with mapping error, given the data that went into training the GAM.

The key assumptions behind using the TPS are that the field of interest is smooth and all of the variability in the estimated quantity of interest can be explained by the covariates used in its GAM [118]. While OM depends on the spatial statistics of the domain over which a field is being estimated, the TPS, in our application, relies on relationships between the bottle data samples (not the spatial statistics), so the TPS estimates and associated errors reflect the physics and chemistry we enter through functional dependencies. If, for instance,

water mass properties according to a temperature-salinity phase diagram vary outside the sampled range, or have an unpredictable relationship, then the TPS estimate would be invalid. The rank, κ, of the reduced rank representation found in the TPS procedure often (except for a few of the parameter estimates such as λ which are on the order of ten) equals one. This does not suggest that multilinear regression would produce equivalent estimates, as addressed below. This merely suggests that a basis representation with a lower rank is sufficient to functionally estimate our fields, given the smoothness penalty. The TPS can also be used to construct hindcasts, but we focus on interpolation and mapping error estimation here.

2.3.1 Specification of a Regression Model

First, to justify our use of a GAM with random effects ("Confidence intervals on the TPS estimates" of Section 2.2.2), we compare an ordinary least squares fit using a generalized linear model (GLM) to two different GAM fits (one with random effects in potential density and one without) for all of the CFC-11 bottle data with quality 2 rating along SR03 in 2001 (Figs. 2.1a-c), A16 in 2003 (Figs. 2.1d-f), and A05 in 1998 (Figs. 2.1g-i). Residuals have clearly non-random structure in Figs. 2.1c, 2.1f, and 2.1h-i, but not in Figs. 2.1a-b, 2.1d-e, or 2.1g, suggesting that the best model fits are those of the TPS with the inclusion of random effects of potential density. The reason for the trend in the residuals shown in Fig. 2.1h is that we have not explained all of the variability with our GAM, leaving out a covariate that affects potential density (which here looks like a smoother function of either temperature or salinity). A05 is different here because [CFC-11] is below the detection limit in a large part of the deep eastern portion of the cross-section, which makes the [CFC-11] field less smooth than the other hydrographic sections. Although different combinations of possible smoothers and tensor products in (2.20) change the estimated fields and associated errors, we were unable to find a different functional specification that eliminates the remaining residual structure. We prefer (2.20) because it is physically motivated.

Also, we justify our use of a GAM with random effects according to the specification in (2.21) for pCFC-ages (specifically, pCFC-11-ages) by comparing a GAM with an alternative

functional form ($\tau_{CFC} = f_{te}(\theta, S) + f_s(CFC)$) to two different GAM fits (each with random effects but one with an AR(1) correlation structure and one with a Gaussian correlation structure). The Gaussian correlation structure uses a worse-case scenario measurement error as a 'nugget,' which is the height of the discontinuity at the origin (for a point correlated with itself as opposed to a point correlated with an adjacent one) of the variogram [107], a function that quantifies spatial correlation. The residuals from the two different GAM fits using (2.21) shown in Figs. 2.2a-b, 2.2d-e, and 2.2g-h are virtually identical, suggesting that incorporating measurement errors in our mapping error estimates is not significant and neither is our choice of correlation structure (between Gaussian and AR(1) on potential density). The residuals have less structure to them using the GAM with functional form given by (2.21) (Figs. 2.2a-b, 2.2d-e, and 2.2g-h) than using the GAM with an alternative functional form (Figs. 2.2c, 2.2f, and 2.2i), but the residuals from the GAM fits for the pCFC-ages here have more structure than the CFC-11 concentration fits shown in Fig. 2.1, particularly at the low and high density ends of the potential density spectrum. The reason for this is that the pCFC-ages can never be less than zero or greater than about sixty years from the raw data. This places a sharp cut-off on the pCFC-age field at both ends as opposed to the pCFC-ages being able to smoothly decrease to values less than zero or increase to values beyond about sixty years.

The functional dependencies of [CFC-11] two covariates at a time show that the standard errors are relatively small (it is difficult to separate out the black and grey hypersurfaces) in Figs. 2.3a, 2.3d, and 2.3g suggests that the GAM specified by (2.20) sufficiently captures the variability in [CFC-11] across each hydrographic cross-section. A feature that should be articulated, demonstrated by the contour magnitudes in Figs. 2.3b-c, 2.3e-f, and 2.3h-i, is that attempting to estimate [CFC-11] outside of the range of values used to train the GAM (e.g., $d \geq 6000$ and either potential temperatures greater than 10^oC or salinities less than 34 PSU along SR03) yields an unrealistic estimate (e.g., greater than any [CFC-11] seen in the ocean: 20 pmol/kg).

Similarly, we inspect the functional dependencies of pCFC-ages two covariates at a time (Fig. 2.4) with standard error estimates also shown for the projection showing only potential temperature and salinity (Figs. 2.4a, 2.4d, and 2.4g). It is evident that the standard errors,

due to the structure in the residuals shown in Figs. 2.2a, 2.2d, and 2.2g, are larger at the extreme ends of the potential density spectrum. The contour magnitudes in Figs. 2.4b-c, 2.4e-f, and 2.4h-i, actually become negative for values of [CFC-11] greater than 3 pmol/kg and potential temperatures greater than 15^oC along A16 (Fig. 2.4e) as well as for values of [CFC-11] greater than 5 pmol/kg and potential temperatures greater than 8^oC along SR03 (Fig. 2.4b). These values for the covariates are outside of the range used to train the GAM for pCFC-ages. Lastly, there is an insignificant amount of variability explained by salinity along A05 (Fig. 2.4i), given the other covariates included in the GAM for pCFC-ages.

2.3.2 Sparsity of Spatial Sampling

To map [CFC-11] along SR03 for the 1991 cruise, using both OM and TPS methods, we use a grid from 44^oS to 67^oS in one degree steps and from 0 dbars (approximately meters) to 4000 dbars in 50 dbar steps. We subsampled the data to test the sensitivity of the results to bottle data distribution (Fig. 2.5). We took out two-thirds of the raw observations (row 1 of Fig. 2.5). Then we took out one-fourth of the raw observations between 1000 and 2000 meters, one-half of the raw observations between 2000 and 3000 meters, and three-fourths of the observations between 3000 and 4000 meters (row 2 of Fig. 2.5). Lastly, we took out all of the raw observations except for the shallowest, the deepest, the northernmost, and the southernmost (row 3 of Fig. 2.5). We compare these results with maps using all of the raw observations (row 4 of Fig. 2.5). Although the TPS smooths to a greater extent than OM, for each subsampling scheme, the TPS estimates preserve many of the essential features, including the bottom water, even when training the GAM with very sparse data. Additionally, as evidenced by row 3 of Fig. 2.5, the standard errors associated with the TPS estimates reflect the sparsity of the bottled samples, becoming larger in regions with few to no observations. The standard errors are such that the 95% confidence intervals based on using one degree of freedom on the TPS estimates contain the objectively mapped fields almost everywhere (outlined in white in column b of Fig. 2.5 where they do not).

2.3.3 Spatial Features of Interpolation Estimates

To demonstrate the importance of using the least biased mapped quantities as covariates, we compare pCFC-ages from OM (column a of Fig. 2.6) and TPS using covariates estimated from OM (column c) and TPS (column b) from two occupations of SR03. The [CFC-11] and pCFC-ages from OM and TPS estimates differ somewhat, especially in the CFC-11 plots (e.g., Fig. 2.5) where Antarctic Bottom Water comes down the continental slope and in the relatively CFC-depleted interior in the pCFC-age plots (columns a-c of Fig. 2.6). Not only do the TPS-estimated pCFC-ages appear to be more similar to the OM-estimated pCFC-ages when covariates gridded from OM are used, but the standard errors shrink as well.

To investigate how well the errors estimated by the TPS method express the features we expect out of a mapping error, we compare the standard error estimates in columns d and e of Fig. 2.6. The TPS-derived errors can be conservative (column d) because they more accurately reflect error that results from attempting to estimate outside of the domain of covariate values used in the training of the statistical model. The relatively smooth maps in column b of Fig. 2.6 have a more limited range of combinations of covariate values used in the GAM because extrema in the [CFC-11], temperature, and salinity fields are attenuated and shifted in their TPS-derived estimates. The bias in the TPS-derived estimates of the GAM covariates influences the magnitudes of the standard errors (column d of Fig. 2.6), causing them to increase at depths below 2000 meters and latitudes south of 55^{o}S, despite how spatial coverage improves going from 1991 to 2008 (not shown). The magnitudes of the standard errors using the OM estimates of the GAM covariates (column e of Fig. 2.6), on the other hand, simply reflect the spatial coverage of the bottle data (not shown).

To investigate the influence of section selection on our conclusions about the general agreement between mapped fields and the utility of the TPS error estimates, we turn our attention to the four occupations of A16 (Fig. 2.8) and the three occupations of A05 (Fig. 2.7). The large gap in observations apparent in 1998 along A16 is reflected in the errors (see Fig. 2.8c3), and the pCFC-age in this gap (at most depths) is underestimated by OM relative to the TPS estimates (not shown). On the other hand, the [CFC-11] that hits a

maximum at intermediate depths just east of the western boundary current along A05 in Fig. 2.7b3 is erroneously shifted deeper by the TPS estimates. A weakness in the TPS estimates (for interpolation purposes) appears in Fig. 2.8b2-b3, demonstrating the extent to which eddies are smoothed out (just south of the Rockall-Hatton Plateau). To further demonstrate how vertical profiles are smoothed out with the TPS method compared to OM, we turn to Fig. 2.9. The local maxima at intermediate depths can be attenuated due to smoothing where extrema occur with TPS, as is shown along A05 for these [CFC-11] profiles. The places where the TPS method smooths over extrema (e.g., Fig. 2.9c) are not reflected in the standard errors, but the error field does reflect greater sparsity in data coverage.

While the large-scale features are always captured by the TPS-estimated fields, the TPS-estimated standard errors only sometimes increase where eddies are smoothed out (e.g., not in 1988 or 2003 along A16), and there are more serious estimate discrepancies where there are extrema in the fields (e.g., near the boundaries of the spatial domains). The OM estimates are not contained within anti-conservative (more than thirty degrees of freedom) TPS-derived 95% confidence intervals on TPS estimates where there are extrema in the OM-estimated fields (not shown), demonstrating that the TPS-estimated standard errors do not properly represent uncertainty due to TPS-based smoothing. On the other hand, using the sum of the effective degrees of freedom (df) on each of the parameters estimated by the GAM, confidence intervals (based on the TPS-estimated standard errors) on the OM estimates contain the raw data everywhere except where there are relatively large vertical gradients (e.g., Fig. 2.9b at 600 meters depth). In Fig. 2.9, $df = 36$ for SR03, $df = 43$ for A16, and $df = 37$ for A05. However, the smallest possible conservative confidence intervals (based on the TPS-estimated standard errors) on the OM estimates that contain the raw data everywhere use one degree of freedom (not shown, but can be inferred from Fig. 2.9). Furthermore, the conservative TPS-derived standard errors behave intuitively based on prior expectation of where the OM estimates are reliable. While we cannot conclude that we should use the sum of the effective degrees of freedom on each of the parameters estimated by the GAM, we can qualitatively conclude that OM is a more reliable vertical interpolator of observed data, and that the uncertainty associated with

mapping can be quantified with TPS-derived standard error estimates.

2.4 Conclusions

The TPS method of mapping discrete tracer samples along oceanographic sections for the purpose of interpolation has distinct strengths and weaknesses when evaluated in comparison to the more standard tool of OM. OM is accurate where spatial coverage is dense, but has difficulty with spatially sparse data coverage. The TPS could be more accurate than OM where data are sparse, but could have difficulty with data coverage that does not span the full range of values the GAM covariates take. Accordingly, the regions where the TPS and OM procedures differ most from one another are where there is a relative dearth of data, and also where there are extrema in the estimated field, including the edges of the maps and just below the thermocline. Differences in these locations are due to the fact that splines are effectively low-pass filters, which move and attenuate extrema when using the TPS method. Extrema are also smeared out across regions that are not well sampled with observations using either mapping method. A set of densely sampled observations, on the other hand, is smoothed out by the TPS relative to OM to prevent "over-training," meaning that the TPS would have a more difficult time estimating a field using covariate values outside of the domain of covariate values used to train the GAM. More observations with the TPS do not serve the same benefit as they would with OM for tracer field estimation. When the distance between samples is larger than the correlation length scale or the signal to noise ratio is small, the TPS can outperform OM because the TPS is a machine learning algorithm trained on available data rather than an interpolation algorithm.

While the TPS estimates with TPS-derived confidence intervals do not always contain the observations, we have found that the TPS maps re-create the essential large-scale features of [CFC-11] and inferred pCFC-ages and that the OM estimates with TPS-derived confidence intervals always contain the observations. The TPS standard errors are large where the data coverage is relatively poor, but do not indicate where smoothing has removed small-scale or boundary extrema. These biases can occur when estimating a quantity using combinations of covariates outside of their domains used in training the statistical model. Because the residuals from the OM estimates are smaller than the TPS estimates,

Table 2.1: Tracer data collected in the listed years along each of the cross-sections of interest.

WOCE Section:	Latitude(s)	Longitude(s)	(T, S, CFC-11)	(δ^3He,Tr,^{4}He)
A16(N)	34^oN to 63^oN	20^oW	1988, 1993, 1998, 2003	1988
A05	24.5^oN	79^oW to 14^oW	1992, 1998, 2004	N/A
SR03	67^oS to 43^oS	135^oE to 147^oE	1991, 1995, 2001, 2008	N/A

the standard errors associated with the TPS estimates are conservative estimates of the OM mapping error. Overall, when data is available and not extremely sparse (e.g., not as sparse as row 3 of Fig. 2.5), we recommend using field estimates from OM, because it is less likely to be biased, and imputing the TPS-derived standard errors, because they more accurately (and conservatively) represent the uncertainty in mapping expected from the sparsity of spatial samples.

The TPS can also be used for hindcasting unobserved tracer quantities, but the physics needs to be included in the associated GAM. If a specific tracer is the field of interest, for example, and its dynamical nature can be explained by other fields, the TPS allows for inclusion of the potential evolution of this tracer. However, the temporal evolution of a tracer such as Tritium, or ^{3}H, can be included in another way such as using the fact that ^{3}H decays exponentially without an additional source (see next chapter). Creating a hindcast using the TPS should not be used as a substitute for observations, but rather is one option that can be used when particular estimable quantities are not included in a suite of sampled quantities. The TPS will always aim to find the optimal balance between bias and variance of an estimate. This is essentially what allows for estimating fields from a range of covariate values within those used to train a GAM, without needing to be explicit about spatial information. Caution must be taken when using the TPS to create a hindcast from a range of covariate values that extend outside of those used to train the associated GAM.

Table 2.2: Measurement errors of quantities of interest. Here, 'PSU' stands for practical salinity units (parts per thousand), 'pmol' stands for pico-moles, 'kg' stands for kilograms, '^{o}C' stands for degrees Celsius, and 'dbar' stands for decibars (approximately equivalent to meters of depth).

Variable:	Units	Measurement error
CFC-11	pmol/kg	0.02× [CFC-11]
Temperature	^{o}C	0.0002
Salinity	PSU	0.0002
Pressure	dbar	0.03

Figure 2.1: Along SR03 in 2001.9 (a-c), along A16 in 2003.5 (d-f), and along A05 in 1998.1 (g-i), shown are the residual CFC-11 concentrations (difference between the CFC-11 concentrations from raw observations and the fitted CFC-11 concentrations) according to the TPS that uses a GAM with AR(1) correlations structure and random effects in potential density according to (2.20) (a, d, and g), the TPS that uses a GAM with AR(1) correlation structure in potential density and no random effects according to (2.20) (b, e, and h), and ordinary least squares that uses a GLM with the same covariates and interactions between covariates (c, f, and i). The GLM is defined as $CFC = b_0 + b_1 d + b_2 \theta + b_3 d + b_4 salt + b_{12} d \times \theta + b_{34} d \times salt$ for θ =potential temperature, d =scaled distance from the equator, and $salt$ =salinity.

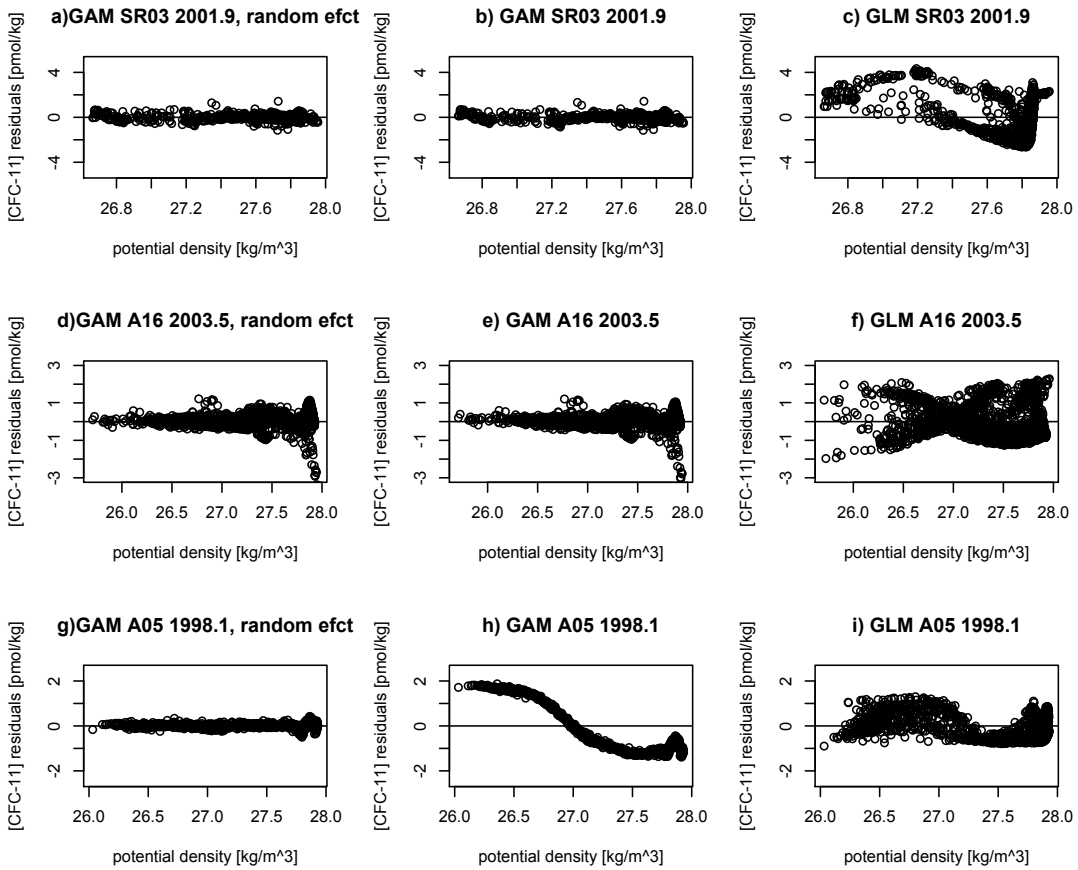

Figure 2.2: Along SR03 in 2001.9 (a-c), along A16 in 2003.5 (d-f), and along A05 in 1998.1 (g-i), shown are the residual pCFC-ages (difference between the pCFC-ages from raw observations and the fitted pCFC-ages) according to the TPS that uses a GAM with AR(1) correlation structure and random effects in potential density according to (2.21) (a, d, and g), the TPS that uses a GAM with Gaussian correlation structure in potential density and a nugget equal to 0.02 times the highest CFC-11 concentration according to (2.21) (b, e, and h), and the TPS that uses a GAM with AR(1) correlation structure and random effects in potential density with the same covariates but different interactions between covariates (c, f, and i). The different GAM is defined as $\tau_{CFC} = f_{te}(\theta, S) + f_s(CFC)$ for θ =potential temperature, CFC =CFC-11 concentration, and $salt$ =salinity.

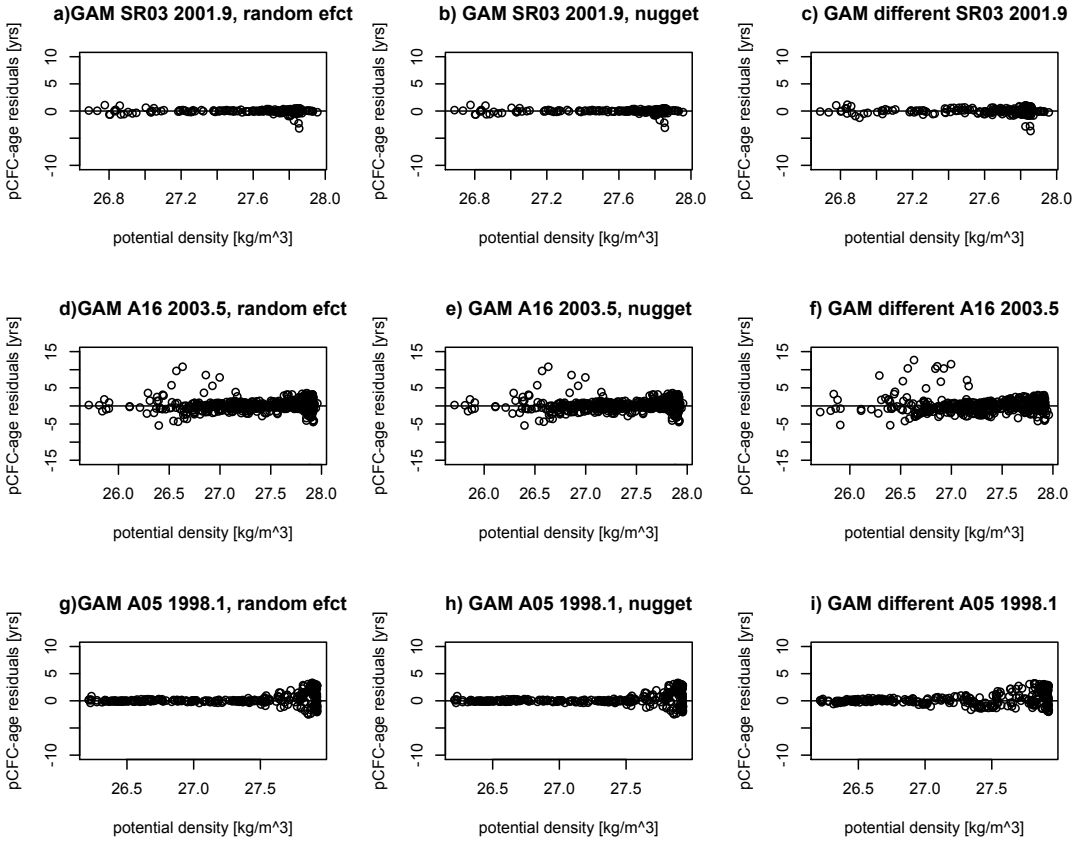

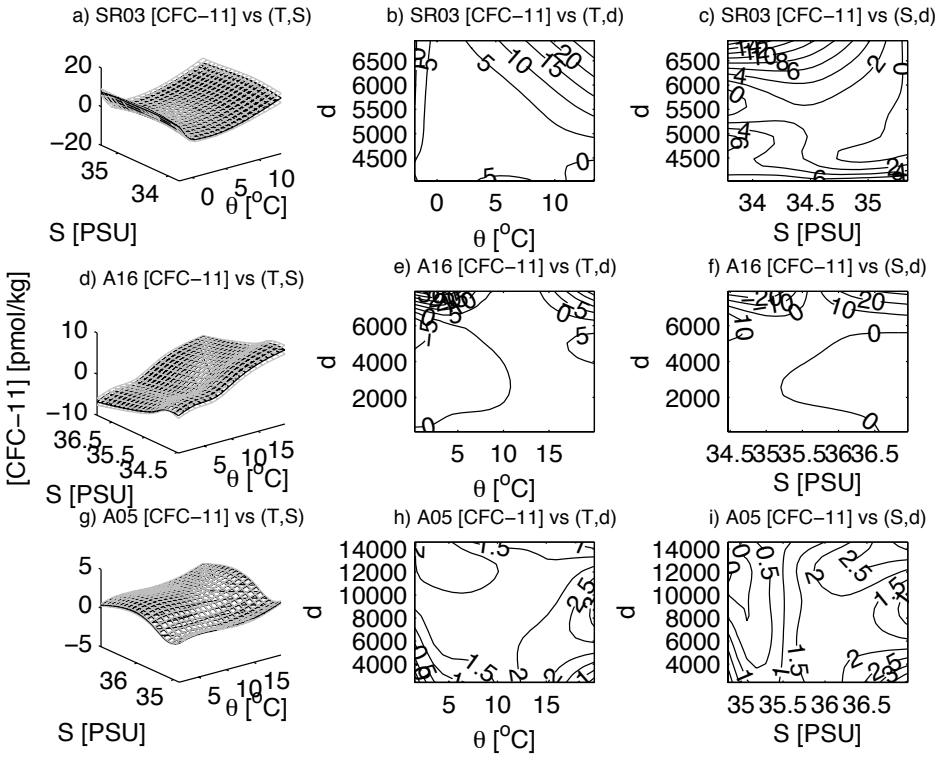

Figure 2.3: Along SR03 in 2001.9 (a-c), along A16 in 2003.5 (d-f), and along A05 in 1998.1 (g-i), shown are the CFC-11 concentrations estimates (black) according to the TPS that uses a GAM with AR(1) correlations structure and random effects in potential density according to (2.20) versus θ =potential temperature and S =salinity (a, d, and g) with upper and lower bounds according to one TPS-derived standard error (grey), CFC-11 (b, e, and h) contours shown against d and θ, and CFC-11 (c, f, and i) contours shown against d and S.

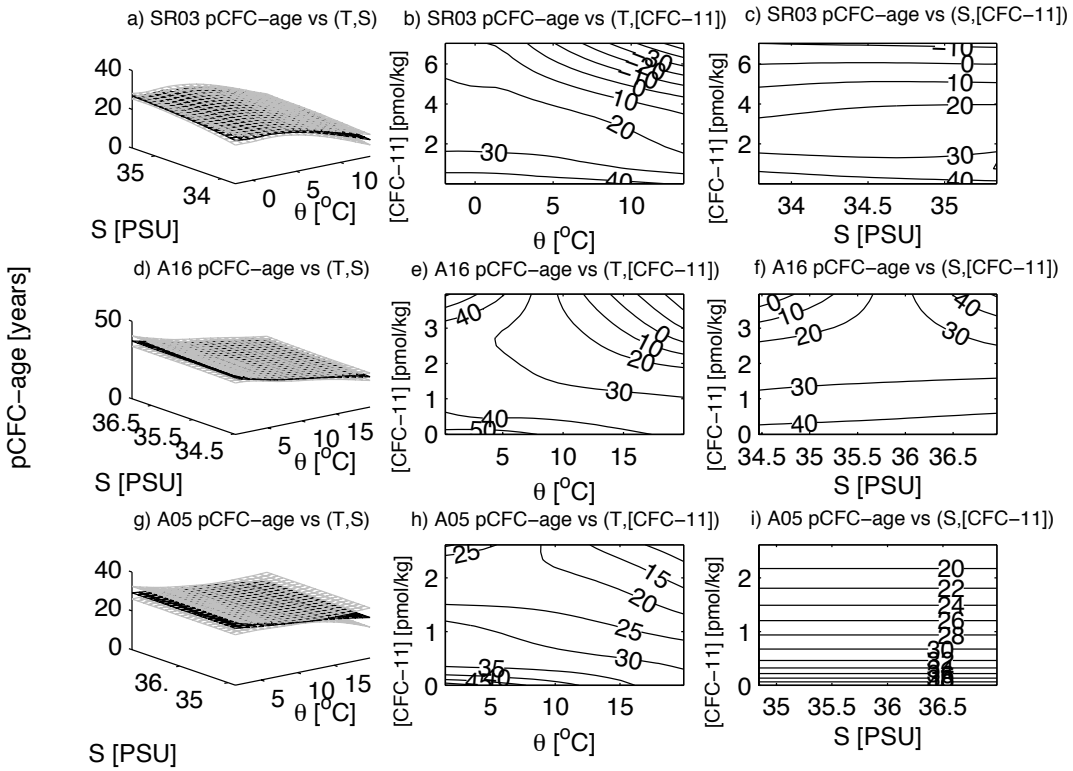

Figure 2.4: Along SR03 in 2001.9 (a-c), along A16 in 2003.5 (d-f), and along A05 in 1998.1 (g-i), shown are the pCFC-age estimates (black) according to the TPS that uses a GAM with AR(1) correlations structure and random effects in potential density according to (2.21) versus θ =potential temperature and S =salinity (a, d, and g) with upper and lower bounds according to one TPS-derived standard error (grey), pCFC-age (b, e, and h) contours shown against CFC-11 and θ, and pCFC-age (c, f, and i) contours shown against CFC-11 and S.

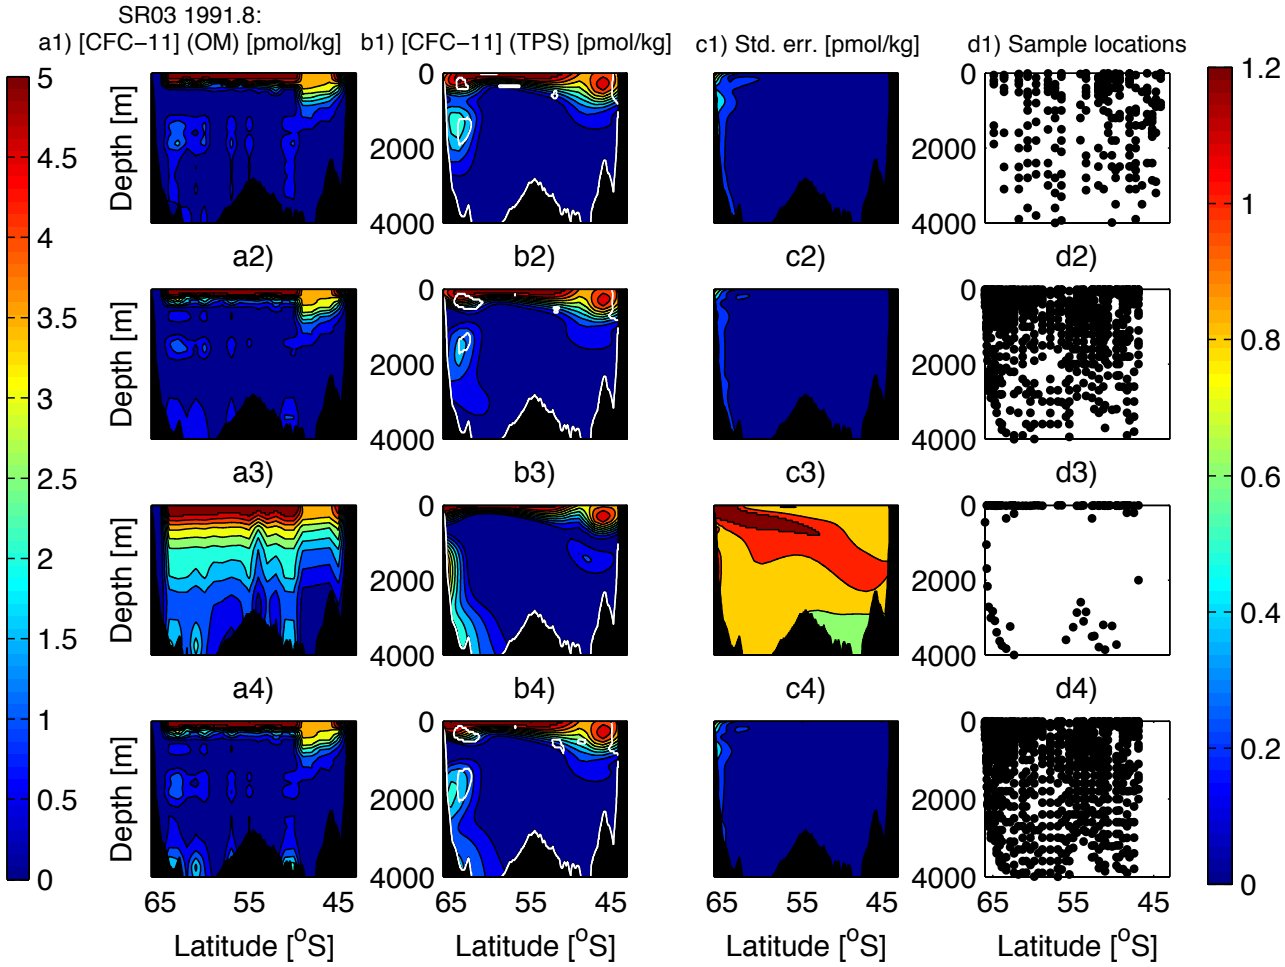

Figure 2.5: The CFC-11 concentrations from objective mapping (column a) and thin-plate spline estimates (column b) and their mapping errors (column c) along SR03 for 1991; also shown are the locations of where the data was sampled (column d) and white contours (column b) bounding the regions where 95% point-wise confidence intervals based on using one degree of freedom on the thin-plate spline estimates do not contain the objective mapping estimates. In row 1, two-thirds of the raw observations were uniformly taken out. Row 2 took out an increasing number of observations with depth. Row 3 took out all of the observations except for those along the boundaries. Row 4 includes all of the raw observations. Columns a and b are on the same scale, as indicated on the left side. Columns c is on the scale indicated on the right side. All columns are over the same domain.

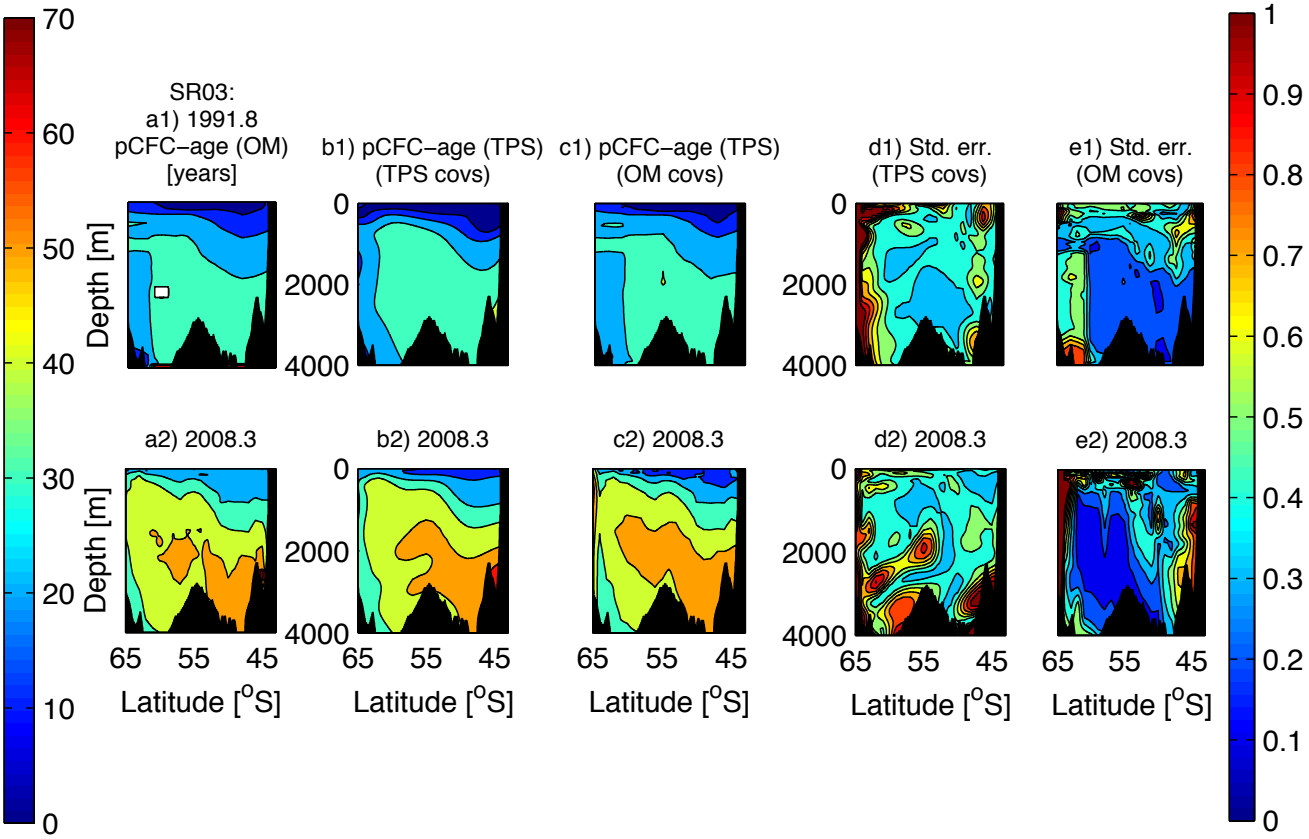

Figure 2.6: The pCFC-ages along SR03 for 1991 (row 1) and 2008 (row 2) from objective mapping (column a) and thin-plate spline estimates using thin-plate spline-derived covariates (column b) and using objectively mapped covariates (column c) and their respective mapping errors (columns d and e) derived from temperature, salinity, and CFC-11 concentrations. Columns a, b, and c are on the same scale, as indicated on the left side. Columns d and e are on the same scale, as indicated on the right side. All columns are over the same domain.

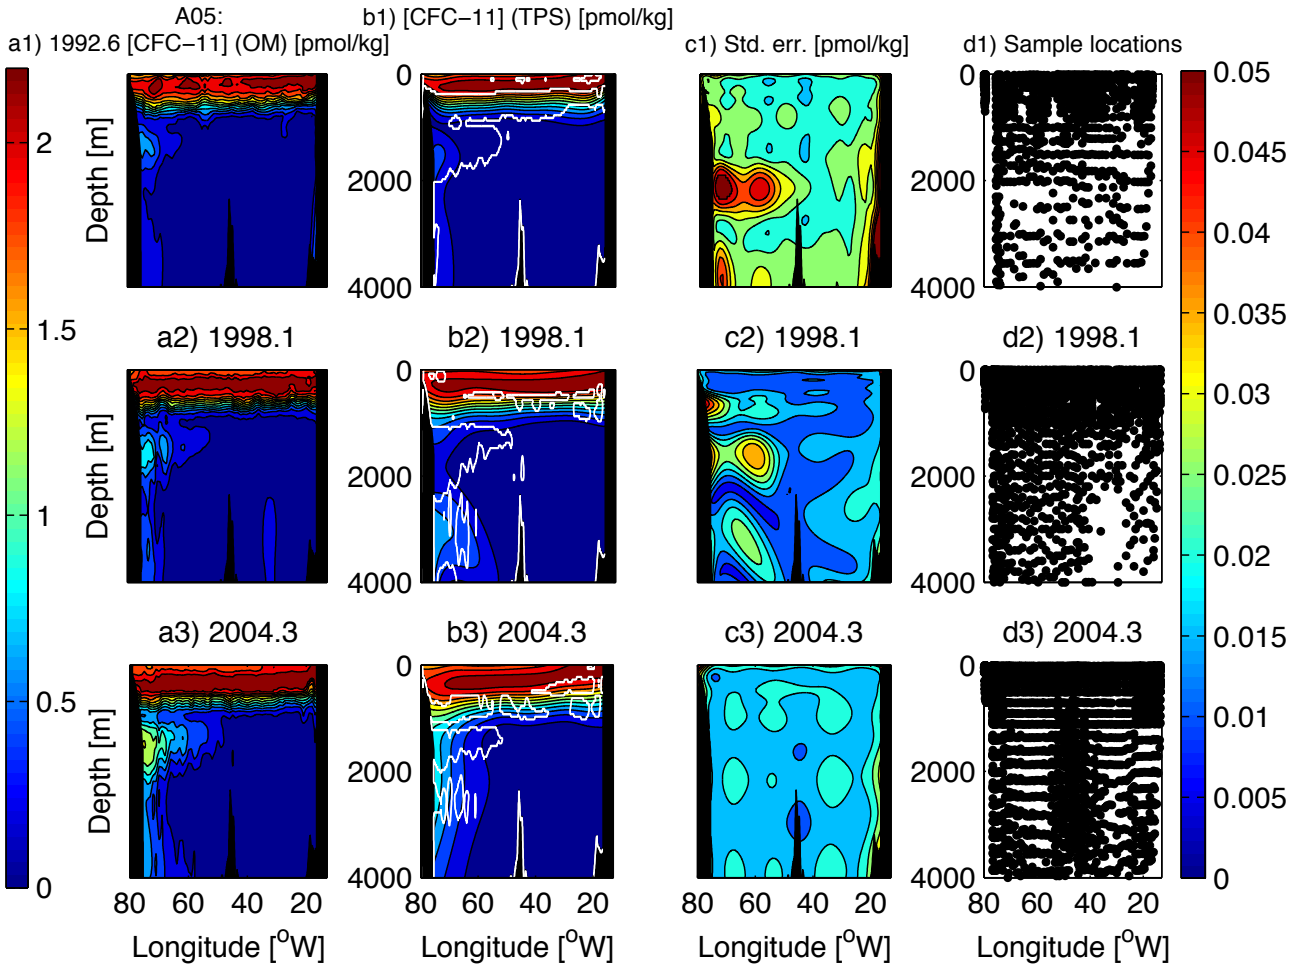

Figure 2.7: The CFC-11 concentrations along A05 for 1992 (row 1), 1998 (row 2), and 2004 (row 3) from objective mapping (column a) and thin-plate estimates (column b) and their mapping errors (column c); also shown are the locations of where the sampled data was bottled (column d) and white contours (column b) bounding the regions where 95% point-wise confidence intervals based on using one degree of freedom on the thin-plate spline estimates do not cover the objective mapping estimates. Columns a and b are on the same scale, as indicated on the left side. Column c is on the scale indicated on the right side. All columns are over the same domain.

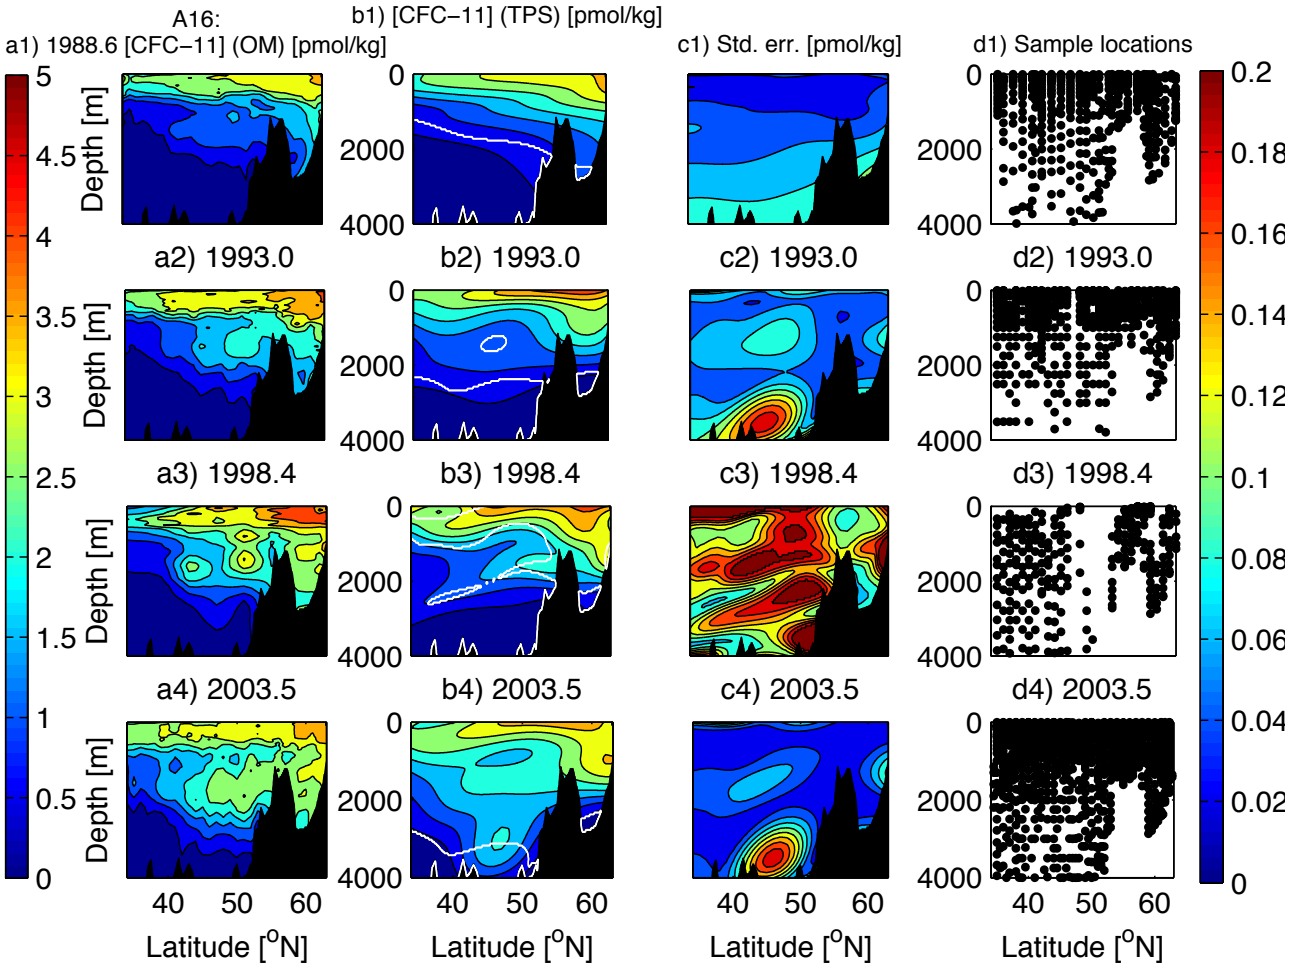

Figure 2.8: The CFC-11 concentrations along A16 for 1988 (row 1), 1993 (row 2), 1998 (row 3), and 2003 (row 4) from objective mapping (column a) and thin-plate estimates (column b) and their mapping errors (column c); also shown are the locations of where the sampled data was bottled (column d) and white contours (column b) bounding the regions where 95% point-wise confidence intervals based on using one degree of freedom on the thin-plate spline estimates do not cover the objective mapping estimates. Columns a and b are on the same scale, as indicated on the left side. Column c is on the scale indicated on the right side. All columns are over the same domain.

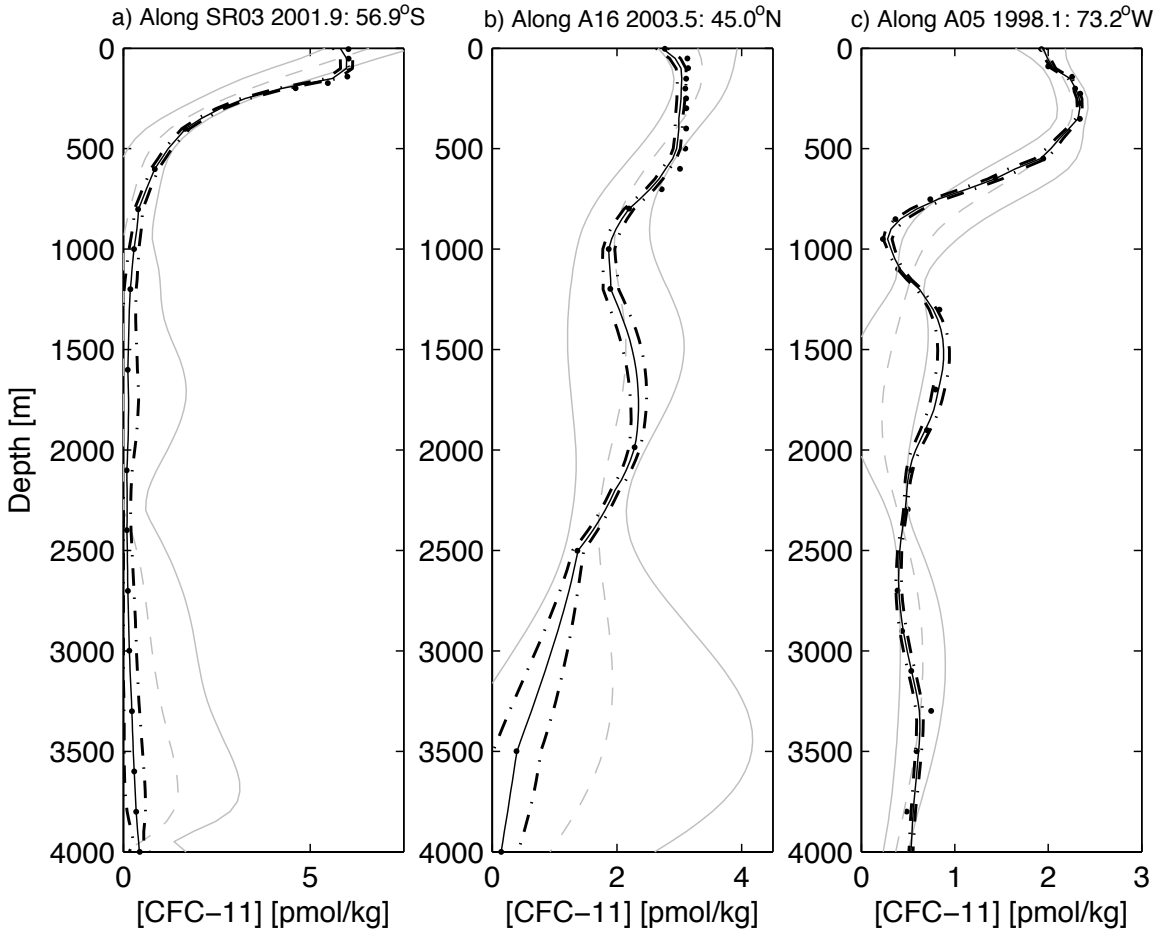

Figure 2.9: The CFC-11 concentration profiles along SR03 at the nearest raw data profile to 57^{o}S in 2001 (a), along A16 at the nearest raw data profile to 45^{o}N in 2003 (b), and along A05 at the nearest raw data profile to 73^oW in 1998 (c) showing the raw data (black dots), objectively mapped data (thin black solid lines) with 95% confidence intervals assuming more than thirty degrees of freedom (thick black dashed lines), and thin-plate spline-estimated data (gray dashed lines) with 95% confidence intervals assuming one degree of freedom (gray solid lines), based on the effective degrees of freedom used at each point by the thin-plate spline.

Chapter 3

BAYESIAN TTD ESTIMATION AND INTERMEDIATE WATER SOURCE INFERENCE

3.1 Introduction

To better understand the ocean's role in storing heat and gases and their exchange with the atmosphere, it is important to quantify the transport of water between the ocean's surface mixed layer and its permanently stratified interior. The ocean's mixed layer and interior are respectively above and below the seasonal thermocline. Both bulk advection and mixing contribute to the transport between the ocean's mixed layer and interior. That is, transport can be along isopycnal layers in which water crosses the mixed layer base while keeping its density (via advection and diffusion) or across isopycnal layers in which water crosses the mixed layer base by changing its density (via mixing).

Many researchers [60, 77, 106] have estimated the subduction rate, the rate at which water crosses the mixed layer base in winter and is not re-entrained within a year. This advective ventilation is found from the vertical velocity and lateral induction of a fluid parcel away from the mixed layer base, assuming this process is purely advective once the parcel leaves the mixed layer permanently. The role of diffusive processes in ventilation, on the other hand, is difficult to separate from a combination of advective and diffusive processes. This combination has been addressed primarily with transient tracer observations and simulations. Some transient tracer concentrations at the surface are determined by the solubility of the gas in water and its saturation, given the atmospheric partial pressure (e.g., CFC-11/CFC-12, or CCl_3F/CCl_2F_2). Other transient tracers are determined by precipitation, runoff, and vapor exchange (like ^3H, or Tritium). (Tritium is the 'T' in HTO that enters the ocean instead of H_2O. We will refer to Tritium as Tr hereafter.) After entering the ocean, a transient tracer is then advected and diffused by the ocean circulation.

The apparent pCFC-ages (the time since the water containing CFCs was last in contact

with the atmosphere) can be derived by taking the measured CFC concentrations, converting them into partial pressures by dividing by their solubility (temperature and salinity dependent via Warner and Weiss, 1985 [111]), and finding the year in the atmospheric time series [110] that matches this partial pressure. ^{3}He/Tr ages can be found from an analytical expression [43] because Tr decays to ^{3}He. These tracer-specific ages are different from the average of the ages of the observed water parcels because of non-linearities in the atmospheric source functions of these tracers and mixing [64]. In both of these cases, the age is a single number, but we know that any given water parcel contains water with a continuous spectrum of ages that originated in many different places. Thus, we use the concept of a transit-time distribution, or TTD, [38] to quantify this spectrum of ages. A TTD can be interpreted as a probability (or mass fraction) of a water parcel starting at a particular surface location getting to a particular interior location over a certain time interval. Because the TTD integrates to unity over all ages, it has units s^{-1}. If it were not for mixing, tracers would have a single 'age' instead of a continuous spectrum of ages. As such, a narrowly peaked TTD corresponds with a steady, advectively dominated flow. Biases in tracer-specific ages can also be understood by noting that their different atmospheric histories weight the TTD differently. For example, a tracer with a history shorter than the oldest ages within a water parcel will give zero weight to the TTD for a finite portion of the age spectrum, which essentially truncates the spectrum and biases inference about the moments of the TTD.

Primeau and Holzer (2006) [76] and Hall et al. (2007) [31] show how ventilation rates, TTDs, and transient tracer budgets are related. Hall et al. (2007) [31] show that a cumulative ventilation-rate distribution, $\Phi(\tau)$, may be defined as the mass flux through surface area, Ω, into the domain, D, of a water mass that resides at least τ amount of time in D before leaving back across Ω, giving $\Phi(\tau)$ units kg$\cdot$ s^{-1}. An equivalent interpretation of $\Phi(\tau)$ is the negative of the rate of change of domain inventory of water as a function of age so that the integral of $\Phi(\tau)$ from ages t_1 to t_2 yields the total inventory gain of water mass that resides between t_1 and t_2 years in D. The cumulative ventilation-rate distribution allows for computation of a ventilation rate because it accounts for both advection and diffusion. The domain-averaged TTD under steady-state transport (i.e., with constant fluid properties at

each point) is directly proportional to this cumulative ventilation-rate distribution, but a time-varying TTD can also be related to a mass flux without the direct proportionality [31].

Of the several established methods to estimate the TTD, only the method of Maximum Entropy [42] does not assume that the ocean transports water following the one-dimensional advection-diffusion equation, but this method has required the circulation to be the same every year in previous studies. Maximum Entropy transforms the distribution estimation problem to a system with an equal number of equations and unknowns. This is accomplished through the use of Lagrangian multipliers of the difference between tracer observations and a model of the transport of transient tracers from the surface ocean to a point in the interior. Holzer et al. (2010) [39] computed TTDs along 53.5^oW in the North Atlantic and Khatiwala et al. (2009) [48] computed TTDs over the entire World Ocean using the Maximum Entropy technique. Using CFC-11 and radioactive carbon-14 data, Holzer et al. (2010) [39] found the mean ages of water and the surface origin of these waters, while Khatiwala et al. (2009) [48] estimated how much anthropogenic carbon is in the ocean.

In this study, we will present a Bayesian method for estimating the TTD based on a combination of high-resolution numerical model output and observations. We will demonstrate the method with a focus on three intermediate water masses along three sections: the SR03 section (from 67^oS to 44^oS, between 135^oE and 147^oE), a transect that goes from Tasmania to Adélie Land of Antarctica; the A16 section that runs southward from the southern side of Iceland (at 63^oN along 20^oW) to a point we chose as 34^oN (because that is the southernmost point of ^{3}He/Tr observations to the depths we consider); and the A05 section that runs along 24.5^oN from about 14^oW to about 79^oW in the North Atlantic. These cross-sections are World Ocean Circulation Experiment (WOCE) transects with bottle data from the Climate Variability and Predictability (CLIVAR)/CO_2 Repeat Hydrography Program (http://cchdo.ucsd.edu/). The data we use along each of these sections, as well as one other (namely, A20) to compare with mean age estimates from [39], are listed in Table 3.1 and shown spatially in Fig. 3.1.

The reason for choosing these sections is that each of them has at least three repeated occupations in which CFCs (and sometimes ^{3}He/Tr as well) were measured from bottle data, which allows us to make inference about mass flux variability. Along the SR03 section,

there are several prominent water masses as described by Rintoul and Bullister (1999) [83] and Tomczak and Liefrink (2006) [104]. Antarctic Intermediate Water (AAIW) can be identified by its prominent salinity minimum. Speich et al. (2002) [100] used an inverse model to suggest that there is both eastward and westward transport of AAIW across the SR03 section. AAIW is formed from the coldest component of water that is ventilated in the southeast Pacific Ocean [94]. After formation, the three documented choke points AAIW travels through (Fig. 3.1) are as follows: The westward transport of AAIW across SR03, south of Tasmania, is referred to as the Tasman Leakage [100]. The westward transport of AAIW through the Indonesian islands is referred to as the Indonesian Throughflow [28]. The eastward transport of AAIW via the Antarctic Cirumpolar Current (ACC) between South America and Antarctica is referred to as the Drake Passage pathway [82]. It has been hypothesized by Rintoul and Bullister (1999) [83] that the East Australian Current and its recirculation [81] in the Tasman Sea carries some of the AAIW that ends up in the Tasman Leakage branch. The second intermediate water mass we examine is the Labrador Sea Water (LSW). Northwest of the A16 section, the North Atlantic Current branches off to feed the Irminger Current [4] which crosses A16, bringing LSW to the section (Fig. 3.1). LSW is a thick, homogeneous layer at intermediate depths, observed to have been formed in the Labrador Sea and Irminger Basin [75], marked by high concentrations of anthropogenic tracers and salinity, potential temperature, and potential vorticity minima [79]. Along the A05 section, North Atlantic Deep Water (NADW) [55] is relatively warm ($< 4^oC$) and nutrient depleted compared to other deep waters, marked by a salinity maximum (> 34.7 PSU) at intermediate depths. Along A05, we will only be investigating the LSW component of NADW, which is formed by mixing LSW with (salty) Mediterranean Outflow Water [55].

This paper addresses the following questions using the above-mentioned tracer data and a previously unused statistical approach to estimating TTDs and its first two moments: 1) What is the minimum number of transient tracers required using Bayesian analysis and what type of prior information (i.e., first-guess without any observational information) needs to be provided to statistically distinguish the first two TTD moments for one water mass from those of another? 2) Is the Green's function solution to the one-dimensional advection-diffusion equation, referred to as the inverse Gaussian (IG) representation, sufficient to

reproduce model output or observations? 3) How well does a high-resolution numerical model simulate the along-isopycnal diffusivities of the ocean? 4) Can we find evidence of a component of AAIW being entrained and/or detrained anywhere other than the southeast Pacific with our method? 5) Can we find evidence of LSW being entrained and/or detrained anywhere other than the Labrador Basin with our method? 6) How does consideration of the full spectrum of water parcel ages influence estimates of spreading rates of a particular water mass, as opposed to considering only the younger components of this water mass? To address these questions, we first describe our statistical technique called Markov Chain Monte Carlo (MCMC) [65]. MCMC has previously been used to reduce the uncertainties on estimates of vertical diffusivities by Goes et al. (2010) [27], for example, but has not been used for the purpose of estimating a TTD and its first two cumulative moments. By comparing the observationally informed estimates with those directly from a high-resolution numerical model, we can perform model evaluation for a few intermediate water masses. Lastly, we provide the volume transports as a function of age for particular water masses across one hydrographic section in the Southern Ocean and two hydrographic sections in the North Atlantic. We conclude with a suggestion for an alternative estimation procedure.

3.2 Data and models

Along the transects listed in Table 3.1, temperature, salinity as well as some transient tracer concentrations (CFC-11, CFC-12, Tr, and ^{3}He) allow us to make estimates of the water age. The water masses of interest, AAIW, LSW, and the LSW component of NADW, along the transects listed in Table 3.1 have distinct properties that can be used to quantitatively define their locations in the water column. For AAIW, we consider each vertical profile of salinity and define AAIW as present by requiring that the salinity, S, be less than $S_{minmax,AAIW}$, and potential densities in the range, $27.05 \leq \rho_\theta \leq 27.40$ kg/m^3, north of 50^{o}S [93]. Here, $S_{minmax,AAIW}$ is defined as the salinity at the shallowest depth where $dS/d|z| < 0$ and $d^2S/d|z|^2 > 0$ (z is depth) so that AAIW is the freshest water north of 50^{o}S. When analyzing AAIW, it will be useful to analyze Subantarctic Mode Water (SAMW) as well. For SAMW,

we consider each vertical profile of large-scale potential vorticity,

$$PV = -\frac{f}{\rho_\theta}\frac{\partial \rho_\theta}{\partial z} \tag{3.1}$$

where f is the Coriolis parameter set to be $2 \times 7.2921 \times 10^{-5} sin(\lambda 180/\pi)$ for latitude λ. We define SAMW as present by requiring that the potential vorticity, PV, be less than $PV_{minmax,SAMW}$, and at potential densities between the mixed layer depth (determined by the $\rho_\theta - \rho_\theta(z = 0) > 0.125$ kg m^{-3} criterion) and 27.05 kg m^{-3}, north of 50^oS [93]. Here, $PV_{minmax,AAIW}$ is defined as the potential vorticity at the shallowest depth where $dPV/d|z| < 0$ and $d^2PV/d|z|^2 > 0$ so that SAMW is the minimum potential vorticity water north of 50^oS. For LSW, we consider each vertical profile of CFC-11 concentrations, [CFC-11], and define LSW as present by requiring that [CFC-11] be greater than $C_{maxmin,LSW}$. Here, $C_{maxmin,LSW}$ is defined to be [CFC-11] at the shallowest depth below 1000 meters where $d[CFC-11]/d|z| > 0$ and $d^2[CFC-11]/d|z|^2 < 0$ so that LSW is the water with the maximum CFC-11 concentration below 1000 meters depth [44]. For the LSW component of NADW, by looking for the shallowest depth, $z_{minmax,NADW}$, where $dS/d|z| < 0$ and $d^2S/d|z|^2 > 0$ and the shallowest depth that is deeper than 1000 meters, $z_{min,NADW}$, where $dS/d|z| > 0$ and $d^2S/d|z|^2 > 0$, we identify the LSW component of NADW as the water at depths, $z_{min,NADW} < z < z_{minmax,NADW}$ [55].

We will compare our MCMC estimates of TTDs (Section 3.4) where these water masses are located as well as our MCMC estimates of mean ages, and Peclet numbers with output from the Parallel Ocean Program (POP) high-resolution model [59]. The POP model has a nominal 0.1^o resolution in longitude with a tripolar grid (with poles over Russia, Canada, and the South Pole) and variable spacing (10 meters near the surface, 250 meters in the deep) in the vertical. This model's atmospheric state was forced by a repeat annual cycle Coordinate Ocean Reference Experiment (CORE) dataset [54] that has averaged 6-hourly fields to monthly fields. The temperature and salinity fields were initialized by interpolating the annual mean WOCE Global Hydrographic Climatology [29]. For both momentum and tracers, vertical mixing coefficients are determined by the K-Profile Parameterization [53], which uses diffusion coefficients of 0.1 m^{2}s^{-1} for resolving gravitational instabilities [58] and background values for vertical tracer diffusion of 10^{-5} m^{2}s^{-1} near the surface to 10^{-4} in the

deep. Five TTDs have been calculated by the POP model, each with the boundary impulse response (BIR) approximation [30] by letting the surface of the ocean be $1/N_s$, for an entire year where N_s is the number of seconds in a year, and keeping track of the values of this tracer. The five different TTDs from POP were calculated starting at different times from one another (0, 2, 4, 30, and 70 years after introducing the first boundary impulse). The five TTDs from POP are primarily different along the margins of shadow zones and ventilation regions, which is attributed to the time-varying eddy fields as discussed in Maltrud et al. (2010) [58]. CFC-11 concentrations from POP are compared with objectively mapped CFC-11 concentrations from bottle data in Section A.2.3 of the Appendix to justify our use of POP. The IG and TTD theory are explained in more detail in the following section.

3.3 Methods

3.3.1 TTDs and the Cumulative Ventilation-Rate Distributions

In general, the tracer concentrations, $C(\vec{x},t)$, in the presence of a source term (Q), velocity field (U), and diffusivity tensor (κ), follow the three-dimensional advection-diffusion equation,

$$\frac{\partial C}{\partial t} + (V(\vec{x},t)\cdot\vec{\nabla} - \vec{\nabla}\cdot\kappa(\vec{x},t)\vec{\nabla})C = Q(\vec{x},t). \tag{3.2}$$

With no source function, $Q = 0$, (except at the surface), we have

$$C(\vec{x},t) = \int_0^t dt_0 \int\int_{\partial D} d^2x_0 C_0(\vec{x_0},t_0)\mathcal{G}(\vec{x},t;\vec{x_0},t_0) \tag{3.3}$$

for t being the current time and t_0 being the time when the tracer starting entering the system through surface, ∂D, of domain, D and $\mathcal{G}$ is the Green's function solution. In particular, the one-dimensional advection-diffusion equation,

$$\frac{\partial C(y,t)}{\partial t} = \frac{\partial}{\partial y}\kappa\frac{\partial C(y,t)}{\partial y} - U\frac{\partial C(y,t)}{\partial y} \tag{3.4}$$

has a tractable analytical Green's function solution [38] assuming that the flow is in steady-state or has a repeating seasonal cycle (which is referred to as cyclostationary by Holzer et al. (2010) [39]). Here, y is the distance from the isopycnal outcrop. Assuming that κ, the

along-isopycnal eddy diffusivity, is a constant,

$$G(y,t) = \sqrt{\frac{y^2}{4\pi\kappa t^3}} exp(-\frac{(y-Ut)^2}{4\kappa t}), \tag{3.5}$$

which is the IG distribution.

At a given location, the Green's function can then be defined by two parameters, the mean age (Γ) and either the half-width (Δ), or equivalently by the mean age and the Peclet number ($Pe = \Gamma^2/\Delta^2$). We define the mean age to be $\Gamma = \int_0^\infty d\xi \xi G_{IG}(\xi)$, the half-variance to be $\Delta^2 = \frac{1}{2}\int_0^\infty d\xi(\xi-\Gamma)^2 G_{IG}(\xi)$, and a drift parameter ν set to be one here so that one year elapses for every unit of time. Then we can write the IG distribution [24] as,

$$G_{IG,1}(t) = \sqrt{\frac{\Gamma^3}{4\pi\Delta^2 t^3}} exp(-\Gamma\frac{(\Gamma-\nu t)^2}{4\Delta^2 t}). \tag{3.6}$$

The two parameters, Γ and Δ, can be inferred by calculating the tracer concentrations using (3.3) and making them consistent with observations [113]. Typically, $\Delta/\Gamma = 1/\sqrt{Pe}$ is between 0.5 and 1.5 [112]. In many places in the World Ocean, (3.6) is inadequate. For example, where water masses of two distinct distributions of ages mix, the IG model should be extended to be a mixture model of IGs,

$$\begin{aligned} G_{IG,2}(t) = \alpha\sqrt{\frac{\Gamma_1^3}{4\pi\Delta_1^2 t^3}} exp(-\Gamma_1\frac{(\Gamma_1-\nu t)^2}{4\Delta_1^2 t}) + \\ (1-\alpha)\sqrt{\frac{\Gamma_2^3}{4\pi\Delta_2^2 (t-t_2^*)^3}} exp(-\Gamma_2\frac{(\Gamma_2-\nu(t-t_2^*))^2}{4\Delta_2^2(t-t_2^*)}), \end{aligned} \tag{3.7}$$

where the new parameters are the relative weight of the two water masses, α, and an offset age at which to begin the second IG, t_2^*. In this case, the mean age here is $\alpha\Gamma_1 + (1-\alpha)\Gamma_2$ and the half-variance is $\alpha^2\Delta_1^2 + (1-\alpha)^2\Delta_2^2$ since both IGs contribute to the overall mean age and half-variance of the distribution. The offset age, t_2^*, is included here because it is needed to fit a second IG to a second TTD peak from the POP model. Since the second peak typically occurs at large ages, by including t_2^* as a parameter, we are allowing for a less diffusive second IG than would be the case if we were to use $t_2^* = 0$. The second IG is used here along SR03 because Antarctic Bottom Water (AABW) and Circumpolar Deep Water (CDW), each with very different age distributions, mix a degree or so off of

the Antarctic continental margin and AAIW forms from SAMW mixing with other waters that have different age distributions.

Hall et al. (2007) [31] showed that the TTD can be used to estimate a distribution of mass flux. To do this, a tracer inventory, $I_{VT}(t, t_0)$, is defined as the mass of water that crosses the boundary between the seasonal thermocline and permanent thermocline (loosely interpreted as the base of the mixed layer) between times t and t_0. Primeau and Holzer (2006) [76] and Hall et al. (2007) [31] show that as the tracer inventory evolves from t_0 to t, the tracer makes contact with surface area Ω (e.g., the base of the mixed layer) and is contained within the volume, V. It follows from these definitions that $I_{VT}(t, t_0)$ can be related to the domain-averaged TTD, $G(t, t')$, by $-dI_{VT}/dt_0 = \rho V G(t, t_0)$. The total mass in V is ρV for ρ being the density of the fluid in V, so the flux of mass through Ω that has resided at least $\tau = t - t_0$ in V is the flux of water mass exiting V that has resided at least τ in V and is given by

$$\Phi_+(t, \tau) = \rho V G(t, t_0) + \rho V \int_0^\tau dt' \frac{\partial}{\partial t} G(t, t - t'). \tag{3.8}$$

At each time, t, in which there was an occupation of a WOCE transect, we can then use (3.8) to evaluate the volume transport spectrum with Φ_+/ρ (units $\mathrm{m}^3 \cdot \mathrm{s}^{-1}$) through the transects. We note that Φ_+ is not a distribution over τ in the classical statistics sense because Φ_+ is not normalized to one.

The above derivation of (3.8) does not require the assumption of steady-state or cyclo-stationary flow [31]. We can also estimate an average spreading rate of a water mass with density ρ over its entire age history,

$$\theta(t) = \frac{1}{\tau_{max} - \tau_{min}} \int_{\tau_{min}}^{\tau_{max}} d\tau \frac{\Phi_+(t, \tau)}{\rho A_{wm}}, \tag{3.9}$$

where τ_{max} is the largest age within the water mass in volume (V) that contributes volume transport above a threshold (defined to be when $G(t, t - \tau) > 0.00125\ \mathrm{s}^{-1}$), τ_{min} is smallest age of the water mass that contributes some volume transport (also defined to be when $G(t, t - \tau) > 0.00125\ \mathrm{s}^{-1}$ but is defined such that $\tau_{min} = 0$ here), and A_{wm} is the cross-sectional area the water mass occupies across a given transect. The quantities, τ_{max} and τ_{min}, within V are slightly uncertain and the density estimates for the TTDs themselves

have some uncertainty associated with them, but these uncertainties are dwarfed by the temporal variability in Φ_+. Our spreading rate estimates, $\theta(t)$, and our volume transport estimates, $\Phi_+(t,\tau)/\rho$, at each hydrographic occupation time, t, are new applications of the TTD. These estimates have not been previously performed because these are relative estimates. We argue that comparing these estimates at each hydrographic occupation time with the other occupations provides us with a way to infer how much $\Phi_+(t,\tau)/\rho$ and $\theta(t)$ change from year to year.

3.3.2 TTDs from CFCs and Tritium

To obtain an estimate of the TTD using a single IG for A16 and A05 and mixture of IGs for SR03, we need to account for the uncertainties in the observations when we check the consistency of our proposed TTD against the data. To simplify the notation, let $\vec{H} = (\Gamma, Pe)$ for A16 and A05 and $\vec{H} = (\alpha, \Gamma_1, \Gamma_2, Pe_1, Pe_2, t_2^*)$ for SR03. Using the IG (or mixture of IGs) as the TTD for a given grid point, the concentration of CFC-11, C, will be given, according to (3.3), by:

$$\frac{C(t)}{(F_{sol,CFC}(\theta,S)f_{sat,CFC})} = \int_{t_T}^{t} dt_0 C_{CFC,atm}(t_0) G_{IG,k}(t-t_0|\vec{H}) \tag{3.10}$$

where k represents the number of inverse Gaussians are used (two for SR03 and one for A16 and A05), $C_{CFC,atm}$ is the atmospheric time series of CFC-11 in partial pressure units, $C(t)$ is a CFC-11 concentration at a given interior point at time t in units of [pmol/kg], $F_{sol,CFC}(\theta,S)$ is the potential temperature- and salinity-dependent solubility function determined by Warner and Weiss (1985) [111], $f_{sat,CFC}$ is the saturation level of CFC-11 (which we treat as unknown but close to unity), $G_{IG,1}$ is our IG approximation ($G_{IG,2}$ our mixture model of IGs approximation) to the TTD, and $t_T = 1945$ is the year that CFC-11 started being measured to have non-zero partial pressure in the atmosphere.

Dreisigacker and Roether (1978) [23] estimated a surface time series of Tr in the North Atlantic Ocean (rather than an atmospheric time series). Thus, for Tr, we use:

$$Tr(t) = \int_{t_T}^{t} dt_0 C_{Tr,ocn}(t_0) G_{IG,k}(t-t_0|\vec{H}) exp[-\frac{log(2)}{12.43}(t-t_0)] \tag{3.11}$$

where $C_{Tr,ocn}$ is the oceanic time series of Tr at a given point estimated by Doney and

Jenkins (1988) [18] (modifying the estimates of Dreisigacker and Roether (1978) [23]), $Tr(t)$ is a Tritium concentration at a given interior point at time t in units of TU (defined in Table 3.2), $t_T = 1700$ to account for primordial Tritium, and the exponential factor takes the radioactivity of Tritium (with a half-life of 12.43 years) into account.

3.3.3 How to estimate IG parameters: A Monte Carlo method

With cross-sections of transient tracer observations and numerical model estimates of TTDs, we develop a technique to estimate the TTD here using a MCMC approach. The method is implemented as follows at each grid point of location and depth: A) We take samples of water parcel ages from the POP TTD, plug these samples into the Expectation-Maximization (EM) algorithm (see Section A.2.1 of the Appendix), and get estimates of the joint distributions of IG parameters with uncertainties associated with how non-IG-like the POP TTD is. Next, $10,000$ iterations of the following sequence of steps (B-F) are performed: B) We propose values for IG parameters, $\vec{H}$, (see "Sampling the IG parameters" of Section 3.3.3) from the joint distributions in A and sample observations (such as CFC-11 concentrations) from a distribution whose width is determined by measurement plus mapping errors and whose correlation between observed quantities is determined by spatial correlation (see "Sampling the observations" of Section 3.3.3). C) We evaluate a product of a first-guess (prior) distribution for the IG parameters and a symmetric function that is maximized when the convolved CFC-11 concentrations at the surface with an IG with proposed parameters from B are closest to the sampled observations, also from B (see "Bayes' Theorem" of Section 3.3.3). D) We evaluate the ratio of the quantity in C from the current iteration of B-F to the quantity in C from the previous iteration of B-F to calculate a probability of accepting the most recently proposed IG parameters (see "Implementing MCMC" of Section 3.3.3). E) We determine whether we save the IG parameters from the previous or current iteration of B-F from the quantity in D as the next value in our samples of IG parameters (see "Implementing MCMC" of Section 3.3.3). The purpose of this is to come up with the smallest uncertainty estimates of the IG parameters that are consistent with the observations as possible, given all of the sources of error with which we inform the procedure. F)

The saved sample of IG parameters from the current iteration defines an IG from which we take a sample of water parcel age at a particular grid point (and time when computing the volume transports). The sample in F is a sample from the Bayesian TTD estimate, which is not an IG if an IG is insufficient for the TTD to be consistent with observations. When $10,000$ iterations of B-F are complete, we have $10,000$ samples of IG parameters and $10,000$ samples of water parcel ages from the Bayesian TTD estimate. G) A kernel, or histogram, estimate of the collected water parcel ages is a numerical representation of the TTD. Enough samples need to be taken to get an accurate numerical representation of the TTD. The TTD derived from one set of samples at a particular time and location should be close enough (i.e., within variability due to sampling, called Monte Carlo error) to the TTD derived from another set of samples at the same time and location.

There are three advantages to using MCMC for TTD estimation: 1) The distribution of uncertainty on physical relevant parameters such as the mean age (given by the first cumulative moment of the TTD) and Peclet number (a ratio of advection to diffusion, which in combination with the mean age gives an estimate of the width of the TTD) can also be estimated. 2) We only need to know the surface concentrations of transient tracers at the surface locally, as opposed to globally, to estimate a TTD at a single location and depth. 3) We can quantify all of the known contributing uncertainties of our estimates. The details for implementing our MCMC procedure are provided in the following Sections ("Sampling the observations" of 3.3.3 to "Implementing MCMC" of 3.3.3) and our assumptions are listed in "Summary of method" of Section 3.3.4.

Sampling the observations

Defining μ as the mean values of the sampled parameters and Σ as the covariance between these parameters (e.g., a standard normal distribution has $\mu = 0$ and $\Sigma = 1$), we sample the observations and saturation levels. Unlike the standard normal distribution, a multivariate normal distribution samples more than one variable and allows for correlation between them. We use objectively mapped bottle data for the mean values and associated uncertainties of the mapping procedure using the thin-plate-spline method (Chapter 2 of the present thesis)

for the diagonals of the covariance matrix in this sampling scheme. The observations that are compared with the model in (3.10) and (3.11) are sampled as follows:

We sample each component of $P(\vec{O})$ ($\vec{O} = (Tr, CFC\text{-}11, T, S, p)$ for A16 and $\vec{O} = (CFC\text{-}11, T, S, p)$ for A05 and SR03) for the MCMC scheme according to a multivariate normal distribution with mean equal to the objectively mapped observed quantities at a given grid point, μ_{obs}, and whose covariance matrix is estimated from the empirical spatial covariance between each component's gridded fields, $\mathbf{\Sigma}_{obs}$. By empirical spatial covariance, we mean that we take the covariance, σ_{ij}, between pairs of components, i and j of $\vec{O}$ (for any $i \neq j$), over the entire field where these components are non-zero, and scale by the ratios of the measurement plus mapping errors of components i and j, Σ_{ii} and Σ_{jj}, to the standard deviations of the ith and jth components of $\vec{O}$ where they are non-zero, σ_{ii} and σ_{jj}:

$$\Sigma_{ij} = \frac{\Sigma_{ii}\Sigma_{jj}}{\sigma_{ii}\sigma_{jj}}\sigma_{ij} = c_{ij}\Sigma_{ii}\Sigma_{jj} \tag{3.12}$$

where c_{ij} is the correlation between the ith and jth components of $\vec{O}$ where these components are non-zero. The diagonals of this covariance matrix, Σ_{ii}, are the sum of the mapping errors (squared) derived from using the thin-plate spline (Chapter 2 of the present thesis) and the measurement errors (squared), as listed in Table 3.2. The multivariate normal distribution to sample $\vec{O}$ is truncated (written as $tMVN(\cdot, \cdot)$) such that the CFC-11 concentrations are non-negative; the salinities, S, are greater than or equal to 30 PSU; the temperatures, T, are greater than or equal to -2^oC; and the pressures, p, are non-negative. Therefore, we can write that the observations are sampled as $\vec{O} \sim tMVN(\mu_{obs}, \mathbf{\Sigma}_{obs})$. Lastly, the saturation, f_{sat}, is sampled from a normal distribution with mean 0.9 and variance 0.1^2. This is one part of step B from Section 3.3.3.

Sampling the IG parameters

We sample all of the parameters, $\vec{H}$, from their prior distributions. We use a prior distribution for the IG parameters either based on the TTDs from the POP model output and the EM algorithm described in the Appendix, Section A.2.1 or a uniformly distributed prior in this sampling scheme. More details on how the prior distributions are specified as follows:

We let $P(\vec{H})$ be a truncated multivariate normal distribution with mean $\mu_{\vec{H}}$ and covariance $\Sigma_{\vec{H}}$ such that each sampled component of $\vec{H}$ is non-negative. Other caveats are listed in the Appendix, Section A.2.2. The mean for $\vec{H}$, $\mu_{\vec{H}}$, is estimated via the EM algorithm (Appendix, Section A.2.1) and the covariance matrix for $\vec{H}$, $\Sigma_{\vec{H}}$, is determined as follows: The covariance matrix associated with taking M samples of the water parcel ages from the POP model TTD is calculated using the Fisher-information matrix, $\mathbf{I}(\cdot)$. The entries of $\mathbf{I}(\vec{H})$ are given by,

$$I(\vec{H})_{i,j} = -\sum_{l=1}^{M} \frac{\partial^2}{\partial H_i \partial H_j} log(G_{IG,k}(t_l - t_0|\vec{H})). \tag{3.13}$$

The covariance matrix for uncertainty associated with $\vec{H}$ is $\mathbf{V}_{\vec{H}} = M/\mathbf{I}(\vec{H})$, calculated using (3.13). In short, $\vec{H} \sim tMVN(\mu_{\vec{H}}, \mathbf{V}_{\vec{H}})$. This truncated multivariate normal prior distribution will be compared with a uniform prior distribution, by which we mean that we uniformly sample over the entire interval of values that come from sampling the truncated multivariate normal distribution. This completes step B from Section 3.3.3.

Bayes' Theorem

To understand our MCMC application, Bayes' Theorem [3] is invoked here. We interpret the TTD as the probability of a water parcel having age, τ, given the (tracer) observations, $\vec{O}$, and the TTD is equivalent to $P(\tau|\vec{O})$. To sample from the TTD, we first sample the IG parameters based on Bayes' Theorem with:

$$P(\vec{H}|\vec{O}) = \frac{P(\vec{O}|\vec{H})P(\vec{H})}{P(\vec{O})} \propto P(\vec{O}|\vec{H})P(\vec{H}), \tag{3.14}$$

where $P(\vec{O})$ is not known for each time and location, the prior ($P(\vec{H})$) is either uniform (constant in $\vec{H}$) or a truncated multivariate normal distribution based on the POP TTDs (as previously mentioned in "Sampling the IG parameters" of Section 3.3.3), and the likelihood ($P(\vec{O}|\vec{H})$) is a product of symmetric distributions in the differences between observed tracer concentrations and (3.10) or (3.11),

$$P(\vec{O}|\vec{H}) = \prod_i \prod_T \frac{1}{\pi s_T} \frac{1}{1 + \frac{(T_{i,obs}(t) - \int_{t_T}^{t} dt_0 C_{T,surf}(t_0) G_{IG,k}(t-t_0|\vec{H}) exp[-\frac{log(2)}{\lambda_T}(t-t_0)])^2}{s_T^2}}. \tag{3.15}$$

Here, T is CFC-11 (for which 'surf' means atmospheric partial pressure or 'atm') and Tritium (for which 'surf' means oceanic concentration or 'ocn') where available, N is the number of repeat observations we have ($N = 3$ for A05 and $N = 4$ for A16 and SR03) ($i = 1, ..., N$), t_T is the time at which tracer T from human activity began the enter the ocean ($t_T = 1945$ for CFC-11 and $t_T = 1700$ for Tritium), s_T is the sum of the measurement and mapping errors on the observed tracer concentration, and λ_T is the half-life of tracer T ($\lambda_{CFC-11} = \infty$ years and $\lambda_{Tr} = 12.43$ years) at each location and depth. In combination with using the prior, whose density is calculated at the sampled IG parameters and whose implementation is described in Section A.2.2 of the Appendix, using (3.15) completes step C from Section 3.3.3. One reason for choosing the likelihood, (3.15), to be symmetric about zero is that we do not want to weight positive differences between observed tracer concentrations and (3.10) or (3.11) differently from negative differences. Also, the distributions in (3.15) have longer tails than normal distributions, so large deviations between the model and observations get a small non-zero weight rather than zero weight (as would be the case with normal distributions). By collecting enough samples of IG parameters, we find that the errors have a Gaussian structure to them.

To understand why we use MCMC instead of deriving $P(\vec{O})$ numerically as a normalization constant, it is important to understand that we are explicitly entering all sources of uncertainty. In order to incorporate these sources of uncertainty in the IG parameters, $\vec{H} = \vec{H}^* + \vec{\epsilon}$, where $\vec{H}^*$ is the vector of the best-fit values for the IG parameters and $\vec{\epsilon}$ is the vector of errors on the IG parameters. The errors are distributed according to the symmetric distribution in (3.15), and we repeat enough iterations of the MCMC procedure by sampling (see Section A.2.2 of the Appendix) the IG parameters, $\vec{H}^*$, to account for the uncertainty due to $\vec{\epsilon}$. By sampling the IG parameters in this fashion, described in the following Sections, 3.3.3 and 3.3.3, we are effectively integrating out the uncertainties associated with the IG parameters with MCMC to yield $P(\tau|\vec{O}) = \int d\vec{H} P(\tau|\vec{H}, \vec{O}) P(\vec{H})$. Thus, by taking a sample from the distribution of IG parameters and then using that sample to construct a distribution that represents one possible TTD estimate, we are sampling from $P(\tau|\vec{O})$.

Implementing MCMC

The Metropolis-Hastings (M-H) Algorithm [65] is a MCMC algorithm that is used here. The M-H algorithm states that if a set of samples from the target distribution (e.g., the distribution of values in which the IG parameters are consistent with observations) has $(n-1)$th sample with value, $\vec{H}^{(n-1)}$, then the value, $\vec{H}^{*(n)}$, is proposed as the nth sample, which is accepted or rejected as the next value in the set of samples from the target distribution according to an acceptance probability,

$$A(\vec{H}^{(n-1)}, \vec{H}^{*(n)}) = min\{1, \frac{P(\vec{H}^{*(n)})P(\vec{O}|\vec{H}^{(n-1)})}{P(\vec{H}^{(n-1)})P(\vec{O}|\vec{H}^{*(n)})}\}. \tag{3.16}$$

This completes step D from Section 3.3.3. Here, $A(\vec{H}^{(n-1)}, \vec{H}^{*(n)}) = A^*$ is compared with a random number, x, where $0 \leq x \leq 1$. We set $\vec{H}^{(n)} = \vec{H}^{*(n)}$ if $A^* \geq x$ and $\vec{H}^{(n)} = \vec{H}^{(n-1)}$ if $A^* < x$. This completes step E and step F follows immediately from Section 3.3.3. After taking enough samples ($n = 1001, ..., 10000$, in our case), the median of these samples is our estimate of $\vec{H}$ and our standard error of this estimate is the standard deviation of these samples.

3.3.4 Summary of method

By fitting the POP model TTD to a mixture model of IGs with the EM algorithm, as discussed in Section A.2.1 of the Appendix, we can estimate a prior distribution for each of the parameters of the IG as in step A from Section 3.3.3. The TTD, using our MCMC technique, is estimated by sampling a set of IG parameters by using additional information from a prior distribution as in part of step B from Section 3.3.3, next sampling from the mixture of IGs with this sampled set of IG parameters as in the other part of step B from Section 3.3.3, defining a first-guess TTD as an IG distribution [24] commonly used to approximate the true TTD [49, 112] as in step C from Section 3.3.3, calculating an acceptance probability and whether the currently proposed IG parameter samples are the next sample in the final collection of IG parameter samples as in steps D and E from Section 3.3.3, plugging the next sample in the final collection of IG parameter samples into the IG and taking a water parcel age sample from that distribution as in step F from Section 3.3.3,

and repeating this procedure thousands of times. We can define the prior distribution used in step D from Section 3.3.3 to be either uniform (a specific type of uninformative) prior or a distribution based on the POP model output (as a specific type of informative prior). By choosing the informative prior, we can constrain our uncertainty estimates to a greater extent.

Having described the method, we list our MCMC model assumptions: 1) CFC-11 (and sometimes Tr) together with the spatial covariance between the IG parameters based on POP model output are sufficient to give a unique TTD. 2) The mixture model of IGs, with uncertainty in its parameters taken into account, can accurately represent the true TTD. 3) The product of solubility and saturation of water carrying CFC-11 was the same at the time of subduction as the time of measurement, schematically shown in Fig. 3.2. We test these three assumptions in concert by using model output as 'data' for SR03 in Section 3.4.1.

3.4 Results

The IG parameters were estimated using observations from all of the occupations across a given section, as in (3.15), instead of estimating the IG parameters for each occupation and then averaging. This is because the number of effective degrees of freedom associated with an estimate that uses a single occupation as data is unknown and it would not be possible to propagate the uncertainties. The volume transports, on the other hand, use observations from only one occupation across a given section, so there is only one distribution as opposed to product of distributions in (3.15) for each volume transport estimate. With the exception of profiles comparing the information provided by additional tracers and additional prior distribution information from POP, we will only be plotting the cross-sections down to a depth at which the mean age approaches 100 years because it is not reasonable to expect CFC-11 or ^{3}He/Tr in combination with model output to provide information on time scales beyond this. The deepest depth shown for SR03 is 1200 meters, for A05 is 2200 meters, and for A16 is 2400 meters. We also plot A20 down to 3000 meters to compare with a previous study [39]. All MCMC runs here take $10,000$ posterior (TTD) samples at each location and depth. Because there is autocorrelation between the Markov Chains starting from the first and bth sample for $b < b_{in} = 1000$ and because it may take several accept/reject iterations

to find the posterior distribution, we only use sample numbers b_{in} to $10,000$ as posterior samples. This makes $9,000$ samples for each location and depth.

3.4.1 SR03

Along SR03, the mean ages and Peclet numbers from POP, estimated from our MCMC technique using observations as data, and estimated from our MCMC technique using POP output as data are shown in Fig. 3.3. The mean age is largest (at about 350 years) according to POP to the south of the topographic ridge around 59^oS at the bottom (not shown). Older water parcels such as CDW mix with AABW along the Antarctic shelves, but this is not apparent in POP (Fig. 3.3a) because the ages along the Antarctic shelves are no younger than their surroundings. The POP model we have used to constrain our mean ages does not have enough AABW export in this region due to a well-known problem with poor representation of formation of dense waters due to brine rejection on the Antarctic shelves [17]. Combining the model with observations via MCMC, on the other hand, the diffusive mixing between AABW and CDW off the shelf around $61 - 64^oS$ is seen below about 1000 meters depth as seen from the Peclet numbers (Fig. 3.3e) relative to the POP output (Fig. 3.3b) with the Peclet numbers in this region estimated to be about $2.0^{\pm}1$ from MCMC as opposed to 3 with POP. The weight, α, has the greatest uncertainty (not shown) where there are more parameters that need to be estimated for the TTD (sum of two IGs rather than one IG) and the uncertainty in α scales proportionally with IG parameter uncertainties; also, the standard errors are larger on the mean ages and/or Peclet numbers in these locations (Figs. 3.3g-j). There are two locations where this occurs: just below the thermocline (about 400 meters depth to the south and about 700 meters depth to the north of 52^oS) which corresponds to entrained and then detrained shallow/intermediate waters in the southeastern Indian Ocean, as portrayed in Fig. 2 of Sallée et al. (2006) [88], and near the continental slope of Adélie Land (between $61 - 63^oS$ below about 400 meters depth) which corresponds to mixing of AABW and CDW. The ratio of the half-variance parameter, Δ, to the mean age parameter, Γ, along SR03, $1/\sqrt{Pe}$, tends to be consistent with the range of values reported by Waugh et al. (2004) [112] since the Peclet numbers in Fig. 3.3e tend

to be between 0.5 and 4. Isochrons approximately follow isopyncals (not shown) and the young age influence of both SAMW between $300 - 500$ meters depth and AAIW between $500 - 1200$ meters depth is plainly seen in the SR03 section (Fig. 3.3c between $45 - 50^o$S).

To address the accuracy of our three MCMC model assumptions jointly ("Summary of method" of Section 3.3.4), we estimated the mean ages and Peclet numbers with our MCMC procedure replacing the CFC-11, temperature, and salinity observations with POP model output of these quantities. The MCMC results (Figs. 3.3d and 3.3f) suggest younger mean ages and larger Peclet numbers than POP (Fig. 3.3a-b). As indicated by the white 'X's in Figs. 3.3d and 3.3f, the mean ages from POP (Fig. 3.3a) are insignificantly different from those estimated by our MCMC method everywhere (Fig. 3.3d), but the Peclet numbers are significantly different (Figs. 3.3b and 3.3f) according to the standard errors estimated from MCMC. The Peclet numbers from MCMC are relatively large at the base of the mixed layer (Fig. 3.3e at ~ 100 meters depth between $56^o - 59^o$S and north of 49^oS) where both the mean age and half-variance are small but finite, suggesting small diffusivities. The discrepancy between the IG parameters derived from POP model TTDs, which are based on cumulative moment definitions, and the IG parameters estimated via MCMC, which are based on parameter definitions of the IG, using POP model output as observations demonstrates that our three joint assumptions are inaccurate ('X' in Fig. 3.3f) at the base of the mixed layer. This suggests that our solubility assumption is not the culprit, but rather our use of the mixture IG to constrain the form of the TTD is could be inaccurate in some locations. On the other hand, the fact that our method infers the water parcels to be more advective and/or less diffusive within the mixed layer but at its base (Figs. 3.3d and 3.3f) implies that our method does not provide information about when the mixed layer deepens and shoals because information provided by the data aliases the seasonal signal. At the locations where there is AAIW in both POP and the observations ('X' in both Figs. 3.3a and 3.3c), the means and standard deviations of the mean age differences (i.e., differences between Figs. 3.3a, 3.3c, and 3.3d) are summarized in Table 3.3. The mean differences in the mean ages are small compared to the mean ages at the locations where there is AAIW ($20 - 50$ years where there is an 'X' in both Figs. 3.3a and 3.3c), suggesting that our three assumptions are adequate for the purpose of estimating TTDs where there is AAIW along

SR03.

3.4.2 North Atlantic Ocean (A20, A16, and A05)

Along the hydrographic sections in the North Atlantic Ocean, we can evaluate the POP model by comparing the IG parameters and considering their estimated uncertainties. We can also compare our mean age estimates (Fig. 3.4) to those from the Maximum Entropy method (Holzer et al., 2010 [39] - Fig. 9) along A20. We only make this comparison down to depths where the time scales of CFC-11 and ^{3}He/Tr are not too short to be able to estimate an unbiased mean age. Our mean ages agree very well in spatial distribution with the bottom-left panel of Fig. 9 in Holzer et al. (2010) [39], but below the depths where the mean ages approach 100 years (at about 1000 meters depth), our results begin to diverge. Because we extrapolated the POP TTDs from about 100 years, we cannot accurately estimate a mean age from the POP model of greater than 100 years. In addition, because we used ^{3}He/Tr instead of $\Delta^{14}C$, we cannot estimate a mean age from the MCMC procedure of much greater than 100 years. It is apparent that there are spatial similarities between the results of the Maximum Entropy technique (Holzer et al., 2010 [39] - Fig. 9) and those of the MCMC technique (Fig. 3.4b). Additionally, our MCMC estimates of the Peclet number at intermediate depths (Fig. 3.4e) are larger than those of the POP model (Fig. 3.4d) with about the same mean ages (Figs. 3.4a-b), which suggests that the POP model is too diffusive at intermediate depths.

The oldest ages along A16 are south of the topographic ridges near the sea floor (Fig. 3.5b) because the youngest water is coming from the north. The limits to which we are confident in our results (i.e., the mean age less than 100 years) are to the sea floor north of 45^{o}N and at depths shallower than 750 m south of 45^{o}N. The subtropical/subpolar gyre boundary near 45^{o}N is apparent from the latitudinal age gradients (Fig. 3.5b) as well as latitudinal Peclet number gradients (Fig. 3.5e). Constant mean age surfaces along A16 approximately follow isopycnals (not shown). LSW is young relative to its surrounding waters along the flanks of the Rockall-Hatton Plateau (near $55 - 59^o$N) at mid-depth. The Peclet numbers are again estimated by MCMC to be larger than the Peclet numbers

from POP with similar mean ages. The Peclet numbers at intermediate depths in the North Atlantic are relatively large (small diffusivities) according to our MCMC estimates, consistent with studies such as Ledwell et al. (1993) [56].

The oldest ages along A05 are east of the Mid-Atlantic Ridge near the sea floor because that water is less well ventilated and is almost entirely devoid of CFCs (Fig. 3.6b). Here, ^{3}He/Tr would provide useful supplementary information. We only trust the magnitudes of our IG parameter estimates down to 1500 meters east of the Mid-Atlantic Ridge (around 50^oW) and at greater depths west of the ridge (where mean ages become much larger than 100 years). The ages are younger near the western boundary where the Deep Western Boundary Current is expected to be found and older where there is NADW. Again, the Peclet numbers estimated by MCMC tend to be larger than the Peclet numbers from POP while the mean ages are more similar.

3.4.3 Posteriors as TTD estimates

Examples of the TTD estimates (posteriors) by using the MCMC technique to sample water parcel ages for each section show the influence of the likelihood relative to the prior distribution (Fig. 3.7) pulling the mean age in one direction or another. The two TTD estimates shown along A16 (Figs. 3.7c-d) and A05 (Figs. 3.7e-f) and one TTD estimate shown along SR03 (Fig. 3.7a) appear to be IG-like, but one estimate along SR03 (Fig. 3.7b) has two peaks that represent two different water masses mixing with each other. While each of these estimates appear to be IGs or a mixture of IGs, we can see that only one (Fig. 3.7c) can be fit precisely with an IG approximation or mixture. When all of the uncertainties on the estimated IG parameters are taken into account, the net estimate for the posterior TTD tends to exhibit non-IG-like properties despite taking samples from IGs or mixtures of IGs in each iteration of the MCMC algorithm. This suggests that a single- or mixture-IG representation of the TTD is less consistent with the observations than the true TTD.

The likelihood was calculated by using (3.15) with a proposed (mixture) IG that has a Peclet number equal to the average of the MCMC-derived samples for the Peclet numbers and varying the mean age. The peak in the likelihood is at the mean age consistent with

observations, assuming the Peclet number is accurate. The (observationally informed) likelihood tugs the posterior distribution away from the (POP) prior, which in most cases is reflected through a shift in the mean age (Figs. 3.7a-c and 3.7f), but the likelihood also influences the Peclet number estimates (not shown). One might expect that because the pCFC- (and ^{3}He/Tr-) ages are (both) biased young, the likelihood would reduce the length of the tail of the TTD. However, the MCMC method is using ages consistent with observations according to (3.10) and (3.11) rather than the biased tracer-derived ages themselves.

The mean age of the posterior is often older than the mean age of the POP output. The likelihood has a large width at some locations (Figs. 3.7a, 3.7c, and 3.7e) since both measurement and mapping error are larger closer to the surface, and has a relatively small width at other locations (Figs. 3.7b, 3.7d, and 3.7f). The mean age is altered less in the locations where the likelihood is closer to the POP mean age. The mean age is about the same in Figs. 3.7d-e because the likelihood is not different enough from the POP mean age, given the estimated Peclet number. The likelihood is much older than the POP mean age and the MCMC-derived mean age in Fig. 3.7b because 1200 meters depth at 62^{o}S is practically devoid of CFCs. The MCMC-derived mean age is much closer to the POP mean age at this location and depth because we do not consider mean ages in our method much older than the maximum of the POP mean age and the time scale on which CFCs have been entering the ocean.

3.4.4 Uncertainties using different tracers and different priors

Next, we demonstrate effects of using a uniform as opposed to POP model-based prior (Fig. 3.8). The width of the confidence intervals on mean age expand the most once a mean age of about 50 years is reached using any combination of tracers with the POP model-based prior (Fig. 3.8). The differences between the widths of the confidence intervals using different priors is apparent whether only CFC-11 (Figs. 3.8a, 3.8c, 3.8d, and 3.8f) or both CFC-11 and ^{3}He/Tr (Figs. 3.8b and 3.8e) are used. The widths of the confidence intervals on the IG parameters are smaller using information from the POP model for a prior. The widths of the confidence intervals on the IG parameters along A16 are almost no different using both

^{3}He/Tr and CFC-11 or only CFC-11 data until 2000 meters depth is reached (not shown). At 2000 meters depth, the Peclet number begins to be distinguishable from zero using the POP model-based prior along A16 when ^{3}He/Tr data from 1988 are included (Fig. 3.8e), while the Peclet number remains indistinguishable from zero using only CFC-11 data (not shown). The most noticeable features that are different between including both ^{3}He/Tr and CFC-11 as opposed to only CFC-11 data along A16 are the estimates of the IG parameters using the uniform prior. The IG parameter estimates using the uniform prior with only CFC-11 along A16 are larger and have larger confidence intervals (not shown).

A combination of the POP model-based prior and CFC-11 provides almost all of the information ^{3}He/Tr data are capable of providing along A16. The information provided in the POP model output can essentially be used as a substitute for an additional transient tracer to add additional information when using Bayesian analysis. Using the uniform prior, the Peclet number overlaps zero almost everywhere along A16 and A05 (Figs. 3.8e-f) and above at least 500 meters depth for mean ages along A16 and A05 (Figs. 3.8b-c). Using the POP model-based prior, in contrast, the mean ages are distinguishable from zero almost everywhere, while the Peclet numbers are distinguishable from zero below 700 meters depth along SR03 (Fig. 3.8d) and below 2000 meters depth along A16 (Fig. 3.8e). Only one transient tracer is necessary along some hydrographic sections if model output is used as a prior in a Bayesian estimation procedure because the POP model-based prior essentially rules out combinations of mean ages and Peclet numbers associated with different pathways. However, we need to be careful about interpreting the uncertainties where a transient tracer used to estimate the IG parameters has a concentration below the detectable limit because the uncertainty could be an underestimate.

Since one of the benefits to using the MCMC technique developed here is that we can quantify all known contributing sources of error on the IG parameters, we break these sources down here: 1) There is measurement error in the observations, whose contribution is small evidenced by Table 3.2, and there is mapping error derived from thin-plate spline estimates, which can be large but not along the profiles shown in Fig. 3.8. 2) There is error associated with the fact that the TTD is not IG-like, which is large primarily due to the uncertainty on the second IG parameters upon inspection of the Fisher's information matrix

derived via the EM algorithm. 3) Another large source of error could be the error associated with the fact that we are only using CFCs (and ^{3}He/Tr) in some of the plotted profiles, which cannot be explicitly accounted for because concentrations below the detectable limit provide no information. By supplementing our estimates with additional transient tracers with different atmospheric histories [113], we would be able to constrain the uncertainties on the IG parameters to a greater extent. 4) There is a small amount of error associated with the fact that we do not know the saturation levels or solubilities of the surface sources, which is reflected through a saturation uncertainty based on the assumption that the solubility function is the same at the time of subduction to the time we make an observation. 5) There is Monte Carlo error, which is negligibly small due to the number of samples we take via MCMC for each grid point, evidenced by the difference between the average sampled age of a water parcel and average estimate for the mean age at each location and depth always being less than two years. 6) There is error associated with the fact that the circulation is not steady, which alters the IG parameters slightly from year to year, as reflected in each volume transport spectrum discussed in the following subsection, and we are only using three or four snapshots to find an average. The differences between the uncertainties on the mean ages and Peclet numbers are due to 2. When the inverse Gaussian, or mixture of inverse Gaussians, representation of the TTD is inappropriate, the uncertainties on the Peclet numbers will never shrink beyond a certain extent regardless of the suite of tracers we use to diminish the errors due to 3 and 4. The fact that we cannot distinguish the Peclet number from zero along A16 or A05 in most locations is partially due to non-IG-likeness and partially due to using as few tracers with little prior information from POP.

3.4.5 Volume transport spectra and spreading rates

We, next, estimate the volume transport via (3.9) as a function of age class for three different estimation methods of the TTD to demonstrate how conclusions made from inference about where particular water masses ventilate would be different with each method. The volume transports are calculated by averaging over the cross-sectional area along a given transect in which a particular water mass is found, so the POP-derived volume transports are different

in each panel of Figs. 3.9, 3.10, 3.11, and 3.12 when bottle data was collected at a different time of the year (when the water mass volume was different from another time of year). In contrast to the spreading rates listed in Table 3.4, the relative magnitudes of volume transport shown in Figs. 3.9, 3.11, and 3.12 are meaningful for comparing the model, IG approximation, and Bayesian estimates using the POP-derived prior, but the absolute magnitudes are volume-dependent. Only a cross-sectional area is known, not the total volume, so the magnitudes of our estimates could be interpreted as $\mathrm{Sv{\cdot}m^{-1}}$ instead of Sv. We do not know the thickness of the water mass in the direction perpendicular to the cross-section. Still, the relative magnitudes of the volume transports across a given section are meaningful because each of the volume transport estimates along a given section are scaled by the same constant from the total volume transport spectrum for a particular water mass.

Along A16, A05, and SR03, the posterior of water parcel age generally estimates lower volume transports (solid lines in Figs. 3.9b, 3.9d, 3.11 and 3.12) and spreading rates for younger water parcels than both the POP estimates (dashed lines) and the IG approximation (dot-dashed lines). Along SR03 and A16, the IG estimates (Figs. 3.9 and 3.11) are different enough from both the posterior- and prior-based estimates to suggest that the IG representation is not sufficient to be consistent with either the observations or POP. We also show the volume transports for SAMW (Fig. 3.10) to suggest that the young AAIW could have come from the SAMW that was entrained and detrained in the southeast Indian Ocean [88].

Our estimates of spreading rates are consistent with those of previous studies (Table 3.4 compared with estimates of Sy et al., 1997 [102] and Sloyan and Rintoul, 2001 [93]). Some possible explanations for any discrepancy with previous studies is that previous estimates rely on mass conservation within the observed domain, which may be violated, and these estimates additionally do not consider the varying transport rates associated with parcels of different ages. In our spreading rate estimates, we are averaging the spreading rates over all water parcel age classes along a particular water mass trajectory from a formation region to a transect rather than an instantaneous advective transport along the transect. We cannot say much about the inter-annual variability or any trend because of seasonal aliasing.

AAIW

The young AAIW spreads at a relatively fast and highly variable rate (Table 3.4), and it would only be possible for such young AAIW (or SAMW that later formed AAIW) to come from upstream in the ACC after being entrained and then detrained the South Indian Ocean thermocline as opposed to coming from the Tasman Leakage. If the young, quick AAIW came from the southeast Pacific without having been ventilated somewhere along the way, spreading rates would need to be > 8 cm/s. This, combined with the double peak in the TTDs just below the thermocline along SR03 (suggested by relative large standard errors in Figs. 3.3g-j), is evidence, consistent with the findings of Sallée et al. (2006) [88], that SAMW is entrained and detrained in the southeast Indian Ocean. We argue this pathway partially contributes to forming AAIW that passes through SR03. Shallow and intermediate waters such as AAIW are entrained and detrained in the South Indian Ocean along the pathway toward the Pacific sector of the Southern Ocean, crossing SR03. The pathways of the older (mean age > 10 years), more sluggish AAIW are less certain. Any of the pathways shown in Fig. 3.1 are possible. The older, more sluggish AAIW could have come from the Indian Ocean after circulating the Indian Gyre, from the South Pacific as part of the Tasman leakage as suggested by Speich et al. (2002) [100], from having recirculated in the Tasman Sea and mixed with older waters as suggested by Rintoul and Bullister (1999) [83], or even from having circulated around the entire ACC as suggested by Sloyan et al. (2010) [94]. The AAIW spreading rates are comparable to previous estimates [93] if only the youngest waters are considered (Table 3.4). The older AAIW diminishes the overall average spreading rate.

There are a few anomalies in the features of the volume transport spectra of AAIW in Fig. 3.9 that are very similar to the volume transport spectra of SAMW in Fig. 3.10. These anomalies are not likely to be due to seasonality because all of the occupation times of SR03 overlap summertime. The posterior estimate in 1991 along SR03 (Fig. 3.9a) is missing the first peak (< 5 years of age). Since the 1991 data are noisy, we find AAIW at deeper (nearly 2000 meter) depths than any of the other occupation years and note that the only non-zero volume transports at young ages (< 5 years) are accounted for entirely

by SAMW (Fig. 3.10a) and above waters. The year 1995 (Figs. 3.9b) has a relatively large secondary peak in the volume transport spectrum compared to 2001 and 2008 (Figs. 3.9c and 3.9d). The same relative magnitudes appear in the volume transport spectra for SAMW (Fig. 3.10), except 1991 (Fig. 3.10a) has a larger young-age (< 5 years) transport than 2008 (Fig. 3.10d). The years 1991 and 1995 (Figs. 3.9a-b) have relatively large secondary peaks in the volume transport spectrum compared to 2001 and 2008 (Figs. 3.9c-d). The same relative magnitudes appear in the volume transport spectra for SAMW, except 1991 has a smaller secondary peak. When there is a non-zero peak at young (< 5 years) ages in the volume transport spectrum for AAIW (Figs. 3.9b-d), it is possible that some of the young (< 5 years) AAIW formed from young SAMW. It is plausible that eddy diffusion is strong enough, as indicated by along- and cross-stream eddy diffusivities in Figs. 3c and 3f in Sallée et al. (2008) [86], east of 80^oE and west of SR03 in the Southern Ocean to support the hypothesis that some AAIW forms from the SAMW that entrains and detrains in the South Indian Ocean.

Where there is AAIW, by geostrophy, the vertical gradient of the zonal transport, $g(d\rho/dy)/f/\rho_0$, is mostly negative but sometimes positive (about $(-0.025^{\pm}0.05)/\rho_0 \mathrm{s}^{-1}$ compared to about $(0^{\pm}0.005)/\rho_0 \mathrm{s}^{-1}$ at deeper depths, not shown). Here, f is the Coriolis parameter, $g = 9.81 \mathrm{m\ s}^{-2}$ is the acceleration due to gravity, ρ_0 is the background density, and ρ is the density. Assuming a depth of no motion at any depth below AAIW suggests that, without a barotropic component to the flow, AAIW is moving eastward. The isopycnals are sloped more in 2001 than in other years (not shown). Since older waters carry lower concentrations of CFCs and the Tasman Sea is relatively CFC-depleted, the older AAIW is likely to have come from mixing in the Tasman Sea. In the year 2008, the Tasman leakage shut off since AAIW close to Tasmania that was present in each of the previous occupations of SR03 is absent (not shown). This could explain the smaller secondary peak magnitude in the volume transport spectrum in 2008 (Fig. 3.9d) relative to 1991 or 1995 (Figs. 3.9a-b) as opposed to 2001 (Fig. 3.9c) which has a smaller secondary peak magnitude because of the eastward flowing AAIW being relatively young.

LSW

LSW also appears to have been ventilated in two different general locations, but this is not as apparent by looking at its volume transport spectra (Fig. 3.11). The separation between the younger and older waters in the volume transport spectrum of LSW along A16 is much less pronounced (in fact, non-existent except for in Fig. 3.11d) than of AAIW along SR03. This can mostly be attributed to the geographic origins of each water mass. Some of the young, faster-spreading LSW could have come from the Labrador Basin because its spreading rate would be consistent with the estimates of Sy et al. (1997) [102] or slower, but the very young (e.g., $\leq 2-3$ years shown in Fig. 1 of Pickart et al., 2003 [75]) LSW must have come from the Irminger Basin. The LSW with mean ages of < 5 years would need to have spread at a rate of at least 4 cm/s to reach A16 from the Labrador Basin, which is greater than the estimates of Sy et al. (1997) [102]. Johnson et al. (2005) [44] found that the time it takes for LSW to get from the Labrador Basin to A16 is about 8 years. The POP model either forms too much LSW in the Irminger Basin or transports LSW very quickly from the Labrador Basin because of its large volume transport at young ages (Figs. 3.11a-d). The North Atlantic Oscillation (NAO) index was positive in 1988 and 1993 and negative in 1998 and 2003, while the volume of LSW transported across A16 in 1988 and 1993 is about half of that in 1998 and 2003 (not shown, but can be inferred from the relative magnitudes shown in Fig. 3.11 and spreading rates tabulated in Table 3.4). The denser LSW spreads more slowly and since LSW in 1988 and 1993 was more dense, we find smaller volume transports during those years than in 1998 and 2003. Dickson et al. (1996) [16] also reported the production of denser LSW when there is a positive NAO index. Furthermore, the volume transport of younger waters (< 5 years) is relatively large in 2003 compared to 1998, which is consistent with the less negative NAO in 2003 than in 1998.

LSW component of NADW

The younger NADW may not have circulated around the North Atlantic Gyre, while the older NADW may have made full revolutions in the gyre to contribute to such a wide

volume transport spectrum for the LSW component of NADW (Fig. 3.12). The larger volume transports according to POP (Figs. 3.12a-c) are primarily due to the larger volume of LSW component of NADW (Figs. 3.6a and 3.6c). The spreading rates estimated from these spectra for each water mass and year in which we have observations is shown in Table 3.4. Because the transport of LSW is steady relative to the LSW component of NADW (Table 3.4), Mediterranean Outflow Water that mixes with LSW to form the LSW component of NADW [55] could have a relatively variable mass flux, but the mass flux of LSW could also be variable in its southward flow due to changes in the North Atlantic Gyre as LSW circulates.

3.5 Conclusions

Using tracer observations mapped from WOCE and CLIVAR repeat hydrography data, we have estimated a TTD two different ways to allow for multiple peaks without assuming steady-state or cyclostationarity: through a mixture of IGs representation and through a Bayesian posterior on water parcel age. Below 500 meters along A05 and below 1000 meters south of 50^oN along A16, the Peclet number is found to be greater than 3, which is the value Hall et al. (2007) [31] found in the shallow (< 1000 m) North Atlantic Ocean, but a value statistically indistinguishable from 3 is commonly found elsewhere. The ratio of Δ to Γ tends to be between 0.5 and 4, consistent with the values found by Waugh et al. (2004) [112]. In addition, using information from the POP model output for a prior distribution in combination with CFC-11 observations, we can constrain the uncertainties in IG parameters to be significantly different from zero in a number of locations without using an additional transient tracer because of our use of a model-based prior distribution and Bayesian analysis. However, it is not possible to infer unbiased mean ages on the order of 100 years using CFC-11 because of how long these tracers have been entering the ocean due to human activity. The POP model output, on the other hand, cannot be used to trustfully infer mean ages well above 100 years due to the length of the model run.

We have estimated the mass flux of three water masses: LSW, AAIW, and the LSW component of NADW through the A16, SR03, and A05 WOCE sections, respectively. We have found that the spreading rates vary considerably more for AAIW along SR03 than for

the LSW component of NADW along A05 which in turn varies more than for LSW along A16. The entrainment and detrainment region of the young, quick AAIW is upstream in the ACC in the South Indian Ocean. The older, more sluggish AAIW is likely to have come from the Tasman Sea recirculation. The large changes in CFC concentrations along the northern end of SR03 where there is AAIW are absent from the POP model output (which uses climatological mixed layer depths). This suggests that the ventilation processes responsible for such large CFC concentration changes relative to other water masses are not represented by POP.

Some of the young, quick LSW could have come from the Labrador Basin, but the very young LSW ($< 2 - 3$ years) must have come from the Irminger Basin. More dense, sluggish LSW corresponds to a positive NAO index and less dense, faster-spreading LSW corresponds to a negative NAO index. It takes at least a hundred years on average for water to make its way from the surface southward to 24.5^oN to form the components of NADW denser than the LSW component.

We have shown an example of how sensitive volume transport spectrum estimates along three WOCE sections are to the discrepancies between model output and IG representations of TTDs. These discrepancies are due to discrepancies in the TTD's IG-likeness and how different the advective-diffusive processes and pathways are according to the POP model as opposed to reality. In the North Atlantic and Southern Ocean, both the IG model TTD representation using our Bayesian estimates of IG parameters and POP model-derived TTD overestimate mass fluxes at younger ages relative to our Bayesian estimates of the TTD. At older ages, both the IG model TTD representation and POP model-derived TTD underestimate mass fluxes relative to our Bayesian estimates of the TTD. Future TTD estimates could improve upon our MCMC model outlined here by using nonparametric Bayesian techniques.

Table 3.1: Tracer data collected in the listed years along each of the cross-sections of interest.

WOCE Section:	Latitude(s)	Longitude(s)	(T, S, CFC-11)	(δ^3He,Tr,^{4}He)
A16(N)	34^{o}N to 63^{o}N	20^{o}W	1988, 1993, 1998, 2003	1988
A05	24.5^{o}N	79^{o}W to 14^{o}W	1992, 1998, 2004	N/A
A20	20^{o}N to 41^{o}N	53.5^{o}W	1997	1997
SR03	67^{o}S to 43^{o}S	135^{o}E to 147^{o}E	1991, 1995, 2001, 2008	N/A

Table 3.2: Measurement errors of quantities of interest. Here, 'PSU' stands for practical salinity units (parts per thousand), 'pmol' stands for pico-moles, 'nmol' stands for nano-moles, 'kg' stands for kilograms, 'TU' stands for Tritium units (which are 10^{18} times the isotopic ratio of ^{3}H to Hydrogen), '^{o}C' stands for degrees Celsius, and 'dbar' stands for decibars (approximately equivalent to meters of depth).

Variable:	Units	Measurement error
Tritium	TU	0.05 (upper ocean, $\leq$ 100 m), 0.005 (elsewhere)
δ^3He	%	1.5 per mille
Helium	nmol/kg	0.005$\times$ [^{4}He]
CFC-11	pmol/kg	0.02$\times$ [CFC-11]
Temperature	^{o}C	0.0002
Salinity	PSU	0.0002
Pressure	dbar	0.03

Table 3.3: The differences between the POP model-derived mean ages based on cumulative moments from TTDs (POP), our MCMC model estimates of the mean age (Γ) using POP model output as observations (model) which only uses CFC-11, and our MCMC model estimates of Γ using bottle data as observations (obs) which only uses CFC-11, each calculated only where there is AAIW (in units of years).

Estimates	mean Γ difference	standard deviation of Γ difference
POP-model	4.97	22.4
POP-obs	3.72	22.9
model-obs	1.25	3.22

Table 3.4: Water mass spreading rates [cm/s] calculated from (3.9) for AAIW along SR03, LSW along A16, and LSW component of NADW along A05; errors are given in terms of a root-mean-squared deviation, R_{dev}, from the average point estimate over a range of ages. Also, given are the water masses' potential densities, ρ_θ, with one standard deviation, σ_{ρ_θ}, intervals around the mean, according to our definitions of these water masses. The estimate, ρ_θ, would be the value for ρ in (3.8), but ρ always cancels out in our calculations. By 'younger,' we are referring to the waters of ages within the first peak in the volume transport spectrum.

Water Mass	Year	ρ_θ ($^{\pm}\sigma_{\rho_\theta}$) [kg/m^3]	Spreading Rate($^{\pm}R_{dev}$) [cm/s] (younger($^{\pm}R_{dev}$))
AAIW	1991.8	27.2$^{\pm}$0.104	0.290$^{\pm}$0.0193 (1.17$^{\pm}$0.0834)
AAIW	1995.0	27.2$^{\pm}$0.0995	0.290$^{\pm}$0.0255 (1.93$^{\pm}$0.120)
AAIW	2001.9	27.2$^{\pm}$0.102	0.290$^{\pm}$0.0341 (1.87$^{\pm}$0.480)
AAIW	2008.3	27.2$^{\pm}$0.101	0.290$^{\pm}$0.0185 (1.22$^{\pm}$0.0666)
LSW	1988.6	27.5$^{\pm}$0.0893	0.330$^{\pm}$0.0376 (2.38$^{\pm}$0.350)
LSW	1993.0	27.5$^{\pm}$0.132	0.330$^{\pm}$0.0387 (2.62$^{\pm}$0.350)
LSW	1998.4	27.6$^{\pm}$0.103	0.330$^{\pm}$0.0380 (2.30$^{\pm}$0.380)
LSW	2003.5	27.5$^{\pm}$0.0740	0.330$^{\pm}$0.0351 (2.51$^{\pm}$0.230)
NADW(LSW)	1992.6	27.3$^{\pm}$0.141	0.220$^{\pm}$0.0124 (0.610$^{\pm}$0.0346)
NADW(LSW)	1998.1	27.3$^{\pm}$0.0803	0.220$^{\pm}$0.0129 (0.610$^{\pm}$0.0380)
NADW(LSW)	2004.3	27.3$^{\pm}$0.101	0.220$^{\pm}$0.0126 (0.610$^{\pm}$0.0341)

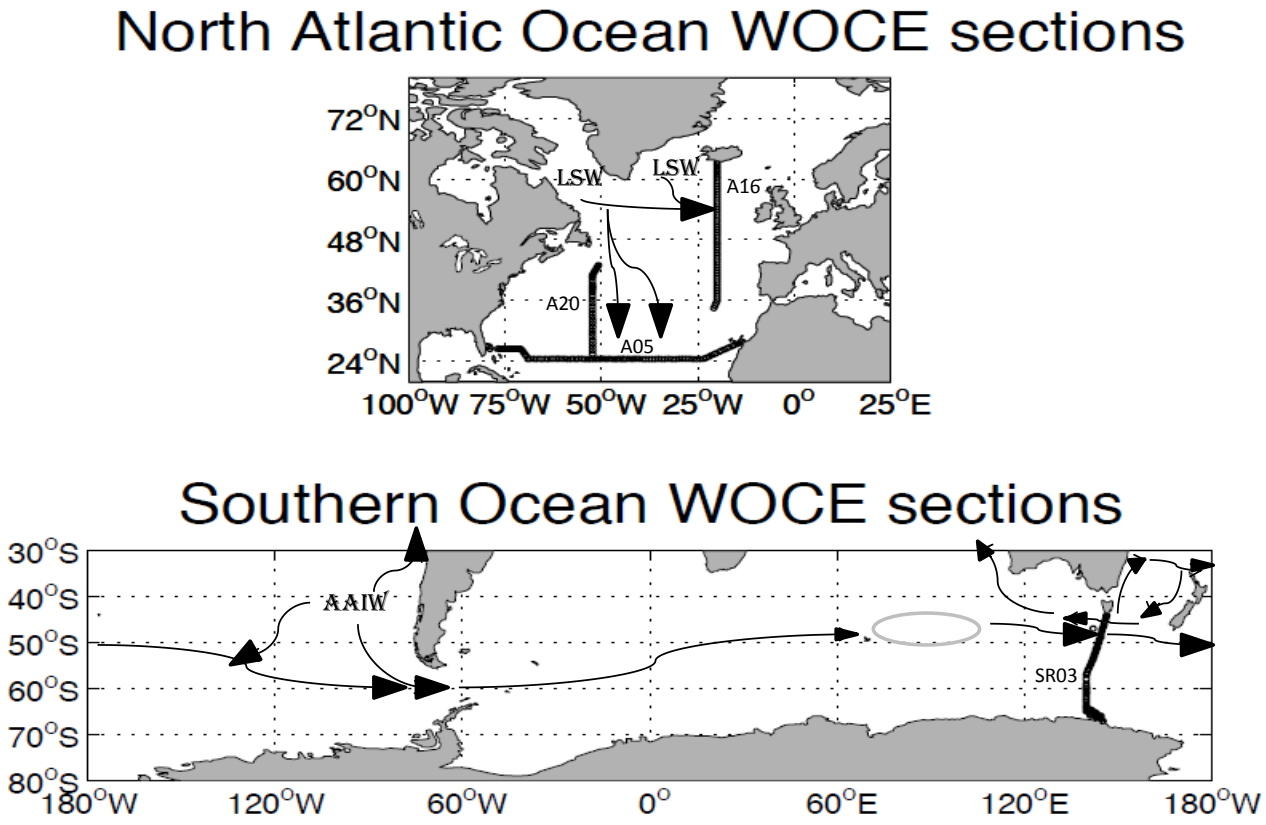

Figure 3.1: Locations of bottle data measurements used in this study taken along WOCE transects (lines of black Os). Also shown are hypothesized pathways of Labrador Sea Water (LSW, thin arrows, top plot) and of Antarctic Intermediate Water (AAIW, thin arrows, bottom plot) from their source regions and a region (light grey oval) in the southeast Indian sector of the Southern Ocean that indicates where AAIW is hypothesized to be entrained and later detrained due to a deepened mixed layer. LSW is labeled where it is formed in the North Atlantic Ocean via deep convective processes and AAIW is labeled where it is formed via turbulent processes in the southeast Pacific Ocean. The only pathway through a choke point for AAIW shown here is the Tasman Leakage around 147^{o}E, 45^{o}S.

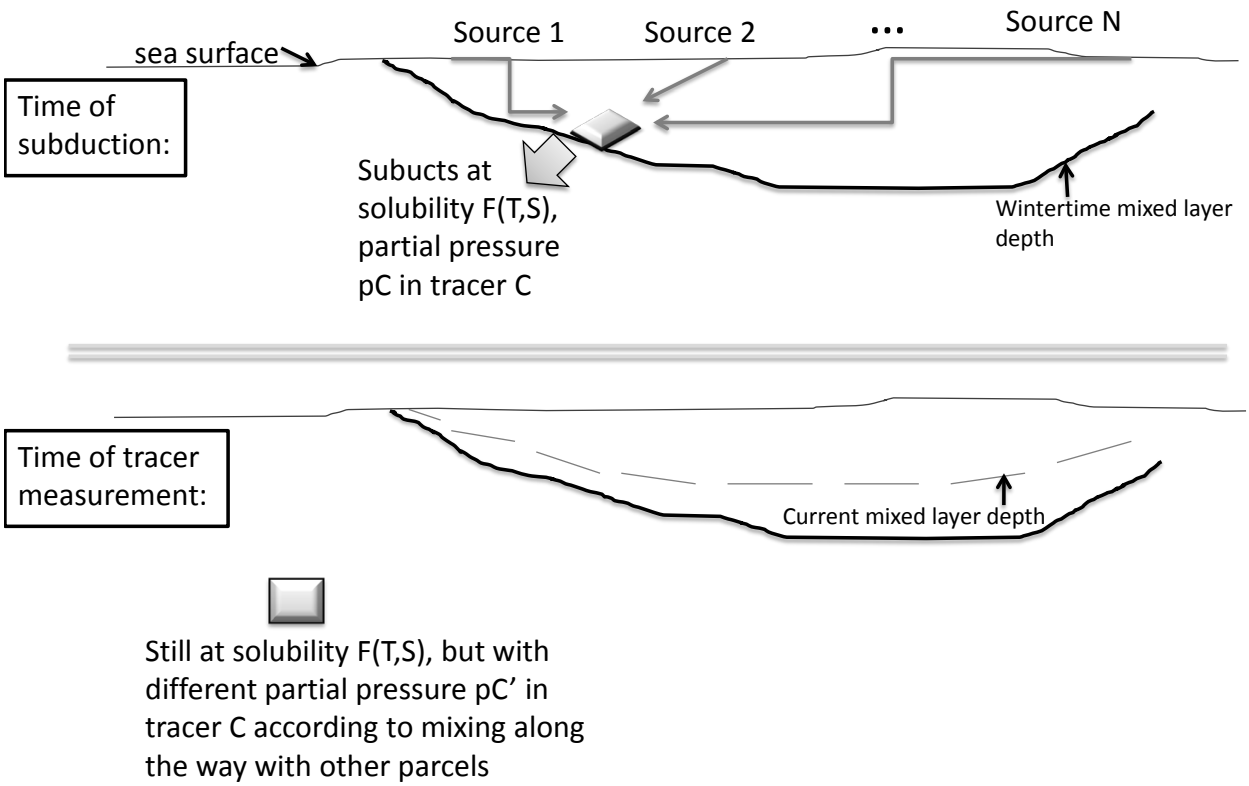

Figure 3.2: Schematic of the time at which a water parcel subducts (top) and the time at which we measure bottle data (bottom) under the assumptions of our MCMC model. That is, the product of the saturation level and solubility of CFC-11 at the time of subduction and the time at which we measure bottle data is the same.

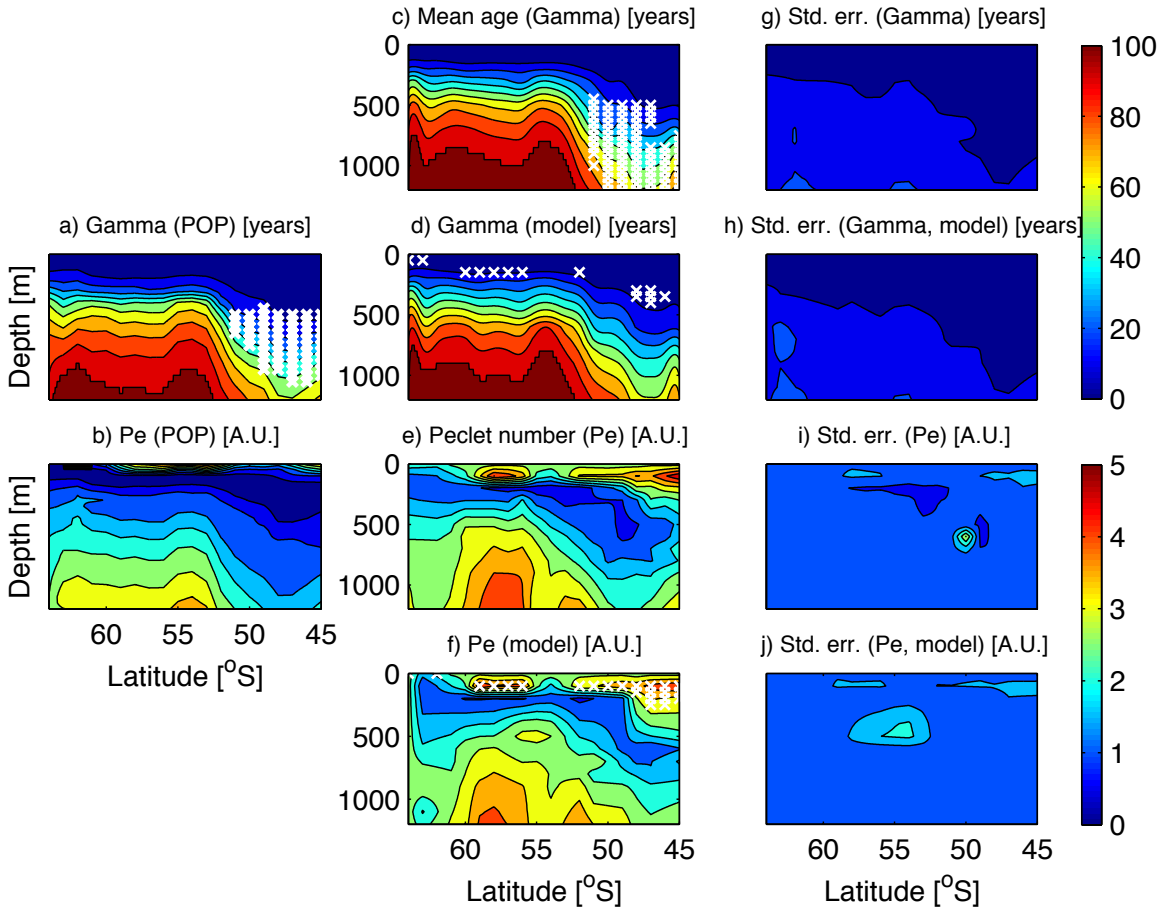

Figure 3.3: Mean ages (a, c, and d) and Peclet numbers (b, e, and f) and their associated errors (g-j) derived via a MCMC scheme (c-f) along the SR03 WOCE section and plotted against depth and latitude and derived via POP TTDs (a and b). CFC-11, T, and S from POP model output were used as 'data' instead of observations (panels d and f), and objectively mapped CFC-11, T, and S bottle data were used as data (panels c and e). The POP-derived TTDs were used as prior distributions in all of the MCMC estimates. AAIW from observations common amongst all four occupations is outlined in panel c with white 'X's and the envelope of all AAIW over each year from POP is outlined in panel a with white 'X's. The white 'X's in panels d and f indicate where the POP-derived TTD parameters do not lie within the 95% confidence intervals based on the MCMC estimates using POP model output as data for the mean age and Peclet number, respectively.

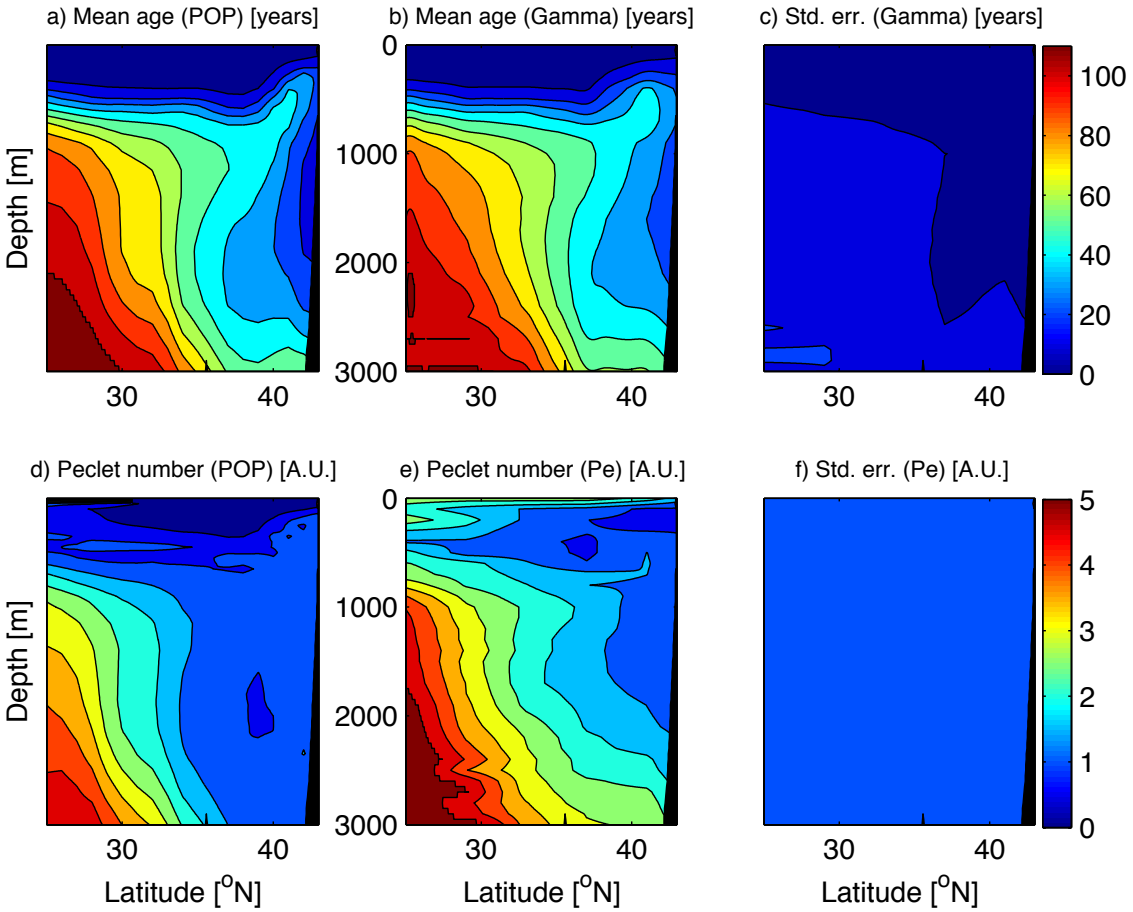

Figure 3.4: Mean ages (a and b) and Peclet numbers (d and e) and their associated errors (c and f) derived via a MCMC scheme (b, c, e, and f) along the A20 WOCE section and plotted against depth and latitude and derived via POP TTDs (a and d). Objectively mapped CFC-11, T, and S and ^{3}He/Tr were used as data, and POP-derived TTD parameters were used as priors.

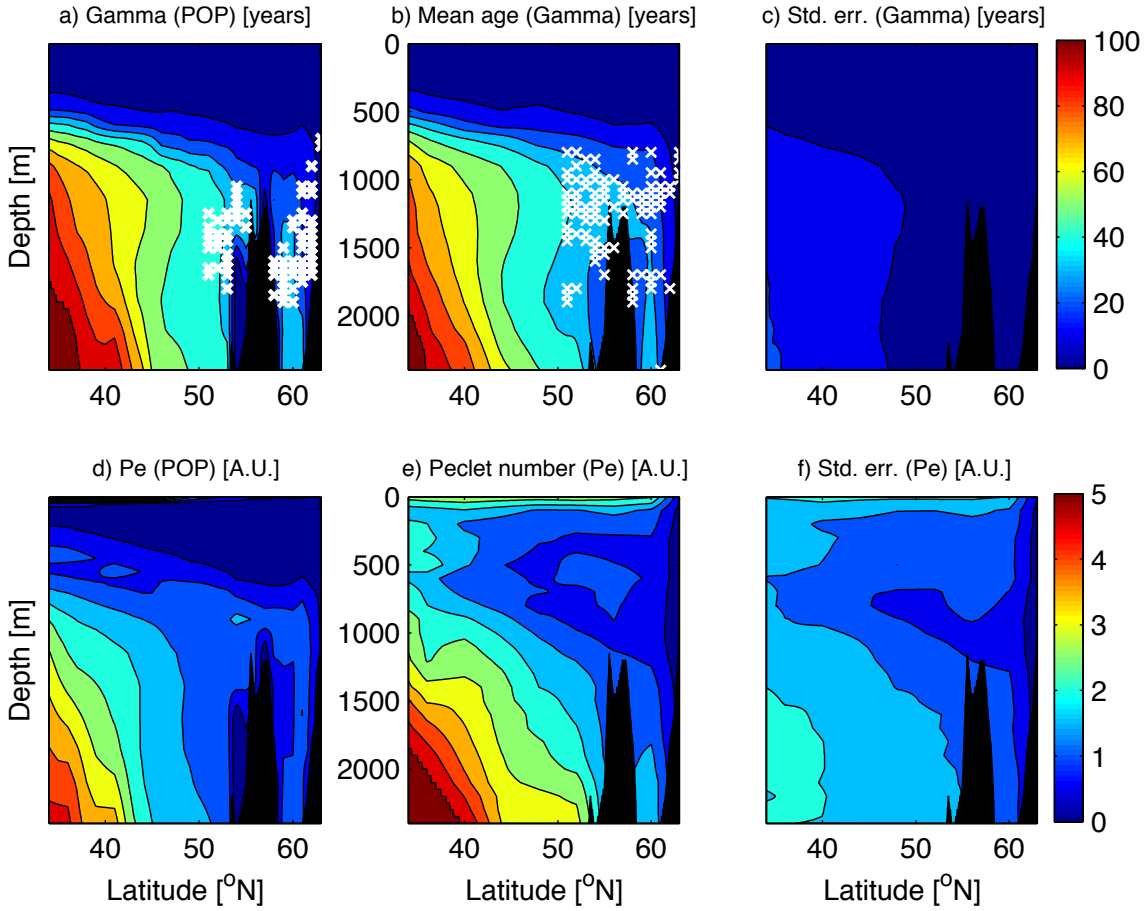

Figure 3.5: Mean ages (a and b) and Peclet numbers (d and e) and their associated errors (c and f) derived via a MCMC scheme (b, c, e, and f) along the A16 WOCE section and plotted against depth and latitude and derived via POP TTDs (a and d). Objectively mapped CFC-11, T, and S and ^{3}He/Tr were used as data, and POP-derived TTD parameters were used as priors. LSW from observations common amongst all four occupations is outlined in panel c with white 'X's and the envelope of all LSW over each year from POP is outlined in panel a with white 'X's.

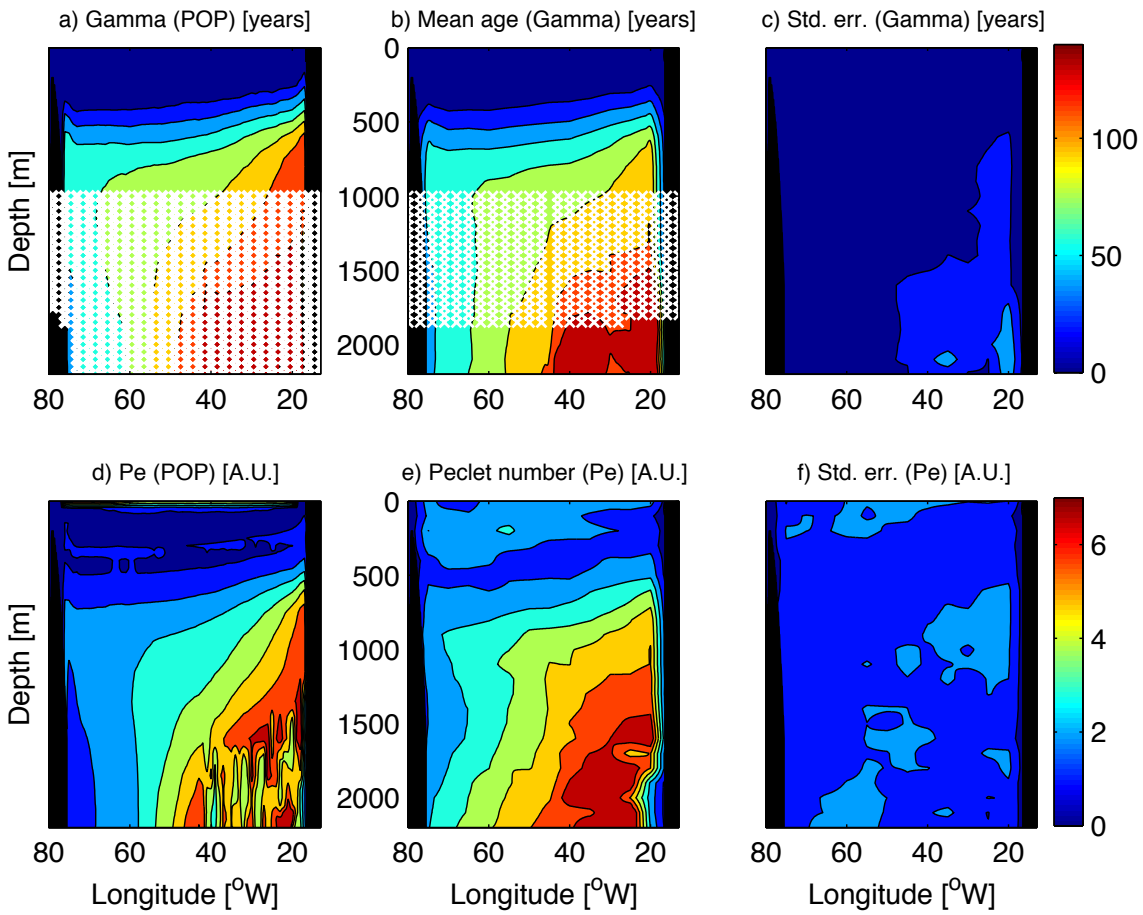

Figure 3.6: Mean ages (a and b) and Peclet numbers (d and e) and their associated errors (c and f) derived via a MCMC scheme (b, c, e, and f) along the A05 WOCE section and plotted against depth and latitude and derived via POP TTDs (a and d). Objectively mapped CFC-11, T, and S were used as data, and POP-derived TTD parameters were used as priors. LSW components of NADW from observations common amongst all three occupations is outlined in panel c with white 'X's and the envelope of all of the LSW components of NADW over each year from POP is outlined in panel a with white 'X's.

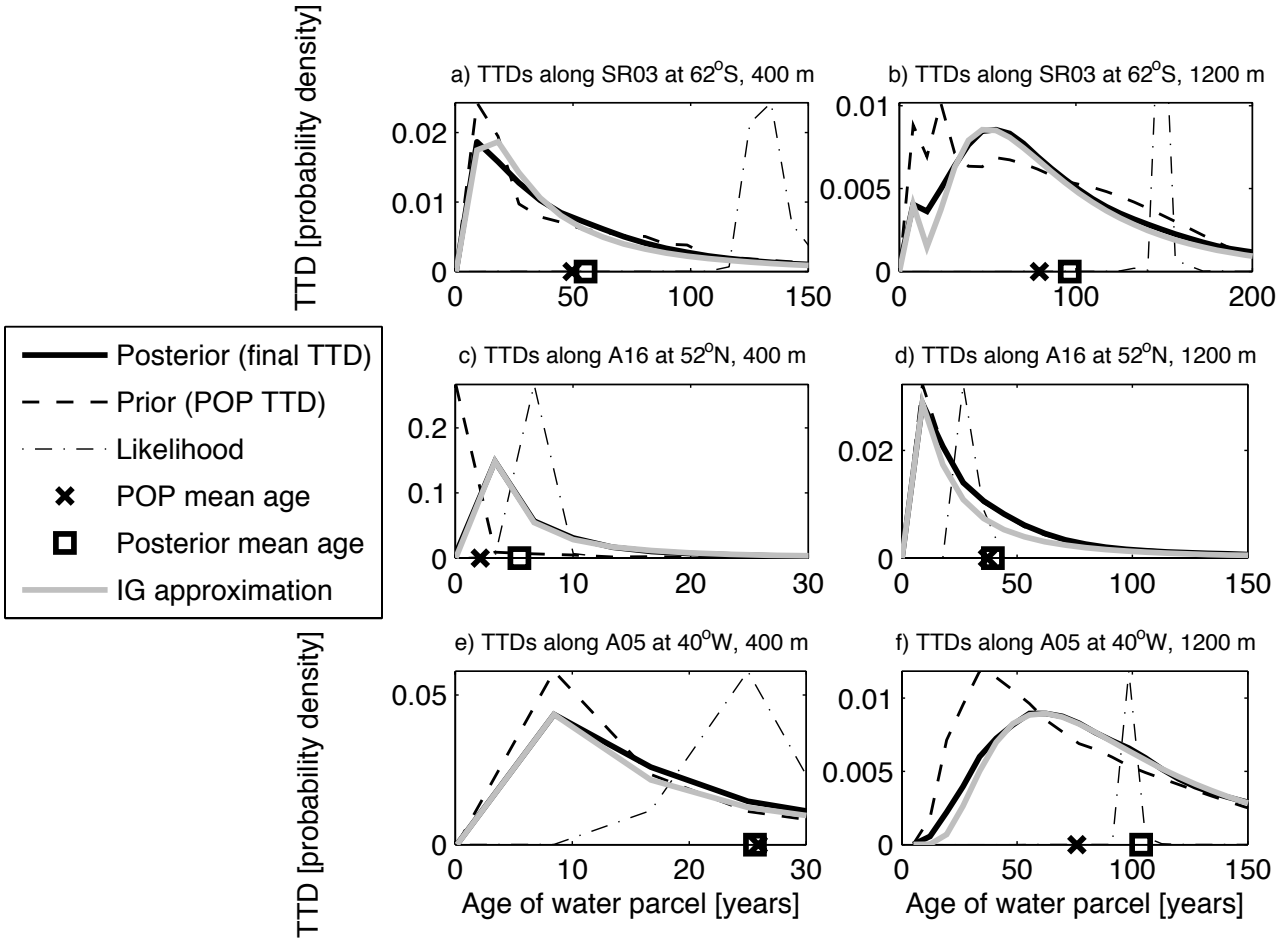

Figure 3.7: TTDs estimated from the BIR method using the POP model (dashed lines) and from our MCMC scheme (solid lines) with a likelihood using a product of symmetric distributions in the variables of (3.10) (and (3.11) for A16) plotted (dot-dashed lines) for SR03 at 62^oS (a-b), A16 at 52^oN (c-d), and A05 at 40^oW (e-f) at 400 meters depth (a, c, and e) and 1200 meters depth (b, d, and f). Also shown is an IG approximation (mixture for SR03) to the posterior TTD estimate entering the MCMC-derived samples into the EM algorithm described in the Appendix (gray lines), the mean age from the posterior TTD estimate (solid square on the abscissa), and the mean age from the POP TTD estimate ('X' on the abscissa).

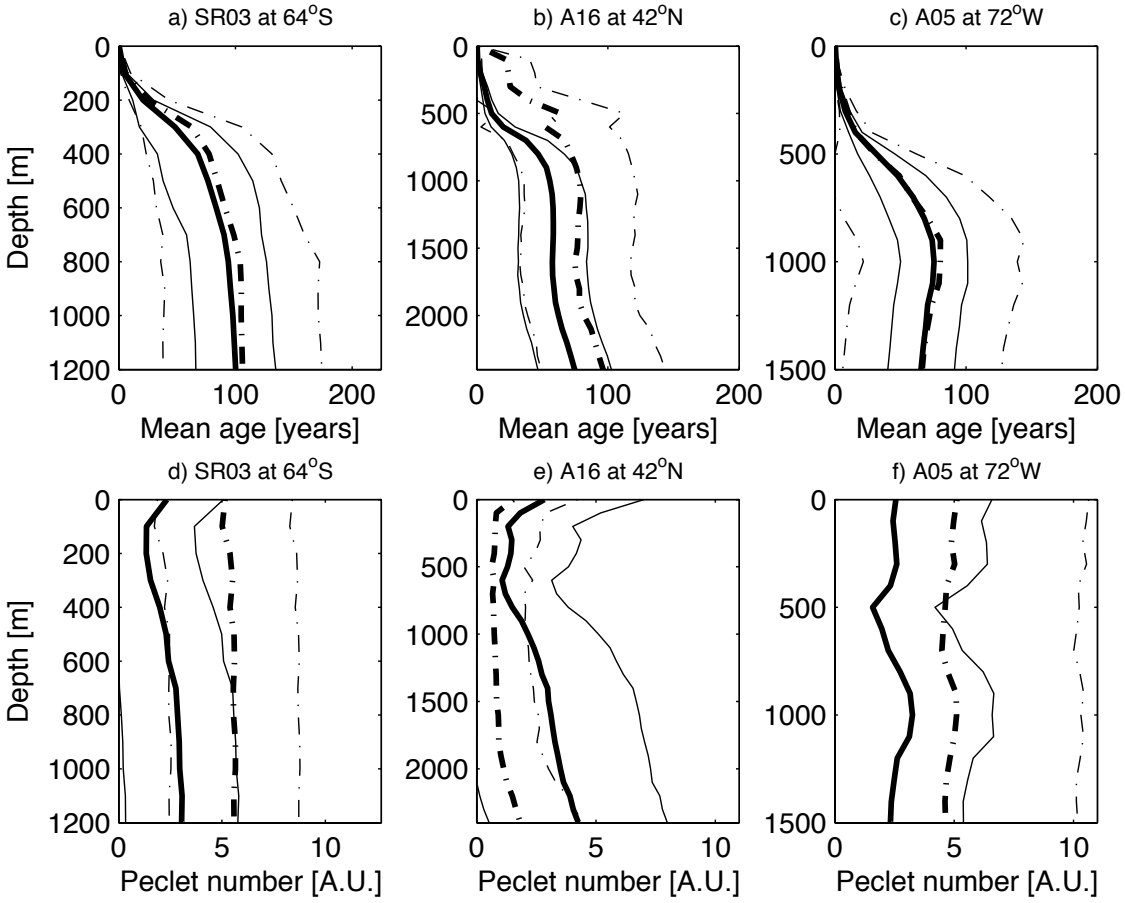

Figure 3.8: The mean ages (a-c) and Peclet numbers (d-f) along SR03 (a and d) at 64^oS, A16 (b and e) at 42^oN, and A05 (c and f) at 72^oW. Shown on each panel is the MCMC-derived estimate of mean age or Peclet number using CFC-11 (and Tritium along A16) with the POP model TTD as a prior (solid lines), and with a uniform prior (dot-dashed lines). A 95% confidence interval (thinner lines) is shown flanking each estimate (thicker lines).

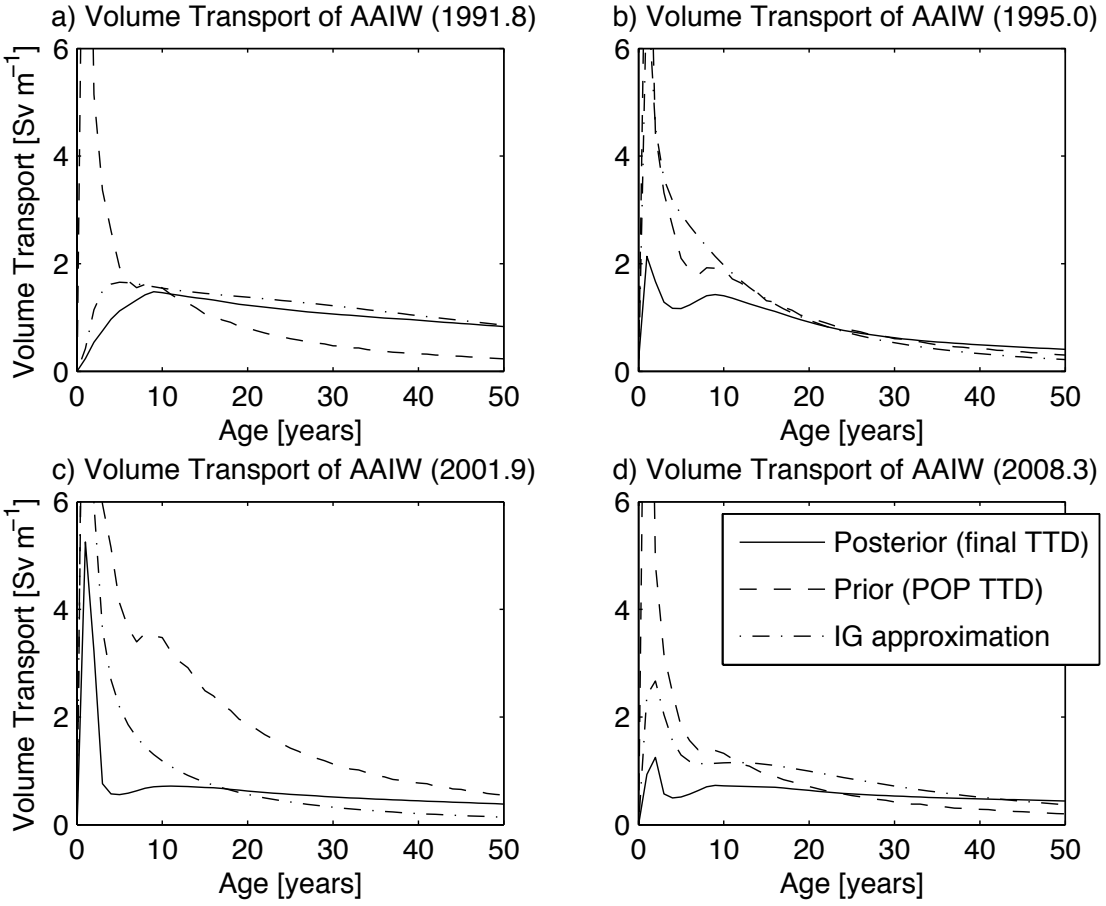

Figure 3.9: The spectra of volume transports [Sv m^{-1}] (proportional to volume-averaged TTDs over the grid boxes that contain AAIW) as a function of age [years] along the SR03 section using the posterior TTD on age (solid), prior TTD from POP (dashed), and posterior mixture model of inverse Gaussians (dot-dashed). Only CFC-11 data were used for a transient tracer. Occupation years of 1991 (a), 1995 (b), 2001 (c), and 2008 (d) are shown. The units, [Sv m^{-1}], of the volume transport spectrum have a m^{-1}, which would be scaled by the meters traveled according to the spreading rate of AAIW to yield the true volume transport spectrum in units, [Sv].

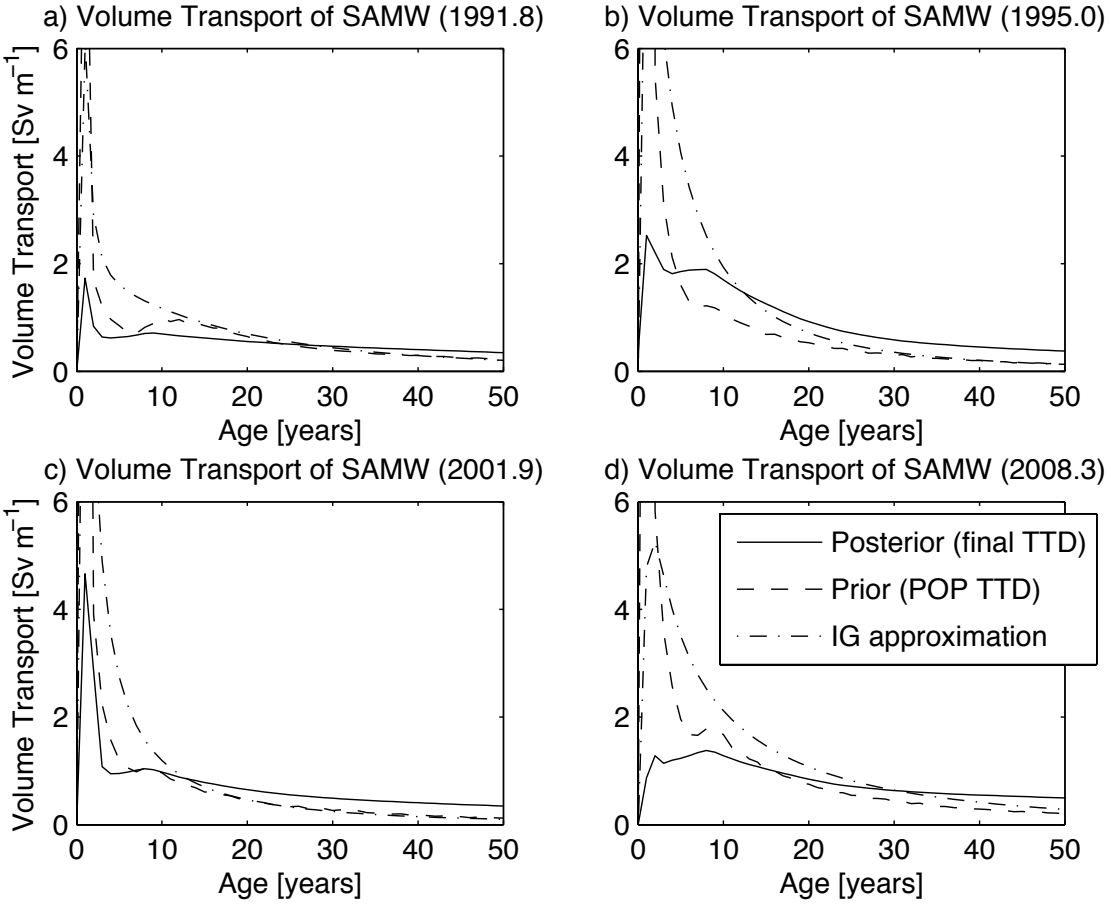

Figure 3.10: The spectra of volume transports [Sv m^{-1}] (proportional to volume-averaged TTDs over the grid boxes that contain SAMW) as a function of age [years] along the SR03 section using the posterior TTD on age (solid), prior TTD from POP (dashed), and posterior mixture model of inverse Gaussians (dot-dashed). Only CFC-11 data were used for a transient tracer. Occupation years of 1991 (a), 1995 (b), 2001 (c), and 2008 (d) are shown. The units, [Sv m^{-1}], of the volume transport spectrum have a m^{-1}, which would be scaled by the meters traveled according to the spreading rate of SAMW to yield the true volume transport spectrum in units, [Sv].

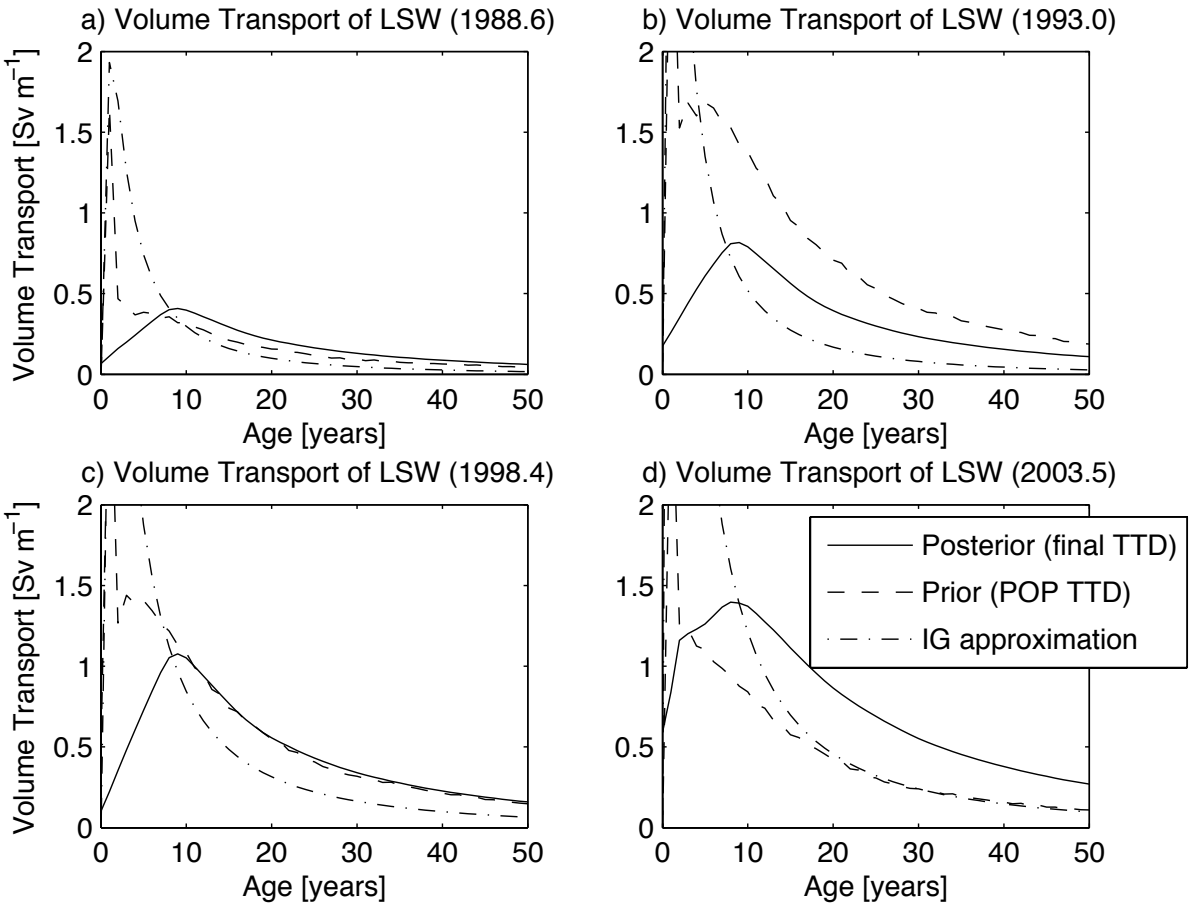

Figure 3.11: The spectra of volume transports [Sv m^{-1}] (proportional to volume-averaged TTDs over the grid boxes that contain LSW) as a function of age [years] along the A16 section using the posterior TTD on age (solid), prior TTD from POP (dashed), and posterior mixture model of inverse Gaussians (dot-dashed). Both CFC-11 and ^{3}He/Tr data were used for transient tracers. Occupation years of 1988 (a), 1993 (b), 1998 (c), and 2003 (d) are shown. The units, [Sv m^{-1}], of the volume transport spectrum have a m^{-1}, which would be scaled by the meters traveled according to the spreading rate of LSW to yield the true volume transport spectrum in units, [Sv].

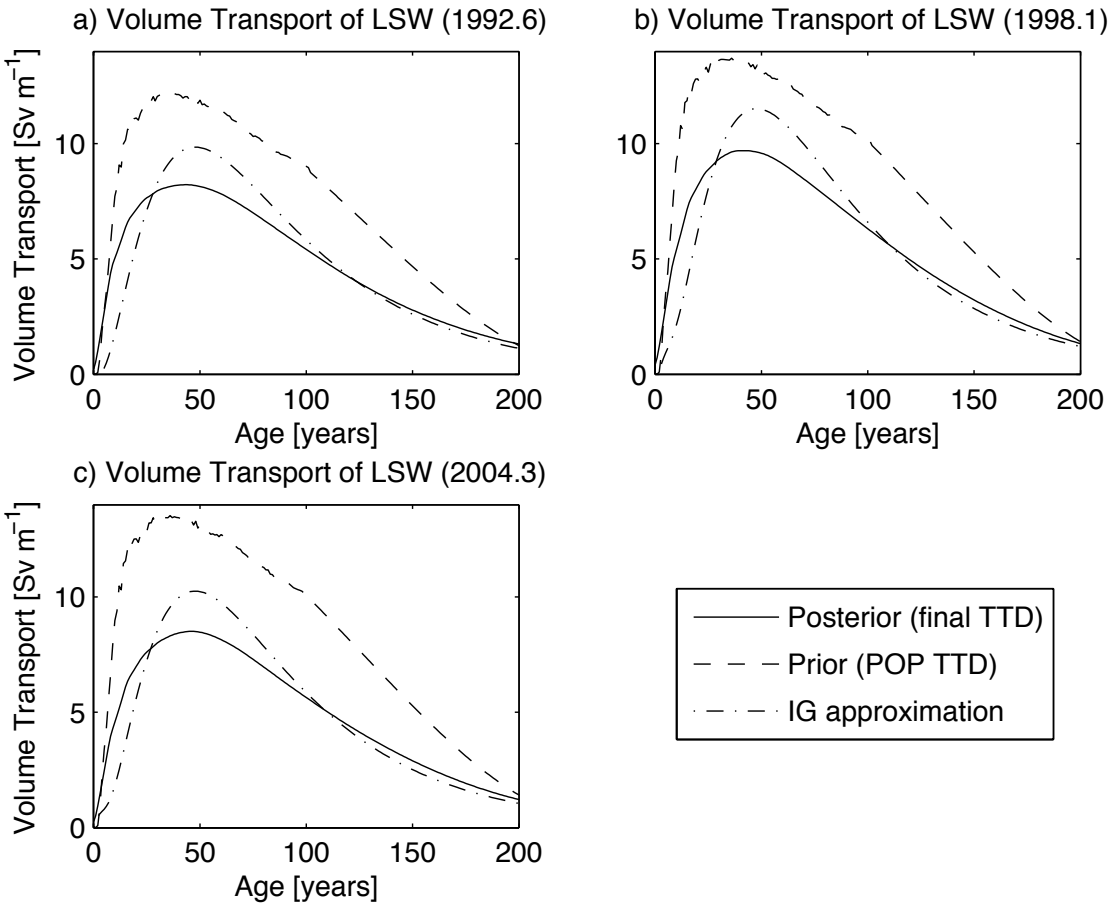

Figure 3.12: The spectra of volume transports [Sv m^{-1}] (proportional to volume-averaged TTDs over the grid boxes that contain the LSW component of NADW) as a function of age [years] along the A05 section using the posterior TTD on age (solid), prior TTD from POP (dashed), and posterior mixture model of inverse Gaussians (dot-dashed). Only CFC-11 data were used for a transient tracer. Occupation years of 1992 (a), 1998 (b), and 2004 (c) are shown. The units, [Sv m^{-1}], of the volume transport spectrum have a m^{-1}, which would be scaled by the meters traveled according to the spreading rate of the LSW component of NADW to yield the true volume transport spectrum in units, [Sv].

Chapter 4

ADVECTIVE VENTILATION RATE ESTIMATES FROM SATELLITE AND PROFILER OBSERVATIONS

This chapter is taken from Trossman et al. (2009) [106] (c)American Meteorological Society. Reprinted with permission.

4.1 Introduction

The upper ocean that separates the atmosphere and the permanent thermocline consists of an Ekman layer, a mixed layer, and a seasonal thermocline. Water parcels gain their identities in the mixed layer and keep their identities to a great extent in the thermocline. To understand climate variability and water mass formation, it is important to know about water mass exchange rates between the surface and the permanent thermocline. The ocean can exchange properties such as heat with the atmosphere and may store these properties for decades or longer through transfer to the thermocline [115].

Subduction is the mass flux from the seasonal thermocline into the permanent thermocline. Obduction is the (irreversible) mass flux into the seasonal thermocline from the permanent thermocline. These processes may also be referred to as ventilation. The top of the seasonal thermocline is defined as the shallowest depth at which the mixed layers base is mixed within a given year. The top of the permanent thermocline is defined as the deepest mixed layer depth within a given year. Seasonally, the entire column of water down to the base of the seasonal thermocline has been mixed, but may not be mixing.

Because temperature and salinity are not conserved when water parcels are mixed, the top of the permanent thermocline is the relevant boundary when considering conservation of physical properties of water parcels. It was believed by Iselin (1939) [41] that Ekman pumping in the late winter (March) forced water to cross the base of the wind-generated Ekman layer and flow along isopycnals thereafter. Stommel (1979) [101] hypothesized that taking all mixed layer parameters from late winter can bypass the complication of the

seasonal cycle and thus the seasonal thermocline can be ignored. Climatological estimates using Stommel's hypothesis assume water parcels subducted from the summer mixed layer are generally re-entrained as the mixed layer deepens the following winter. This is usually referred to as Stommel's mixed layer demon.

In addition to the sign difference, obduction and subduction are different in three ways, according to Qiu and Huang (1995) [77]: 1) Different physical processes govern them. Temperature and salinity properties, as determined by the mixed layer processes, will be carried by the subducted water to the permanent thermocline. On the other hand, water loses its identity via convective mixing when water is obducted. 2) Subduction and obduction have different annual cycles. Subduction occurs after late winter when the mixed layer restratifies and obduction occurs between late fall and early winter when the mixed layer is deepening. 3) A mean annual subduction rate is an estimate of how much water enters the permanent thermocline, while a mean annual obduction rate is an estimate of how much water comes from the permanent thermocline.

Local maxima in the subduction of water occur where there are 'mode' waters, where the water column is nearly vertically homogeneous in density. The mode waters of interest here are also characterized by temperatures between 17 and 19^{o}C [119]. Because of this property, its alternate name is Eighteen Degree Water (EDW). This well-defined water mass is important for the storage of heat and chemical constituents in the North Atlantic. While surface heat fluxes largely determine the sea surface temperature (SST) variability in most of the World Ocean, heat transport by geostrophic flow makes as comparable contribution as surface heat fluxes in the northwestern North Atlantic [22]. Dong and Kelly (2004) [22] have also shown that the interannual heat transport is primarily driven by geostrophic currents, not by Ekman currents.

EDW is a link between the ocean and atmosphere. Production of EDW is believed to be connected to the North Atlantic Oscillation (NAO), sea surface density anomalies near the Gulf Stream, and deep water formation in the North Atlantic subpolar regions [11]. The EDW typically forms just south of the Gulf Stream, where it cools (to the lower end of the 17 to 19^{o}C range) and recirculates. Years with greater EDW production tend to have EDW closer to 17^{o}C while years with lower production have warmer EDW, as the EDW mixes

with warmer recirculating water [50]. EDW is also an important component of the near surface overturning circulation and for transferring biogeochemical properties between the near-surface and the ocean interior since EDW is depleted of nutrients and rich in carbon content [73].

Joyce et al. (2000) [45] found that the Gulf Stream (northern latitude) position, NAO index, and EDW potential vorticity (or lack of EDW production) are highly correlated with one another with a zero lag time. The EDW potential vorticity is also highly correlated with large-scale atmospheric forcing. Hence, three factors have been thought to influence EDW variability: (1) Gulf Stream position, (2) ocean-atmosphere heat fluxes, and (3) temperature and salinity anomaly advection. However, Dong et al. (2007) [21] have found that both the oceanic advective convergence and the product of heat fluxes and the area of the outcropping 17 to 19^oC isotherms, not variability of surface heat flux, play the most important roles in EDW variability. The oceanic advective convergence indicates the importance of (1) and (3) while the area of the outcropping 17 to 19^oC isotherms indicates the importance of preconditioning. Since the NAO affects the SSTs, a link between the NAO and EDW may be drawn if it may be shown that the SSTs affect EDW volume.

Kwon and Riser (2004) [51] have shown that EDW volume reaches its annual peak in the winter. Using profile estimates, they show that the formation rate of EDW is on the order of $4-5$ Sv from profile estimates, as opposed to the $15-20$ Sv Speer and Tziperman (1992) [99] estimated from air-sea fluxes using the Walin (1982) [109] formulation. The Walin (1982) [109] formulation diagnoses the EDW formation rates by using diabatic forcings as a function of outcropping isopycnals. No circulation or mixing is included so mixing of density within the interior is ignored. Nurser (1999) [69] found both deepening of the Winter mixed layer and lateral mixing within the mixed layer (including entrained denser waters into lighter ones) to be potentially significant at mid-latitudes, albeit less significant than mixing in sustained upwelling regions in the tropics. The discrepancy between the estimates of Kwon and Riser (2004) [51] and Speer and Tziperman (1992) [99] may be caused by an error in the Walin (1982) [109] formulation or inaccurate air-sea fluxes, but most probably can be accounted for by the lateral mixing, or eddy processes, that result in lateral transport in the mixed layer near the Gulf Stream. These processes have patchy spatial distributions to

which bulk parameters may not do justice.

Here we build on the work of Qiu and Huang (1995) [77] who made subduction rate estimates for the North Atlantic (and Pacific) with climatological estimates for MLDs, climatological estimates for currents, and Hellerman and Rosenstein (1983) [35] wind stress data (which has been noted to overestimate the wind stress by 30 % [33]). We recalculate Ekman pumping, as Xie (2004) [120] did with QuikSCAT wind stress data from 1999 to 2003, but including the impact of upper ocean geostrophic currents on Ekman velocities. We include winds, currents, observationally-derived MLDs, and vertical shear to calculate (EDW) ventilation rates and compare with previous estimates. Section 4.2 describes the data we use in our Ekman pumping and (EDW) ventilation rate estimates. We end with results in Section 4.3 before summarizing.

4.2 Data and methods

Owing to recent realizations of satellite and profiler systems, we are able to investigate not only interannual variability in a few selected locations but also intra-annual variability across the Gulf Stream and Sargasso Sea. QuikSCAT winds are available daily from July 1999 and mapped as in Kelly et al. (1999) [47] (more information available online at `http://airsea.jpl.nasa.gov/DATA/QUIKSCAT/wind/`); Archiving, Validation, and Interpretation of Satellite Oceanographic data (AVISO) sea surface heights (SSHs) are available weekly from October 1992 (`http://www.aviso.oceanobs.com/en/data/products/sea-surface-height-products/global/msla/index.html`); Argo temperatures and salinities are available every 10 days from the year 2000
(`http://www.usgodae.org/ftp/outgoing/argo/geo/atlantic_ocean`); and World Ocean Circulation Experiment/Atlantic Climate and Circulation Experiment (WOCE/ACCE) mixed layer depths and temperatures are from 1997 through 2003. World Ocean Atlas 1994 (WOA) temperature and salinity climatologies are available for each month of the year (`http://www.cdc.noaa.gov/data/gridded/data.nodc.woa94.html`). Our data are summarized in Table 4.1.

4.2.1 Surface currents

The AVISO SSH product, on a $1/3 \times 1/3^o$ grid, also includes error maps, which indicate significant errors located between the tracks of Ocean Topography Experiment (TOPEX)/Poseidon and Jason-1. Because of these errors, surface geostrophic flow calculated from the SSH anomalies shows signals that reflect the satellite track patterns. The anomalies have high-frequency noise and coverage errors in them while the mean does not, as the mean was derived from historical and hydrographic float (surface) data following Singh and Kelly (1997) [92]; more information available online at
`http://kkelly.apl.washington.edu/matl/meanssh.html`. Because more structural detail would be lost if the sum of the anomalies and mean SSHs is filtered, we first filter the anomalies and then add on the mean SSHs (not filtered by us). We use spatial low-pass filtering on the SSH anomaly data after temporally averaging over three-week blocks of time. We use a half power of six longitudinally (between 80^o and 30^oW) and a half power of five latitudinally (between 20^o and 55^oN), low-pass filtering once in each direction with a Butterworth filter. In regions where the errors are greater than their corresponding SSH anomalies, we assume that these SSH anomalies are zero, which is typically when errors exceed 20 cm. The only places with greater errors are along the coast (with errors sometimes as high as 50 cm). We set our final SSH to be the sum of these low-passed anomalies and the mean AVISO SSH. The resulting currents are derived from the geostrophic relationships

$$u_g = -\frac{g}{f}\frac{\partial \eta}{\partial y} \tag{4.1}$$

$$v_g = \frac{g}{f}\frac{\partial \eta}{\partial x}, \tag{4.2}$$

where $g = 9.81$ m s^{-2}; f is the Coriolis parameter, which we allow to vary with latitude; and h is the SSH. These currents do not show a track pattern. Additional filtering was applied to mask over the jump between land and nonzero currents. If the difference between the currents between two grid points was 0.25 m s^{-1} or larger, the currents from that longitude to the westernmost grid point and within a box of three latitudinal grid points are set to zero. The mean annual geostrophic surface current structure maps out the Gulf Stream (lhs of Fig. 4.1). This objective map gives smaller SSH anomalies in regions of larger errors,

which in turn gives slower current speeds.

4.2.2 Shear

To obtain the vertical structure of the currents, we low-pass filter the potential densities averaged monthly from January 2002 to December 2006 using the same procedure as in low-pass filtering the SSH: once with a half power of six longitudinally (between 80^o and 30^oW) and a half power of five latitudinally (between 208 and 558N). We find these potential densities using WOA (T, S) data when the errors associated with objectively mapped Argo (T, S) data exceed 50 m as calculated according to the objective mapping method of Le Traon (1990) [105]. Then we use the thermal wind equations,

$$\frac{\partial u_g}{\partial z} = \frac{g}{f\rho_o}\frac{\partial \rho_\theta}{\partial y} \tag{4.3}$$

$$\frac{\partial v_g}{\partial z} = -\frac{g}{f\rho_o}\frac{\partial \rho_\theta}{\partial x}, \tag{4.4}$$

where ρ_θ is the potential density and ρ_o is the mean density over our domain.

We grid these current shears with 50-m vertical resolution, vertically integrate, and then use the AVISO surface currents from the previous section as a reference. The ratio of the currents at some depth to currents at the surface is capped at one and filled with the spatially averaged ratios for this depth if greater than one (which occurs at less than 10% of grid points). These currents are low-pass filtered once in both horizontal directions with a half power of four. Last, we fit logarithmic curves to the profiles of these ratios and low-pass filter with a half power of four again for a small shear profile.

4.2.3 Ekman pumping

One way that water parcels may obduct or subduct is by vertical advection or vertical pumping through Ekman currents (Appendix, Section A.1). Ekman pumping is defined by $-w_e \equiv -\vec{\nabla} \cdot \vec{U_e}$, as a convergence of Ekman transports $\vec{U_e}$, which leads to a downward motion. An upward motion results from positive Ekman suction, w_e. Usually, the Ekman transports are approximated with wind momentum input uncoupled from currents:

$$U_e = \frac{1}{\rho}\frac{\tau^y}{f} \tag{4.5}$$

$$V_e = -\frac{1}{\rho}\frac{\tau^x}{f}, \tag{4.6}$$

here called the simple Ekman balance (Appendix, Section A.1). This analytical approximation, which Qiu and Huang (1995) [77] and Xie (2004) [120] have used for w_e, is valid in most regions, but the assumptions fail in the Gulf Stream. If surface geostrophic currents are strong, the vertically averaged flow in the Ekman layer satisfies (Appendix, Section A.1)

$$\frac{\partial U_e}{\partial t} + U_e\frac{\partial u_g}{\partial x} + u_g\frac{\partial U_e}{\partial x} + V_e\frac{\partial u_g}{\partial y} + v_g\frac{\partial U_e}{\partial y} - fV_e = \frac{\tau^x}{\rho} - \tau_{damp}U_e \tag{4.7}$$

$$\frac{\partial V_e}{\partial t} + U_e\frac{\partial v_g}{\partial x} + u_g\frac{\partial V_e}{\partial x} + V_e\frac{\partial v_g}{\partial y} + v_g\frac{\partial V_e}{\partial y} + fU_e = \frac{\tau^y}{\rho} - \tau_{damp}V_e. \tag{4.8}$$

We have used the approximation of Qiu and Huang (1995) [77] and Xie (2004) [120] for a first guess to the Ekman pumping. We also impose the simple Ekman balance at the coasts. Such regions serve as boundaries to the regions in which the Qiu and Huang (1995) [77] approximation is not valid and help define a well-posed problem. To numerically solve for a converged solution, we have applied a time difference method with time steps, $\Delta t = 1/2$ hr. The geostrophic currents are held fixed for each week of Aviso SSH data and QuickScat wind data. To find the Ekman solutions, we time step, including accelerations as well as damping terms, $\tau_{damp}U_e$ and $\tau_{damp}V_e$. The damping factor, τ_{damp}, should be short relative to the temporal resolution of our data, a week, and long relative to the time steps, Δt. Running for $2/\tau_{damp}$ days, the solution reaches steady state. This yields the Ekman currents, and from this, we calculate w_e. The root-mean-squared differences between using $\tau_{damp} = 1/1$ day^{-1} and $\tau_{damp} = 1/1.675$ day^{-1} for u_e, v_e, and w_e, respectively are 0.0020 +/- 0.0014 m/s, 0.0016 +/- 0.0009 m/s, and 2.55×10^{-7} +/- 1.19×10^{-7} m/s. The solution does depend on τ_{damp}, but we choose $\tau_{damp} = 1/1.675$ day^{-1} to ensure enough time for the Ekman layer to set up.

4.2.4 Subduction rate approximation

Subduction is defined as positive when fluid descends through the thermocline, so subduction is positive when w_e is negative. This can be seen in the expression for approximate mean annual subduction rates [60, 77]:

$$S_{ann} = -\overline{(w_e - \frac{\beta}{f}\int_{-\delta_e}^{0} v_g dz)} + \frac{1}{\tau}(h_{m,0} - h_{m,1}). \tag{4.9}$$

Here, τ is one year and $h_{m,0}$ and $h_{m,1}$ are the March MLDs for the year in which we want to compute the mean subduction rate and the year after that, respectively. The first term represents vertical pumping and the second term represents lateral induction. Average rates of obduction may be computed similarly but with MLD information from one year and the year previous. Over a period, τ, an average subduction rate may be computed via:

$$S_{avg} = -\frac{1}{\tau}\int_{t_0}^{t_1} w_{tr}dt - \frac{1}{\tau}\int_{t_0}^{t_1}(\frac{\partial h_m}{\partial t} + u_{tr}\frac{\partial h_m}{\partial x} + v_{tr}\frac{\partial h_m}{\partial y})dt \tag{4.10}$$

where t_0 and t_1 are the initial and final times (functions of space and year in our parcel tracking model to be discussed later), $\tau = t_1 - t_0$, w_{tr} and (u_{tr}, v_{tr}) are the vertical and horizontal velocities at the base of the mixed layer, h_m is the maximum depth of the mixed layer in the given year, and the $\partial h_m/\partial t$ term averages to zero over a sufficiently long τ [77].

4.2.5 MLDs

Water parcels may also obduct or subduct by movement of the MLDs relative to a water parcel that allows the MLD to entrain (obduct) or detrain (subduct) this parcel (4.9). The MLDs are calculated using Argo and WOCE/ACCE data. When available, we use adjusted temperatures and salinities and only measurements with quality "1" ratings (`http://www.coriolis.eu.org/cdc/argo_rfc.htm`). The floats with World Meteorological Organization identification number 852 from 2004 onward have been excluded owing to a nonconstant cold bias discovered in February of 2007 (`http://argo.jcommops.org`). These floats were fairly uniformly distributed in the North Atlantic, leaving coverage as shown in Fig. 4.2. Thus, while we expect the MLDs to be biased low including these data and to be higher excluding these data, the MLDs should be better with their exclusion.

We use the shallowest pressures at which the criterion, $\Delta T > 0.3^o$C, is satisfied unless stated otherwise. Here, ΔT is the difference between the surface temperature and the *in situ* temperature at some depth. Alternatively, $\Delta\rho_\theta > 0.125$ kg/m^3 may be used as a criterion where $\Delta\rho_\theta$ is the difference between the surface density and the potential density at some depth. Using the $\Delta T > 0.3$ criterion to estimate the MLDs yields results that look qualitatively different from those obtained using $\Delta\rho_\theta > 0.125$ as a criterion along the coasts and north of 45^oN. Despite how we should be using the $\Delta\rho_\theta > 0.125$ criterion, our

WOCE/ACCE data lacks salinity so we settle with the $\Delta T > 0.3$ criterion.

For each grid point, we calculate a simple two-sample t statistic with

$$t = (\overline{MLD_T} - \overline{MLD_{\rho_\theta}})\sqrt{(\frac{\sigma_T^2}{n_T} + \frac{\sigma_{\rho_\theta}^2}{n_{\rho_\theta}})^{-1}}, \tag{4.11}$$

with degrees of freedom

$$df = (\frac{\sigma_T^2}{n_T} + \frac{\sigma_{\rho_\theta}^2}{n_{\rho_\theta}})^2\sqrt{(\frac{1}{n_T - 1}(\frac{\sigma_T^2}{n_T})^2 + \frac{1}{n_{\rho_\theta} - 1}(\frac{\sigma_{\rho_\theta}^2}{n_{\rho_\theta}})^2)^{-2}}, \tag{4.12}$$

σ^2 representing the sample temporal variance where the MLDs are not zero, $n = 48$ representing the number of months in our sample, T subscripts corresponding to estimates derived from the $\Delta T > 0.3$ criterion, and ρ_θ subscripts corresponding to estimates derived from the $\Delta\rho_\theta > 0.125$ criterion [97]. (A more accurate test to use in this case would be a paired difference test [97], but the conclusion of this hypothesis test is identical.) We reject the null hypothesis that the (averages of the) MLDs are the same along the coasts and in the northeast region of our domain using the two criteria to 95 % confidence. In these regions, using the $\Delta\rho_\theta > 0.125$ criterion yields deeper MLDs. This is consistent with where salinity profiles are relatively homogeneous or compensate for the temperature change in the vertical.

The MLDs we estimate based on each of the two criteria are objectively mapped according to the method of Le Traon (1990) [105], building on the work of Bretherton et al. (1976) [8]. The number of grid points that had valid estimates (non-zero MLD anomalies) was largest in 2004 because that year had fewer number 852 floats than subsequent years. The number of points our objective mapping routine uses on average in a given month to describe the data are 1572 in 2002, 1605 in 2003, 1614 in 2004, 1597 in 2005, 1604 in 2006, and 1625 in 2007. We set our large scale basis function to be $1 + x + y + x^2 + y^2$, use a Poissonian covariance function, and use a decorrelation length scale of 150 kilometers.

The boundary between the seasonal and permanent thermoclines is defined to be the maximum MLD, h_m, over one year time. Here, one year is the time from when the MLD is deepest one year to when the MLD is deepest in the next year, at each horizontal grid point (e.g. Fig. 4.3). This can be shorter or longer than one year and this variability of time between maximum MLDs greatly influences the result. A concern is whether our results will

be an artifact of better data coverage in later years. The means and standard deviations of the time-averaged MLDs and their time-averaged associated objective mapping errors calculated according to the method of Le Traon (1990) [105] are shown in the top and bottom half of Figure 4.4 respectively. Because the spatial variance from objective mapping is constant, on the order of $10,000$ m^2, the errors are relatively small.

4.2.6 Parcel Tracking

Following the method of Qiu and Huang (1995) [77], fluid parcels are tracked from a given initial position and time until they are obducted or run into the coast, with time steps of seven days. We track the parcel within the upper 1000 meters, and within the domain of $27 - 53^o$N and $78 - 32^o$W. The parcels are placed on a 2^o latitude by 2^o longitude by 100 meters depth grid system and tracked for each horizontal and vertical grid point. The geostrophic velocities of these water parcels are calculated subject to the vertical shear at each of these grid points. The purposes of this $2^o \times 2^o \times 100$ m grid are to: (1) resolve the vertical shear in the geostrophic flow and (2) start following water parcels from each of these grid points.

To track the parcels, a fourth-order Runge-Kutta integration was used. The numerical time intergrations have been performed with Matlab's `ode45` (fourth-order Runge-Kutta) routine at a relative tolerance of 1%. The velocities were integrated along the trajectory of each water parcel, $u_{tr} = (u_g + u_e)/(111200 \times cos(\pi\lambda/180))$, $v_{tr} = (v_g + v_e)/111200$, and

$$w_{tr} = \begin{cases} w_e z/\delta_e + \frac{\beta}{f}\int_{-z}^{0} v_g dz` & \text{if } 0 > z > \delta_e \\ w_e + \frac{\beta}{f}\int_{-z}^{0} v_g dz` & \text{if } \delta_e > z \end{cases}$$

where λ is degrees latitude, $\delta_e = 40$ meters is the Ekman layer depth (below which u_e and v_e are 0), 111200 is the conversion factor from m/s to degrees/s, f is the Coriolis parameter allowed to vary with latitude as $sin(\pi\lambda/180)$, and β is the latitudinal derivative of f. Finally, we assume that flows just below the mixed layer are geostrophic.

We find the depths at which water columns cross h_m to find the thickness of each water parcel that gets transferred to or from the thermocline over τ, the amount of time it takes for the mixed layer to be deepest from one year to the next. Here, τ is the interval of

time between releases of water parcels. A subducted (Fig. 4.5, top) or obducted (Fig. 4.5, bottom) water parcel is defined by its top and bottom. For subducted parcels, the bottom is the maximum MLD in its first year of being tracked and the top is determined by tracking the parcel. Likewise, for obducted parcels, the top is the maximum MLD in its first year of being tracked and the bottom is determined by tracking the parcel. We interpolate from a $2^o \times 2^o \times 100$ m grid onto a $1/2^o \times 1/2^o \times 20$ m grid to find the top of subducted parcels and the bottom for obducted parcels. The 20 m vertical resolution comes from our approximate error estimates in regions where ventilation is predicted to be non-zero. We have performed the same calculation with a 10 meter resolution and obtained results that are different by at most 10 m yr^{-1}, less than our error. We assume that water parcels that go outside of our model domain neither subduct if they are above the MLD nor obduct if they are below the MLD when this occurs.

4.2.7 Parcel Identities

Water parcel properties (temperature, salinity, etc.) are constant below the mixed layer and only change in the mixed layer. A parcel's properties, once in the permanent thermocline, are identical to those when and where they subduct in our Lagrangian trajectory calculation so a parcel's properties are set once out of the mixed layer. Due to mixing at the surface, if a parcel comes back into the mixed layer, its properties are no longer assumed to be constant.

The waters that are EDW: 1) have temperatures between 17 and 19^{o}C, 2) are West of 40^{o}W and North of 25^{o}N, and 3) have a vertical temperature gradient $dT/dz < 0.006^o$C/m. Kwon and Riser (2004) [51] impose 2) onto their region to calculate EDW volume. This makes at most a negligible 1% difference in volume given the two other properties of EDW. We estimate that the equivalent criterion of (3) in terms of potential vorticity, fN^2, for Brunt-Väisälä frequency $N^2 = -g/\rho(\partial\rho/\partial z)$, Coriolis parameter f, and $g = 9.81$ m/s^2 to be $|fN^2| < 1 \times 10^{-7}$ m^{-1} s^{-1}.

We calculate the vertical thickness of EDW and multiply by the area of a grid cell at each grid point to find the volume of EDW in a given month from objectively mapped thicknesses of Argo and WOCE/ACCE temperature data. This objective mapping procedure is

the same as the one for the MLDs. We calculate EDW ventilation rates by using the objectively mapped temperatures from Argo and WOCE/ACCE data to identify the EDW as we track it. We do not use WOA data substitutions here because this reduces the interannual variability in the record. By tracking how much EDW enters the permanent thermocline each year, we can infer the mean annual EDW subduction rates. We do the same with EDW exiting the permanent thermocline to get the mean annual EDW obduction rates. The EDW ventilation rates are dependent upon the aforementioned mixing assumption of our model, but the EDW volume is not since it is calculated directly from data.

4.2.8 Error

Using an estimated variance, σ_h^2 on the order of $10,000$ m^2 for the MLDs, we may find an upper bound on the error in ventilation rate between two consecutive years. To do this, we set $(\Delta h_{err})/\tau = e_m$. Here, $(\Delta h_{err})^2 = \varepsilon_{m,0}^2 + \varepsilon_{m,1}^2$ is the sum of squared MLD errors for two consecutive years . Both $\varepsilon_{m,0}^2$ and $\varepsilon_{m,1}^2$ are computed by multiplying the associated errors from objective mapping in percent of variance by σ_h^2. We may also estimate the ventilation rates using the lateral induction term in the Qiu and Huang (1995) [77] formulation (4.9) with $(\Delta h_m)/\tau$.

The spatial distribution of $(\Delta h_{err})/\tau$ is similar to that of ε_m in Figure 4.4 but patchier. The errors are largest around 60^oW and 37.5^oW in the south and off the coast of Newfoundland. By multiplying each $(\Delta h_{err})/\tau$ and $(\Delta h_m)/\tau$ ('analytical' estimates tabulated in Table 4.3) by the area of a grid cell at each grid cell point and summing all of the grid points where ventilation is non-zero, we may get volume transport estimates for error and ventilation. We find that the volume transport error estimates can be as large as 50 % of our volume transport ventilation estimates.

4.3 Results

4.3.1 Ventilation analysis

By taking advantage of new datasets, we are able to make new estimates for the quantities needed to estimate the ventilation of the thermocline in the North Atlantic. The

ageostrophic currents are small in most places relative to the geostrophic ones (right-hand panel of Fig. 4.1). The first-order corrected ageostrophic currents only get up to 15.8% of the geostrophic currents at any given place or time and almost never get above 10%. The greatest geostrophic current contributions to the Ekman transports come from terms with derivatives in the geostrophic flow (e.g., $U_e \partial u_g / \partial x$). The contributions from the derivatives in the Ekman transports are negligible, both by scaling arguments and direct calculation, especially outside of the Gulf Stream.

Accounting fully for the geostrophic currents numerically, w_e is noticeably different when calculating w_e directly from the wind (Fig. 4.6). The advection of Ekman velocities by geostrophic currents decreases w_e by 20% where the Gulf Stream makes its sharpest turns (consistent with our ageostrophic velocities being largest there). The advection of Ekman velocities by geostrophic currents increases w_e by about a 10% elsewhere. Comparing our w_e with Xie (2004) [120], who also used QuikSCAT winds, we see similar structure and magnitude, whereas w_e of Qiu and Huang (1995) [77] is only half as big in the Gulf Stream. There is less pumping south of the Gulf Stream in the Sargasso Sea and less suction north of the Sargasso Sea than in Xie (2004) [120].

Parcels travel at most 5^o latitudinally but may travel as many as 35^o longitudinally in a year. The means and standard deviations of longitudinal and latitudinal distances water parcels travel each year in our model are given in Table 4.2. Lateral induction, $(\Delta h_m)/\tau$, dominates both subduction and obduction away from the coast, contributing a few hundreds of meters per year while vertical transport contributes at most 75 m yr^{-1} in the Qiu and Huang (1995) [77] formulation (4.9). However, vertical transport via w_e becomes quite significant close to the coasts.

With shear in weekly currents and monthly MLDs for individual years from 2002 to 2005, our ventilation rate estimates are shown in Fig. 4.7. This may be compared to our use of the Qiu and Huang (1995) [77] formulation with associated errors (Table 4.3). Using the Qiu and Huang (1995) [77] formulation, the errors are much less than ventilation rate estimates everywhere except the coast of Newfoundland and two southern regions of our domain. Because of these errors, the estimated magnitudes of obducted water parcels from along the coast of Newfoundland are questionable.

By taking the monthly average of the Argo-derived MLDs over 2002-06, annual average of the currents with shear over 1999-2006, and annual average of the vertical velocities over 1999-2006, we may compute climatological ventilation rates (Fig. 4.8). Vertical shear in the geostrophic flow has a second-order effect on the ventilation rates. This shear may cause up to a 10% decrease in subduction but is not clearly out of our range of error. To see this, an average monthly MLD climatology has been constructed over 2002-2006 and used to calculate ventilation both with and without shear (Fig. 4.8, top and second from top).

Comparing the use of our $\Delta T > 0.3$ criterion (Fig. 4.8, top) with the $\Delta\rho_\theta > 0.125$ criterion (Fig. 8, second from bottom), we find up to a 30% decrease in obduction using the former because h_m is smaller in the northeast region of our domain where we estimate more obduction than subduction. The difference between total subduction and obduction is closer to zero using the $\Delta\rho_\theta > 0.125$ criterion. Comparing our estimates with and without Ekman currents (Fig. 4.8, top and bottom), we find up to a 15% decrease in ventilation when $w_e = 0$.

We may also compare the climatology using the $\Delta T > 0.3$ criterion with shear to the ventilation rates that we get by using the monthly Argo-derived MLDs (not averaged over 2002-06), average of currents over 1999-2006 with shear, and average of vertical velocities over 1999-2006, and finally taking the average of these ventilation rates over 2002-06 (Fig. 4.9). We find that the ventilation rates calculated using mean currents are about 5% smaller and using mean MLDs are about 45% smaller. The lateral induction term of (4.9) is mostly responsible for this (nonlinear) difference in ventilation using means instead of weekly currents and monthly MLDs.

More specifically, the MLDs are not deepest in March every year (Fig. 4.3, bottom). For example, one year may have a deepest MLD in April and the next may have a deepest MLD in February at a given horizontal grid point, as governed by convective cooling and stratification strength. Then the ratio of Δh_m to τ may not stay constant from year to year. Here, t is the time from when h_m is deepest one year to when h_m is deepest in the following year. The other nonlinear component to the ventilation rate estimation pertains to spatial distributions of currents and MLDs. The MLDs within the Gulf Stream are not as deep as they are to its south and east; as stratification is strong within the Gulf Stream, the

negative buoyancy forcings are sufficiently counteracted by advected heat. The subduction rates are nonzero south and east of the Gulf Stream (Fig. 4.7) where the MLDs deepen most (Fig. 4.3, top). Their magnitudes are consistent with Qiu and Huang (1995) [77].

On the other hand, the imbalance of subduction and obduction over our time series warrants further attention: as Table 4.4 shows, over our four years there is a mean net subduction (subduction minus obduction) of 0.6 Sv of water with more subduction than obduction only in 2004 and 2005. A mean net subduction greater than zero would imply a shoaling of MLDs over the time period, but this is sensitive to the domain over which this is estimated. It should be noted that, if the $\Delta\rho_\theta > 0.125$ criterion were used, the MLDs in the northeast region would be deeper. There would be less subduction and more obduction. We have not tracked parcels north of 50^oN where MLDs get deep enough to also yield more obduction that may account for the remaining discrepancy.

Last, we cannot estimate how long water parcels reside in the thermocline after subduction because we only have four years of data that can be tracked. However, we see that some of the parcels that reside in the thermocline for more than a year can be identified as EDW. Table 4.4 shows that there is a net subduction of EDW greater than zero for each of the four years, while the EDW volume stays relatively stable. The EDW is not likely to travel north of 50^oN and obduct. Any EDW that travels south of 25^oN is unlikely to obduct. We may hypothesize that, by properly taking mixing into account, this EDW ventilation imbalance will diminish, but we cannot test this with our model.

4.3.2 Density class analysis

The volume transport rates in our domain for σ_θ up to 28.2 kg/m^3 are in Table 4.4. The subducted waters come from regions south and southeast of the Gulf Stream, while obducted waters come from regions along the Gulf Stream and closer to shore. Most of the subducted waters are slightly less dense than the obducted waters (Fig. 4.10), consistent with Qiu and Huang (1995) [77]. There is generally a right-skewed Gaussian-shaped distribution of ventilation as a function of density and a secondary peak in subduction for some years. The reason for this double peak in subduction is the local subduction minimum in the

$27.2\sigma_\theta$ density class (seen barely in 2002, in 2003, and markedly in 2005), consistent with observations that there is a local minimum in oxygen concentration there [80]. It is unclear why this local minimum disappears in 2004.

Waters denser than EDW, of 27.0 kg/m^3 and above, are subducted where the ageostrophic currents are relatively large on an annual average as well as on a winter average, accelerating the subduction process into the Gulf Stream. Waters less dense than 22.2 kg/m^3 subduct in areas where MLDs are relatively shallow and Ekman pumping is relatively strong; in particular, along the coast (Fig. 4.11). While subduction of the lightest waters in the slope region is consistent with the model of Speer and Tziperman (1992) [99], this is the region where the errors in the Argo data are largest. The presence of water below 22.2 kg/m^3 may be an exaggeration of how light this water is. Lastly and also consistent with the model of Speer and Tziperman (1992) [99], we found that Ekman pumping makes a negligible contribution to ventilation where EDW subducts (Fig. 4.12).

4.3.3 EDW formation characteristics

Applying the density-based criterion, $\Delta\rho_\theta > 0.125$, makes no qualitative difference in our volume estimates. The t-statistics for the differences between the mean EDW thicknesses using the density-based and temperature-based criteria over 2002 to 2006 where the EDW thicknesses are not zero do not show that there is a significant difference in the mean EDW thicknesses to 95% confidence. Because there is no EDW along the coasts and north of 45^oN, this is consistent with where the MLD estimates are significantly different. The density range between 25.2 and 27.8 kg m^3 accounts for most of the ventilated waters, including newly subducted EDW in the Sargasso Sea. Most of the ventilated EDW lies within the range of σ_θ between 26.0 and 26.8 kg/m^3, consistent with the EDW densities of McCartney and Talley (1982) [62]. In 2005, the distribution of densities of ventilated EDW is wider than in previous years with a longer tail toward lower densities. This may be explained by the subduction and obduction of a wider spatial distribution of water parcels formed over a range of years.

EDW is found to form just south of the Gulf Stream. This subtropical mode water

follows the North Atlantic recirculation gyre toward the south and subducts toward the edges of the recirculation gyre with a higher population of water between 17 and 19^{o}C during springtime (Fig. 4.12). We find that there is weaker stratification at the top of the thermocline where there is EDW, consistent with the definition of EDW. There is virtually no EDW north of the Sargasso Sea because older EDW mixes and is modified by heat fluxes after entering the Gulf Stream. The rest of the water between 17^o and 19^{o}C is of subpolar origin.

We find EDW subduction rates closer to that of Kwon and Riser (2004) [51] than Speer and Tziperman (1992) [99]. The Walin (1982)[109] formation equates the diapycnal volume flux with the difference between the air-sea buoyancy and diffusive fluxes over outcrop areas. Since the air-sea buoyancy fluxes are spatially and temporally variable, there may be substantial error in the air-sea buoyancy flux term, propagating error to the annual formation rate estimate of EDW. Additionally, neglecting mesoscale variability that contributes to diapycnal volume flux may propagate error to the annual formation rate estimate of EDW. According to our estimates, as well as those of Kwon and Riser (2004) [51], the Walin (1982) [109] formulation appears to equate wintertime formation with annual formation. More water between 17^o and 19^{o}C is formed than subducts (Table 4.4). The differences between the EDW formation rates of Kwon and Riser (2004) [51] and Speer and Tziperman (1992) [99] also appear to the explained by the differences between the mean annual volume of water between 17^o and 19^{o}C and the EDW volume. Our EDW subduction and obduction rates are also tabulated in Table 4.4. An increase in net EDW subduction is consistent with an increase in net subduction of all waters. However, it is apparent that a net positive subduction rate of EDW does not necessarily indicate an increase in EDW volume and similarly a net positive obduction rate of EDW does not dictate a decrease in EDW volume. The missing EDW from this picture is a) EDW that is modified in the main thermocline, b) newly formed EDW in the mixed layer that has not obducted, and c) EDW that is not modified after obducting.

The spatial distribution of the errors in EDW volume is the same as that of the MLDs because the same objective mapping procedure was applied to the EDW thicknesses themselves. By estimating the thickness of water between 17 and 19^{o}C via objectively mapped

Argo temperatures, we find similar volumes to the thickness objective mapping procedure until 2005 and 2006 when the thickness objective mapping procedure yields EDW volumes that are twice as large. The spatial distributions of EDW and seasonal variability of its volume between the order of 1×10^{14} m^3 and about 10×10^{14} m^3 are robust. Using a combination of objectively mapped temperatures from both Argo and WOA data yields EDW volumes with a depressed seasonal variation.

Subducted waters are about the same density on average as obducted waters. However, while obducted water may generally be colder than subducted water, obducted waters are generally warmer relative to their surroundings than subducted waters. Similarly, since EDW is cold relative to surrounding waters in the mixed layer, we reason that the mixed layer should be coldest in 2002, warm due to an increase in net subduction of EDW from 2002 to 2004, and cool due to a decrease in net subduction of EDW in 2005. Consistent with this, the winter (January through March average) mixed layer (potential) temperature means and standard deviations respectively for our domain are 16.3^{o}C and 4.76^{o}C (4023 data points) for 2002, 16.7^{o}C and 4.82^{o}C (4134 data points) for 2003, 16.8^{o}C and 5.02^{o}C (4010 data points) for 2004, 17.6^{o}C and 4.77^{o}C (3885 data points) for 2005, and 17.0^{o}C and 5.02^{o}C (3987 data points) for 2006.

4.4 Conclusions

The wind stresses are significantly dependent upon the relative velocities between the winds and the geostrophic currents. This affects the Ekman divergences by about 10%-20% as well as the Sverdrup transports. Even so, vertical pumping is increasingly negligible relative to lateral induction farther away from the coast. Including a small vertical shear can make about a 10% difference in ventilation rates, albeit within our error.

Because the ratio of Δh_m to τ is not constant, 'analytical' estimates of ventilation rates may be inaccurate. The use of climatologies in our Lagrangian trajectory method may underestimate ventilation by similar reasoning. Nevertheless, the length of our time series may also explain the underestimates of ventilation. By how much averaging MLDs matters remains uncertain but appears to be marginally larger than our error: at most 50 %. Winter mixed layer temperatures change in concert with net (EDW) ventilation rates, which

Table 4.1: The data we used in our calculations

Quantity	Winds	SSH	MLD, σ_θ, EDW	shear
Source	QuikSCAT	Aviso	Argo, WOCE/ACCE	WOA
Dates	Jan 02-Aug 06	Jan 02-Aug 06	Jan 02-Jun 07	pre-1994
Frequency	daily	weekly	every 10 days	monthly

suggests that mixed layer temperature variability may be closely linked to ventilation rate variability. However, the criterion by which we calculate the MLDs may affect obduction estimates because of a difference in h_m where salinity profiles have significantly different gradients than elsewhere.

Subducted waters tend to come from south of the Gulf Stream, while obducted waters come from along the Gulf Stream. The lightest waters subduct at the coast where vertical velocities can dominate MLD variability. Water heavier than EDW is aided by the ageostrophic currents in subduction more than other waters. Both lateral induction and vertical pumping are nonnegligible in the Gulf Stream. MLD variability matters most for ventilation elsewhere.

WOA data have been used in related studies (e.g., [45, 50]), but owing to the relatively unpronounced seasonal variability in WOA data the temporal resolution made possible by the Argo system allows us to investigate variability at higher frequencies. The importance of intra-annual variations in MLDs, for instance, must be emphasized in estimating EDW ventilation rates. These rates are greatest during the springtime when the mixed layer shoals. Some of the subducted EDW resides in the thermocline for a period longer than the length of our time series: at least 3.5 yr. This warrants future investigations because the EDW residence time may play a crucial role in the heat budget and a potential NAO-EDW coupled oscillation.

Table 4.2: Means and standard deviations (std) of the longitudinal and latitudinal distances water parcels travel in a given year through our model

Distances	2002	2003	2004	2005
mean of longitudinal distance [^{o}]	6.24	6.34	6.40	6.19
std of longitudinal distance [^{o}]	6.87	7.07	7.55	7.48
mean of latitudinal distance [^{o}]	1.11	0.989	0.963	1.07
std of latitudinal distance [^{o}]	1.48	1.24	1.37	1.32

Table 4.3: Spatial mean and standard deviation (std) of subduction and obduction, and root mean square (rms) error: 'analytical' ($(\Delta h)/\tau$ from Qiu and Huang (1995) [77]) versus modeled ['-' is insufficient data]

estimation	2002	2003	2004	2005	2006
Analytical rms error [m/yr]	74.7	57.6	41.2	48.3	44.1
Analytical subduction mean [m/yr]	45.1	58.7	67.3	46.5	46.3
Analytical subduction std [m/yr]	50.5	68.5	69.8	40.2	51.6
Modeled subduction mean [m/yr]	51.8	57.5	52.8	54.9	-
Modeled subduction std [m/yr]	59.3	57.5	49.9	56.3	-
Modeled obduction mean [m/yr]	64.6	63.9	49.6	47.1	-
Modeled obduction std [m/yr]	79.9	73.2	57.9	53.9	-

Table 4.4: Ventilation and EDW rates and EDW volume from 2002 to 2006 ['-' is insufficient data]

mean annual quantity	2002	2003	2004	2005	2006
$17 - 19^oC$ water volume [10^{14} m^3]	4.44	8.78	10.3	13.5	12.8
$17 - 19^oC$ water subduction [Sv]	3.72	8.17	6.10	6.20	-
$17 - 19^oC$ water obduction [Sv]	1.30	1.70	2.03	2.30	-
EDW volume [10^{14} m^3]	1.58	1.63	1.69	1.35	1.89
EDW subduction [Sv]	0.80	1.86	1.74	3.65	-
EDW obduction [Sv]	0.43	0.72	1.00	1.12	-
Subduction [Sv]	17.4	19.0	18.1	18.2	-
Obduction [Sv]	20.4	20.5	15.5	13.9	-

Figure 4.1: Fig. 1 from Trossman et al. (2009) [106] (c)American Meteorological Society. Reprinted with permission. Average of the geostrophic and ageostrophic (as estimated from the first-order quasi-geostrophic equations described in the Appendix, Section A.1) currents from Aviso in the Gulf Stream region over 1999 to 2006.

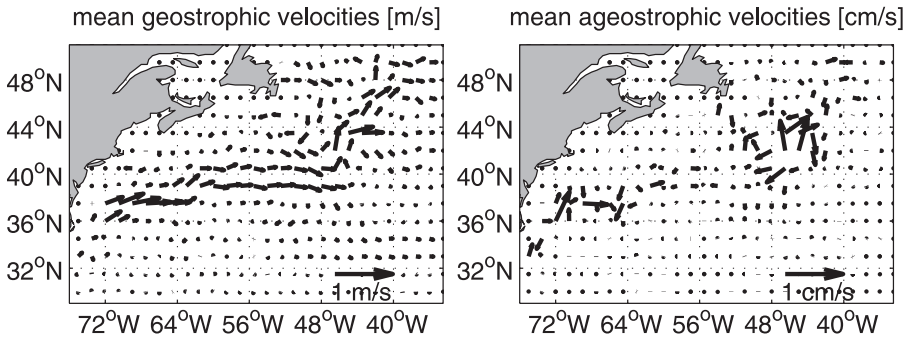

Figure 4.2: Fig. 2 from Trossman et al. (2009) [106] (c)American Meteorological Society. Reprinted with permission. Locations of non-852 Argo floats that reported data in the month of March from 2002 to 2005.

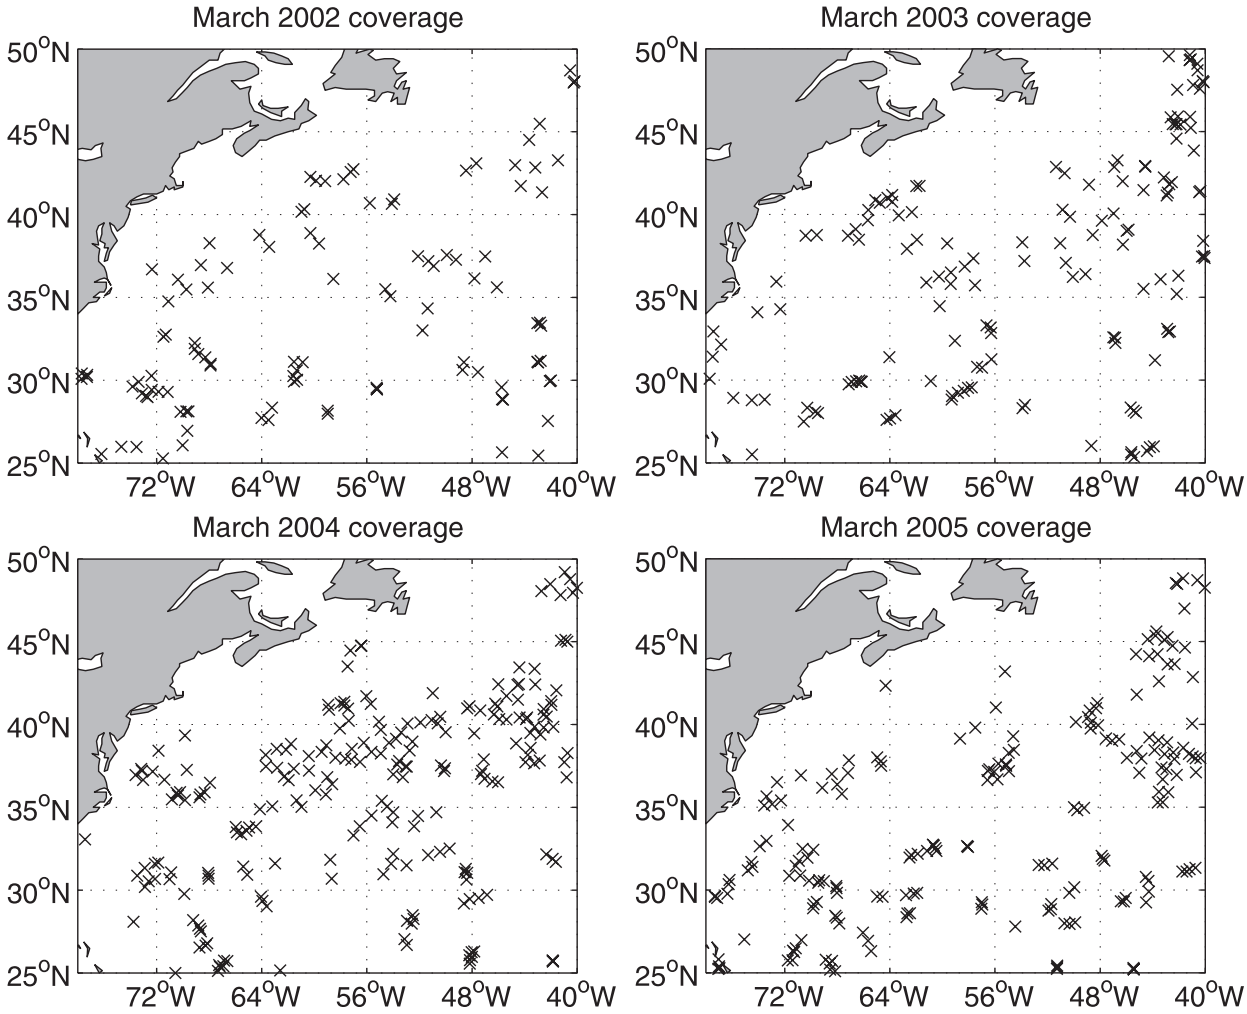

Figure 4.3: Fig. 3 from Trossman et al. (2009) [106] (c)American Meteorological Society. Reprinted with permission. The MLDs in March (top) and the month numbers (e.g. March = 3) when the MLDs reach a maximum in late winter/early spring of 2004 (left) and 2007 (right). These are the years with the most and least spatial coverage by non-WMO 852 Argo floats from when WMO 852 Argo floats errors were introduced.

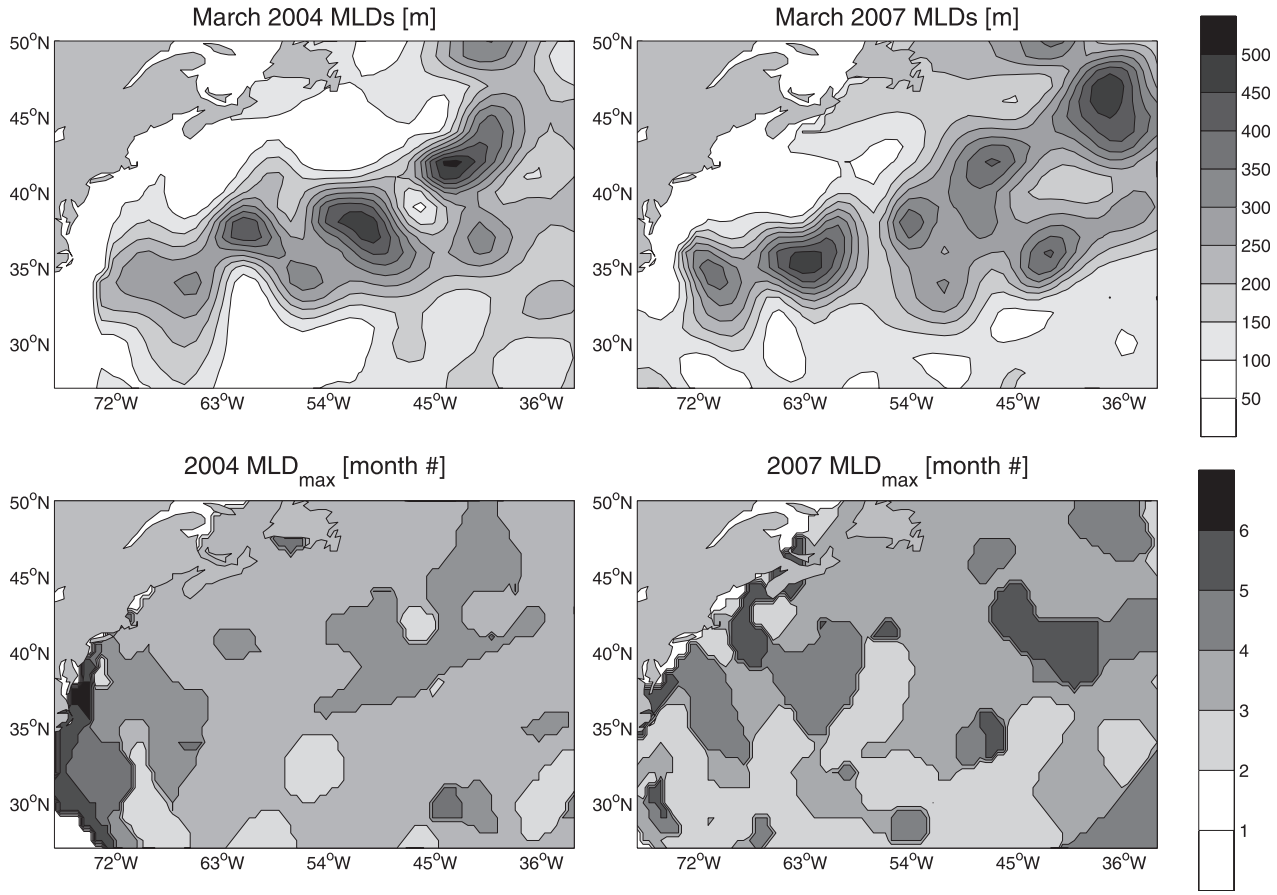

Figure 4.4: Fig. 4 from Trossman et al. (2009) [106] (c)American Meteorological Society. Reprinted with permission. Top-left: Mean of the MLDs over 2002 to 2007 in meters. Top-right: Standard deviation of the MLDs over 2002 to 2007 in meters. Bottom-left: Mean of the associated objective mapping errors over 2002 to 2007 in meters. Bottom-right: Standard deviation of the associated objective mapping errors over 2002 to 2007 in meters.

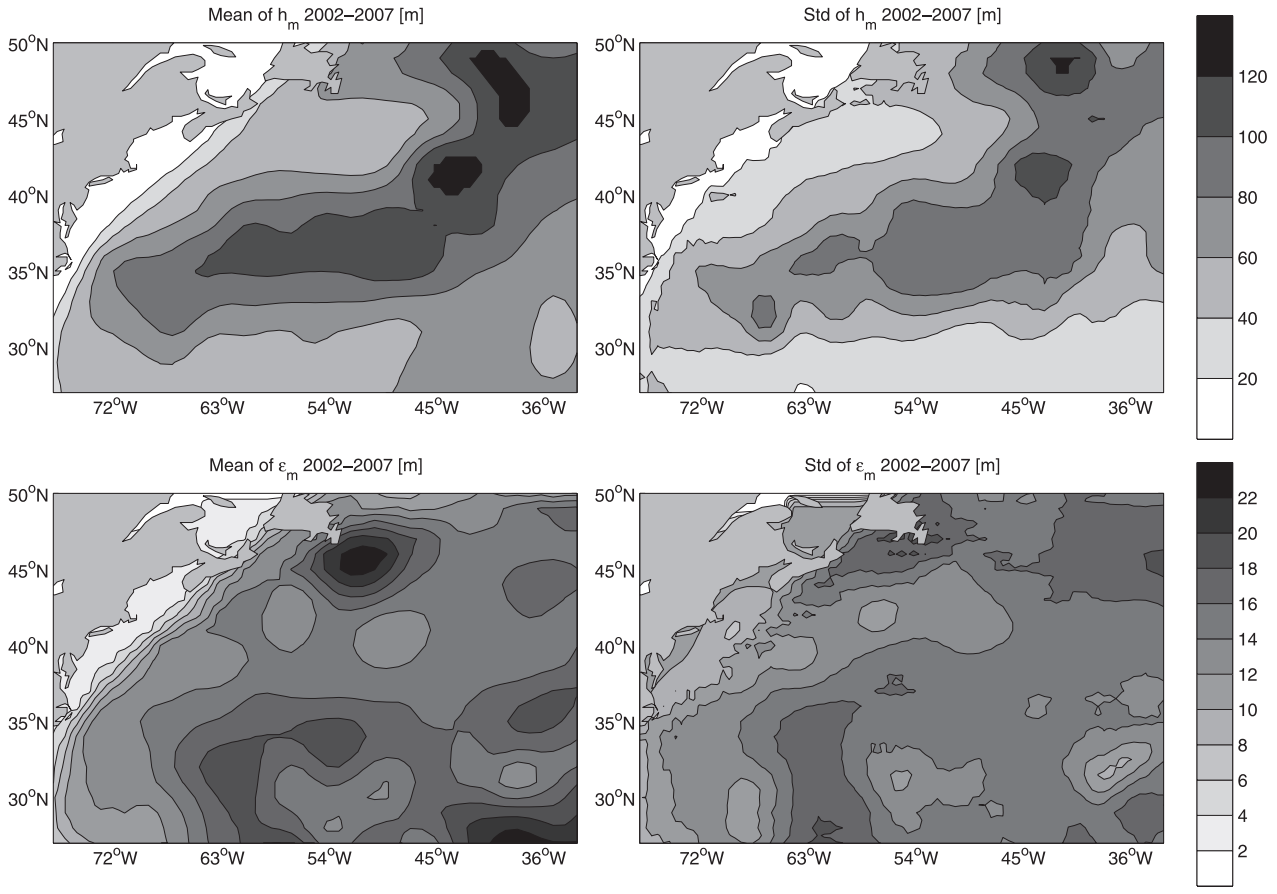

Figure 4.5: Fig. 5 from Trossman et al. (2009) [106] (c)American Meteorological Society. Reprinted with permission. The atmosphere, surface Ekman layer, mixed layer, seasonal thermocline, and permanent thermocline are shown to represent an idealized cross-section across a strong density gradient. The gray parallelograms are approximate representations of the water columns that either subduct or obduct. Here, h_m is the boundary between the seasonal and permanent thermocline. In the top case, the water column subducts and stays in the permanent thermocline. In the bottom case, the water column obducts. The dashed arrow indicates time evolution, the interval it takes for the MLD to be deepest from one year to the next. For simplicity, h_m is the same in each year.

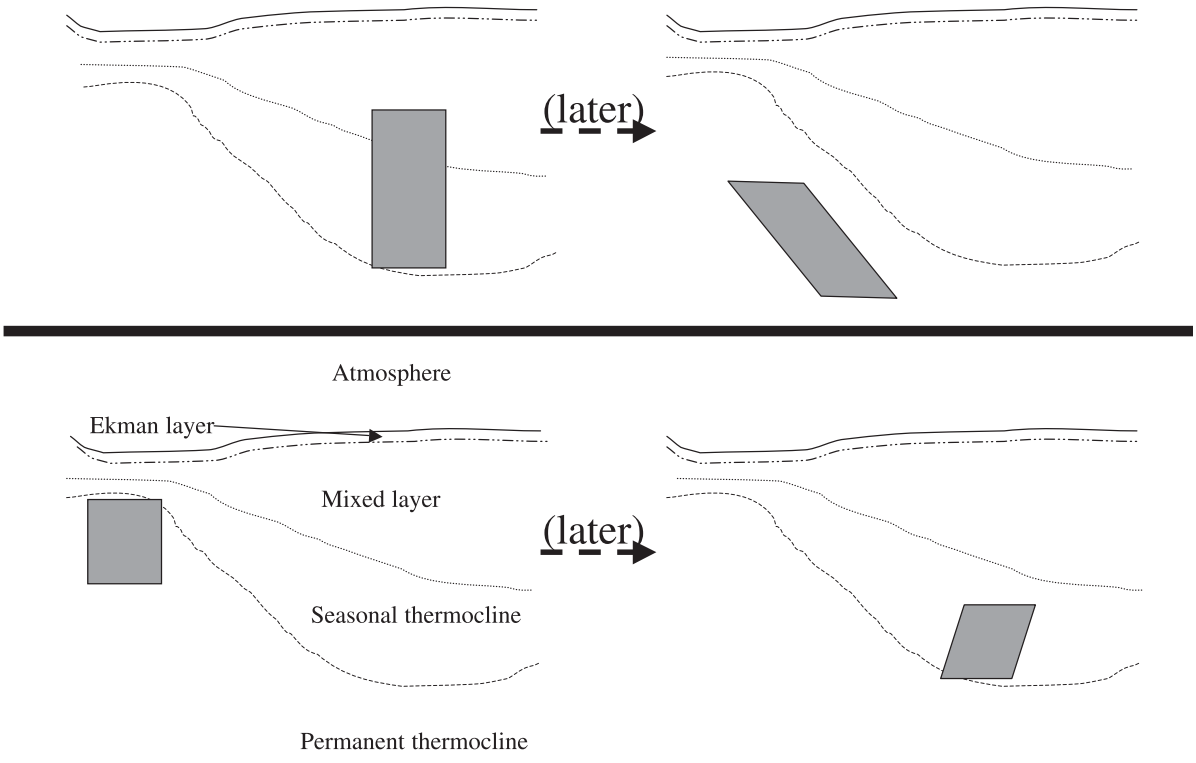

Figure 4.6: Fig. 6 from Trossman et al. (2009) [106] (c)American Meteorological Society. Reprinted with permission. The average of Ekman pumping (m yr^{-1}) in the Gulf Stream and Sargasso Sea regions, including (top) currents (solved numerically) over 1999 to 2006 with mean SSH contours in white and (bottom) the difference between the average of w_e (m yr^{-1}) including currents and that without currents in the same region over the same time period.

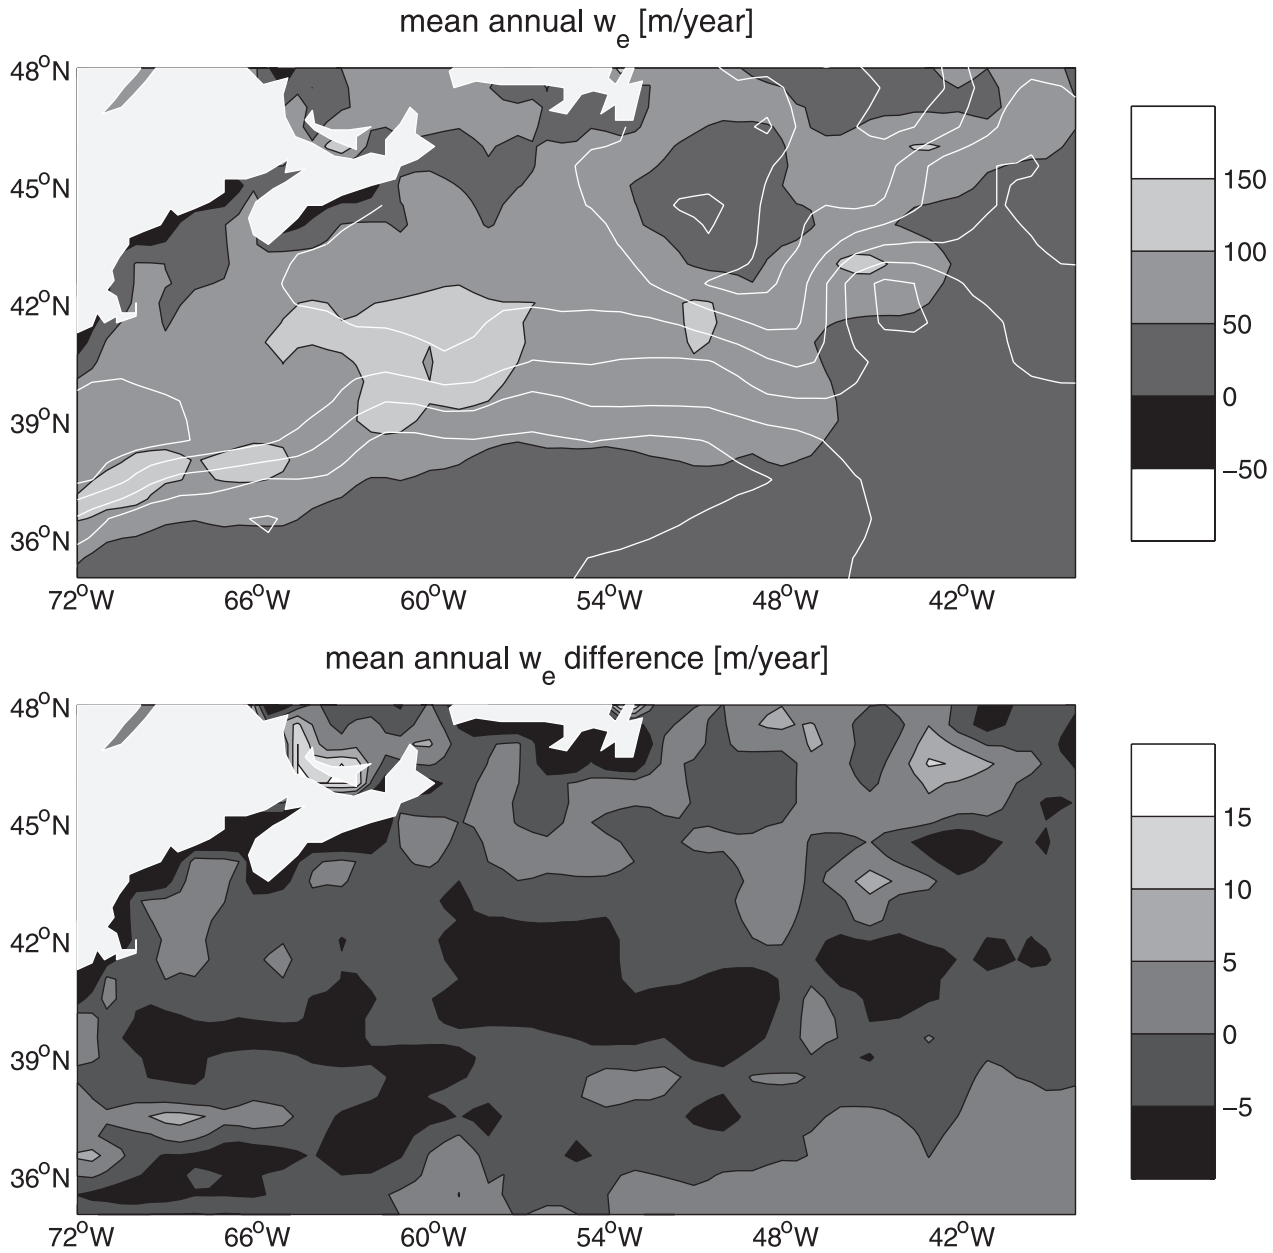

Figure 4.7: Fig. 7 from Trossman et al. (2009) [106] (c)American Meteorological Society. Reprinted with permission. The mean annual subduction [in m yr^{-1}] (left) and obduction [in m yr^{-1}] (right) rates for 2002-2003 (top), 2003-2004 (second to top), 2004-2005 (second to bottom), and 2005-2006 (bottom) using monthly-varying MLDs and weekly-varying currents with shear. The corresponding volume transport [in Sv] is in the lower right-hand corner box of each plot.

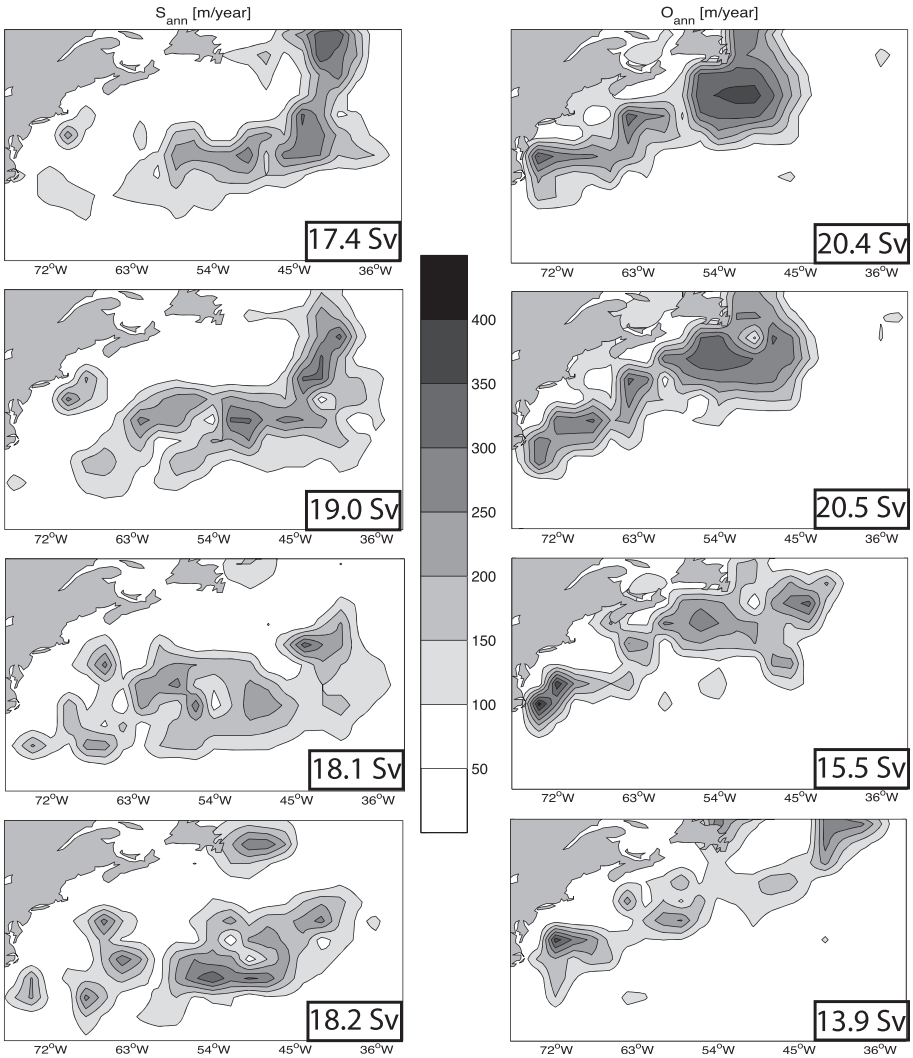

Figure 4.8: Fig. 8 from Trossman et al. (2009) [106] (c)American Meteorological Society. Reprinted with permission. The mean annual subduction [in m yr^{-1}] (left) and obduction [in m yr^{-1}] (right) rates using 1999-2006 average currents and a climatology of 2002-2005 average monthly-varying MLDs. The top corresponds to using the $\Delta T > 0.3$ criterion with vertical current shear. The second to top corresponds to using the $\Delta T > 0.3$ criterion with no vertical current shear. The second to bottom corresponds to using the $\Delta \rho_\theta > 0.125$ criterion with vertical current shear. The bottom corresponds to using the $\Delta T > 0.3$ criterion and vertical shear, but no Ekman currents ($w_e = 0$). The corresponding volume transport in and out of the permanent thermocline [in Sv] is in the lower right-hand corner box of each plot.

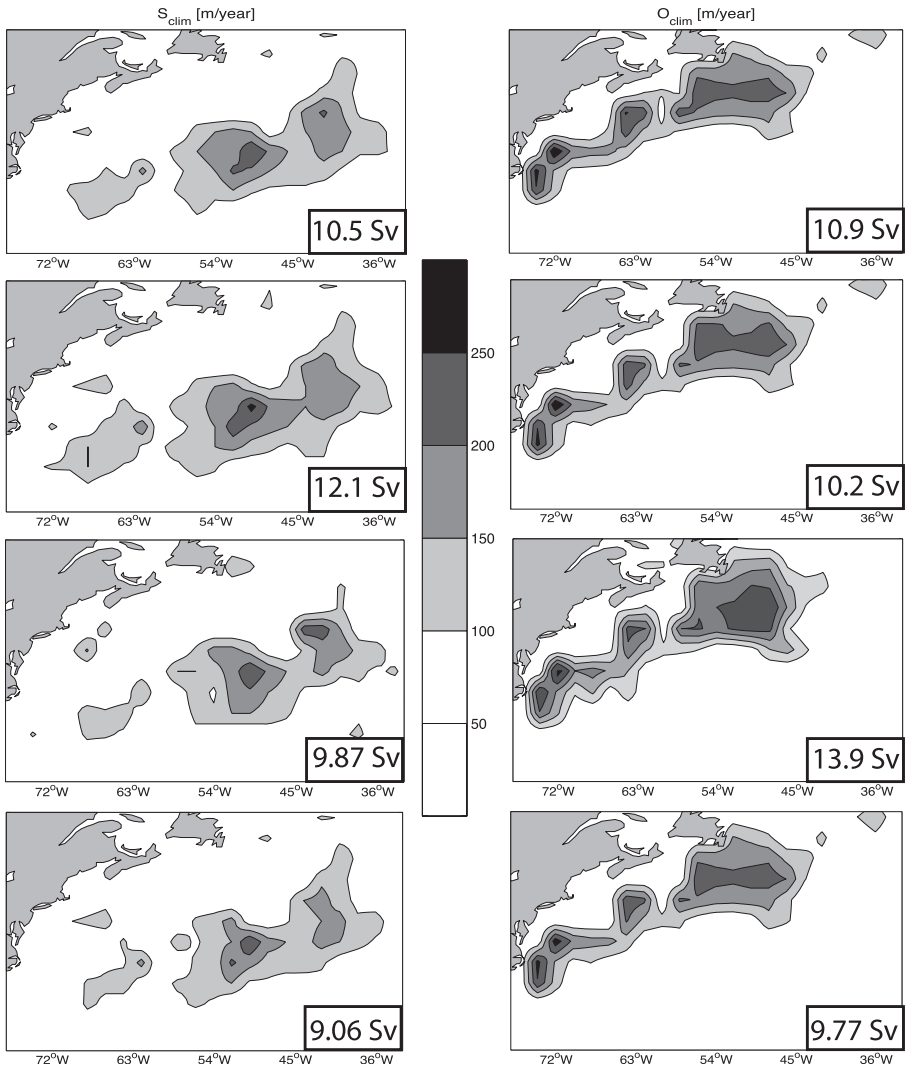

Figure 4.9: Fig. 9 from Trossman et al. (2009) [106] (c)American Meteorological Society. Reprinted with permission. The mean (left) subduction (m yr^{-1}) and (right) obduction (m yr^{-1}) rates over 2002-2006 using (top) mean velocities and mean MLDs, (middle) mean velocities and estimated monthly MLDs, and (bottom) estimated weekly velocities and estimated monthly MLDs, all with shear. The corresponding volume transport in and out of the permanent thermocline (Sv) is in the lower right-hand corner box of each plot.

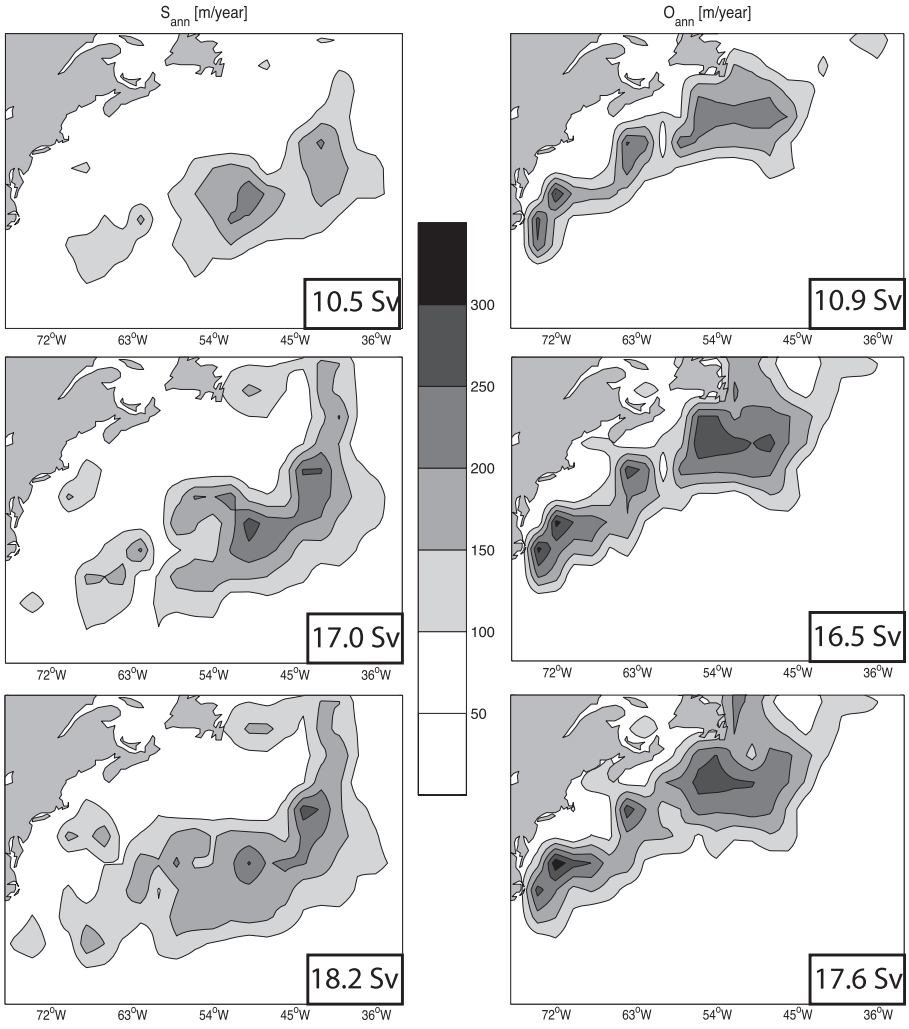

Figure 4.10: Fig. 10 from Trossman et al. (2009) [106] (c)American Meteorological Society. Reprinted with permission. Volume transports [in Sv] categorized into subduction (gray positive bars), EDW subduction (black positive bars), obduction (gray negative bars), and EDW obduction (black negative bars), broken into intervals of 0.2 kg/m^3 for (a) 2002, (b) 2003, (c) 2004, and (d) 2005.

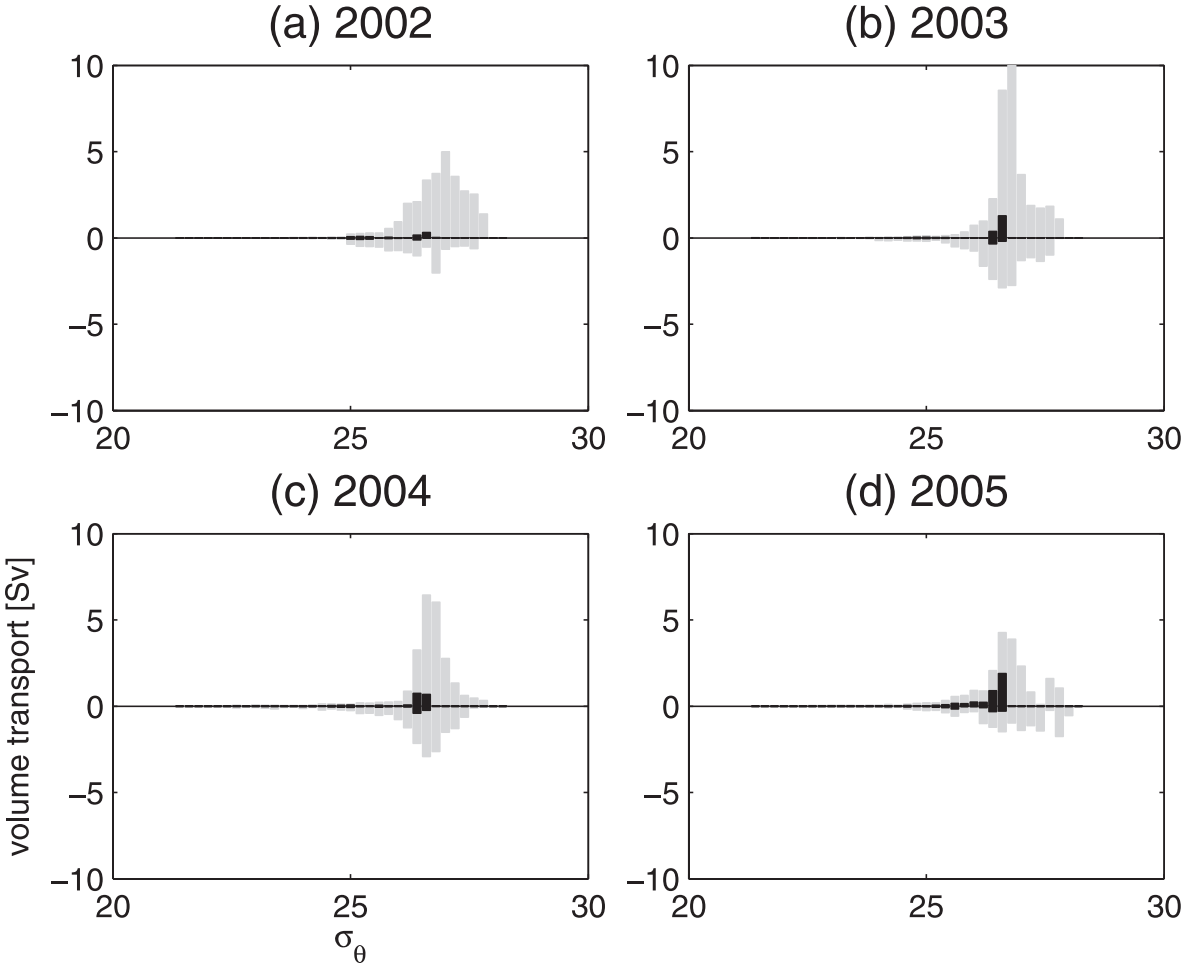

Figure 4.11: Fig. 11 from Trossman et al. (2009) [106] (c)American Meteorological Society. Reprinted with permission. Locations of where anomalously large subduction rates of heavy water, 27.0 kg/m^3 and above, [x] and light water, 22.0 kg/m^3 and below, [o] occur for (a) 2002, (b) 2003, (c) 2004, and (d) 2005.

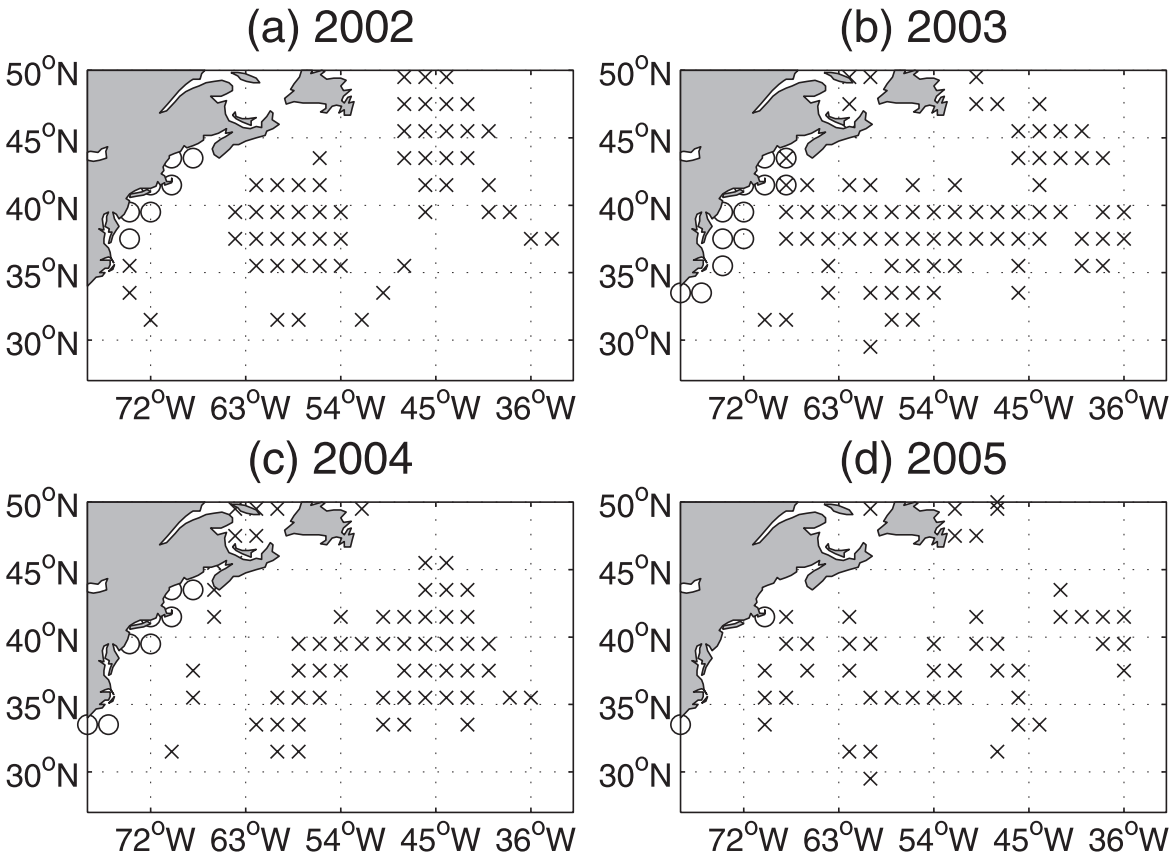

Figure 4.12: Fig. 12 from Trossman et al. (2009) [106] (c)American Meteorological Society. Reprinted with permission. Thickness of waters between 17 and 19^{o}C [in meters] in springtime (April-June) with locations where greater than average subduction rates of EDW occur in our domain [x] in (a) 2002, (b) 2003, (c) 2004, and (d) 2005.

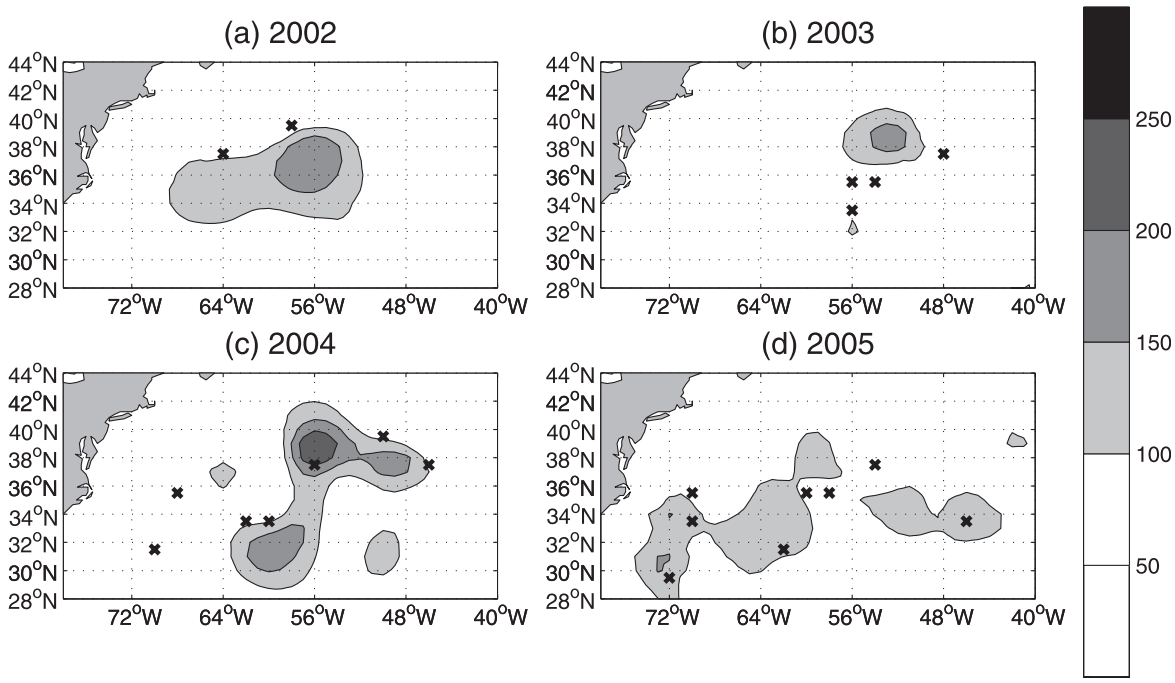

Chapter 5

THE ROLE OF ALONG-ISOPYCNAL DIFFUSION IN VENTILATION

5.1 Introduction

Mode waters are the manifestation of air-sea interaction and subduction and serve to ventilate the thermocline [88]. We focus on the ventilation of the thermoclines of two oceans here: the North Atlantic and Southern Oceans. Specifically, we look at the subtropical mode water (STMW) and subpolar mode water (SPMW) of the North Atlantic, and the subantarctic mode water (SAMW) of the Southern Ocean. Mode waters have minima in large-scale potential vorticity [62] in the vertical, making them the most homogeneous water masses within a given gyre above the thermocline. STMW has observed potential density anomalies between 26.0 and 26.6 kg m^{-3} [55], SPMW has observed potential density anomalies mostly between 27.2 and 27.5 kg m^{-3} [62, 6], and SAMW has observed potential density anomalies between 26.5 and 27.05 kg m^{-3} [88] (26.6 and 26.8 kg m^{-3} in the Indian Sector [93]). Below the thermocline, the most homogeneous water mass in the North Atlantic is Labrador Sea Water (LSW), which also has a minimum in large-scale potential vorticity in the vertical. LSW has observed potential density anomalies between 27.4 and 27.7 kg m^{-3} [44] and we will distinguish from SPMW by with salinity minimum [79].

The North Atlantic's subtropical gyre has been extensively studied and provided the motivation for the ventilation theory developed by Luyten et al. (1983) [57]. Primarily following isopycnals but also influenced by lateral mixing under a field of Ekman currents, water from the surface layer of the ocean is said to be ventilated when this water is advected and diffused into the permanent thermocline. Trossman et al. (2009) [106] have used satellites and profiler observations to estimate ventilation with an advective parcel tracking model and were able to estimate a rate of STMW ventilation to be $2 - 8$ Sv, consistent with previous estimates [50]. SPMW transport has been estimated to be $7.5 - 20.5$ Sv [6], most of which is not attributed to advection from the mixed layer because SPMW is formed

primarily through mixing processes below the mixed layer, incorporating water upwelled from the interior.

Ventilation of the Southern Ocean has been less extensively studied. We will be focusing on a subset of the Southern Ocean. From 20^o in the South Indian Ocean to 180^o E in the South Pacific Ocean, there are several distinguishable water masses, one of which is SAMW, but there are potentially many different origins of SAMW that give rise to different types of SAMW in the South Indian Ocean that are denser the farther east they form. SAMW is formed by deep mixing on the northern side of the ACC, where water (north of 50^oS) is downwelled due to convergence of wind stress [88]. Studies such as Sloyan and Rintoul (2001) [93] have estimated the ventilation rate of SAMW to be about $6.5 - 8.5$ Sv in the South Indian Ocean (23.6 Sv over the entire Southern Ocean) based on an inverse model and observations.

The main goals in this paper are to provide a precise constraint on the range of mode water ventilation (or more specifically, subduction) rates in the North Atlantic and Southern Oceans and to quantify the relative contributions of advection and diffusion to the ventilation rates. The World Ocean Circulation Experiment (WOCE) and Climate Variability and Predictability (CLIVAR) & Carbon observations (Table 5.1) are utilized along with outputs from several isopycnal coordinate model runs. These model runs are ranked according to which simulates output that is closer to observations. Instead of selecting one 'best' model or averaging the model outputs together with equal weights, we can provide more realistic estimates by averaging the model outputs together with weights that reflect their relative rankings. We calculate the subduction rates of SAMW (Indian Ocean), STMW (North Atlantic), and SPMW (North Atlantic) using hydrographic sections that run across the Southern, Subtropical North Atlantic, and Subpolar North Atlantic Oceans, respectively, with a statistical technique called Bayesian model averaging (BMA). BMA incorporates observations by treating model runs as random variables themselves, each with an associated posterior probability (roughly interpreted to be the degree of our belief that a model simulates the observations better than other models) that we use to weight the outputs to come up with a final estimate.

In order to evaluate the contribution of diffusion in each basin according to our technique,

utilizing a numerical model and BMA, we estimate a transit-time distribution, or TTD, [38] at each point according to a method evaluated by Haine et al. (2008) [30], called the boundary impulse response, or BIR, method. A TTD is a Green's function solution to the three-dimensional advection-diffusion equation that quantifies the fraction of mass that gets from a boundary (e.g., the surface of the ocean) to some other location (e.g., the interior of the ocean) from the time a tracer enters the ocean at the boundary to some other time. A purely advective water parcel would have a delta spike for a TTD, so the half-variance of the TTD essentially is a measure of diffusivity. We compare the TTD-based ventilation estimates, which take advection and diffusion into account, with diapycnal velocity-based ventilation estimates, which do not take diffusion once the water has been subducted into account.

The main topics of this paper are arranged as follows: We review the theory of how a TTD is related to mass flux when the ocean is not in steady-state as well as the method of BMA. By varying the time at which we spike the surface with a tracer in the model TTD estimation, we can estimate a time-varying TTD which can in turn be related to a time series of ventilation rates of a particular density class of water. We describe a set of model runs that we then average using varying weights according to BMA. When CFC-11 concentrations in the model output are closer to the objectively mapped CFC-11 concentrations from bottle data, the model's output from that particular simulation is more heavily weighted than others in the ensemble of simulations. All of the model simulations are weighted to find final estimates of the mode water ventilation rates and uncertainty on these estimates are quantified. We thus use TTDs and BMA together to end up with ventilation rate estimates of the following mode waters in the North Atlantic and Southern Oceans: SAMW, STMW, and SPMW.

5.2 Methods

5.2.1 Data

In order to evaluate a TTD according to our numerical model, we use CFC-11 concentrations along repeat hydrographic sections that include CFC-11 observations (see Table 5.1). These

sections are as follows: The SR03 section (from 67^{o}S to 44^{o}S, between 135^{o}E and 147^{o}E) is a meridional transect that goes from Tasmania to Adélie Land of Antarctica. The I09 section (from approximately 32^{o}S to 63^{o}S between 81^{o}E and 97^{o}E) is a meridional transect that crosses the Indian and Southern Oceans west of Australia. The A16 section (from 63^{o}N to approximately 34^{o}N along 20^{o}W and southward to approximately 55^{o}S between 20^o and 36^{o}W) is a meridional transect that goes southward in the eastern North Atlantic basin and crosses the equator going southward in the western North Atlantic basin. The A05 section is a zonal transect that follows 24^{o}N from about 20^{o}W to about 79^{o}W in the North Atlantic. The A20 section (from 21^{o}N to 41^{o}N along 53.5^{o}W) is a meridional transect from Brazil to Newfoundland. The years in which the CFC-11 data was taken for each section are listed in Table 5.1 and a map of these sections is shown in Fig. 5.1.

Model Runs

An offline isopycnal coordinate model is ideal for this study because we can vary the background diffusivity without changing the circulation and we can vary the circulation without changing the background diffusivity. The numerical model is an offline advection diffusion code in isopycnal coordinates called Offtrac [15]. Offtrac uses mass flux fields based on output from the Hallberg Isopycnal Model (HIM) [32]. HIM is a 42 layer, isopycnal coordinate, nominally 1^o longitudinal model on a tripolar grid that has been forced by the repeat annual cycle Coordinate Ocean Reference Experiment (CORE) dataset [54] which contains monthly fields that were averaged from 6-hourly fields.

Climatological fields were constructed from the last 120 months of a 250 year spin-up of HIM, using the CORE fields as input. When we use the climatological fields, we use the average field of each month of the last 120 months of the spin-up, Jan.-Dec., as input to Offtrac and repeat these twelve month-long fields every year. Hindcast fields were constructed based on using CORE-based precipitation, radiation, relative humidity, sea level pressure, temperature at 10 meters above sea level, and wind velocity fields at 10 meters above sea level January 1958 to December 2004 [54] as input for HIM. When we use the hindcast fields, the twelve month-long climatological fields are first repeated each year for

28 years since we start the model at the calendar year, 1931. In the calendar year, 1959, we start using the HIM-derived hindcast fields of variable wind stress, density layer thickness, temperature, salinity, zonal and meridional transport, and diapycnal velocity fields. The hindcast fields are repeated every 47 years thereafter when estimating the TTDs because the last calendar year of hindcast fields is 2005 and the tail of the TTDs at most depths below several hundred meters is beyond $28 + 47$ years. Gas exchange parameterization follows the the Ocean Carbon-cycle Model Intercomparison Project (OCMIP-2) protocol [20].

We construct an ensemble of Offtrac model runs that consists of 75 year monthly CFC-11 concentration output to get the weights for BMA. Here, the Offtrac model is forced with climatological fields and three different values of the along-isopycnal diffusivity: 5×10^2, 1×10^3, and 1×10^4 $\mathrm{m^2 s^{-1}}$. We also perform runs with the Offtrac model forced with hindcast fields and the same three different values of the along-isopycnal diffusivity. The range of diffusivities was chosen based on the work of Shuckburgh et al. (2011) [91], Abernathey et al. (2010) [1], and Mecking et al. (2004) [64].

Identifying Water Masses

Since our task is to find the ventilation rate of mode waters, we need a definition of SAMW, STMW, and SPMW according to the HIM hindcast and climatological fields. The properties of these water masses are summarized in Table 5.2. The way we identify SAMW, STMW, and SPMW is to use the large-scale potential vorticity [52],

$$PV_k = \frac{f}{\rho_k} \frac{\frac{1}{2}(\rho_{k+1} + \rho_k) - \frac{1}{2}(\rho_k + \rho_{k-1})}{h_k} \tag{5.1}$$

where ρ_k is the potential density at the kth isopycnal layer, f is the Coriolis parameter set to be $2 \times 7.2921 \times 10^{-5} sin(\lambda 180/\pi)$ for latitude λ, and h_k is the thickness of the kth isopycnal layer in the model. (To calculate the derivatives, we ignore the isopycnal layers that are less than a centimeter thick.) We find the depths at which it is the case that the magnitude of large-scale potential vorticity, $|PV|$, is near its minimum in the vertical. That is, we look for depths at which either $\partial|PV|/\partial z < 0$ and $\partial^2|PV|/\partial z^2 > 0$ or $\partial|PV|/\partial z > 0$ and $\partial^2|PV|/\partial z^2 > 0$ at each horizontal location in the domains of the Southern and North Atlantic Oceans. We identify SAMW, STMW, and SPMW as the waters where this

condition is satisfied with $PV < 1 \times 10^{-9}$ $\mathrm{m}^{-1}\mathrm{s}^{-1}$, and the potential densities are within a range such that SAMW can be found in the southeast Indian Ocean (northern Southern Ocean), STMW can be found southeast of the Gulf Stream, and SPMW can be found over the southwest and southeast coasts of Greenland as well within as most of the rest of the Subpolar North Atlantic Ocean to the south of Greenland. We limit identification of the mode waters to the longitudes and latitudes listed in Table 5.2 with the additional criterion for SPMW that the salinity, S, is not near its minimum in order to distinguish SPMW from LSW. That is, we add the additional criterion for SPMW that the depths we find SPMW also satisfy either either $\partial S/\partial z > 0$ and $\partial^2 S/\partial z^2 < 0$ or $\partial S/\partial z < 0$ and $\partial^2 S/\partial z^2 < 0$.

5.2.2 Two Methods to Estimate Ventilation Rates

Using only the locations where there is SAMW, STMW, or SPMW, we calculate the subduction rates of these three water masses in two different ways: one using a TTD approach and one using the diapycnal velocities at the base of the mixed layer directly. As described in Primeau and Holzer (2006) [76] and Hall et al. (2007) [31], the total water mass in volume, V, is ρV where ρ is the density of the fluid in V, so the flux of water mass through a surface, Ω, that has resided at least $\tau = t - t_0$ in V is the flux of water mass exiting V that has resided at least τ in V and equals

$$\Phi_+(t,\tau) = \rho V G(t,t_0) + \rho V \int_0^\tau dt' \frac{\partial}{\partial t} G(t, t-t'). \tag{5.2}$$

Here, the TTD averaged over the grid points at which there is a water mass with density ρ is $G(t,\cdot)$. The residence time of the water mass of density ρ exiting V that has resided in V for τ amount of time before eroding can then be calculated. The total mass of water that is eroded after residing in V between τ and $\tau + \Delta\tau$ amount of time is

$$(\frac{M_e}{M_{wm}})(t,\tau) = -\tau d\Phi_+(t,\tau)/d\tau, \tag{5.3}$$

and the rate of input of water that resides in V for τ amount of time is $d\Phi_+(t,\tau)/d\tau$ [76]. The residence time of a water mass is the ratio of the total mass to rate of input [76]. The volume transport spectrum, $\Phi_+(t,\tau)/\rho$ via (5.2), and the mass fraction erosion spectrum, M_e/M_{wm} via (5.3), are both functions of the residence time, τ.

In order to the estimate subduction rate as $\int d\tau \Phi_+(t,\tau)/\rho$, we need an ensemble of model simulations for the TTD. Each of the six members in our ensemble of (three climatological and three hindcast) simulations starts the age clock at given time, t_0, so in order to integrate the derivative and derive a time series of the TTD that resolves decadal variability, we start the TTD at $t_0 = 0$, 5,..., 75 years after the model begins. The TTD is the same for each t_0 when we use the climatological circulation fields, but the TTD is different for each t_0 when we use the hindcast circulation fields. This makes 3×16 hindcast and 3 climatological, or 51, model runs. For each of the six members in our ensemble of model simulations with a given t_0, we compute the TTD from t_0 to $t_0 + 200$ years. The TTD is estimated with the BIR approximation [30] by letting the surface of the ocean be $1/N$ for an entire year where N is the number of time steps in a year and keeping track of the values of this tracer. We do not need to extrapolate the TTD because 200 years of model run time is sufficient to capture the tails of older ages for the water masses we consider here.

Our estimate of the subduction rate from (5.2) will be compared to a diapycnal velocity-based subduction rate estimate [52], which is calculated as

$$V_{ann} = \int \int_{outcrop} dA(w_{de} - w_{en}). \tag{5.4}$$

Here, w_{de} is the detrainment velocity (negative diapycnal velocity just below the base of the mixed layer) and w_{en} is the entrainment velocity (positive diapycnal velocity just below the mixed layer), each averaged over the time when the winter mixed layer is deepest in one year to the time when the winter mixed layer is deepest in the following year. The integral in (5.4) is over the area of the ocean surface (outcrop) in which the density of the surface waters lies within the range of densities for a given mode water (Table 5.2). By applying BMA, as described in the following subsection, we can determine the uncertainties associated with our subduction rate estimates from each of the two different methods.

5.2.3 How to Choose Which Model Output is Best: Bayesian Model Averaging

Since no single model simulation is perfect, we would like to come to a better estimate of the quantity of interest by weighting model outputs based on comparisons with observations. To do this, we need to decide, *a priori*, which model parameters we are uncertain about

(which here are κ_H and circulation fields) but affect the quantity of interest (which here is a model output, the subduction rates of mode waters using two different methods). We need to vary these model parameters to span a realistic range of values they can take and choose to use CFC-11 concentrations as a surrogate performance measure of each model simulation. Using BMA, we can estimate a quantity of interest (which here is a subduction rate) with uncertainties that indicate how much we trust these estimates. BMA basically treats our model simulations as random variables and places a probability density on these random variables with a peak located at the model simulation that reproduces the observations most closely.

Two comments on estimating a quantity with a physical model are worth emphasizing: 1) we need a good physical model of the quantity of interest and 2) we need to incorporate the uncertainty associated with all of our most imperfectly known model parameters. The uncertainty of an estimated quantity attributed to the model parameters that are most imperfectly known needs to overwhelm the uncertainty attributed to the model parameters that are least imperfectly known. As the top panel of Fig. 5.2 suggests, running a model that is not very accurate in the first place does not necessarily have point-wise uncertainty estimates that contain the 'truth' (i.e., what a perfectly accurate measurement would be). Furthermore, the bottom panel of Fig. 5.2 suggests that failing to take all sources of uncertainty into account using a model that does well on average over periodic intervals of time will lead to point-wise uncertainty estimates that do not contain the truth either. Because the error associated with using the climatological or hindcast fields in Offtrac and the error associated with using a single along-isopycnal diffusivity, κ_H, should dominate the uncertainty in our model parameters, we vary these parameters and weight each model output according to the probability that this output is the truth given observations of CFC-11 along several transects in the years in which there were measurements. CFC-11 concentrations from Offtrac and objectively mapped observations along each of the WOCE transects listed in Table 5.1 are shown and described in Section A.4 of the Appendix.

To be more mathematically precise, we have different input parameter configurations for HIM/Offtrac runs, $M_1, ..., M_k$, where here, $k = 6$, and we want to weight each of these configurations by the relative degree to which we believe in them. Each model has

a probability density for some random variable, X, (which can be any output variable or parameter of the model, but here are the TTD- and diapycnal velocity-based estimates of subduction rates) and unknown parameters (which here are κ_H and the climatological or hindcast HIM circulation fields). We can find the posterior probability of the jth model configuration being the true model for the data amongst all possible models, $M_1, ..., M_k$, which we denote by m_j. Here, we let the CFC-11 concentration according to the jth model simulation at a location and time be $D_{model,j} = D^*_{obs} + \epsilon_{model,j}$, where $D^*_{obs} = D_{obs} + \epsilon_{obs}$ is the observed CFC-11 concentration, D_{obs} is the actual CFC-11 concentration, ϵ_{obs} is the error is the observed CFC-11 concentration from objectively mapped bottle data at a location and time (which we model with a Gaussian distribution), and $\epsilon_{model,j}$ is the error of the jth model simulation relative to the observation at a location and time (which we model with a different Gaussian distribution from ϵ_{obs}). Then we calculate the posterior probability with

$$m_j = P(M_j|D_{obs}) = \frac{P(D_{obs}|M_j)P(M_j)}{\sum_{i=1}^{k} P(D_{obs}|M_i)P(M_i)} \propto P(D_{obs}|M_j)P(M_j), \tag{5.5}$$

where $\sum_{j=1}^{k} m_j = 1$ and $P(M_j)$ is a prior belief, or first-guess distribution without incorporating information from the observations on the jth model, by Bayes' Theorem [3]. Bayes' Theorem expresses the relationship between marginal and conditional probabilities of random variables. Leaving out a conditional probability such as $P(D_{obs}|M_3)$, for instance, would bias inference about m_j for all $j = 1, ...k$ because each probability is dependent upon the others by the fact that probabilities sum to one.

To estimate each $P(D_{obs}|M_j)$, we need to incorporate the uncertainty associated with the CFC-11 observations. Here, D^*_{obs} are sampled from a normal distribution, $N(\mu_{obs}, \epsilon^2_{obs})$, within the Markov Chain Monte Carlo, or MCMC, algorithm (e.g., explained in an oceanographic context in Chapter 3 of the present thesis in which a TTD and its first two cumulative moments were estimated with uncertainties via MCMC). The mean of the normal distribution, μ_{obs}, equals the objectively mapped value for CFC-11 concentration and the variance, ϵ^2_{obs}, equals $\sqrt{(0.02 \times D^*_{obs})^2 + (\sigma_{objm})^2}$, where the measurement error is $0.02 \times D^*_{obs}$ and error due to mapping the bottle data, estimated with the thin-plate spline method described in Chapter 2 of the present thesis, is σ_{objm}. By incorporating the uncertainty

associated with the CFC-11 observations, we find the likelihood given by

$$P(D_{obs}|D^*_{obs}, M_j) = \prod_{x,t,s} \frac{1}{\sqrt{2\pi(\sigma_{CFC-11,j})^2}} exp[-\frac{(D^{*(s(x,t))}_{obs} - D^{(s(x,t))}_{model,j})^2}{2(\sigma_{CFC-11,j})^2}], \tag{5.6}$$

of the jth model output is interpreted to be the probability that the data (here, the observed CFC-11 concentrations, D_{obs}, at the times, t, and locations, x, along sections, s, in which we have them) is perfectly simulated by the jth model's output (here, its CFC-11 concentrations, $D_{model,j}$, at the same times and locations), and the width of the likelihood,

$$\sigma_{CFC-11,j} = \sqrt{\frac{1}{N-1} \sum_{x,t,s} (D^{*(s(x,t))}_{obs} - D^{(s(x,t))}_{model,j})^2}. \tag{5.7}$$

The uncertainties associated with D_{obs} in the likelihood, (5.6), for each model simulation are integrated out by performing MCMC to evaluate the integrated likelihood, given by

$$P(D_{obs}|M_j) = \int dD^*_{obs} P(D_{obs}|D^*_{obs}, M_j) P(D^*_{obs}|M_j), \tag{5.8}$$

where $P(D^*_{obs}|M_j)$ is a prior belief on the CFC-11 observations given the CFC-11 concentrations from the jth model.

Finally, we calculate m_j by taking $10,000$ samples from $P(M_j|D_{obs})$ via MCMC that take on the value one (when m_j is found to be a better model than the previous iteration) or zero (when not). The way we take a sample of $P(M_j|D_{obs})$ at each MCMC iteration is as follows: Let the $(n-1)$th sample of the set (Markov Chain) of samples whose proportion of ones determines $P(M_j|D_{obs})$ be one for model, $M^{*(n-1)}_j$, and zero for all other models. Then letting the model, $M^{*(n)}_j$, be the proposed model that is one for the nth sample in the Markov Chain, we accept or reject $M^{*(n)}_j$ as the next value in the Markov Chain according to an acceptance probability,

$$A(M^{*(n-1)}_j, M^{*(n)}_j) = min\{1, \frac{P(M^{*(n)}_j) P(D_{obs}|M^{*(n-1)}_j)}{P(M^{*(n-1)}_j) P(D_{obs}|M^{*(n)}_j)}\}. \tag{5.9}$$

Here, $A(M^{*(n-1)}_j, M^{*(n)}_j) = A^*$ is compared with a random number, x, where $0 \leq x \leq 1$. We set the nth sample in the Markov Chain for model $M^{*(n)}_j$ equal to one if $A^* \geq x$ and the nth sample in the Markov Chain for model $M^{*(n-1)}_j$ is one if $A^* < x$. The posterior probability, m_j, is set to be the proportion of ones for model M_j.

Using the above posterior probabilities, m_j, we can find an estimate for the expected value and variance of this estimate for an output of our models. Specifically, our final model estimate of X (which here is the subduction rate) is given by

$$\hat{X} = \sum_{j=1}^{k} m_j X_j \tag{5.10}$$

where X_j is the value of X according to model M_j. Also, for $Var[X|M_j, D]$ being the variance of the estimator of X according to model M_j, the variance on our estimator of X is

$$Var[\hat{X}] = \sum_{j=1}^{k} m_j Var[X|M_j, D] \tag{5.11}$$

where $Var[X|M_j, D]$ is just the variance of the Markov Chain of ones and zeroes found with MCMC for the jth model simulation multiplied by the square of the TTD- and diapycnal velocity-based estimates of subduction rates inferred from the simulation of model j.

5.3 Results

Subduction rates from (5.2) and (5.4) are computed for each of the six model runs with a given t_0 (in five year increments with $t_0 = 0$ corresponding to 1931,..., $t_0 = 75$ corresponding to 2006). These six estimates, for a given t_0, are used to estimate a subduction rate and its uncertainty with (5.10) and (5.11) using CFC-11 observations from all of the occupations across a given basin to compute the BMA weights. For the Southern Ocean, each occupation of SR03 and I09 at each location and depth above 1000 meters was used in the likelihood. For the Subtropical North Atlantic Ocean, each occupation of A05 and A20 at each location south of 30^oN and depth was used in the likelihood, (5.6). For the Subpolar North Atlantic Ocean, each occupation of A16 and A20 at each location north of (and including) 30^oN and depth was used in the likelihood, (5.6).

All MCMC runs here take $10,000$ posterior samples of $P(M_j|D)$ (Table 5.3) in each basin to calculate the weights, m_j, equal to the mean value of the Markov Chain of values for M_j (0 or 1). Because there is autocorrelation between the Markov Chains starting from the first and bth sample for $b < b_{in} = 1,000$ and because it may take several accept/reject iterations to find the posterior distribution, we only use sample num-

bers b_{in} to $10,000$ as posterior samples. This makes $9,000$ samples for each location and depth. We also tried using two priors: $[P(M_1), P(M_2), P(M_3), P(M_4), P(M_5), P(M_6)] = [0.050, 0.125, 0.150, 0.500, 0.050, 0.125]$ and $[1/6, 1/6, 1/6, 1/6, 1/6, 1/6]$ for M_1, M_3, and M_5 using the climatological circulation fields; M_2, M_4, and M_6 using the hindcast fields; M_1 and M_2 using $\kappa_H = 5.0 \times 10^2$ m^{2}s^{-1}; M_3 and M_4 using $\kappa_H = 1.0 \times 10^3$ m^{2}s^{-1}; and M_5 and M_6 using $\kappa_H = 1.0 \times 10^4$ m^{2}s^{-1}. We found no distinguishable difference by using either prior distribution.

5.3.1 Water mass thicknesses by basin

According to the climatological and hindcast HIM outputs, we identify SAMW, STMW, and SPMW with the potential vorticity and density ranges as listed in Table 5.2. The thicknesses of SAMW, STMW, and SPMW according to the model climatology (Figs. 5.3a, 5.3d, and 5.3g) and hindcast (Figs. 5.3b, 5.3e, and 5.3h) are similar in spatial distribution to previous estimates [88, 31, 6]. Because there is no trend in the mode water thicknesses, we assume that mode water is eroded where mode water thicknesses become smaller and find that SAMW primarily is eroded in the ACC region and the Tasman Sea (Fig. 5.3c); STMW is eroded in the outer regions of the Subtropical North Atlantic gyre (Fig. 5.3f); and SPMW is eroded in the Labrador Sea, the Irminger Sea, and most of the central Subpolar North Atlantic (Fig. 5.3i).

5.3.2 Volume transport and mass fraction erosion spectra

The volume transport spectra of SAMW, STMW, and SPMW (Figs. 5.4a, 5.4c, and 5.4e) suggest that there is a small fraction of mode waters in the Southern and North Atlantic Oceans that exceed about 20 years of age since having been at the surface. The mode of the volume transport spectrum of SAMW is at 2 years (Fig. 5.4a) with a mean residence time of $14.5(14.0, 15.0)$. The mode of the volume transport spectrum of STMW is at 1 year (Fig. 5.4c) with a mean residence time of $11.9(11.3, 12.5)$. The mode of the volume transport spectrum of SPMW is at 0 years (Fig. 5.4e) with a mean residence time of $3.41(2.70, 4.11)$. The temporal variability, used as the standard errors for the confidence intervals, of STMW

is greater than that of SAMW or SPMW in percent of volume transport magnitudes. The absolute temporal variability and magnitudes of volume transport, however, is greatest for SPMW.

The mass fractions of SAMW, STMW, and SPMW eroded as a function of residing beneath the mixed layer for some amount of time (Figs. 5.4b, 5.4d, and 5.4f) suggest that mode waters are partially composed of waters that get recycled and stay beneath the mixed layer because there is a significant amount of mode water with residence times that exceed 20 years. The mode of the mass fraction erosion spectrum of SAMW is at 8 years with a secondary peak at 13 years (Fig. 5.4b) and the largest mass fraction eroding at large residence times of the mode waters considered here. The mode of the mass fraction erosion spectrum of STMW is at 2 years with the secondary and tertiary peaks at 5 and 12 years (Fig. 5.4d). The mode of the mass fraction erosion spectrum of SPMW is at 1 year with a secondary peak at 11 years (Fig. 5.4f) and the smallest mass fraction eroding at large residence times of the mode waters considered here. The temporal variability of the mass fraction erosion spectrum for SAMW is largest with STMW second-to-largest.

The secondary peaks in the mass fraction erosion spectra of SAMW (Fig. 5.4b), STMW (Fig. 5.4d), and SPMW (Fig. 5.4f) are consistent with erosion processes with at least two distinguishable time scales. One erosion process is entrainment into the mixed layer, corresponding to one of the peaks in the mass fraction erosion spectrum, and another involves upwelling and mixing [62]. The third distinguishable time scale for STMW at about 5 years is not pronounced for SPMW and absent for SAMW, which does not necessarily suggest a third physical process by which mode water erodes. Instead, this could indicate the response of STMW to some atmospheric forcing, such as the North Atlantic Oscillation (NAO) [45, 44], but this is beyond the scope of the present study.

5.3.3 Subduction rate estimates

Based on the BMA-derived posterior probabilities using a uniform prior (Table 5.3) and using the TTD method, estimates of the subduction rates and their estimates' variances were computed for years 1931 through 1958 using only the climatological fields and 1959 through

2005 using both the climatological and hindcast fields. Thus, when we label the abscissa with years 1931 through 1958, these should not be literally interpreted as calendar years as opposed to model years. Additionally, a calculation based on the diapycnal velocities via (5.4) was performed and is represented by the grey lines in Fig. 5.5. The diapycnal velocity-based estimates represent the formation of water in the mixed layer with subsequent advective subduction, but does not allow for modification of the water mass once it is in the ocean interior (by diffusive processes).

The time series of subduction rates and their estimates' 95% confidence intervals are shown in Fig. 5.5, which suggests an increasing trend for SAMW and SPMW without the uncertainties and an oscillatory pattern for STMW. Investigation of the physical mechanisms responsible for the changes in subduction rates is beyond the scope of this manuscript, but possible culprits include an increasingly stratified upper water column from greater warming rates closer to the surface than in the interior and changes in the winds such as those associated with changes in the Southern Annular Mode (SAM) [87] or NAO. The uncertainties on the TTD-based estimates suggest that the increasing trends in mode water subduction rates are not significant until after the 1970's for SAMW (Figs. 5.5a) and the 2000's in the case of SPMW (Fig. 5.5c).

The trends in Fig. 5.5 are only significant if the circulation fields and κ_H are the dominant sources of uncertainty within each model configuration. The diapycnal velocity-based estimates of the subduction rates do not have a significant trend. In contrast to SAMW, STMW has a peak in its subduction rates around 1970. It is likely to be a coincidence that this is when the NAO index switched from being more frequently negative to more frequently positive because there is no NAO regime change in our model. STMW volume should be greater during the post-1970 period according to the correlation found by Kwon and Riser (2004) [51], consistent with the implied greater STMW formation after 1970 found here. In the Southern Ocean, mixed layer depth anomalies are primarily determine by the SAM [87]. Just like it is possible that the hindcast circulation fields are consistent with the changes in the NAO, it is possible that the hindcast circulation fields are consistent with the changes in the SAM observed in the real atmosphere, but this is beyond the scope of the present study. Lastly, SPMW subduction rates could be increasing for a reason different

from that of SAMW and SPMW formation could be experiencing a different trend since SPMW subduction contributes only a small fraction to SPMW formation [6].

The diapycnal velocity-based estimates of subduction rates of SAMW are never larger than the TTD-based estimates (albeit, the estimates differ by a statistically insignificant amount). Whatever distinguishable difference there is between the TTD-based and diapycnal velocity-based estimates could be attributed to diffusion enhancing the ventilation process and/or diffusion eroding SAMW in the thermocline. The assumption that the properties of water masses are approximately conserved in the thermocline after being ventilated is not particularly good one in the Southern Ocean because SAMW is known to be altered and turned into AAIW [94]. The subduction rates of SAMW derived from the TTD-based and diapycnal velocity-based estimates could also be slightly different because the TTD-based estimates include all SAMW in the domain shown in Figs. 5.3a-b, which could have been ventilated elsewhere, while the diapycnal velocity-based estimates only include SAMW ventilated in the domain shown in Figs. 5.3a-b. However, this discrepancy would not get worse with time, so the differences in the two time series estimates cannot be attributed to a domain boundary problem. This is not a problem with the estimates for STMW or SPMW because there is no STMW or SPMW outside the domains shown in Figs. 5.3d-e and 5.3g-h.

5.4 Conclusions

SAMW (Indian Ocean), STMW (North Atlantic), and SPMW (North Atlantic) have different distributions of residence times below the mixed layer. SAMW is the mode water with the longest residence times. Each of the ventilated mode waters erode by processes with at least two distinguishable time scales. Our estimates provide evidence that SAMW (Indian Ocean) and SPMW (North Atlantic) subduction rates increase over time. Since the diapycnal velocity-based method underestimates the subduction rates relative to the TTD-based method where diffusion is relatively important for water mass modification, along-isopycnal mixing is inferred to have become more important for SAMW formation over time.

While it is possible that the increasing trend we find in SAMW subduction rates is real, we cannot conclude that we have found a trend because we have likely underestimated the

mode water subduction rates ($1-4$ Sv for STMW, $1-4$ Sv for SPMW, and $4-10$ Sv for SAMW). For example, Trossman et al. (2009) [106] estimated at least a few years in which STMW (North Atlantic) subduction was about $6-8$ Sv, albeit with 50% error, which could possibly be attributed to the fact that HIM/Offtrac does not resolve eddies. Future studies are needed to establish the possible link between the SAM and SAMW subduction, the possible link between the NAO and STMW and/or SPMW subduction, whether SAMW subduction has been increasing over time, and whether along-isopycnal diffusion is a cause of the increase in SAMW subduction over time.

Even if mode water volume formed approximately balances mode water volume erosion, mode water formation rates do not need to equal mode water erosion rates because averaging sums of rates is not the same as averaging sums of volume (which is the approximately conserved quantity). Also, data processing error propagation is missing from our error analysis, which could account for some of the imbalance in formation versus erosion of mode waters. Lastly, the distribution of ventilated waters according to density class could be more uniform and less peaked at the mode, which could explain changes in mode water formation versus erosion as opposed to other water mass formation versus erosion. More formation than erosion of mode waters due to greater stratification of the surface ocean (due to larger warming rates at the surface than just below the surface) would have consequences for the stratification of the interior and the feedback the ocean has on the atmosphere via release of heat and carbon stored in mode waters. Further studies are needed to establish the extent to which an increasing subduction rate trend will persist.

Studies such as Glecker et al. (2008) [26] have attempted to construct performance metrics for climate models. We argue here that BMA is another way in which a climate model can be judged with respect to itself or others, which may be an enhancement of the weighted mean squared error approach of Glecker et al. (2008) [26] because the posterior probabilities serve as the weights themselves. When a model is judged against itself, there is a straightforward interpretation of what the "best" parameter configurations should be. When judged against other models, with possibly different parameters altogether, the interpretation of the parameter configurations is less straightforward, but the posterior probability of a model given the observations used in the BMA calculation still has the same

Table 5.1: Tracer data collected in the listed years (before 2006) along each of the cross-sections of interest.

WOCE Section:	Latitude(s)	Longitude(s)	CFC-11
A16(N)	34^oN to 63^oN	20^oW	1988, 1993, 1998, 2003
A05	24.5^oN	79^oW to 14^oW	1992, 1998, 2004
A20	20^oN to 41^oN	53.5^oW	1997
SR03	67^oS to 43^oS	135^oE to 147^oE	1991, 1995, 2001
I09(S)	63^oS to 35^oS	115^oE	1995

interpretation. One warning here would be that, in our case, we only used a set of sparse, aliased tracer concentrations, which may have been observed at the times when the model performed better or worse than at other times. In the case of future climate modeling, there are no future observations with which to use in BMA, but all available observations relevant for the quantity one would like to calculate for future times should be used. Lastly, if different physical models (with potentially different parameters) are used, then a random sample of all possible physical models should be taken rather than the suite of all models available because there may be (e.g., as is the case with the IPCC models [71]) duplication of more similar models than others which would bias a final estimate of an output quantity from these models via BMA.

Table 5.2: The potential vorticities (PV) [$20 \times 10^{-10}\mathrm{m}^{-1}\mathrm{s}^{-1}$] and potential densities (minimum, $\rho_{\theta,min}$, to maximum, $\rho_{\theta,max}$) [kg/m^3] of Subantarctic Mode Water (SAMW), North Atlantic Subtropical Mode Water (STMW), and North Atlantic Subpolar Mode Water (SPMW) according to either the hindcast and climatological circulation fields.

Water Mass	circulation fields	longitudes	latitudes	PV	$\rho_{\theta,min}$-$\rho_{\theta,max}$
SAMW	climatology	20-180^{o}E	65-20^{o}S	6.85($^{\pm}$4.64)	26.7($^{\pm}$0.0434)-26.8($^{\pm}$0.0716)
SAMW	hindcast	20-180^{o}E	65-20^{o}S	6.83($^{\pm}$4.65)	26.7($^{\pm}$0.0430)-26.8($^{\pm}$0.0756)
STMW	climatology	80-0^{o}W	10-40^{o}N	8.49($^{\pm}$3.57)	26.2($^{\pm}$0.1112)-26.5($^{\pm}$0.0431)
STMW	hindcast	80-0^{o}W	10-40^{o}N	8.49($^{\pm}$3.57)	26.2($^{\pm}$0.1075)-26.5($^{\pm}$0.0436)
SPMW	climatology	80-0^{o}W	40-70^{o}N	7.37($^{\pm}$4.40)	27.3($^{\pm}$0.0221)-27.4($^{\pm}$0.0256)
SPMW	hindcast	80-0^{o}W	40-70^{o}N	7.34($^{\pm}$4.41)	27.3($^{\pm}$0.0211)-27.4($^{\pm}$0.0245)

Table 5.3: The posterior probabilities, $P(M_j|D)$, using a uniform prior, $P(M_j) = 1/6$ ($j = 1, ..., 6$), for the six configurations of along-isopycnal diffusivity, κ_H, and hindcast/climatological fields for both the Southern Ocean (SO) using I09 and SR03, the Subtropical North Atlantic (STNA) using A05 and A20, and the Subpolar North Atlantic (SPNA) using A20 and A16.

climatology/hindcast	κ_H	$P(M_j\|D)$: (SO)	(STNA)	(SPNA)
climatology	5.0×10^2	0.159	0.205	0.248
hindcast	5.0×10^2	0.175	0.130	0.0837
climatology	1.0×10^3	0.155	0.203	0.246
hindcast	1.0×10^3	0.180	0.133	0.0901
climatology	1.0×10^4	0.155	0.199	0.244
hindcast	1.0×10^4	0.176	0.129	0.0873

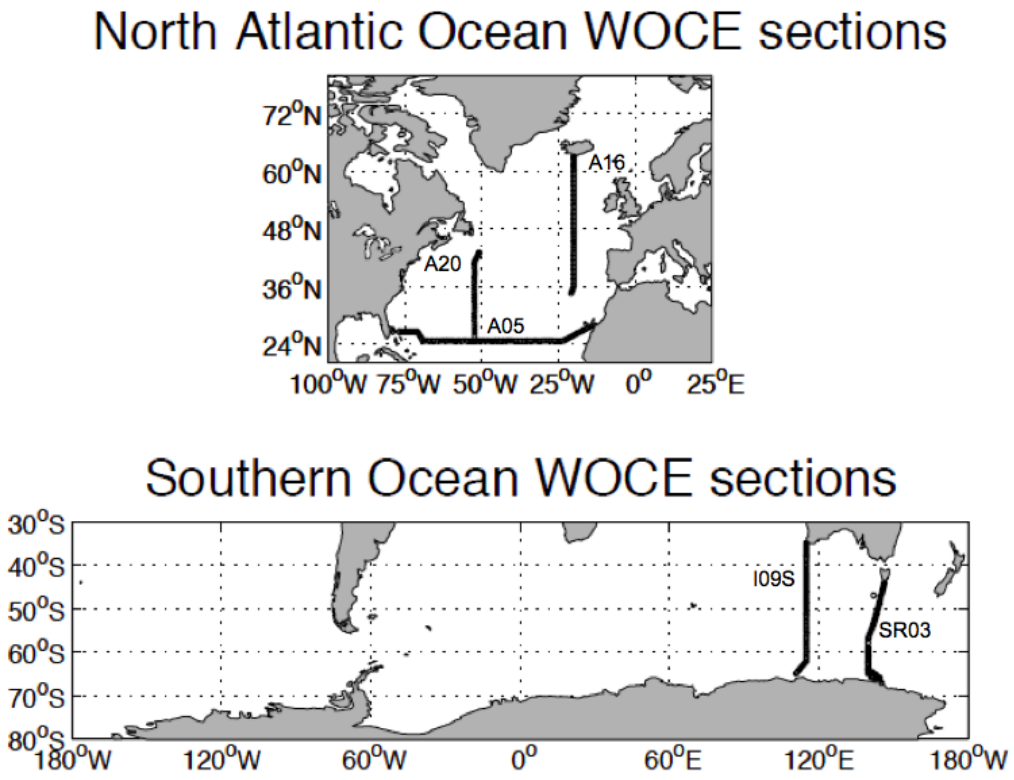

Figure 5.1: Locations of bottle data measurements used in this study taken along WOCE transects (lines of black Os), listed in Table 5.1.

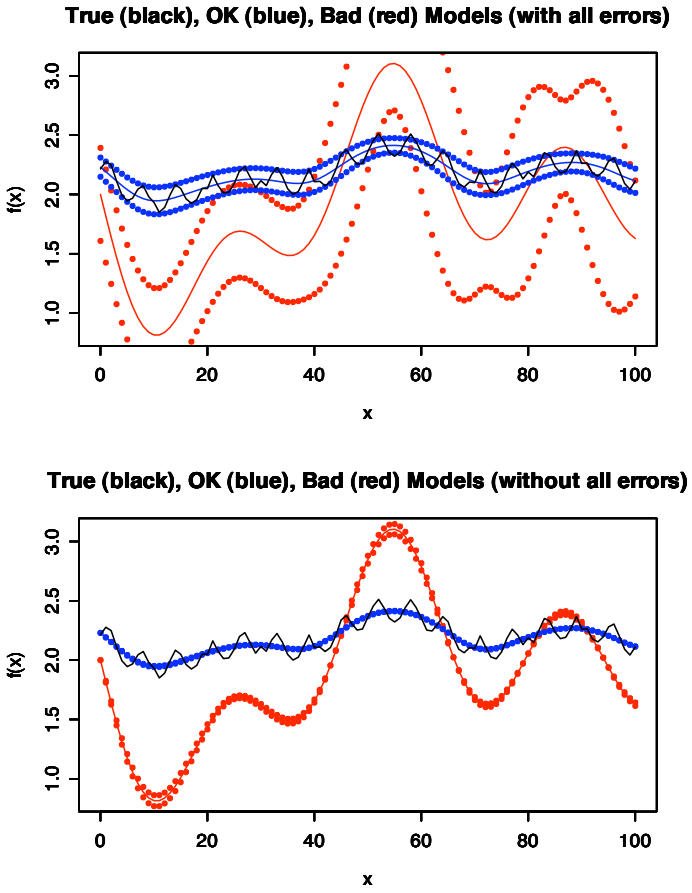

Figure 5.2: Using a deterministic function that we call 'truth' (solid black), an approximation that filters out the higher frequencies (solid blue), and a crude approximation that intersects the truth least often (solid red), the top panel shows 95% confidence intervals (dots) taking all sources of uncertainty (parameters in the deterministic function) into account and the bottom panel show 95% confidence intervals (dots) taking only one of the sources of uncertainty into account.

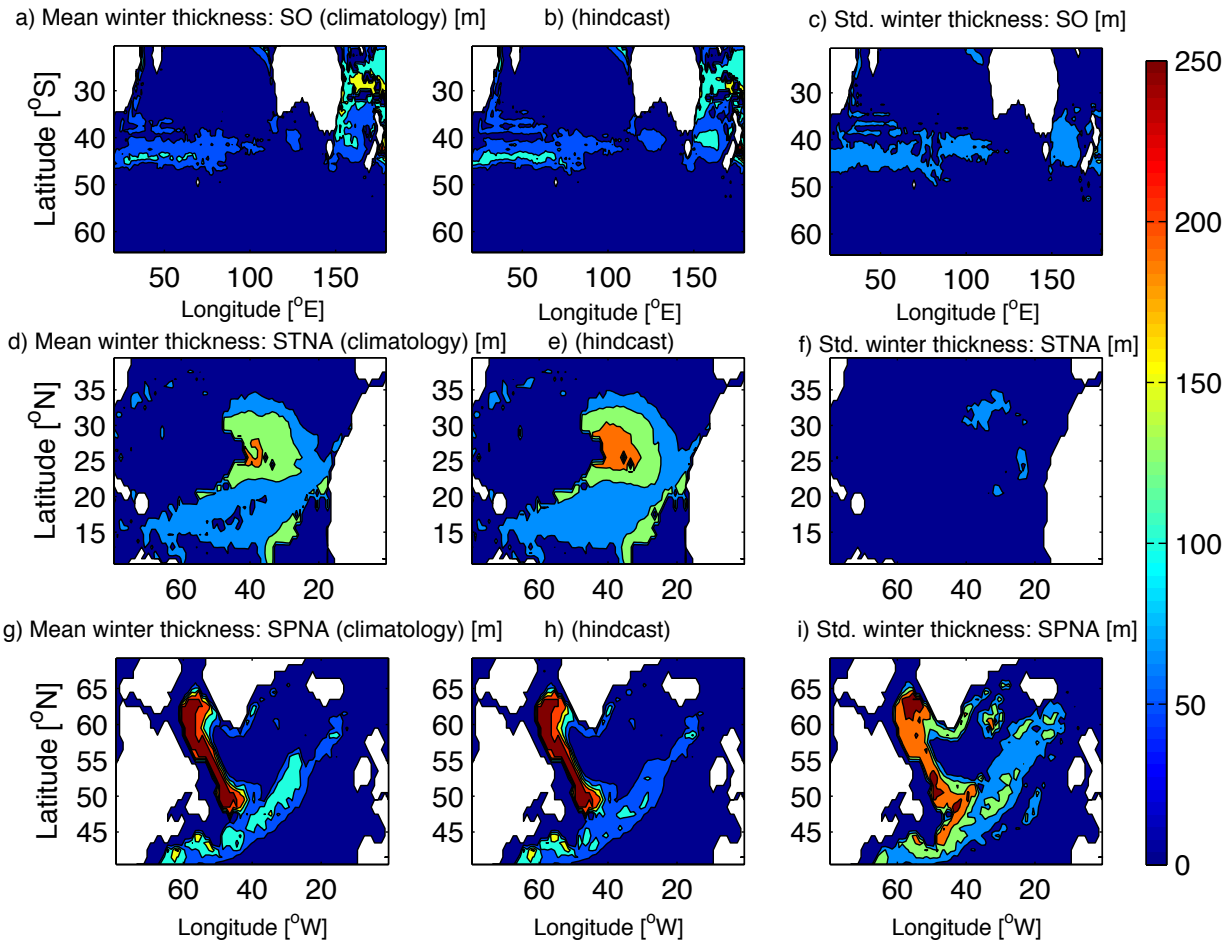

Figure 5.3: The average mode water thickness in the water column at each horizontal location over the last 120 months of the 250 year climatology (panels a, d, and g) and the entire 47 years of the hindcast (panels b, e, and h) in the Southern (panels a, b, and c), Subtropical North Atlantic (panels d, e, and f), and Subpolar North Atlantic (panels g, h, and i) Oceans. Also plotted are the standard deviations about the average mode water thickness in the water column at each horizontal location over the entire 47 years of the hindcast (panels c, f, and i) in the Southern (panel c), Subtropical North Atlantic (panel f), and Subpolar North Atlantic (panel i) Oceans.

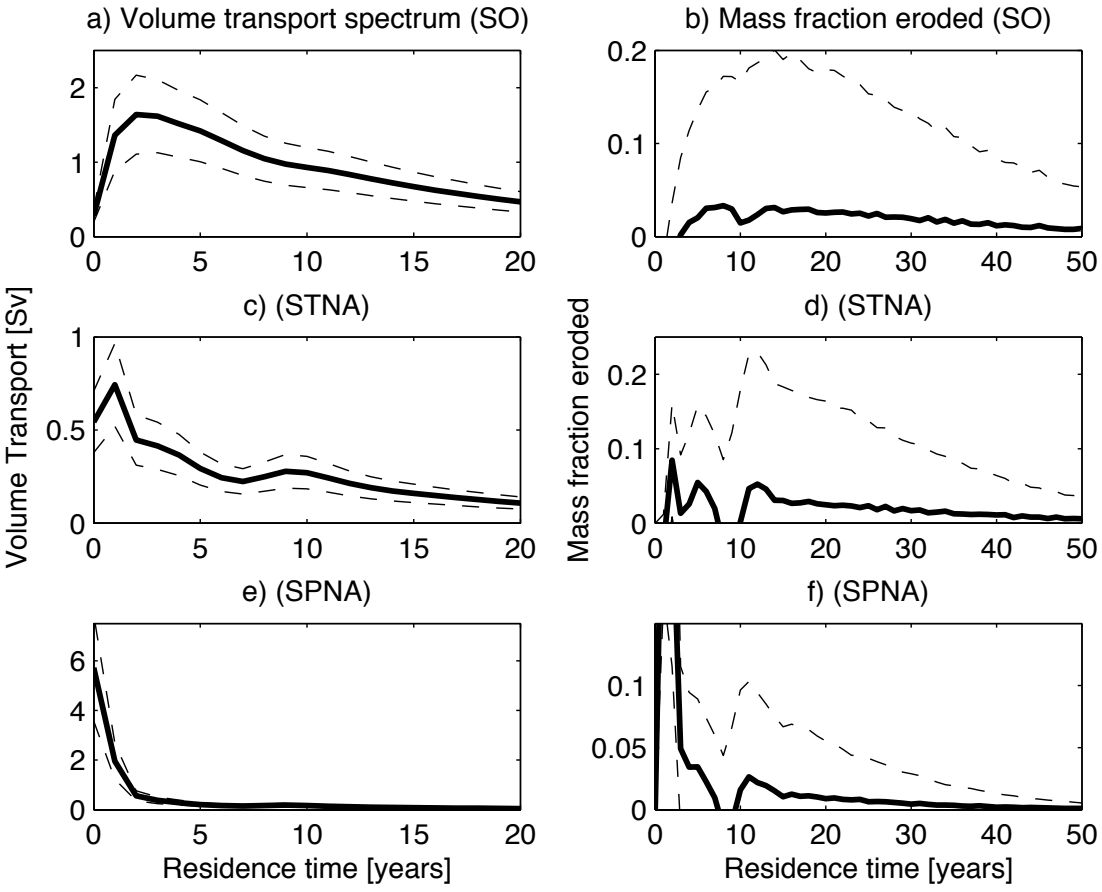

Figure 5.4: The volume transport spectra [Sv= 10^6 m^3s^{-1}] (a, c, and e) and mass fractions eroded (b, d, and f) (solid lines) and their variability over time (dashed lines) for Subantarctic Mode Water (a-b), Subtropical Mode Water (c-d), and Subpolar Mode Water (e-f) for years 1931 through 2006.

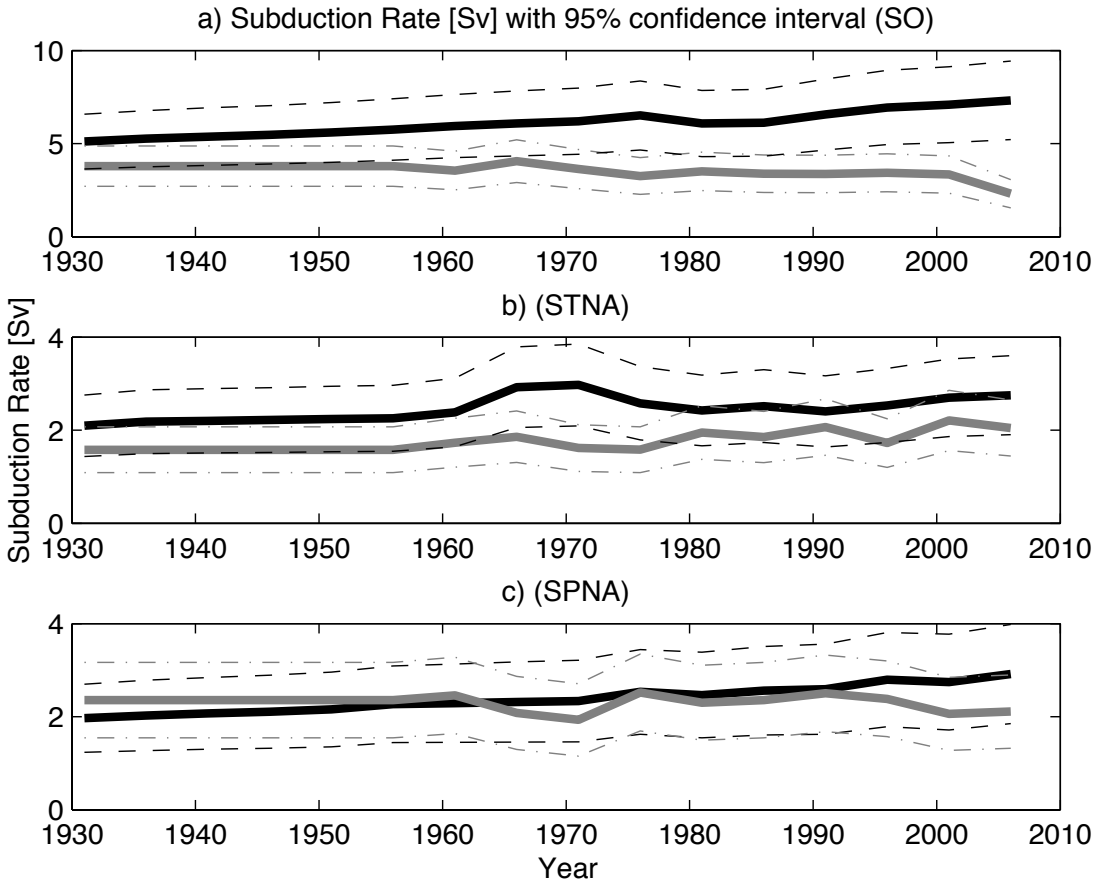

Figure 5.5: The subduction rates [Sv= 10^6 m^3s^{-1}] (TTD-based method in solid black lines and vertical velocity-based method in grey lines) and their uncertainties derived from BMA estimate variances (dashed black lines for the TTD-based-method and dot-dashed grey lines for the vertical-velocity-based method) for Subantarctic Mode Water (panel a), Subtropical Mode Water (panel b), and Subpolar Mode Water (panel c) for years 1931 through 2006.

Chapter 6

POTENTIAL DENSITY-DEPENDENCE OF ALONG-ISOPYCNAL DIFFUSIVITY

6.1 Introduction

While studies from Ledwell et al. (1993) [56] to Goes et al. (2010) [27] have constrained the effective diffusivity in the World Ocean across constant density surfaces (diapycnal diffusivity $\sim 1 - 2 \times 10^{-5}$ m^2s^{-1} in pelagic regions), the diffusivity (along isopycnals) is at least seven orders of magnitude larger and is only known to about an order of magnitude. Studies such as Mecking et al. (2004) [64] and Abernathey et al. (2010) [1] suggest that the along-isopycnal diffusivity should have some density- and spatial-dependence as opposed to being a fixed background value. Mecking et al. (2004) [64] showed, in the North Pacific Ocean, that along-isopycnal diffusivity needs to be 5×10^3 m^2s^{-1} in the Kuroshio Extension (2×10^3 m^2s^{-1} elsewhere) to simulate observed CFC concentrations. Abernathey et al. (2010) [1] showed, in the Southern Ocean, that the effective diffusivity can be as high as 2×10^3 m^2s^{-1} near the steering levels of baroclinic waves and as low as 5×10^2 m^2s^{-1} near the surface in the Antarctic Circumpolar Current (ACC) using the method of Nakamura (1996) [68]. Shuckburgh et al. (2011) [91] further suggest that the mixed layer lateral eddy diffusivities are a function of both the associated physical processes (either eddy stirring and small-scale mixing or surface damping by air-sea interaction) that give rise to the distribution of surface tracers and the type of tracer, ranging from $\sim 5 \times 10^2$ to 1×10^4 m^2s^{-1}.

The average along-isopycnal diffusivity of a water parcel along its path from where it was formed to its current position is expected to be different for a parcel near the surface where there is more mechanical energy to stir the water than in the abyss. Furthermore, it is possible that the path-averaged along-isopycnal diffusivity is distinguishable by water mass because water masses form and are altered by different physical processes. For water

masses that form by deep convection, their properties are set near their formation site and then they are subsequently advected and diffused. For water masses that are formed by mixing distinguishable water masses, on the other hand, diffusion acts in the formation process itself, potentially leading to a higher average along-path diffusivity. The water masses of interest here are shallow and intermediate water masses that are spatially and temporally sampled along several hydrographic cross-sections in the North Atlantic and Southern Oceans. In the North Atlantic Ocean, we will be focused on in Subpolar Mode Water (SPMW), Subtropical Mode Water (STMW), and Labrador Sea Water (LSW). North Atlantic SPMW is formed through cooling and freshening of surface waters in upwelling regions that then enter the Gulf Stream and North Atlantic Current and move to the northeast subpolar regions. North Atlantic STMW is formed through subduction in the western boundary current extension [60]. LSW is formed through deep convective processes in the Labrador and Irminger Seas [75]. In the Southern Ocean, we will be interested in Subantarctic Mode Water (SAMW) which forms on the northern side of ACC after enhanced vertical diffusivity below the mixed layer in austral summers and autumns alters the interior stratification and preconditions the water column to form deep mixed layers [94], and Antarctic Intermediate Water (AAIW) which forms in the southeast Pacific and has been argued to form from the coldest component of SAMW via a combination of convective and Ekman processes at the Antarctic Polar Front (between $50 - 60^o$S) to sink to depths below 1000 meters south of 50^oS [61].

The ultimate goals of our efforts here are to estimate the path-averaged along-isopycnal diffusivities of water parcels in different regions and test whether water masses that form from different mechanisms require different path-averaged along-isopycnal diffusivities to be consistent with observations in the North Atlantic and Southern Oceans. The World Ocean Circulation Experiment (WOCE) and Climate Variability and Predictability (CLIVAR) & Carbon observations (Table 5.1) are utilized along with outputs from several isopycnal coordinate model runs. These model runs are ranked according to which simulates output that is closer to observations. Instead of averaging the model outputs together with equal weights, we can provide more realistic estimates by averaging the model outputs together with weights that reflect their relative rankings. We estimate the point-wise along-isopycnal

diffusivities and their uncertainties along each hydrographic section with a statistical technique called Bayesian model averaging (BMA). BMA incorporates observations by treating model runs as random variables themselves, each with an associated posterior probability (roughly interpreted to be the degree of our belief that a model simulates the observations better than other models) that we use to weight the outputs to come up with a final estimate.

For comparison, we find a numerical model's path-averaged diffusivity at each point along each hydrographic section with our technique, utilizing BMA. To do this, we first estimate a transit-time distribution, or TTD, [38] at each point according to a method evaluated by Haine et al. (2008) [30], called the boundary impulse response (BIR) method. A TTD is a Green's function solution to the three-dimensional advection-diffusion equation that quantifies the fraction of mass that gets from a boundary (e.g., the surface of the ocean) to some other location (e.g., the interior of the ocean) from the time a tracer enters the ocean at the boundary to some other time. A purely advective water parcel would have a delta spike for a TTD, so the half-variance of the TTD essentially is a measure of the path-averaged diffusivity.

The main topics of this paper are arranged as follows: The observational data (CFC-11 concentrations) and model simulations (of CFC-11 concentrations and TTDs) we have are outlined. We design a set of model simulations that are eventually averaged with weights derived from the observational data according to BMA. We then review the method of BMA. Using the idea that when CFC-11 concentrations in our model output are closer to those of the observations, the model's output from that particular run is more heavily weighted than others in the ensemble of runs. With this weighting, we determine an estimate of the along-isopycnal diffusivity and the uncertainty on this estimate. We compare the BMA-derived weights over entire ocean basins to quantify the importance of along-isopycnal diffusion in the North Atlantic and Southern Oceans. Lastly, we use BMA to ascertain a density dependence of the along-isopycnal diffusivity.

6.2 Methods

6.2.1 Data and Model Runs

In order to evaluate the path-averaged along-isopycnal diffusivity at each point according to our numerical model, we use CFC-11 concentrations from five repeat hydrographic sections listed in Table 5.1. These sections were described in the previous chapter and shown in Fig. 5.1. The offline isopycnal coordinate model (Offtrac), the Offtrac parameter specifications (either climatological or hindcast circulation fields and three different values for the background along-isopycnal diffusivity κ_H), and the mass flux fields Offtrac uses based on output from another isopycnal coordinate model (HIM) were also described in the previous chapter. To justify our use of HIM/Offtrac, comparison plots are shown and discussed in Section A.4 of the Appendix.

What is different in this chapter from the previous chapter is that only six model simulations are performed for 75 years of monthly CFC-11 output and six model simulations are performed for 200 years of TTD output. Each of the six members in our ensemble of simulations starts the age clock at given time, $t_0 = 0$, and we compute the TTD from that time, t_0, to t_0+200 years. The TTD is estimated with the BIR approximation [30] by letting the surface of the ocean be $1/N$ for an entire year where N is the number of time steps in a year, keeping track of the values of this tracer. We need to extrapolate the model TTD simulations to capture the tails beyond 200 years. To perform this extrapolation, following the method described in the Section A.2.2 of the Appendix, we use a cubic spline fit with knots at every three months, average over windows of eight months, and force the TTD to vanish one-third of the way between the peak of the TTD and the year 1200 via log-linear extrapolation. We then reassign the raw TTD to the first 200 years, and numerically integrate the resulting 1200 year-long TTD to get its first two cumulative moments.

6.2.2 How to Choose Which Model Output is Best: Bayesian Model Averaging

Since no single model simulation is perfect, we want to come to a better estimate of a quantity of interest by weighting model outputs based on comparisons with observations. To do this, we need to decide, *a priori*, which model parameters affect the quantity of

interest (which here is a model parameter, κ_H) but that we are uncertain about (which here are κ_H and circulation fields) (see previous chapter). We need to vary these model parameters to span a realistic range of values and choose to use CFC-11 concentrations as a surrogate performance measure of each model simulation. Using BMA, we can estimate a quantity of interest (which here is κ_H) with uncertainties. BMA basically treats our model simulations as random variables with a probability density that has a peak located at the model simulation that reproduces the observations most closely.

To be more mathematically precise, we have different input parameter configurations for HIM/Offtrac runs, $M_1, ..., M_k$, where here, $k = 6$, and we want to weight each of these configurations by the relative degree to which we believe in them. Each model has a probability density for some random variable, X, (which can be any output variable or parameter of the model, but here is κ_H) and unknown parameters (which here are κ_H and the climatological or hindcast HIM circulation fields). We can find the posterior probability of the jth model configuration being the true model for the data amongst all possible models, $M_1, ..., M_k$, which we denote by m_j. The background for implementing BMA here was developed in the previous chapter, but in the current chapter, we will use two different forms of the likelihood, $P(D_{obs}|M_j)$, following the notation of the previous chapter. By incorporating the uncertainty associated with the CFC-11 observations, we find one form of the likelihood given by

$$P(D_{obs}|D^*_{obs}, M_j) = \prod_{x,t,s} \frac{1}{\sqrt{2\pi(\sigma_{CFC-11,j})^2}} exp[-\frac{(D^{*(s(x,t))}_{obs} - D^{(s(x,t))}_{model,j})^2}{2(\sigma_{CFC-11,j})^2}], \tag{6.1}$$

of the jth model output is interpreted to be the probability that the data (here, the observed CFC-11 concentrations, D_{obs}, at the times, t, and locations, x, along sections, s, in which we have them) is perfectly simulated by the jth model's output (here, its CFC-11 concentrations, $D_{model,j}$, at the same times and locations), where the width of the likelihood is given by

$$\sigma_{CFC-11,j} = \sqrt{\frac{1}{N_{tot} - 1} \sum_{x,t,s} (D^{*(s(x,t))}_{obs} - D^{(s(x,t))}_{model,j})^2} \tag{6.2}$$

for N_{tot} being the number of unique x, t, and s. Another form of the likelihood used here applies this method to individual grid points, in which case the sums over s, x, t become

sums over t only. In this form, we find the likelihood to be given by

$$P(D_{obs}|D^*_{obs}, M_j) = \prod_t \frac{1}{\sqrt{2\pi(\sigma_{CFC-11,j})^2}} exp[-\frac{(D^{*(s(x,t))}_{obs} - D^{(s(x,t))}_{model,j})^2}{2(\sigma_{CFC-11,j})^2}], \tag{6.3}$$

where, the observed CFC-11 concentrations, D_{obs}, are at a given location and depth during occupation times, t, and the width of the likelihood is given by

$$\sigma_{CFC-11,j} = \sqrt{\frac{1}{N}\sum_t (D^{*(s(x,t))}_{obs} - D^{(s(x,t))}_{model,j})^2} \tag{6.4}$$

where N is the number of occupations along the given section.

Using the posterior probabilities, m_j, equal to the proportion of the Markov Chain with value j formed by performing MCMC, we can find an estimate for the expected value and variance of this estimate for an output of our models. Specifically, our final model estimate of κ_H is given by

$$\hat{\kappa_H} = \sum_{j=1}^{k} m_j \kappa_{H,j} \tag{6.5}$$

where $\kappa_{H,j}$ is the value of κ_H according to model M_j. Also, for $Var[\kappa_H|M_j, D]$ being the variance of the estimator of κ_H according to model M_j, the variance on our estimator of κ_H is

$$Var[\hat{\kappa_H}] = \sum_{j=1}^{k} m_j Var[\kappa_H|M_j, D] \tag{6.6}$$

where $Var[\kappa_H|M_j, D]$ is just the variance of the Markov Chain of ones and zeroes found with MCMC for the jth model simulation multiplied by $\kappa^2_{H,j}$.

6.3 Results

Using (6.5) and (6.6), κ_H and its uncertainty were estimated two different ways: A) using CFC-11 observations from all of the occupations across a given basin and B) using CFC-11 observations from all of the occupations along a given section at a given location and depth. For method A), the number of points used in the likelihood, (6.1), is the number of locations, depths, and occupations along the sections that are in a given ocean basin. For the Southern Ocean in method A), each occupation of SR03 and I09 at each location and depth was used in the likelihood. For the Subtropical North Atlantic Ocean in method

A), each occupation of A05 and A20 at each location south of 30^oN and depth was used in the likelihood. For the Subpolar North Atlantic Ocean in method A), each occupation of A16 and A20 at each location north of (and including) 30^oN and depth was used in the likelihood. For method B), the number of occupations at each ($1^o \times 50$ meter depth) grid point is equal to the number of data points used in the likelihood, (6.3), given by Table 5.1 as $N = 4$ pre-2006 occupations along A16, $N = 3$ pre-2006 occupations along A05 and SR03, and $N = 1$ pre-2006 occupation along I09 and A20. In the plots of profiles, the deepest depth shown for each of the hydrographic sections is 2000 meters because κ_H is estimated to be constant but uncertain below that depth. We do not expect CFC-11 observations to give much information below 2000 meters because CFC-11 has only been entering the ocean since 1945.

All MCMC runs here take $10,000$ posterior samples of $m_j = P(M_j|D)$, either in each basin in method A) or at each location and depth in method B). Because there is autocorrelation between the Markov Chains starting from the first and bth sample for $b < b_{in} = 1,000$ and because it may take several accept/reject iterations to find the posterior distribution, we only use sample numbers b_{in} to $10,000$ as posterior samples. This makes $9,000$ samples for each location and depth. We also tried using two priors: $[P(M_1), P(M_2), P(M_3), P(M_4), P(M_5), P(M_6)] = [0.050, 0.125, 0.150, 0.500, 0.050, 0.125]$ and $[1/6, 1/6, 1/6, 1/6, 1/6, 1/6]$ for M_1, M_3, and M_5 using the climatological circulation fields; M_2, M_4, and M_6 using the hindcast fields; M_1 and M_2 using $\kappa_H = 5.0 \times 10^2$ m^{2}s^{-1}; M_3 and M_4 using $\kappa_H = 1.0 \times 10^3$ m^{2}s^{-1}; and M_5 and M_6 using $\kappa_H = 1.0 \times 10^4$ m^{2}s^{-1}. We found no distinguishable difference by using either prior distribution.

6.3.1 BMA by basin

We calculated estimates of the posterior probabilities, m_j, via MCMC in the Southern and both the Subtropical and Subpolar North Atlantic Oceans (Table 6.1). The CFC-11 concentrations along SR03 and I09 in the Southern Ocean are closest to the observations using the climatological circulation fields and any value for the along-isopycnal diffusivity between 5×10^2 and 10^4 m^{2}s^{-1}. The model simulations of CFC-11 concentrations are slightly

better using the smaller value for κ_H, but by an insignificant amount due to Monte Carlo error. Similarly, the CFC-11 concentrations along A20 (as far south as 30^oN) and A16 in the Subpolar North Atlantic Ocean, and A20 (as far north as 30^oN) and A05 in the Subtropical North Atlantic Ocean are closest to the observations using the climatological circulation fields and any value for the along-isopycnal diffusivity between 5×10^2 and 10^4 m^{2}s^{-1}. Because of the more extensive historical coverage of the North Atlantic Ocean relative to the Southern ocean and because the hindcast mixed layer depths and winds, in particular, are likely to be more variable in ways that are not quite exact, it is not surprising that the model runs using the climatological circulation fields rank higher in the North Atlantic Ocean. The posterior probabilities calculated using method A) with a non-uniform prior (not shown) or a uniform prior (Table 6.1) are about the same because hundreds of grid points of CFC-11 observations are used in the likelihood, (6.1), to dwarf the influence of the prior. The difference between the posterior probabilities is essentially the Monte Carlo error.

6.3.2 BMA by section

To address whether the rankings tabulated in Table 6.1 have smaller-scale detail than on the basin scale, we calculated estimates of the posterior probabilities, m_j, via MCMC for individual sections by averaging over the weights calculated at each ($1^o \times 50$ m depth) grid point (Table 6.2). Only N data points here were used in the likelihood at each point as opposed to hundreds or thousands. Using these rankings by grid point and averaging over individual sections, we find that the hindcast circulation fields rank higher for the Southern Ocean sections with κ_H closer to 10^3 m^{2}s^{-1} than to 5×10^2 along SR03 and not distinguishably closer to 10^4 than to 10^3 m^{2}s^{-1} for I09. The hindcast also ranks higher along A05 with κ_H not distinguishably closer to 10^3 than to 5×10^2 m^{2}s^{-1}. The climatology ranks higher along A20 with κ_H not distinguishably closer to 10^3 than to 10^4 m^{2}s^{-1} and the climatology ranks higher along A16 with κ_H closer to 10^4 than to 10^3 m^{2}s^{-1}. The numerical model with greater along-isopycnal diffusivities along A16 ranks higher than the other sections.

To investigate whether using data across an entire basin using method A) obscures smaller-scale detail, we use method B) and average the weights themselves calculated at each grid point over entire basins (Table 6.3). The rankings are different between Tables 6.1 and 6.3, which suggests that the along-isopycnal diffusivities should vary across the basins to be more consistent with observations. It is also possible that information provided to the weights via the BMA procedure at each grid point can be extracted by surrounding grid points. To demonstrate how the BMA weights vary in space, an example of the weights and derivatives in objectively mapped dynamical tracer fields derived from observations for each grid point are shown along SR03 in Fig. 6.1. The structure of the weights and derivatives in the tracer fields along SR03 (as well as the other hydrographic sections not shown here) suggest that the weights are a function of density class, which are reflected in the final estimates of κ_H at each grid point. Also suggested by Fig. 6.1 is that the weights at each grid point are different but have the same general spatial structure using either the climatological or hindcast circulation fields.

The spatial structure suggested by the weights in Fig. 6.1 implies a density-dependence rather than a water mass dependence. SAMW (north of 50^oS at 500-1000 meters depth in Fig. 6.1) has a distinct relatively constant large-scale potential vorticity (defined by (5.1)) signature. The vertical gradient of large-scale potential vorticity (Fig. 6.1a) and the second horizontal derivative in salinity are the only spatial derivatives taken in potential vorticity, temperature, or salinity that vaguely outline the density-dependent structure we find in κ_H. The vertical gradient in salinity (not shown) only outlines where there is AAIW (north of 50^oS at 1000-1200 meter depth in Fig. 6.1) as opposed to the entire density class that AAIW belongs to (extending as far south as 55^oS at 1000 meters depth in Fig. 6.1).

6.3.3 North Atlantic estimates

To investigate whether there is a density-dependent structure to κ_H in the North Atlantic Ocean, we plot the estimates of κ_H based on the BMA-derived weights at each grid point. The spatial structure in Fig. 6.2 suggests that, using only CFC-11 and assuming that we have accounted for all sources of uncertainty, we can constrain the uncertainty to an extent

that can distinguish the path-averaged along-isopycnal diffusivity in one density class from another almost everywhere. The uncertainties on these estimates are greater than 1×10^3 $\mathrm{m^2 s^{-1}}$ in density classes where the estimates are greater than 5×10^3 $\mathrm{m^2 s^{-1}}$. These large uncertainty regions appear at about 1000 meters depth in Figs. 6.2a-b, between 1000 and 2000 meters depth in Figs. 6.2d-e, and between both 500-1000 meters depth across $20-80^o$W and 1000-2000 meters depth across $40-60^o$W in Figs. 6.2g-h. AAIW, along A20, AAIW along A16, AAIW along A05, and Classical LSW (CLSW) along A05 can be identified in these large along-isopycnal diffusivity regions, respectively. We see smaller along-isopycnal diffusivities in other regions (about 300 meters depth in Fig. 6.2a, at most depths between $60-62^o$N in Fig. 6.2d, and between 1000 and 2000 m across $20-40^o$W in Fig. 6.2g). STMW along A20, SPMW along A16, LSW along A16, and Upper LSW (ULSW) along A05 can be identified in these small along-isopycnal diffusivity regions, respectively. Again, if we have accounted for all of the uncertainties, we can distinguish these water mass from the others by their along-isopycnal diffusivities. Finally, the medium diffusivity water mass directly below the ULSW and CLSW ($\sim 3-4 \times 10^3$ $\mathrm{m^2 s^{-1}}$) can be identified as Iceland-Scotland Overflow Water (ISOW) (Fig. 2b of LeBel et al., 2008 [55]) with densities greater than those of the LSW components.

The estimated κ_H here is a function of time and space. Since κ_H represents an average along-isopycnal diffusivity along the water mass's path since formation, we elaborate on how events during the life of the parcel could alter κ_H. If the along-isopycnal diffusivities are high when LSW is first formed and low thereafter, then Fig. 6.2g suggests that the ULSW has had a lower along-isopycnal diffusivities for longer than CLSW. If the along-isopycnal diffusivities are low when LSW is first formed and high sometime after that, then Fig. 6.2g suggests that the CLSW has had a higher along-isopycnal diffusivities for longer than ULSW. Both of these stories are consistent with the relative magnitudes of κ_H shown in Fig. 6.2. The path-averaged along-isopycnal diffusivity of CLSW is a little larger than that of ULSW or LSW where CLSW is observed around 22^oN at 2000 meters depth along A20 (Fig. 6.2a) and between $40-60^o$W around 1500-2000 meters depth along A05 (Fig. 6.2g). The path-averaged along-isopycnal diffusivity of ULSW is a little smaller than that of CLSW where ULSW is observed around 25^oN at 1500 meters depth along A20 (Fig.

6.2a) and between $20 - 40^o$W around 1000-2000 meters depth along A05 (Fig. 6.2g). The path-averaged along-isopycnal diffusivity of LSW is smaller than that of CLSW where LSW is observed north of 60^oN (Fig. 6.2d). To deduce where κ_H begins to differ along the paths of ULSW and CLSW, we would need to estimate κ_H at the source region of LSW, where LSW begins to mix with Mediterranean Outflow Water (MOW) to form CLSW, and where eastward traveling LSW separates from southward traveling LSW. While we see that the along-isopycnal diffusivities differ between ULSW and CLSW as far north as about 25^oN from Figs. 6.2a and 6.2g, we do not know whether the diffusivity of CLSW is turned up sharply around 25^oN to yield a larger path-averaged diffusivity than ULSW or the diffusivity of CLSW is turned up less sharply farther north.

The magnitudes of κ_H may be biased here because of incorporation of only local information about grid points (each of which have assumed to be independent) and the three background values of κ_H in our numerical model simulations probably do not adequately sample all possible κ_H. However, the spatial structure of κ_H (Fig. 6.2) is robust as long as the uncertainties have been sufficiently constrained to distinguish water masses from each other. Looking at several profiles of κ_H versus depth (Fig. 6.3) and versus potential density (Fig. 6.4), we can see the distinction in κ_H between different density classes are significant to 90% and 95% confidence. The path-averaged along-isopycnal diffusivities in the North Atlantic Ocean are different where there is SPMW (Figs. 6.3-6.4d and 6.2d at 100-300 meters depth across $58 - 62^o$N), LSW (Figs. 6.3-6.4d and 6.2d below 750 meters depth north of 59^oN) and ULSW (Figs. 6.3-6.4e and 6.2g at 500-2000 meters depth east of the ridge at $47 - 48^o$W and 1000-1500 meters depth west of the ridge).

6.3.4 Southern Ocean estimates

Previous studies [1] have estimated, using the method of Nakamura (1996) [68] and climatological circulation fields in the Southern Ocean, that a higher along-isopycnal diffusivity at middepth and lower along-isopycnal diffusivity near the surface is needed in models of the Southern Ocean. Our results suggest that more variable (hindcast) circulation fields do not model tracer concentrations as well, given one background along-isopycnal diffusivity.

This does not imply that either the along-isopycnal diffusivity parametrization or the circulation fields used are unimportant relative to each other, but rather, given this result and that of Abernathey et al. (2010) [1], that having both more variable circulation fields and more density-dependent along-isopycnal diffusivity parametrizations could be necessary for accuracy in simulated tracer fields at middepths and above.

There are two water masses for which we can distinguish κ_H in one region from another in the Southern Ocean: SAMW (Figs. 6.5a and 6.5d between 700-1000 meters depth north of 55^oS) and AAIW (Figs. 6.5a and 6.5d between 1000-1200 meters depth north of 50^oS) with AAIW having relatively large along-isopycnal diffusivities relative to SAMW. Profiles of κ_H versus both depth and potential density demonstrate how the along-isopycnal diffusivities of SAMW (Figs. 6.3-6.4a between 700-1000 meters depth or at about 27.23 kg m^3 and 6.3-6.4b at about 900 meters depth or about 27.65 kg m^3) and AAIW along SR03 and I09 (Figs. 6.3-6.4a between 1000-1200 meters depth or at about 27.29 kg m^3 and 6.3-6.4b between 1000-1300 meters depth or about 27.70 kg m^3) differ. Also along SR03, we see smaller along-isopycnal diffusivities at depths just below the thermocline between the surface and 1000 meters south of 60^oS (Fig. 6.5a), which could be Antarctic Bottom Water (AABW) flowing off the shelf or the slope current. Since the formation of AABW is not expected to be modeled particularly well by Offtrac, given the circulation fields, short spin-up time, and low resolution, we should be careful not to interpret this region as anything significantly different from the background diffusivity.

We find a maximum in κ_H at about 1000 meters depth. Smith (2007) [95] found that the steering level of the thermocline-level baroclinic instabilities associated with interior potential vorticity gradients is found at about 1000 meters depth, where the potential vorticity gradients change sign. Smith and Marshall (2009) [96] found that the maximum in baroclinic wave energy conversion is found that about 1000 meters as well. We will elaborate on whether these observations are coincidental or our estimates of κ_H represent the path-averaged along-isopycnal diffusivities of the real ocean in the following subsection.

6.3.5 Physical significance of κ_H

A water mass with a given density forms due to either deep convective processes (e.g., mode waters) or the mixing of several different water masses together (e.g., LSW mixing with MOW to form CLSW). The source of energy for the mixing that takes places between several different water masses can come from growing baroclinic waves due to thermal wind shear, which can result in maximum path-averaged diffusivities near the steering level [96]. This has been demonstrated in the Southern Ocean by Abernathey et al., 2010 [1] and is argued here to be the case for both CLSW in the North Atlantic Ocean and AAIW in both the Southern Ocean and Atlantic Ocean. The remaining question of interest in detecting differences in along-isopycnal diffusivities for particular water masses is whether there is something peculiar about the Offtrac model that would require the local κ_H to be smaller for ULSW, LSW, and SAMW or whether this represents something physically real. Since ULSW, LSW, and SAMW form by way of deep convective processes, it is possible that Offtrac simulates the physical processes of deep convection much better or worse than the physical processes by which mixed water masses are formed. This cannot be ruled out by qualitative inspection, as discussed in Section A.4 of the Appendix.

To address the interpretation of κ_H, Figs. 6.5c, 6.5f, 6.2c, 6.2f, and 6.2i show the half-variance, Δ^2, of the model TTD at each grid point along each hydrographic section. We compare magnitudes of Δ^2 and κ_H relative to surrounding locations because we know that Δ^2 reflects the path-averaged diffusivity according to Offtrac, but we do not know whether κ_H reflects only the path-averaged (along-isopycnal) diffusivity according to Offtrac or whether κ_H accounts for inaccuracy in the HIM circulation fields. The half-variance is estimated by using BMA (not at each grid point, but over the entire spatial domain of each hydrographic section) on the six Offtrac runs with either climatological or hindcast fields and different κ_H. Here, Δ^2 levels off below some depth because running Offtrac for 200 years is not sufficient to capture the entire TTD, even with a log-linear extrapolation, and the peak of the TTD was not reached after 200 years. Nonetheless, the depths we are interested in analyzing are shallower than these depths ($<$ 2000 meters depth). The uncertainties associated with the half-variance from BMA (not shown) are on the same

order of magnitude as the Δ^2 estimates where there are intermediate waters but are much smaller at depths above that (< 1000 m).

We find estimates of κ_H relative to surrounding regions to be consistent with Δ^2 relative to surrounding regions in some locations and inconsistent in other locations from BMA. There is consistency in the relative diffusivity magnitudes where there is SAMW and AAIW along SR03 and I09 (500-1000 and 1000-1200 meters depth in Figs. 6.5c and 6.5f), AAIW and LSW components of North Atlantic Deep Water (NADW) along A05 (500-1000 and 1000-2000 meters depth in Fig. 6.2i), SPMW along A16 (northern end between 100-300 meters depth in Fig. 6.2f), and STMW along A20 (1000 meters depth in Fig. 6.2c), but there is inconsistency in the relative diffusivity magnitudes where there is LSW along A16 (2000 meters depth at 60^oN in Fig. 6.2f). In the Offtrac model (Fig. 6.2i) ULSW appears at deeper depths than in observations (Figs. 6.2g), but the diffusivity magnitudes for ULSW from both the estimated Δ^2 (Fig. 6.2i) and κ_H (Fig. 6.2g) are consistently smaller than surrounding water masses. If the BMA-derived weights were exclusively indicating which density classes have more accurately modeled physical processes than other density classes, then we should see a difference in the estimated κ_H between where both Offtrac (1200-2000 meters depth in Fig. 6.2i) and observations (Fig. 2b of LeBel et al., 2008 [55]) place ULSW along A05 and where only one or the other places ULSW along A05. Because this is not the case (Fig. 6.2g), the question arises: is it possible that 100 % of the differences in κ_H in one density class from another can be explained by inaccuracies with Offtrac as opposed to something physically real? Either the diffusivities we estimate for LSW and the ULSW have some degree of physical reality to them, or the Offtrac model treats LSW as a highly diffusive water mass and ULSW as a water mass with relatively little diffusion sometime between when ULSW and LSW diverge and when they are found along A16.

6.4 Conclusions

It is found that the climatological circulation fields are more likely than the hindcast circulation fields to be accurate (when compared with CFC-11 observations) for the Southern Ocean and both North Atlantic Ocean gyres, and that there is evidence for a density-dependent structure to the along-isopycnal diffusivity. This suggests that, while previous

studies have found that allowing for a κ_H that varies in space can mask the importance of variable circulation, using variable circulation can mask the importance of a κ_H that varies in space as well in a numerical model. The uncertainties in the estimates of the model parameter, κ_H, incorporate how surprising it would be for CFC-11 concentrations to be very different from the model runs, given the observations, considering both within- and between-model differences when κ_H is varied.

Water masses that are ventilated after forming via deep convective processes such as (U)LSW, STMW, SPMW, and SAMW have smaller along-isopycnal diffusivities ($\kappa_H \sim 1 \times 10^3$ m^{2}s^{-1}) than water masses that are formed by turbulent processes that mix distinguishable water masses together such as AAIW and CLSW ($\kappa_H \sim 6 \times 10^3$ m^{2}s^{-1}) and all other water masses ($\kappa_H \sim 3 \times 10^3$ m^{2}s^{-1}). There are several regions where we can estimate a value of κ_H to be distinguishably different from other regions to 95% confidence, like in the Southern Ocean for both SAMW and AAIW along SR03 (500-1000 and 1000-1200 meters depth in Figs. 6.5a-c) and I09 (500-1000 and 1000-1200 meters depth in Fig. 6.5d-f); for LSW, SPMW, and AAIW along A16 in the North Atlantic Ocean (northern end between 1000-2000, northern end between 100-300, and southern end between 500-1000 meters depth in Figs. 6.2d-f); for AAIW and the LSW components of North Atlantic Deep Water (NADW) (both ULSW and CLSW) along A05 in the North Atlantic Ocean (500-1000 and 1000-2000 meters depth in Figs. 6.2g-i); and for STMW along A20 in the North Atlantic Ocean (1000 meters depth in Figs. 6.2a-c).

The water masses that have higher κ_H could reflect the mixing of dynamical tracers, but this is a topic for future investigation. The conditions set by the surrounding water masses could be forcing the interior of a high-diffusivity water mass to adjust in whichever way it can. Since an interior water mass cannot adjust through advection to a greater extent than a converting potential energy into kinetic energy via baroclinic instability would dictate and diapycnal diffusivities are small relative to along-isopycnal diffusivities, the along-isopycnal diffusivities adjust. If this is real, then water masses with lower diffusivities would have more information about surface forcings and the inputs at the surface than water mass with higher diffusivities, which is consistent with what we find here. If the HIM circulation fields for a given density class, such as that of LSW, accurately simulate the real ocean along

one pathway but not another, then it is possible that our estimates of κ_H are not just the path-averaged along-isopycnal diffusivities. It could be the case that κ_H also includes an additional diffusivity that accounts for discrepancies between the HIM circulation fields and the circulation fields of the real ocean.

There are some unresolved issues that need to be addressed before we claim to have demonstrated that the path-averaged along-isopycnal diffusivity of the real ocean is a function of density class. First, a test to see if a model run with circulation fields and κ_H estimated from BMA is weighted more than any of our model runs here via another BMA calculation is left for a future study. A separate test to see if a model run with circulation fields estimated from BMA using observed currents rather than CFC-11 concentrations yields different estimates of κ_H via another BMA calculation is also left for a future study. This is not so straightforward because the circulation fields specified for Offtrac runs from inline runs of HIM conserve mass. However, using another model could also prove to be a problem because of volume conservation. This problem could be approached by using Lagrangian multipliers, minimizing the difference between the circulation fields and observations at locations, depths, and times where observations are available, where not conserving mass or volume would penalize the weighted combination of model runs.

Addressing the interpretation of κ_H could involve an evaluation how well a weighted average of the HIM circulation fields using the m_j from BMA simulates the real ocean's circulation. The differences in the diffusivities between LSW in the northeast Atlantic Ocean and NADW after their paths diverge would essentially need to be differences in circulation patterns that are inadequately represented by the climatological or hindcast HIM circulation fields if our estimates of κ_H do not represent the path-averaged along-isopycnal diffusivities of the real ocean. Presumably, if the circulation fields were this inadequate in simulating the real ocean, we would not see LSW or NADW in approximately the same places as in observational data, but there are multiple paths water parcels can take to get from one location and another. We have found that Offtrac should have a density-dependent κ_H to be more consistent with observations, but interpretation of κ_H as a path-averaged along-isopycnal diffusivity has not been established here.

It is possible that, because our experimental design included such a large diffusivity

Table 6.1: Using a uniform prior, $P(M_j) = 1/6$ (j=1,...,6), the posterior probabilities, $P(M_j|D)$, for the six configurations of along-isopycnal diffusivity, κ_H, and hindcast/climatological fields for both the Southern Ocean (SO) using I09 and SR03, the Subtropical North Atlantic (STNA) using A05 and A20 (south of 30^oN), and the Subpolar North Atlantic (SPNA) using A20 (north of 30^oN) and A16.

climatology/hindcast	κ_H	$P(M_j\|D)$: (SO)	(STNA)	(SPNA)
climatology	5.0×10^2	0.241	0.203	0.246
hindcast	5.0×10^2	0.0924	0.132	0.0862
climatology	1.0×10^3	0.237	0.200	0.244
hindcast	1.0×10^3	0.0987	0.136	0.0924
climatology	1.0×10^4	0.236	0.197	0.242
hindcast	1.0×10^4	0.0952	0.132	0.0894

($\kappa_H = 10^4$ m^{2}s^{-1}), our uncertainties may be conservative, but it is also possible that our uncertainties are underestimates because we have not accounted for the uncertainties associated with all model parameters, for instance. To determine the point-wise spatial structure of κ_H, using information along each water mass' pathway and additional tracers may constrain the uncertainties to a greater extent. We leave this work for future analyses, potentially with the statistical framework that Markov Random Fields [5] provide or by finding a way to include nearby observations but down-weight them relative to the observations at a particular point of interest for an estimate.

Table 6.2: Using a uniform prior, the point-wise posterior probabilities, $P(M_j|D)$, for each of the six configurations of along-isopycnal diffusivity, κ_H, and hindcast/climatological fields averaged across each of the five hydrographic sections in Table 5.1: SR03, I09, A05, A20, and A16.

climatology/hindcast	κ_H	$P(M_j\|D)$: (A20)	(A16)	(A05)	(SR03)	(I09)
climatology	5.0×10^2	0.167	0.154	0.166	0.171	0.161
hindcast	5.0×10^2	0.169	0.135	0.169	0.174	0.171
climatology	1.0×10^3	0.172	0.164	0.162	0.167	0.164
hindcast	1.0×10^3	0.162	0.161	0.174	0.181	0.170
climatology	1.0×10^4	0.168	0.193	0.162	0.150	0.166
hindcast	1.0×10^4	0.162	0.193	0.168	0.157	0.169

Table 6.3: Using a uniform prior, the point-wise posterior probabilities, $P(M_j|D)$, for each of the six configurations of along-isopycnal diffusivity, κ_H, and hindcast/climatological fields averaged over each of the three ocean gyres: Southern Ocean (SO), Subtropical North Atlantic (STNA) Ocean, and Subpolar North Atlantic (SPNA) Ocean.

climatology/hindcast	κ_H	$P(M_j\|D)$: (SO)	(STNA)	(SPNA)
climatology	5.0×10^2	0.165	0.165	0.165
hindcast	5.0×10^2	0.171	0.166	0.166
climatology	1.0×10^3	0.164	0.167	0.167
hindcast	1.0×10^3	0.173	0.163	0.163
climatology	1.0×10^4	0.166	0.166	0.166
hindcast	1.0×10^4	0.169	0.164	0.163

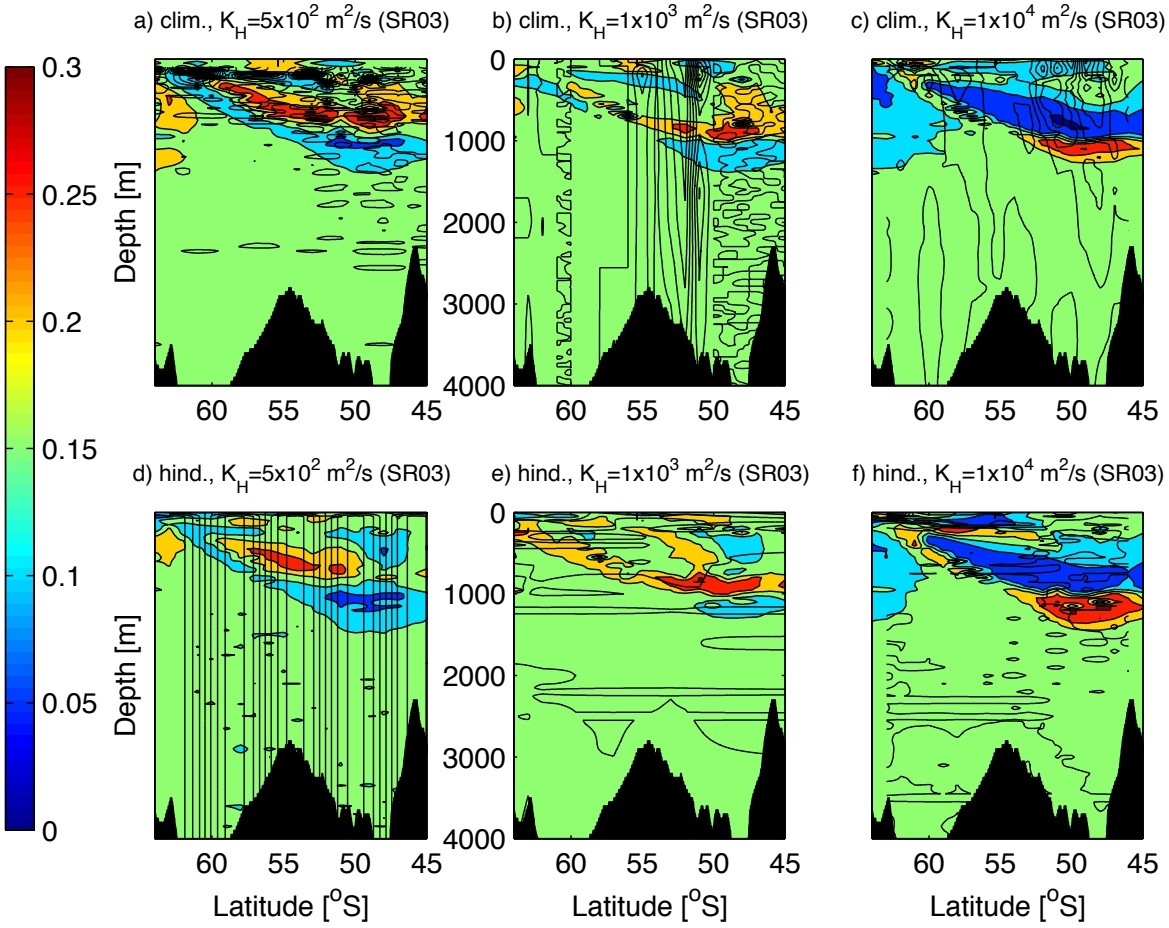

Figure 6.1: The BMA weights, m_j, along SR03 for the climatological HIM circulation fields (a-c) and hindcast HIM circulation fields (d-f) using along-isopycnal diffusivities, κ_H, 5×10^2 m^{2}s^{-1} (a and d), 1×10^3 m^{2}s^{-1} (b and e), and 1×10^4 m^{2}s^{-1} (c and f). Black solid contours indicate the average (over all repeat hydrographic occupations) of the vertical gradient of the large-scale potential vorticity field, $\partial PV/\partial z$, (z is depth) from observations (a); the average of the second vertical derivative, $\partial^2 T/\partial z^2$, of the temperature fields from observations (b); the average of the second vertical derivative, $\partial^2 S/\partial z^2$, of the salinity fields from observations (c); the average of the horizontal gradient of the large-scale potential vorticity field, $\partial PV/\partial \lambda$, ($\lambda$ is the latitude) from observations (d); the average of the second horizontal derivative, $\partial^2 T/\partial \lambda^2$, of the temperature fields from observations (e); and the average of the second horizontal derivative, $\partial^2 S/\partial \lambda^2$, of the salinity fields from observations (f).

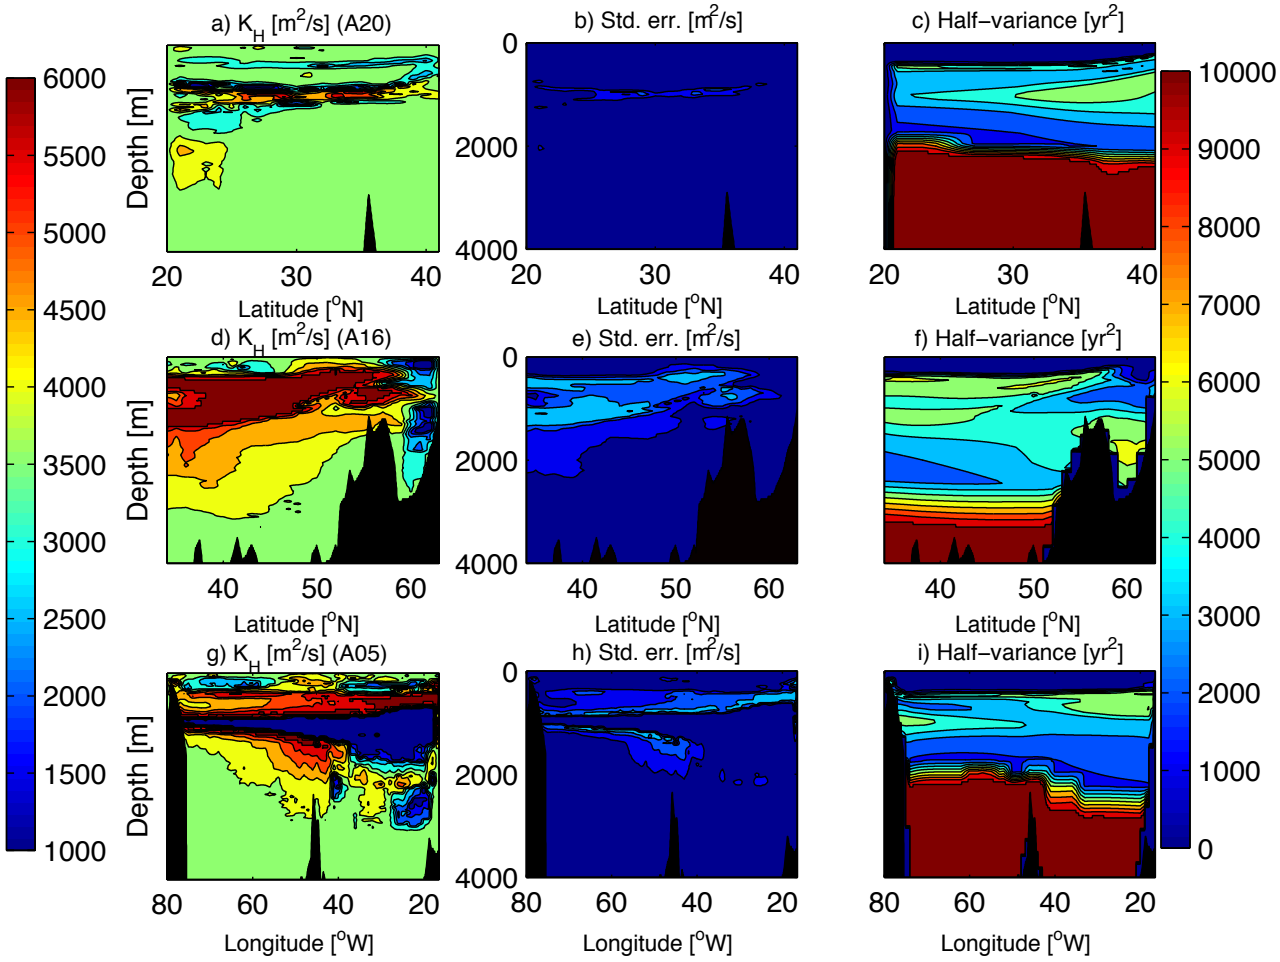

Figure 6.2: The along-isopycnal diffusivities, κ_H, (a, d, and g) and their standard errors (b, e, and h, which when multiplied by 1.96 are the confidence interval half-widths) along A20 (a-c), A16 (d-f), and A05 (g-i) performing BMA at each grid point using CFC-11 concentrations as data for the likelihood with colorbar shown on the left. Also shown are the TTD-derived half-variances, Δ^2, (c, f, and i) with colorbar shown on the right.

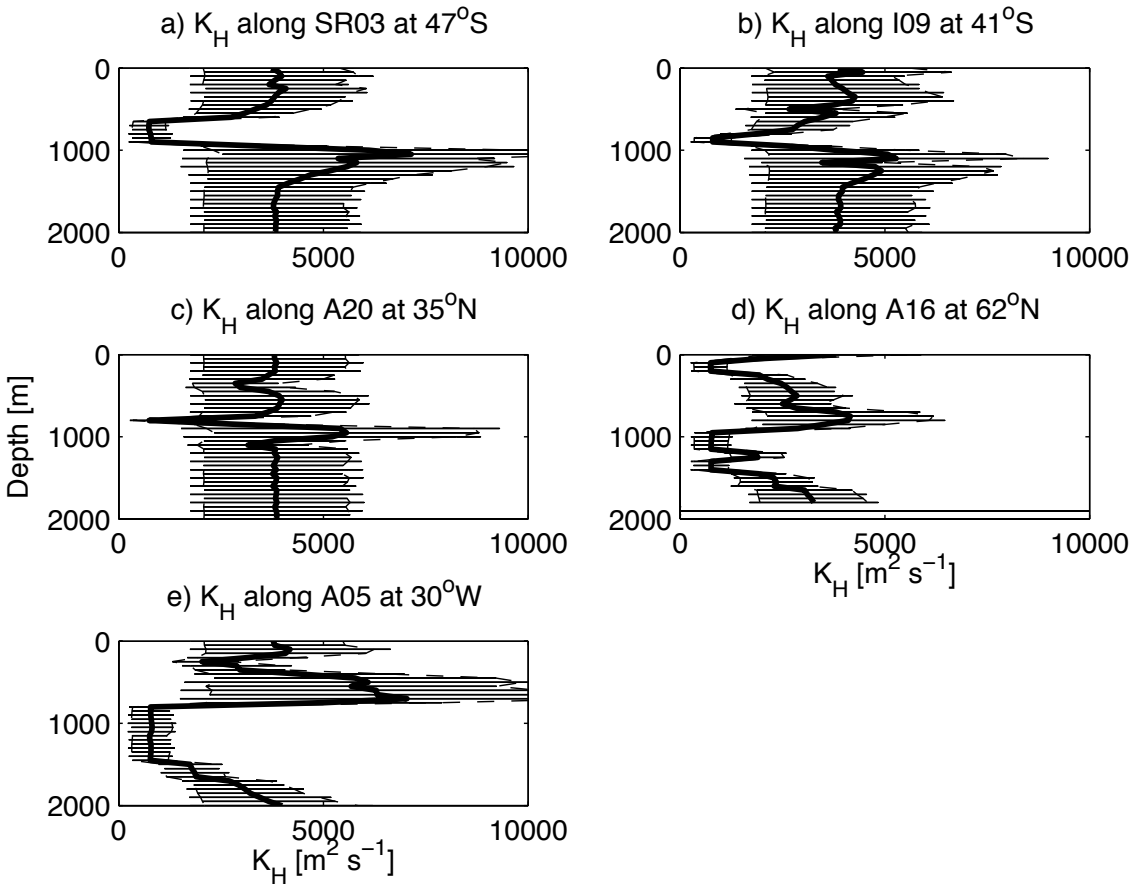

Figure 6.3: The along-isopycnal diffusivities, κ_H, (thick solid lines) versus depth [m] along SR03 at 47^{o}S (a), I09 at 41^{o}S (b), A20 at 35^{o}N (c), A16 at 62^{o}N (d), and A05 at 30^{o}W (e) and their 90% (shaded interior regions) 95% (extra length indicated by lines outward) confidence intervals.

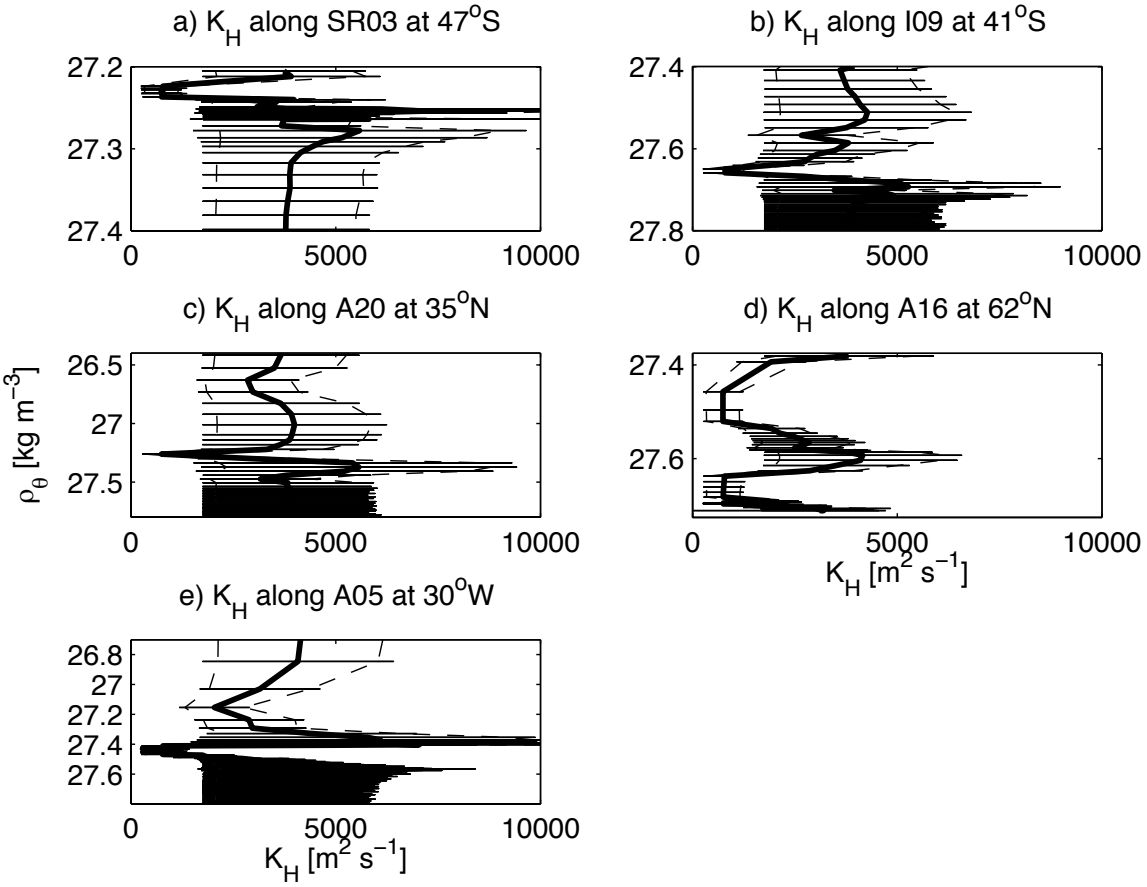

Figure 6.4: The along-isopycnal diffusivities, κ_H, (thick solid lines) versus potential density [kg m^3] along SR03 at 47^{o}S (a), I09 at 41^{o}S (b), A20 at 35^{o}N (c), A16 at 62^{o}N (d), and A05 at 30^{o}W (e) and their 90% (shaded interior regions) 95% (extra length indicated by lines outward) confidence intervals.

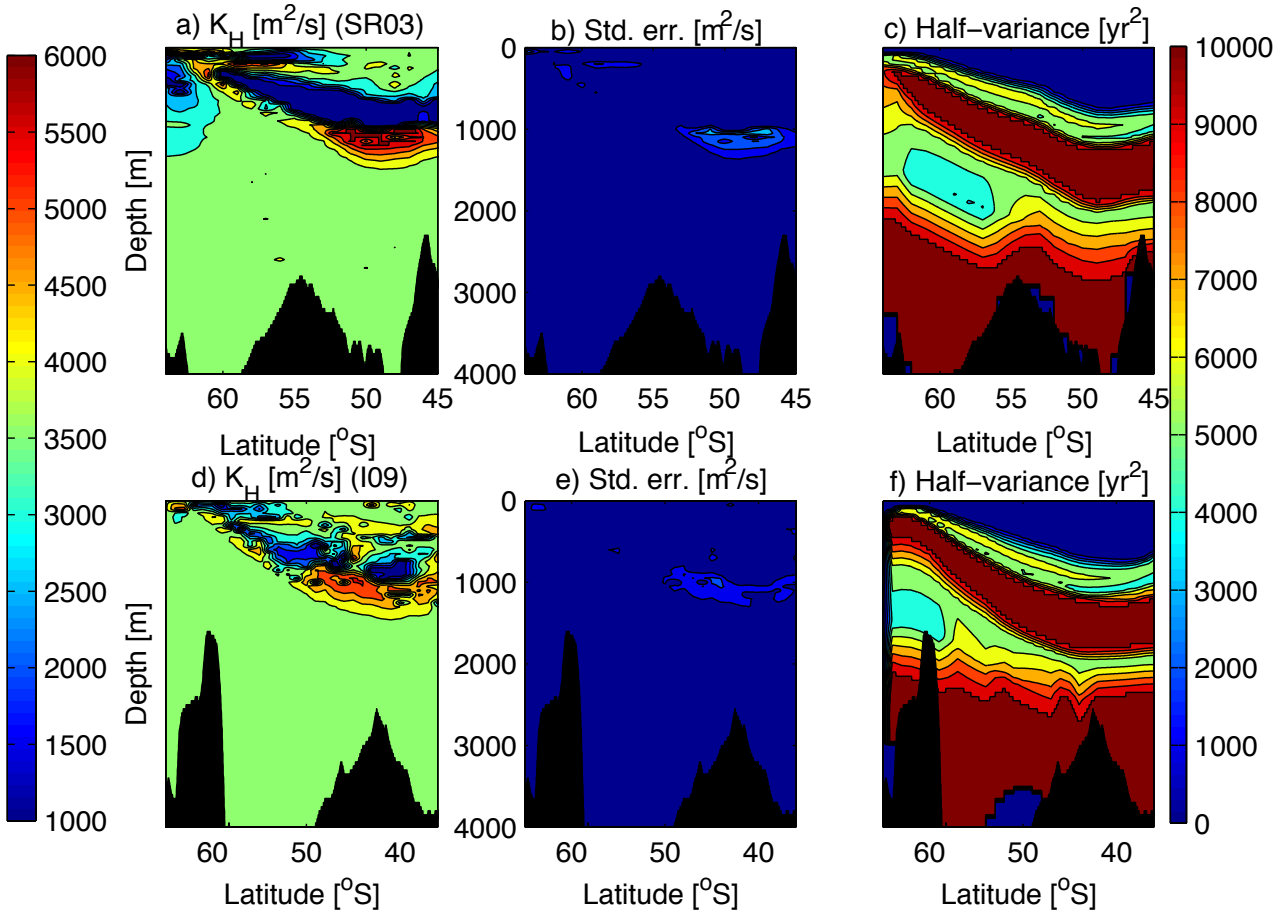

Figure 6.5: The along-isopycnal diffusivities, κ_H, (a and d) and their standard errors (b and e, which when multiplied by 1.96 are the confidence interval half-widths) along SR03 (a-c) and I09 (d-f) performing BMA at each grid point using CFC-11 concentrations as data for the likelihood with colorbar shown on the left. Also shown are the TTD-derived half-variances, Δ^2, (c and f) with colorbar shown on the right.

Chapter 7

CONCLUSIONS

This thesis has developed ways to apply statistical methods such as TPS, MCMC, and BMA to oceanographic contexts. The contexts presented here have primarily been concerned with how tracers that enter the ocean at its surface get to the oceanic interior. We have estimated mapping errors (Chapter 2) and used them in application of MCMC to estimate TTDs (Chapter 3) and model weights (Chapters 5-6). TTDs hold information about how long it has been since a water parcel was at the surface, the ratio of advection to diffusion in the process of getting to its current location, and the mass flux between well-defined volumes within a particular density class. Relating a mass flux to a ventilation rate, we have found a comparable magnitude to a ventilation rate estimated more directly from information about MLDs and diapycnal velocities (Chapter 5). We have used model weights to estimate uncertainties on these ventilation rates (Chapter 5) as well as along-isopycnal diffusivities (Chapter 6).

As demonstrated in Chapter 2, the TPS standard errors are large where the data coverage is relatively poor, but do not necessarily indicate where smoothing has removed small-scale or boundary extrema. These errors reflect the attempt to estimate a quantity using combinations of covariates outside of their domains used in training the statistical model. Overall, when data is available and not extremely sparse, using estimates from OM and imputing the TPS-derived standard errors for a guide on how certain we are about oceanic tracer field estimates are recommended.

As demonstrated in Chapter 3, based on CLIVAR Repeat Hydrography data and error estimates from Chapter 2, TTDs were estimated to be different from IGs. It was inferred that young (< 5 years) AAIW is likely comprised of the SAMW that is entrained and detrained in the southeast Indian Ocean upstream in the ACC relative to SR03. The older, more sluggish AAIW is likely to have come from the Tasman Sea recirculation. Some of

the young, quick LSW could have come from the Labrador Basin, but the very young LSW ($< 2 - 3$ years) must have come from the Irminger Basin. More dense, sluggish LSW corresponds to a positive NAO index and less dense, faster-spreading LSW corresponds to a negative NAO index. It takes at least a hundred years on average for water to make its way from the surface southward to 24.5^oN to form the components of NADW denser than the LSW component.

Ventilation rates from profilers and satellites during the 2002-2006 period were estimated in Chapter 4 for all water and STMW and ventilation rates from using CFC-11 observations during the 1959-2006 period were estimated in Chapter 5 with BMA using a numerical model runs for STMW, SPMW, and SAMW. A trend was not apparent in the five year time series directly from observations, considering the 50% error we found by propagating error assuming a covariance function that only depends on distance rather than location with objective mapping. An increase in the SAMW subduction rates, and in particular, an increase in the contribution due to mixing at isopycnal outcrops, was suggested over the 1959-2006 model time interval if all dominating sources of uncertainty are encapsulated in the background along-isopycnal diffusivities and circulation fields of our model runs. An increase in the stratification of the waters in the South Indian Ocean and/or changes in the winds due to changes in the SAM are possible causes. Further investigation is needed to tell whether this trend pertains to the real ocean and what causes this trend. A steady, nearly significant difference between the STMW subduction rate time series from 1959-2006 using two different estimation techniques was found with no significant trend. No significant difference between the two SPMW subduction rate time series estimates over the same interval was found with a nearly significant trend. These North Atlantic results suggest that diffusion plays more of a role in STMW subduction than SPMW subduction, and whether the winds as indicated by the NAO index effect SPMW subduction differently from STMW subduction requires further investigation.

As demonstrated in Chapter 4, using the 2002-2006 time series with only observations, wind stresses were found to contribute about $10 - 20$ % of the Ekman divergences and Sverdrup transports, but vertical pumping was found to be increasingly negligible relative to lateral induction farther away from the coast. The lightest waters subduct at the coast

where vertical velocities can dominate MLD variability, but MLD variability matters most for ventilation elsewhere. A residence time of at least 3.5 years was found for STMW with the shorter time series with only observations (Chapter 4) and mean residence times of $11.3 - 12.5$ years for STMW, $2.70 - 4.11$ years for SPMW, and $14.0 - 15.0$ years for SAMW were found with BMA using numerical model runs and CFC-11 observations (Chapter 5). Each of these ventilated mode waters were inferred to be eroded by processes with at least two distinguishable time scales.

As suggested in Chapter 6, allowing for a density-dependent κ_H can mask the importance of more variable circulation fields, but using more variable circulation fields can mask the importance of a density-dependent κ_H as well. Water masses that are ventilated after forming via deep convective processes have smaller along-isopycnal diffusivities ($\sim 10^3$ m s^{-1}) than water masses that are formed by turbulent processes that mix distinguishable water masses together ($\sim 5 - 6 \times 10^3$ m s^{-1}). The density classes that have higher κ_H appear at the steering level of baroclinic instabilities associated with interior potential vorticity gradients in the Southern Ocean. The suggested density-dependence of path-averaged along-isopycnal diffusivity here warrants future investigation.

Bayesian statistical analysis provides a way to perform model evaluation of physically interpretable quantities. While an isopycnal-coordinate non-eddy-resolving numerical model needs larger along-isopycnal diffusivities to be consistent with transient tracer observations at intermediate depths in the North Atlantic Ocean (Chapter 6), a depth-coordinate eddy-resolving numerical model has path-averaged along-isopycnal diffusivities that tend to be too large to be consistent with transient tracer observations at intermediate depths in the North Atlantic Ocean (Chapter 3). However, both the isopycnal-coordinate non-eddy-resolving numerical model with $\kappa_H = 10^3$ m^2 s^{-1} and depth-coordinate eddy-resolving numerical model have path-averaged along-isopycnal diffusivities that are too large for LSW in the North Atlantic Ocean to be consistent with transient tracer observations. The physical validity of the findings presented in this thesis will be subject to further scrutiny, but the methods presented here are readily available to improve the ability of models to simulate the real ocean.

BIBLIOGRAPHY

[1] R. Abernathey, J. Marshall, M. Mazloff, and E. Shuchburgh. Enhancement of mesoscale eddy stirring at steering levels in the southern ocean. *J. Phys. Ocean.*, 40:170–184, 2009.

[2] S. L. Barnes. A technique for maximizing details in numerical weather map analysis. *J. Appl. Meteor.*, 3:396–409, 1964.

[3] R. Bayes and R. Price. An essay towards solving a problem in the doctrine of chance. by the late rev. mr. bayes, communicated by mr. price, in a letter to john canton, m. a. and f. r. s. *Philosophical Transactions of the Royal Society of London*, 53:370–418, 1763.

[4] M. Bersch. On the circulation of the northeastern north atlantic. *Deep Sea Res. I*, 42:1583–1607, 1995.

[5] J. Besag. On the statistical analysis of dirty pictures. *J. Royal Stat. Soc. Series B (Methodological)*, 48(3):259–302, 1986.

[6] E. Brambilla, L. D. Talley, and P. E. Robbins. Subpolar mode water in the northeastern atlantic: 2. origin and transformation. *J. Geophys. Res.*, 113:C04026, 2008.

[7] J.-M. Brankart and P. Brasseur. Optimal analysis of in situ data in the western mediterranean using statistics and cross-validation. *J. Atmos. Ocean. Tech*, 13(2):477–491, 1996.

[8] F. P. Bretherton, R. E. Davis, and C. B. Fandry. A technique for objective analysis and design of oceanographic experiments applied to mode-73. *Deep Sea Res.*, 23:559–582, 1976.

[9] R. C. H. Cheng and N. A. K. Amin. Maximum likelihood estimation of parameters in the inverse gaussian distribution with unknown origin. *Technometrics*, 23(3):257–263, 1981.

[10] P. C. Chu, L. M. Ivanov, T. P. Korzhova, T. M. Margolina, and O. V. Melnichenko. Analysis of sparse and noisy ocean current data using flow decomposition. part i: Theory. *J. Atmos. Ocean. Tech.*, 20:478–491, 2003.

[11] R. G. Curry and M. S. McCartney. Ocean gyre circulation changes associated with the north atlantic oscillation. *J. Phys. Ocean.*, 31:3374–3400, 2001.

[12] R. E. Davis. Preliminary results from directly measuring middepth circulation in the tropical and south pacific. *J. Geophys. Res.*, 103(C11):24619–24639, 1998.

[13] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *J. Royal Stat. Soc. B (Meth.)*, 39(1):1–38, 1977.

[14] K. L. Denman and H. J. Freeland. Correlation scales, objective mapping and a statistical test of geostrophy over the continental shelf. *J. Marine Res.*, 43:517–539, 1985.

[15] C. Deutsch, H. Brix, T. Ito, H. Frenzel, and L. Thompson. Climate forcing of ocean hypoxia. *Science*, page submitted, 2011.

[16] R. Dickson, J. Lazier, J. Meincke, P. Phines, and J. Swift. Long-term coordinated changes in the convective activity of the north atlantic 1968-1982. *Prog. Ocean.*, 38:241–295, 1996.

[17] S. C. Doney and M. W. Hecht. Antarctic bottom water formation and deep-water chlorofluorocarbon distributions in a global climate model. *J. Phys. Ocean.*, 32:1642–1666, 2002.

[18] S. C. Doney and W. J. Jenkins. The effect of boundary conditions on tracer estimates of thermocline ventilation rates. *J. Marine Res.*, 46:947–965, 1988.

[19] S. C. Doney, W. J. Jenkins, and J. L. Bullister. A comparison of ocean tracer dating technique on a meridional section in the eastern north atlantic. *Deep Sea Res. I*, 44(4):603–626, 1997.

[20] S. C. Doney, K. Lindsay, K. Caldeira, J.-M. Campin, H. Drange, J.-C. Dutay, M. Follows, Y. Gao, A. Gnanadesikan, N. Gruber, A. Ishida, F. Joos, G. Madec, E. Maier-Reimer, J. C. Marshall, R. J. Matear, P. Monfray, A. Mouchet, R. Najjar, J. C. Orr, G.-K. Plattner, J. Sarmiento, R. Schlitzer, R. Slater, I. J. Totterdell, M.-F. Weirig, Y. Yamanaka, and A. Yool. Evaluating global ocean carbon models: The importance of realistic physics. *Global Biogeochem. Cycles*, 18:GB3017, 2004.

[21] S. Dong, S. L. Hautala, and K. A. Kelly. Interannual variations in upper ocean heat content and heat transport convergence in the western north atlantic. *J. Phys. Ocean.*, 37(11):2682–2697, 2007.

[22] S. Dong and K. A. Kelly. Heat budget in the gulf stream region: The importance of heat storage and advection. *J. Phys. Ocean.*, 34(5):1214–1231, 2004.

[23] E. Dreisigacker and W. Roether. Tritium and ^{90}sr in north atlantic surface water. *Earth and Planet. Sci. Lett.*, 38:301–312, 1978.

[24] J. L. Folks and R. S. Chhikara. The inverse gaussian distribution and its statistical application - a review. *J. R. Statist. Soc. B*, 40(3):263–289, 1978.

[25] G. Fuchs, W. Roether, and P. Schlosser. Excess ^{3}he in the ocean surface layer. *J. Geophys. Res.*, 92:2599–2611, 1987.

[26] P. J. Glecker, K. E. Taylor, and C. Doutriaux. Performance metrics for climate models. *J. Geophys. Res.*, 113:D06104, 2008.

[27] M. Goes, N. M. Urban, R. Tonkonojenkov, M. Haran, A. Schmittner, and K. Keller. What is the skill of ocean tracers in reducing uncertainties about ocean diapycnal mixing and projections of the atlantic meridional overturning circulation? *J. Geophys. Res.*, 115:C12006, 2010.

[28] A. L. Gordon. Interocean exchange of thermocline water. *J. Geophys. Res.*, 91:5037–5046, 1986.

[29] V. V. Gouretski and K. P. Kolterman. Woce global hydrographic climatology. *A Technical Report, Berichte des Bundesamtes fr Seeschiffahrt un Hydrographi*, 35, 2004.

[30] T. W. N. Haine, H. Zhang, D. W. Waugh, and M. Holzer. On transit-time distributions in unsteady circulation models. *Ocean Modelling*, 21:35–45, 2008.

[31] T. M. Hall, T. W. N. Haine, M. Holzer, D. A. LeBel, F. Terenzi, and D. W. Waugh. Ventilation rates estimated from tracers in the presence of mixing. *J. Phys. Ocean.*, 37:2599–2611, 2007.

[32] R. Hallberg. *Some Aspects of the Circulation in Ocean Basins with Isopycnals Intersecting the Sloping Boundaries.* University of Washington-Seattle, PhD Thesis, 1995.

[33] D. E. Harrison. On climatological monthly mean wind stress and wind stress curl fields over the world ocean. *J. Climate*, 2:57–70, 1989.

[34] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning.* Springer Verlag, 2001.

[35] S. Hellerman and M. Rosenstein. Normal monthly wind stress over the world ocean with error estimates. *J. Phys. Ocean.*, 13:1093–1104, 1983.

[36] L. Herraiz-Borreguero and S. R. Rintoul. Subantarctic mode water variability influenced by mesoscale eddies south of tasmania. *J. Geophys. Res.*, 115:C04004, 2010.

[37] J. R. Holton. *An Introduction to Dynamic Meteorology, 4th Edition.* Academic Press, 2004.

[38] M. Holzer and T. M. Hall. Transit-time and tracer-age distributions in geophysical flows. *J. Atmos. Sci.*, 57:3539–3558, 2000.

[39] M. Holzer, F. W. Primeau, W. M Smethie Jr., and S. Khatiwala. Where and how long ago was water in the western north atlantic ventilated? maximum entropy inversions of bottle data from woce line a20. *J. Geophys. Res.*, 115:C07005, 2010.

[40] H. Inoue. A least-squares smooth fitting for irregularly spaced data: Finite-element approach using the cubic. b-spline basis. *Geophysics*, 51:2051–2066, 1986.

[41] C. O. D. Iselin. The influence of vertical and lateral turbulence on the characteristics of the waters at mid-depths. *Trans. Amer. Geophys. Union*, 28(20):1–26, 1939.

[42] E. T. Jaynes. *Maximum-Entropy and Bayesian Methods in Science and Engineering: The Relation of Bayesian and Maximum Entropy Methods*, volume 1. Kluwer Academic Publishers, 1988.

[43] W. J. Jenkins. Tritium and he-3 in the beta triangle and observation of gyre ventilation and oxygen utilization rates. *J. Phys. Ocean.*, 17:763–783, 1987.

[44] G. C. Johnson, J. L. Bullister, and N. Gruber. Labrador sea water property variations in the northeastern atlantic ocean. *Geophys. Res. Lett.*, 32:L07602, 2005.

[45] T. M. Joyce, C. Deser, and M. A. Spall. The relation between decadal variability of subtropical mode water and the north atlantic oscillation. *J. Climate*, 13:2550–2569, 2000.

[46] R. H. Karsten and J. C. Marshall. Constructing the residual circulation of the acc from observations. *J. Phys. Ocean.*, 32:3315–3327, 2002.

[47] K. A. Kelly, S. Dickinson, and Z.-J. Yu. Nscat tropical wind stress maps: Implications for improving ocean modeling. *J. Geophys. Res.*, 104(11):11291–11310, 1999.

[48] S. Khatiwala, F. Primeau, and T. M. Hall. Reconstruction of the history of anthropogenic co_2 concentrations in the ocean. *Lett. Nature*, 462:346–349, 2009.

[49] S. Khatiwala, M. Visbeck, and P. Schlosser. Age tracers in an ocean gcm. *Deep Sea Res. I*, 48:1423–1441, 2001.

[50] Y.-O. Kwon. *Observation of General Circulation and Water Mass Variability in the North Atlantic Subtropical Mode Water Region.* University of Washington-Seattle, PhD Thesis, 2003.

[51] Y.-O. Kwon and S. Riser. North atlantic subtropical mode water: A history of ocean-atmosphere interaction 1961-2000. *Geophys. Res. Lett.*, 31(10):L19307, 2004.

[52] C. Ladd and L. Thompson. Water mass formation in an isopycnal model of the north pacific. *J. Phys. Ocean.*, 31:1517–1537, 2001.

[53] W. G. Large, J. C. McWilliams, and S. C. Doney. Oceanic vertical mixinga review and a model with a nonlocal boundary layer parameterization. *Rev. Geophys.*, 32:363–403, 1994.

[54] W. G. Large and S. G. Yeager. Diurnal to decadal global forcing for ocean and sea-ice models: the datasets and flux climatologies. *NCAR Technical Note TN-460+STR*, 2004.

[55] D. A. LeBel, W. M. Smethie Jr., M. Rhein, D. Kieke, R. A. Fine, J. L. Bullister, D.-H. Min, W. Roether, R. F. Weiss, C. Andrie, D. Smythe-Wright, and E. P. Jones. The formation rate of north atlantic deep water and eighteen degree water calculated from cfc-11 inventories observed during woce. *Deep Sea Res. I*, 55:891–910, 2008.

[56] J. R. Ledwell, A. J. Watson, and C. S. Law. Evidence for slow mixing across the pycnocline from an open-ocean tracer-release experiment. *Nature*, 364:701–703, 1993.

[57] J. Luyten, J. Pedlosky, and H. Stommel. The ventilated thermocline. *J. Phys. Ocean.*, 13:292–309, 1983.

[58] M. Maltrud, F. Bryan, and S. Peacock. Boundary impulse response functions in a century-long eddying global ocean simulation. *Env. Fluid. Mech.*, 10:275–295, 2010.

[59] M. E. Maltrud and J. L. McClean. An eddy resolving global 1/10 degree ocean simulation. *Ocean Modelling*, 8:31–54, 2005.

[60] J. C. Marshall, A. J. G. Nurser, and R. G. Williams. Inferring the subduction rate and period over the north atlantic. *J. Phys. Ocean.*, 23(7):1315–1329, 1990.

[61] M. McCartney. The subtropical recirculation of mode waters. *J. Mar. Res.*, 40:427–464, 1982.

[62] M. S. McCartney and L. D. Talley. The subpolar mode water of the north atlantic ocean. *J. Phys. Ocean.*, 12:1169–1188, 1982.

[63] P. McIntosh. Oceanographic data interpolation: objective analysis and spline. *J. Geophys. Res.*, 95:13529–13541, 1986.

[64] S. Mecking, M. J. Warner, C. E. Greene, S. L. Hautala, and R. E. Sonnerup. The effects of mixing on cfc concentrations and cfc-derived age distributions in the north pacific thermocline. *J. Geophys. Res.*, 109:C07014, 2004.

[65] N. Metropolis and S. Ulam. The monte carlo method. *J. Amer. Statist. Assoc.*, 44(247):335–341, 1949.

[66] D. L. Musgrave. A numerical study of the roles of subgyre-scale mixing and the western boundary current on homogenization of a passive tracer. *J. Geophys. Res.*, 90(C4):7037–7043, 1985.

[67] D. L. Musgrave. Numerical studies of tritium and helium-3 in the thermocline. *J. Phys. Ocean.*, 20:344–373, 1990.

[68] N. Nakamura. Two-dimensional mixing, edge formation, and permeability diagnosed in an area coordinate. *J. Atmos. Sci.*, 53:1524–1537, 1996.

[69] A. J. G. Nurser, R. Marsh, and J. C. Marshall. Diagnosing water mass formation from air-sea fluxes and surface mixing. *J. Phys. Ocean.*, 29:1468–1487, 1999.

[70] D. Nychka. Bayesian confidence intervals for smoothing splines. *Journal of the American Statistical Association*, 83:1134–1143, 1988.

[71] Intergovernmental Panel on Climate Change. Climate change 2007 - the physical science basis: Contribution of working group i to the fourth assessment report of the ipcc. *Cambridge University Press*, 4, 2007.

[72] A. H. Orsi, T. Whitworth III, and W. D. Nowlin Jr. On the meridional extent and fronts of the antarctic circumpolar current. *Deep Sea Res. I*, 42(5):641–673, 1995.

[73] J. B. Palter, M. S. Lozier, and R. T. Barber. The effect of advection on the nutrient reservoir in the north atlantic subtropical gyre. *Nature*, 437:687–692, 2005.

[74] C. Piani, D. J. Frame, D. A. Stainforth, and M. R. Allen. Constraints on climate change from a multi-thousand member ensemble of simulations. *Geophys. Res. Lett.*, 32:L23825, 2005.

[75] R. S. Pickart, F. Straneo, and G. W. K. Moore. Is labrador sea water formed in the irminger basin? *Deep Sea Res. I*, 50:23–52, 2003.

[76] F. W. Primeau and M. Holzer. The ocean's memory of the atmosphere: Residence-time distributions and water-mass ventilation. *J. Phys. Ocean.*, 36:1439–1456, 2006.

[77] B. Qiu and R. X. Huang. Ventilation of the north atlantic and north pacific: Subduction versus obduction. *J. Phys. Ocean.*, 25(10):2374–2390, 1995.

[78] A. E. Raftery, F. Balabdaoui, T. Gneiting, and M. Polakowski. Using bayesian model averaging to calibrate forecast ensembles. *Technical Report, University of Washington-Seattle*, 440, 2003.

[79] J. F. Read. Convex-91: water masses and circulation of the northeast atlantic subpolargyre. *Prog. Ocean.*, 2:182–189, 1955.

[80] F. A. Richards and A. C. Redfield. Oxygen-density relationships in the western north atlantic ocean. *Deep Sea Res.*, 2:182–189, 1955.

[81] K. R. Ridgway and J. R. Dunn. Mesoscale structure of the mean east australian current system and its relationship with topography. *Prog. Ocean.*, 56(2):189–222, 2003.

[82] S. R. Rintoul. South atlantic interbasin exchange. *J. Geophys. Res.*, 96:2675–2692, 1991.

[83] S. R. Rintoul and J. L. Bullister. A late winter hydrographic section from tasmania to antarctica. *Deep Sea Res.*, 46:1417–1454, 1999.

[84] M. Rixen, J. M. Beckers, J.-M. Brankart, and P. Brasseur. A numerically efficient data analysis method with error map generation. *Ocean Modelling*, 2:45–60, 2000.

[85] D. Roemmich. Optimal estimation of hydrographic station data and derived fields. *J. Phys. Ocean.*, 13:1544–1549, 1983.

[86] J.-B. Sallee, K. Speer, R. Morrow, and R. Lumpkin. An estimate of lagrangian eddy statistics and diffusion in the mixed layer of the southern ocean. *J. Marine Res.*, 66:441–463, 2008.

[87] J.-B. Sallee, K. Speer, and S. R. Rintoul. Zonally asymmetric response of the southern ocean mixed-layer depth to the southern annular mode. *Nature Geoscience*, 3:273–279, 2010.

[88] J.-B. Sallee, N. Wienders, K. Speer, and R. Morrow. Formation of subantarctic mode water in the southeastern indian ocean. *Ocean Dynamics*, 56:525–542, 2006.

[89] P. Schlosser, M. Stute, H. Dorr, C. Sonntag, and K. O. Munnich. Tritium/^{3}he dating of shallow groundwater. *Earth and Planet. Sci. Lett.*, 89:353–362, 1988.

[90] R. Seager, Y. Kushnir, M. Visbeck, N. Naik, J. Miller, G. Krahmann, and H. Cullen. Causes of atlantic ocean climate variability between 1958 and 1998. *J. Climate*, 13:2845–2862, 2000.

[91] E. Shuckburgh, G. Maze, D. Ferreira, J. Marshall, H. Jones, and C. Hill. Mixed layer lateral eddy fluxes mediated by air-sea interaction. *J. Phys. Ocean.*, 41:130–144, 2011.

[92] S. Singh and K. A. Kelly. Monthly maps of sea surface height in the north atlantic and zonal indices for the gulf stream using topex/poseidon altimeter data. *Technical Report, WHOI*, 1997.

[93] B. M. Sloyan and S. R. Rintoul. Circulation, renewal, and modification of antarctic mode and intermediate water. *J. Phys. Ocean.*, 31:1005–1030, 2001.

[94] B. M. Sloyan, L. D. Talley, T. K. Chereskin, R. Fine, and J. Holte. Antarctic intermediate water and subantarctic mode water formation in the southeast pacific: The role of turbulent mixing. *J. Phys. Ocean.*, 40:1558–1574, 2010.

[95] K. S. Smith. The geography of linear baroclinic instability in earths oceans. *J. Mar. Res.*, 65:655–683, 2007.

[96] K. S. Smith and J. Marshall. Evidence for enhanced eddy mixing at middepth in the southern ocean. *J. Phys. Ocean.*, 39:50–69, 2009.

[97] G. W. Snedecor and W. G. Cochran. *Statistical Methods (Eighth Edition)*. Iowa State University Press, 1989.

[98] S. Sokolov and S. R. Rintoul. Some remarks on interpolation of nonstationary oceanographic fields. *J. Atmos. Ocean. Tech.*, 16:1434–1449, 1999.

[99] K. Speer and E. Tziperman. Rates of water mass formation in the north atlantic ocean. *J. Phys. Ocean.*, 25(1):93–104, 1992.

[100] S. Speich, B. Blanke, P. de Vries, S. Drijfhout, and K. Doos. Tasman leakage: A new route in the global ocean conveyor belt. *Geophys. Res. Lett.*, 29(10):1416, 2002.

[101] H. Stommel. Determination of water mass properties of water pumped down from the ekman layer to the geostrophic flow below. *Collected Works of Henry Stommel*, 2(1):207–211, 1979.

[102] A. Sy, M. Rhein, J. R. N. Lazier, K. P. Koltermann, J. Meincke, A. Putzka, and M. Bersch. Surprisingly rapid spreading of newly formed intermediate waters across the north atlantic ocean. *Nature*, 386:675–679, 1997.

[103] M. Tomczak and S. Liefrink. Interannual variations of water mass volumes in the southern ocean. *J. Atmos. and Ocean Sci.*, 10(1):31–42, 2005.

[104] M. Tomczak and S. Liefrink. nterannual variations of water mass properties and volumes in the southern ocean. *Ocean Sci. Discuss.*, 3:199–219, 2005.

[105] P. Y. Le Traon. A method for optimal analysis of fields with spatially variable mean. *J. Geophys. Res.*, 93(8):13543–13547, 1990.

[106] D. S. Trossman, L. Thompson, K. A. Kelly, and Y.-O. Kwon. Estimates of north atlantic ventilation and mode water formation for winters 2002-06. *J. Phys. Ocean.*, 39(10):2600–2617, 2009.

[107] H. Wackernagel. *Multivariate Geostatistics (Third Edition)*. Springer-Verlag, 2003.

[108] G. Wahba and J. Wendelberger. Some new mathematical methods for variational objective analysis using splines and cross validation. *Monthly Weather Review*, 108:1122–1143, 1980.

[109] G. Walin. On the relation between sea-surface heat flow and thermal circulation in the ocean. *Tellus*, 34:187–195, 1982.

[110] S. J. Walker, R. F. Weiss, and P. K. Salameh. Reconstructed histories of the annual mean atmospheric mole fractions for the halocarbons cfc-11, cfc-12, cfc-113 and carbon tetrachloride. *J. Geophys. Res.*, 105(C6):14285–14296, 2000.

[111] M. J. Warner and R. F. Weiss. Solubilities of chlorofluorocarbons 11 and 12 in water and sea water. *Deep Sea Res. A*, 32:1485–1497, 1985.

[112] D. W. Waugh, T. W. N. Haine, and T. M. Hall. Transport time and anthropogenic carbon in the subpolar north atlantic ocean. *Deep Sea Res.*, 51:1475–1491, 2004.

[113] D. W. Waugh, T. M. Hall, and T. W. N. Haine. Relationships among tracer ages. *J. Geophys. Res.*, 108(5):C001325, 2003.

[114] R. F. Weiss. Solubility of helium and neon in water and seawater. *J. Chemical and Engineering Data*, 16(2):235–241, 1971.

[115] R. G. Williams. Ocean subduction. *Science*, 109(10):1982–1983, 2001.

[116] S. N. Wood. Thin plate regression splines. *J. R. Statist. Soc. B*, 65(1):95–114, 2003.

[117] S. N. Wood. Confidence intervals for generalized additive models based on penalized regression splines. *Aust. N. Z. J. Statist.*, 48(4):445–464, 2006.

[118] S. N. Wood. *Generalized Additive Models: An Introduction with R.* Chapman and Hall/CRC Press, 2006.

[119] L. Worthington. The 18^o water in the sargasso sea. *Deep Sea Res.*, 5:297–305, 1959.

[120] S.-P. Xie. Satellite observations of cool ocean-atmosphere interaction. *Bull. Amer. Meteor. Soc.*, 85(2):195–208, 2004.

Appendix A

APPENDICES

A.1 Appendix for Chapter 2

The R language is available from http://cran.r-project.org/ and the two additional downloads required to use the thin-plate spline with possible mixed model random effects are the `mgcv` and `nlme` libraries. Let the one-density index lagged autocorrelation for each tracer of quantity of interest be u. As an example of implementation, to estimate potential density, temperature, or salinity (y), with latitude (x), depth (z), and scaled distance from the equator (d), ($y = f_{te}(x, z) + f_s(d)$) we use the command:

```
gam.out<-gamm(y~te(x,z)+s(d),correlation=corAR1(value=u))
```

Within this call, the first argument is equivalent to the specification in (2.19) using R functions, `te`, which specifies a tensor product function, and `s`, which specifies a smoothing spline (or smoother) function. To put the TPS estimates on a grid, we reference the GAM object, `gam.out$gam`, and use the `predict.gam` function to evaluate our fits at the desired grid points (specified with `newdata` =). The `gamm` function can be slow and is not recommended when one is merely making an estimate and approximation of its error for large ($\gg$ 1000 data points) amounts of data. In such a case, one should use either use the `bam` or `gam` function:

```
bam(y~te(x,z)+s(d),method="ML")
gam(y~te(x,z)+s(d),method="REML")
```

If only the mapped field estimates are needed, then using the above call to `bam` would be optimal (at the expense of potentially biased standard errors), but if a first-order estimate of the standard errors is required, assuming no correlation between different grid points, then using the above call to `gam` would be optimal.

If both values of the field and its standard errors are important, then there is no substitute for the `gamm` function, but a call to a mixed model function (namely, `lme`) by this function with too many data points can fail to converge (using a penalized iteratively-reweighted least squares method). A couple of tactics exist for this situation, one of which is to alter the model. For example, a smoothed mapped field may not be desirable, so adding the option, $\texttt{sp} = \texttt{c}(\texttt{rep}(-99999, \texttt{n}))$, with n being the number of covariates in the model (not distinct covariates, but the number of times covariates are referred to in the functional expression), will force the smoothing parameters to be as small as possible (log of 0 is ≈ -99999). Another example of a way to alter the model in a way that may allow the mixed model subroutine called by the `gamm` subroutine to converge is to specify the number of knots to be used. One such call would be:

```
gamm(y~te(x,z,k=4)+s(d,k=10))
```

where four knots (the minimum allowed) are used for the tensor product fit and ten are used for the smoother fit. This option is generally not recommended because it alters the optimality properties of the fit. One last way to alter the model would be to put in random effects by specifying the option, $\texttt{random} = \texttt{list}(\texttt{x} = \ 1 + \texttt{x})$, for example, which estimates a different intercept and slope parameter for each unique x. The way in which we have used random effects in potential density in this study is by specifying $\texttt{random} = \texttt{list}(\texttt{pd} = \ 1)$, which incorporates a random intercept parameter.

A.2 Appendix for Chapter 3

A.2.1 How to Represent a TTD with Two Peaks Analytically: Expectation-Maximization Algorithm

The POP TTD has two peaks in various locations (e.g., AABW mixing with CDW near the continental shelf of Antarctica). To represent these two peaks, we use a mixture model of IGs, (3.7). Because we only have one equation (mixture model integrates to one) and six unknowns: α, t_2^*, Γ_1, Γ_2, Δ_1^2, and Δ_2^2, we cannot simply solve for each of these parameters using our small number of observations, N, at each grid point. Instead, using samples

from a numerically represented distribution (POP model TTD estimates) as data, we can estimate each of these six parameters using the EM algorithm [13].

The EM algorithm is an iterative computational technique in which parameters are estimated, given data. (In our application, the 'data' are samples from the POP-derived TTD.) It guarantees that the likelihood of sampling the data, given the parameters we are trying to estimate, cannot decrease with each successive iteration. (In our application, the likelihood in the context of this algorithm is the mixture model of IGs.) The expectation part of this algorithm takes an expectation over the (log of the) likelihood evaluated at the data with parameters that are believed to maximize the likelihood. (In our application, the parameters believed to maximize the likelihood are first guesses of the moments of the POP TTD, mentioned following (3.7) as a scaling of the weights, α, times the numerically integrated TTD-derived moments.) The maximization part of this algorithm updates the parameter values such that the likelihood is closer to being maximized. If we observe a subset of independent and identically distributed data, $\{t_1, ..., t_M\}$ (e.g., M is the number of samples we take from the POP TTD), and want to estimate the six parameters (namely, α, t_2^*, Γ_1, Δ_1, Γ_2, and Δ_2) associated with a mixture model of two IGs,

$$p(t_i;\alpha,\Gamma_1,\Delta_1,\Gamma_2,\Delta_2) = \alpha\sqrt{\frac{\Gamma_1^2}{4\pi\Delta_1^2 t_i^3}}exp(-\Gamma_1\frac{(\Gamma_1-\nu t_i)^2}{4\Delta_1^2 t_i}) + \tag{A.1}$$
$$(1-\alpha)\sqrt{\frac{\Gamma_2^2}{4\pi\Delta_2^2(t_i-t_2^*)^3}}exp(-\Gamma_2\frac{(\Gamma_2-\nu(t_i-t_2^*))^2}{4\Delta_2^2(t_i-t_2^*)}),$$

the algorithm becomes 1) evaluating the expectation of weight for each IG (for $i = 1, ..., M$):

$$p_1(t_i) = \frac{\alpha\sqrt{\frac{\Gamma_1^2}{4\pi\Delta_1^2 t_i^3}}exp(-\Gamma_1\frac{(\Gamma_1-\nu t_i)^2}{4\Delta_1^2 t_i})}{p(t_i;\alpha,\Gamma_1,\Delta_1,\Gamma_2,\Delta_2)} \tag{A.2}$$

$$p_2(t_i) = \frac{(1-\alpha)\sqrt{\frac{\Gamma_2^2}{4\pi\Delta_2^2(t_i-t_2^*)^3}}exp(-\Gamma_2\frac{(\Gamma_2-\nu(t_i-t_2^*))^2}{4\Delta_2^2(t_i-t_2^*)})}{p(t_i;\alpha,\Gamma_1,\Delta_1,\Gamma_2,\Delta_2)} \tag{A.3}$$

and 2) maximization of likelihood using the IG distribution:

$$\alpha = \frac{1}{M}\sum_{i=1}^{M} p_1(t_i) \tag{A.4}$$

$$\Gamma_1 = \frac{\sum_{i=1}^{M} t_i p_1(t_i)}{\sum_{i=1}^{M} p_1(t_i)} \tag{A.5}$$

$$\Gamma_2 = \frac{\sum_{i=1}^{M} t_i p_2(t_i)}{\sum_{i=1}^{M} p_2(t_i)} \tag{A.6}$$

$$\Delta_1^2 = \frac{\Gamma_1^3}{2} \frac{\sum_{i=1}^{M} (\frac{1}{t_i} - \frac{1}{\Gamma_1}) p_1(t_i)}{\sum_{i=1}^{M} p_1(t_i)} \tag{A.7}$$

$$\Delta_2^2 = \frac{\Gamma_2^3}{2} \frac{\sum_{i=1}^{M} (\frac{1}{t_i} - \frac{1}{\Gamma_2}) p_2(t_i)}{\sum_{i=1}^{M} p_2(t_i)}, \tag{A.8}$$

which is repeated until each of these values do not change beyond some threshold value (which we set to be 10^{-4}) with a successive iteration. Lastly, during each iteration, an estimate for t_2^* is evaluated with the method of Cheng and Amin (1981) [9] in which t_2^* is initially set to $\tau_0 = t_1^* - (2(\sum_{i=1}^{M} t_i^2/M - (\sum_{i=1}^{M} t_i/M)^2) log M)^{-1} (\sum_{i=1}^{M} t_i - t_1^*)^3$, where t_1^* is the minimum of our samples of ages, which here is set to be $(1-\alpha)\tau_{1,max} + \alpha\tau_{2,max}$ with $\tau_{1,max}$ and $\tau_{2,max}$ the ages at which the distribution has its first and second maximum, respectively. Then t_2^* is set to be the value that the following iterative scheme converges to:

$$\tau_{m+1} = \tau_m + [\frac{3}{M} \sum_{i=1}^{M} (t_i - \tau_m)^{-1} + \lambda_2 \{\Gamma_2^{-2} - \frac{1}{M} \sum_{i=1}^{M} (t_i - \tau_m)^{-2}\}] \times \{3\Gamma_2^{-1} \lambda_2^{-2} + 12\lambda_2^{-2}\} \tag{A.9}$$

with $\lambda_2 = \Gamma_2^3/(2\Delta_2^2)$ being the current EM iteration's values for these second IG parameters. If at some stage, it happens to be that $\tau_{m+1} > t_1^*$, then the iterations are continued with $\tau_{m+1} = (t_1^* + \tau_m)/2$, as suggested by Cheng and Amin (1981) [9].

A.2.2 Non-uniqueness of solutions, sampling from distributions, and extrapolating model output

We should define what we mean by 'sampling' and clarify a few problems we have resolved by having defined the sampling schemes (Section 3.3.3) for taking observational uncertainty into account and the two possible prior distributions (the truncated multivariate normal distribution using POP model output or a uniform distribution) that affect the probability of accepting a given sample (given by (3.16)). Here, by sampling, we are referring to drawing a random number between zero and one and finding the value in the domain at which the cumulative distribution function of the unknown distribution equals this randomly drawn number.

1) We need to be careful about the non-uniqueness of a Green's function solution. That is, different combinations of values of $\vec{H}$ will yield the same concentrations (different advection-diffusion pathways) and multiple maxima in the likelihood (3.15). By utilizing multivariate distributions with non-zero off-diagonal entries for covariances estimated from a model and observations, we have essentially made the combinations of IG parameters that do not appear in the model or observations less likely to be proposed solutions.

2) What matters most in the MCMC scheme is the value of the ratios of the likelihoods and priors in the current and previous iterations, so while we have the likelihoods explicitly from (3.15), we need to be careful not to calculate the prior values using a kernel density estimator such as `ksdensity` in Matlab because we sampled the parameters jointly. (Not only is a kernel density sometimes a crude approximation to a probability density, but assuming that $P(H_1)P(H_2)\cdots = P(\Gamma_1)P(\Gamma_2)P(Pe_1)P(Pe_2)P(\alpha)P(t_2^*) = P(\Gamma_1, \Gamma_2, Pe_1, Pe_2, \alpha, t_2^*) = P(\vec{H})$ is only valid if each component of $\vec{H}$ is independent of each of the other components (e.g., Γ is independent of Pe).) Instead of `ksdensity` on each parameter independently, we use Matlab's `mvnrnd` function to sample the parameters and use Matlab's `mvnpdf` to find the multivariate normal probability density values to plug into the acceptance probability ratio, (3.16), along with the likelihoods. The multivariate normal prior densities do not need to be normalized due to truncation of the multivariate normal distribution because that renormalization factor cancels in the acceptance probability ratio, (3.16). In the uniform prior case, the ratio of prior distributions in the acceptance probability ratio, (3.16), cancels out.

3) For each TTD sample, we sample one of the five TTD members to account for the variability due to the eddy activity of the ocean's circulation in our prior distribution. Similar to a uniform prior, sampling one of the five ensemble TTD members from the POP model with equal probability cancels out in the acceptance probability ratio, (3.16). The reason this does not change the acceptance probability calculation for our MCMC procedure is that $P(\vec{H})$ can be replaced with $P(\vec{H}|m = j)P(m = j)$ where the probability of using member, $m = j$, from the ensemble of TTD outputs is $P(m = j) = 1/5$. This uniform probability of $1/5$ is just another constant of proportionality that cancels out in (3.16).

4) As discussed in Maltrud et al. (2010) [58], we need to extrapolate the POP model

TTD estimates to capture the tails of older ages than the time over which the BIR method was applied ($\leq$ 106 years). To perform this extrapolation, we use a cubic spline fit with knots at every three months, average over windows of eight months, force the TTD to vanish one-third of the way between the peak of the TTD and the year 1200 via log-linear extrapolation since only T_m (106 for $m = 1$, 104 for $m = 2$, 102 for $m = 3$, 30 for $m = 4$, and 70 for $m = 5$) years of the run are available for the five TTDs in the ensemble. We then reassign the raw TTD to the first T_m years, and numerically integrate the resulting 1200 year-long TTD to get its first cumulative moment. Since TTDs $m = \{4, 5\}$ are so much shorter than the others, we assign the average of the $m = \{1, 2, 3\}$ TTDs from year 31 to year 102 for $m = 4$ and from year 71 to year 102 for $m = 5$.

A.2.3 CFC-11 comparisons between POP and observations

When applying Bayesian analysis with an informative prior (e.g., the one derived at each location and depth from POP model output in this study), an implicit assumption is made that the information provided by the prior is not going to bias the estimate, and this is guaranteed to be satisfied if the prior is closer to reality. If, for instance, in the current study, the TTD at each location and depth derived from POP were all peaked at 100 years, then our Bayesian TTD estimates would have more weight in the larger age classes, which would have skewed our estimates in the wrong direction at shallow ($<$ 1000 meters) depths. One way to qualitatively address whether POP is modeling the ocean well enough to use its five TTD estimates is to look at maps of a transient tracer at various times. Here, we show maps of CFC-11 concentrations according to POP and from objectively mapping observations along each hydrographic section and at each of their occupation times (Table 3.1).

In the North Atlantic Ocean, the POP output and observations are compared along A20 (Fig. A.1), A05 (Fig. A.2), and A16 (Fig. A.3). It is apparent from Figs. A.1 and A.2 that POP is too diffusive and/or has a deep western boundary current (DWBC) that is too strong because the CFC-11 concentrations according to POP penetrate much farther into the interior of the North Atlantic Ocean and POP simulates larger concentrations in

the core of the DWBC along both A20 and A05. In the current study, our mean age estimates suggest that the DWBC in POP is not too strong, so as long as water parcel pathways in POP are consistent with water parcel pathways in the real ocean, POP is too diffusive at intermediate depths in the North Atlantic Ocean. Nevertheless, the magnitude changes in CFC-11 concentrations from one occupation time to another and the general spatial structures such as implied mixed layer depths are broadly consistent between POP and observations. Where there is LSW (Fig. A.3) and where there is the LSW component of NADW (Fig. A.2), POP is consistent enough with observations to have its output used for prior information, but is not consistent enough (particularly, near the DWBC) to be interpreted as reality.

Similarly, in the Southern Ocean, the POP output and observations are compared along SR03 (Fig. A.4). CFC-11 concentrations are modeled by POP to be very similar to observations in the same ways as those in the North Atlantic Ocean except for one feature unique to this geographic region: POP does not model AABW very well [17], which is most apparent in Fig. A.4 in the southern portion of the domain (south of 60^oS) at bottom depths. The CFC-11 concentrations should be elevated in locations where there is AABW because AABW is relatively young, but POP simulates very small CFC-11 concentrations in these locations. Where there is AAIW and SAMW (Fig. A.4), POP is consistent enough with observations to have its output used for prior information, but is not consistent enough (particularly, near AABW) to be interpreted as reality.

A.3 Appendix for Chapter 4

Starting from the momentum-balance equations for an incompressible fluid and neglecting viscosity, we have the set of coupled partial differential equations (PDEs),

$$\frac{\partial u}{\partial t} + u\frac{\partial u}{\partial x} + v\frac{\partial u}{\partial y} + w\frac{\partial u}{\partial z} - fv = \frac{1}{\rho}\frac{\partial \tau^x}{\partial z} - \frac{1}{\rho}\frac{\partial p}{\partial x} \tag{A.10}$$

$$\frac{\partial v}{\partial t} + u\frac{\partial v}{\partial x} + v\frac{\partial v}{\partial y} + w\frac{\partial v}{\partial z} + fu = \frac{1}{\rho}\frac{\partial \tau^y}{\partial z} - \frac{1}{\rho}\frac{\partial p}{\partial y} \tag{A.11}$$

where the zonal, meridional, and vertical currents are $u = u_e + u_g$, $v = v_e + v_g$, and $w = w_e$, respectively,the sums of their Ekman and geostrophic contributions. Ageostrophic

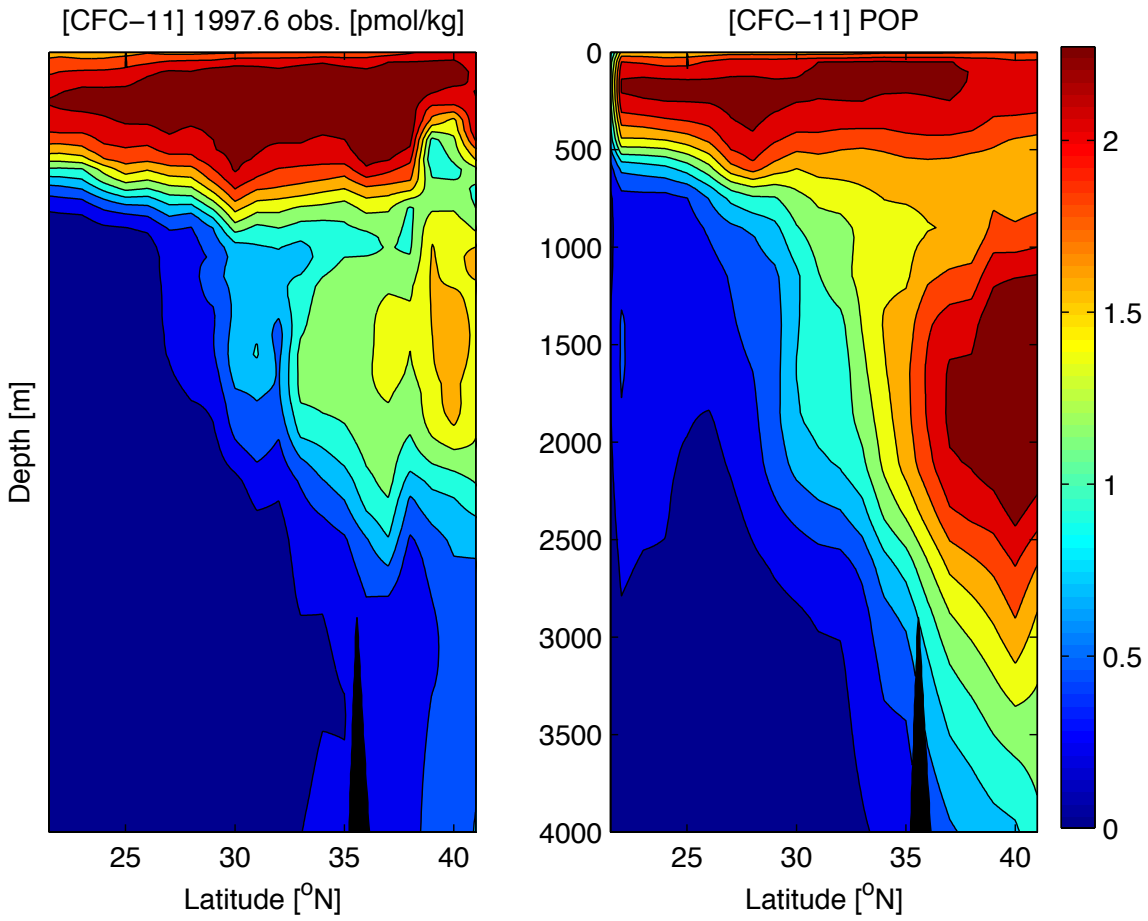

Figure A.1: Comparison plots of [CFC-11] in pmol/kg along A20 from objectively mapped observations (left) and POP (right) in 1997.

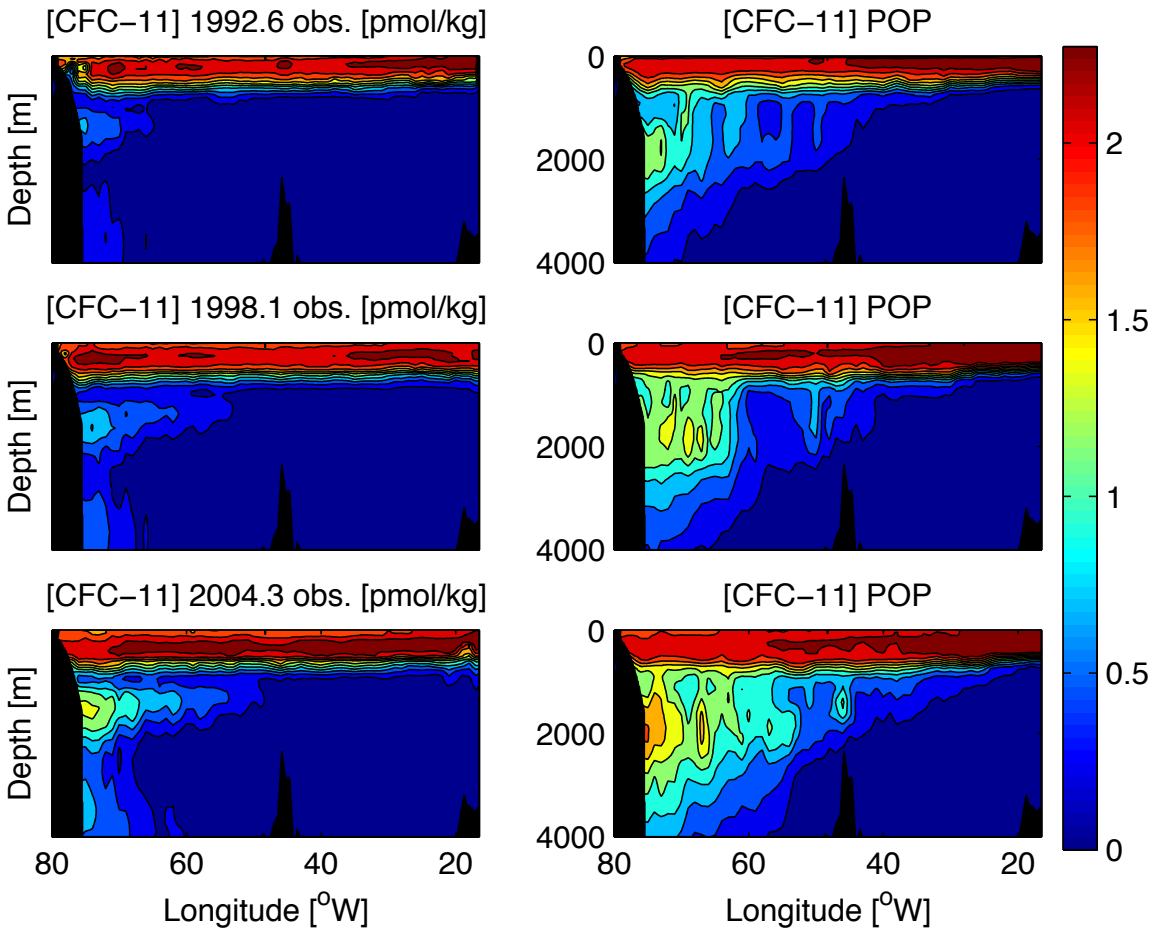

Figure A.2: Comparison plots of [CFC-11] in pmol/kg along A05 from objectively mapped observations (left) and POP (right) in 1992 (top), 1998 (middle), and 2004 (bottom).

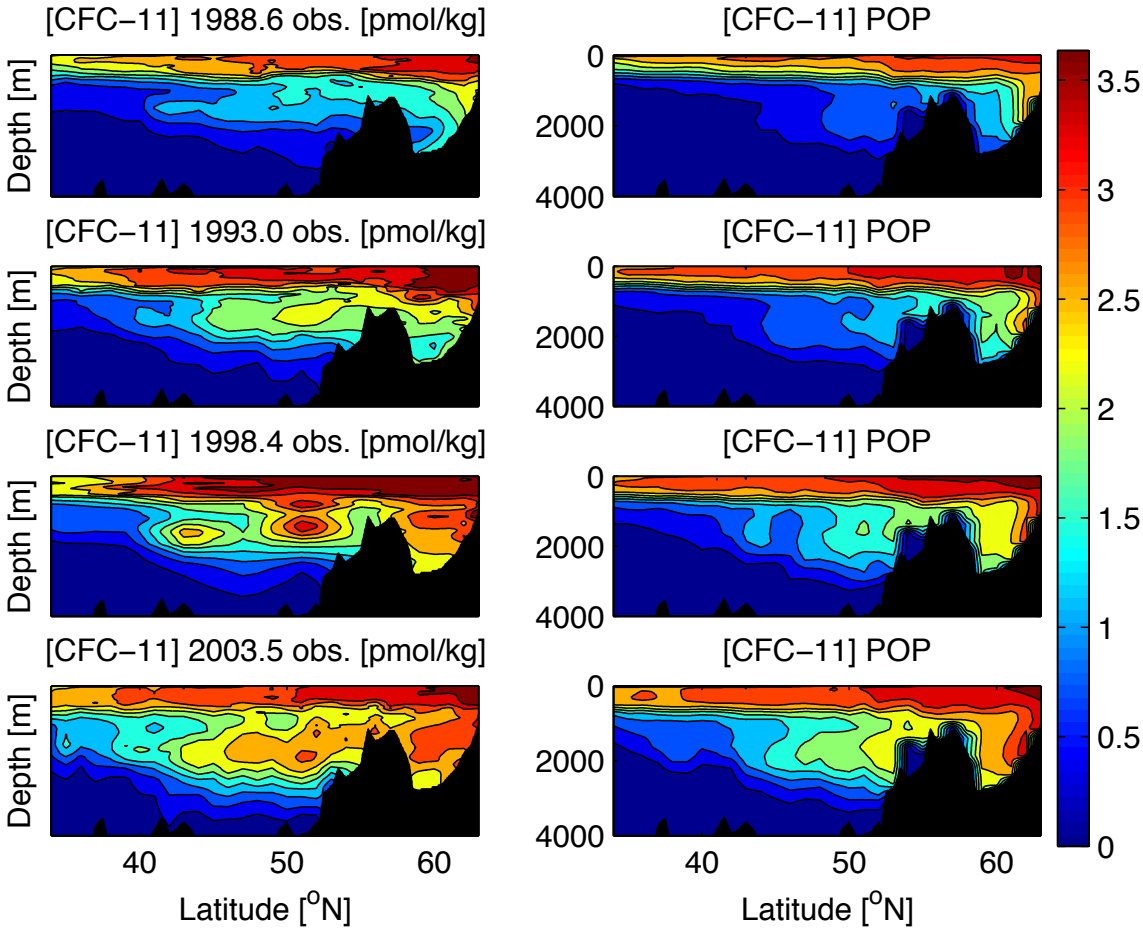

Figure A.3: Comparison plots of [CFC-11] in pmol/kg along A16 from objectively mapped observations (left) and POP (right) in 1988 (top), 1993 (second-to-top), 1998 (second-to-bottom), and 2003 (bottom).

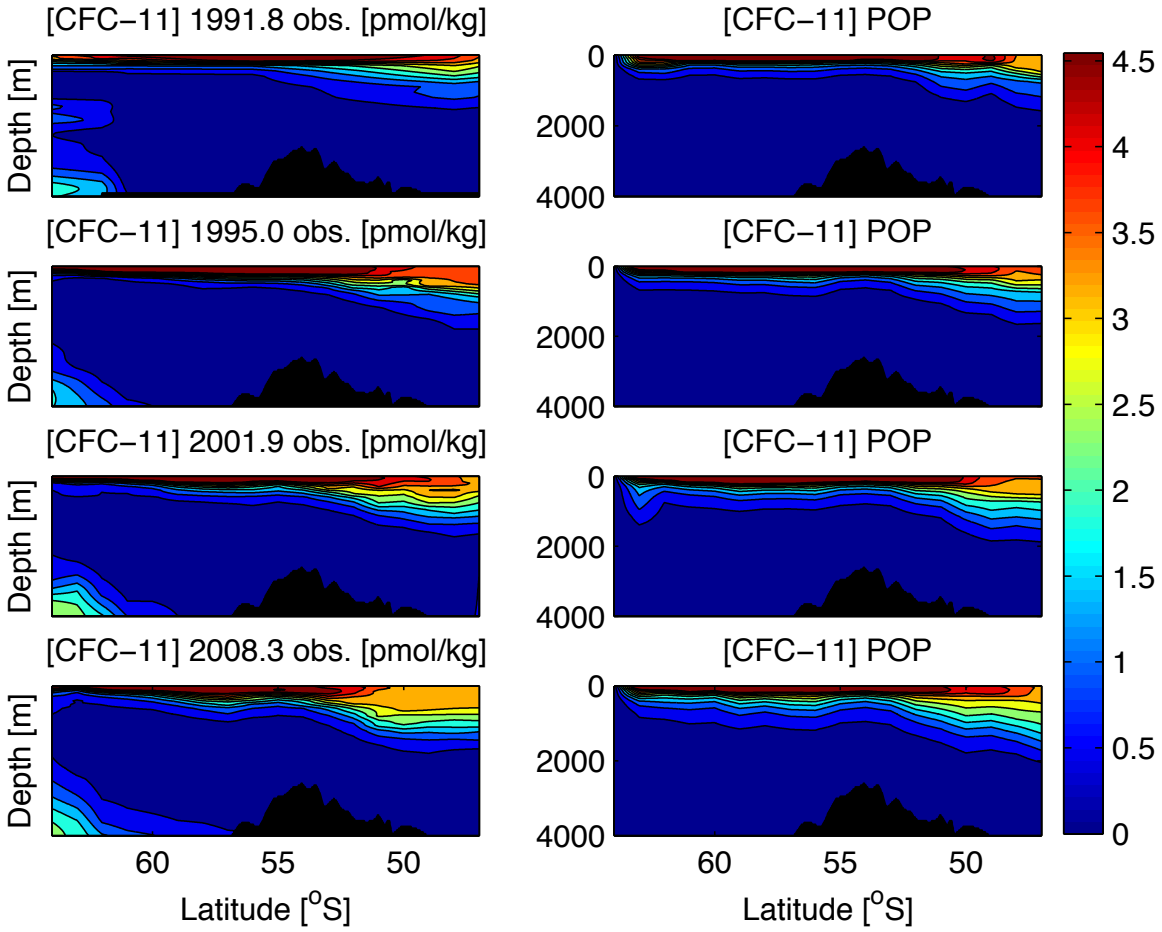

Figure A.4: Comparison plots of [CFC-11] in pmol/kg along SR03 from objectively mapped observations (left) and POP (right) in 1991 (top), 1995 (second-to-top), 2001 (second-to-bottom), and 2008 (bottom).

currents, which include Ekman currents, are defined to be the difference between the total currents and the geostrophic currents. Here, non-Ekman ageostrophic currents and vertical geostrophic currents are ignored. Furthermore, because we have defined the geostrophic current in terms of the gradient of the total pressure field, the pressure gradients cancel with the geostrophic currents without derivatives [37].

Now we apply the following assumptions: i) The geostrophic currents are vertically uniform in the Ekman layer. This is reasonable if we assume that the Ekman layer depth is finer than our vertical resolution. ii) The geostrophic currents obey quasigeostrophic dynamics. That is, $\partial u_g/\partial t + u_g \partial u_g/\partial x + v_g \partial u_g/\partial y - f v_a = 0$ and $\partial v_g/\partial t + u_g \partial v_g/\partial x + v_g \partial v_g/\partial y + f u_a = 0$ to first-order. Here, 'a' denotes 'ageostrophic' but non-Ekman. This essentially assumes that Rossby number, $Ro = U/fL$, is much less than unity. The ageostrophic current terms are of order Ro compared to the geostrophic velocities. Also, the vertical velocity, w, is of the order of Ro compared to aspect ratio scaling. Thus, by vertically integrating over the Ekman layer with depth, δ_e, we get the set of coupled PDEs,

$$\begin{aligned} \frac{\partial (U_e + \delta_e u_g)}{\partial t} + \int_{-\delta_e}^{0} u_e \frac{\partial u_e}{\partial x} dz + U_e \frac{\partial u_g}{\partial x} + u_g \frac{\partial U_e}{\partial x} + \delta_e u_g \frac{\partial u_g}{\partial x} + \\ \int_{-\delta_e}^{0} v_e \frac{\partial u_e}{\partial y} dz + V_e \frac{\partial u_g}{\partial y} + v_g \frac{\partial U_e}{\partial y} + \delta_e v_g \frac{\partial u_g}{\partial y} + \int_{-\delta_e}^{0} w_e \frac{\partial u}{\partial z} dz - f V_e - f V_a &= \frac{\tau^x}{\rho} \end{aligned} \quad \text{(A.12)}$$

$$\begin{aligned} \frac{\partial (V_e + \delta_e v_g)}{\partial t} + \int_{-\delta_e}^{0} u_e \frac{\partial v_e}{\partial x} dz + U_e \frac{\partial v_g}{\partial x} + u_g \frac{\partial V_e}{\partial x} + \delta_e u_g \frac{\partial v_g}{\partial x} + \\ \int_{-\delta_e}^{0} v_e \frac{\partial v_e}{\partial y} dz + V_e \frac{\partial v_g}{\partial y} + v_g \frac{\partial V_e}{\partial y} + \delta_e v_g \frac{\partial v_g}{\partial y} + \int_{-\delta_e}^{0} w_e \frac{\partial v}{\partial z} dz + f U_e + f U_a &= \frac{\tau^y}{\rho} \end{aligned} \quad \text{(A.13)}$$

The integrated zonal Ekman currents are U_e and the integrated meridional Ekman currents are V_e. The Ekman currents are so smooth that, in fact, all first-order derivatives of them turn out to be negligible. This may also be seen by a scaling argument. Thus, with one further assumption, that iii) terms to second order in u_e and v_e (e.g. $\int_{-\delta_e}^{0} dz u_e \partial u_e/\partial x$) are negligible, we find:

$$U_e \frac{\partial u_g}{\partial x} + u_g \frac{\partial U_e}{\partial x} + V_e \frac{\partial u_g}{\partial y} + v_g \frac{\partial U_e}{\partial y} - f V_e = \frac{\tau^x}{\rho} \quad \text{(A.14)}$$

$$U_e \frac{\partial v_g}{\partial x} + u_g \frac{\partial V_e}{\partial x} + V_e \frac{\partial v_g}{\partial y} + v_g \frac{\partial V_e}{\partial y} + f U_e = \frac{\tau^y}{\rho}. \quad \text{(A.15)}$$

After introducing τ_{damp} to iteratively find a numerical solution, we get Equations 4.7 and 4.8.

We may, alternatively, solve for the Ekman transports, U_e and V_e, explicitly if we assume i), ii), iii), and make further assumptions about the strengths of the Ekman and geostrophic currents and their derivatives in space with respect to one another: If iv) $(\partial u_e/\partial x)/(\partial u_g/\partial x) \ll 1$, v) $u_g/u_e \ll 1$ (similarly for derivatives in y and for the meridional currents), and vi) products of derivatives in any currents are negligible, then we get,

$$U_e \approx \frac{1}{\rho}\left(\frac{\tau^x \frac{\partial v_g}{\partial y}}{\left(f\left(f + \frac{\partial v_g}{\partial x} - \frac{\partial u_g}{\partial y}\right)\right)} + \frac{\tau^y}{\left(f + \frac{\partial v_g}{\partial x}\right)}\right) \tag{A.16}$$

$$V_e \approx \frac{1}{\rho}\left(\frac{\tau^y \frac{\partial u_g}{\partial x}}{\left(f\left(f + \frac{\partial v_g}{\partial x} - \frac{\partial u_g}{\partial y}\right)\right)} - \frac{\tau^x}{\left(f - \frac{\partial u_g}{\partial y}\right)}\right). \tag{A.17}$$

From these analytical results, we can find $w_e = \partial U_e/\partial x + \partial V_e/\partial y$ to get,

$$w_e \approx \frac{1}{\rho}\frac{1}{\left(f + \frac{\partial v_g}{\partial x} - \frac{\partial u_g}{\partial y}\right)}\left(\frac{\partial \tau^y}{\partial x} - \frac{\partial \tau^x}{\partial y} - \frac{\beta}{f}\tau^x - \frac{\beta}{\left(f + \frac{\partial v_g}{\partial x} - \frac{\partial u_g}{\partial y}\right)}\right.$$
$$\left.\times\left((\tau^y - \tau^x)\left(\frac{\partial v_g}{\partial x} - \frac{\partial u_g}{\partial y}\right) + 2\left(f(\tau^y - \tau^x) - \tau^x\left(\frac{\partial v_g}{\partial x} - \frac{\partial v_g}{\partial y}\right) + \tau^y\left(\frac{\partial u_g}{\partial x} - \frac{\partial u_g}{\partial y}\right)\right)\right)\right). \tag{A.18}$$

It is assumption v) that fails almost everywhere, but this is our first guess in numerically solving for the Ekman currents. Note that a negative value for we represents water transport from the Ekman layer down to the mixed layer because of how we have defined this quantity. A positive Ekman suction is a positive we and a positive Ekman pumping is a negative w_e.

A.4 Appendix for Chapters 5 and 6

When applying Bayesian analysis in which a 'best' model is weighted more than the others, an implicit assumption is made that there is at least one model in the suite of possible models (combinations of climatological and hindcast circulation fields with three different values for κ_H in the current study) that simulates reality reasonably well. If, for instance, in the current study, the CFC-11 concentrations simulated by Offtrac were nothing like the observed CFC-11 concentrations and were approximately constant as functions of latitude and longitude, linearly increasing as a function of depth, then our Offtrac simulations would not provide information about important features such as the structure of the mixed layer and/or spatial distribution of water masses. One way to qualitatively address whether Offtrac is modeling the ocean well enough to use its simulated CFC-11 concentrations is

to look at maps of these concentrations at various times. Here, we show maps of CFC-11 concentrations according to Offtrac and from objectively mapping observations along each hydrographic section and at each of their occupation times (Table 5.1).

In the North Atlantic Ocean, the Offtrac output and observations are compared along A20 (Fig. A.5), A05 (Fig. A.6), and A16 (Fig. A.7). It is apparent from Figs. A.5 and A.6 that Offtrac has a problem with simulating the deep western boundary current because the CFC-11 concentrations according to Offtrac do not penetrate to deeper depths than about 1000 meters or very far into the interior of the North Atlantic Ocean. However, the magnitude changes in CFC-11 concentrations from one occupation time to another and general shallow-water spatial structures such as implied mixed layer depths are broadly consistent between Offtrac and observations. Nonetheless, the BMA estimates in Figs. A.6 and A.7 are mostly improved directly beneath the mixed layer and the BMA estimates in Fig. A.5b improve the CFC-11 concentration simulations from any individual model run (Figs. A.5c-h) the most beneath the mixed layer, just above where there is NADW. Within the mixed layer and where there is LSW (Fig. A.7), Offtrac is consistent enough with observations to have its output used, but is not consistent enough (particularly, near the DWBC) to provide much information elsewhere.

Similarly, in the Southern Ocean, the Offtrac output and observations are compared along SR03 (Fig. A.8) and I09(S) (Fig. A.9). CFC-11 concentrations are modeled by Offtrac to be very similar to observations in the same ways as those in the North Atlantic Ocean except for one feature unique to this geographic region. Offtrac does not model AABW very well, which is most apparent in Fig. A.8 in the southern portion of the domain (south of 60^oS) at bottom depths. The CFC-11 concentrations should be elevated in locations where there is AABW because AABW is relatively young, but Offtrac simulates very small CFC-11 concentrations in these locations, similar in magnitude to surrounding waters. Nonetheless, the BMA estimates in Fig. A.8 are mostly improved directly beneath the mixed layer (closer to the less steep vertical concentration gradient in observations) and the BMA estimates in Fig. A.9b improve the CFC-11 concentration simulations from any individual model run (Figs. A.9c-h) the most beneath the mixed layer, down to where there is AAIW. Within the mixed layer and where there is SAMW (Fig. A.8), Offtrac is

consistent enough with observations to have its output used, but is not consistent enough (particularly, near AABW) to provide much information elsewhere.

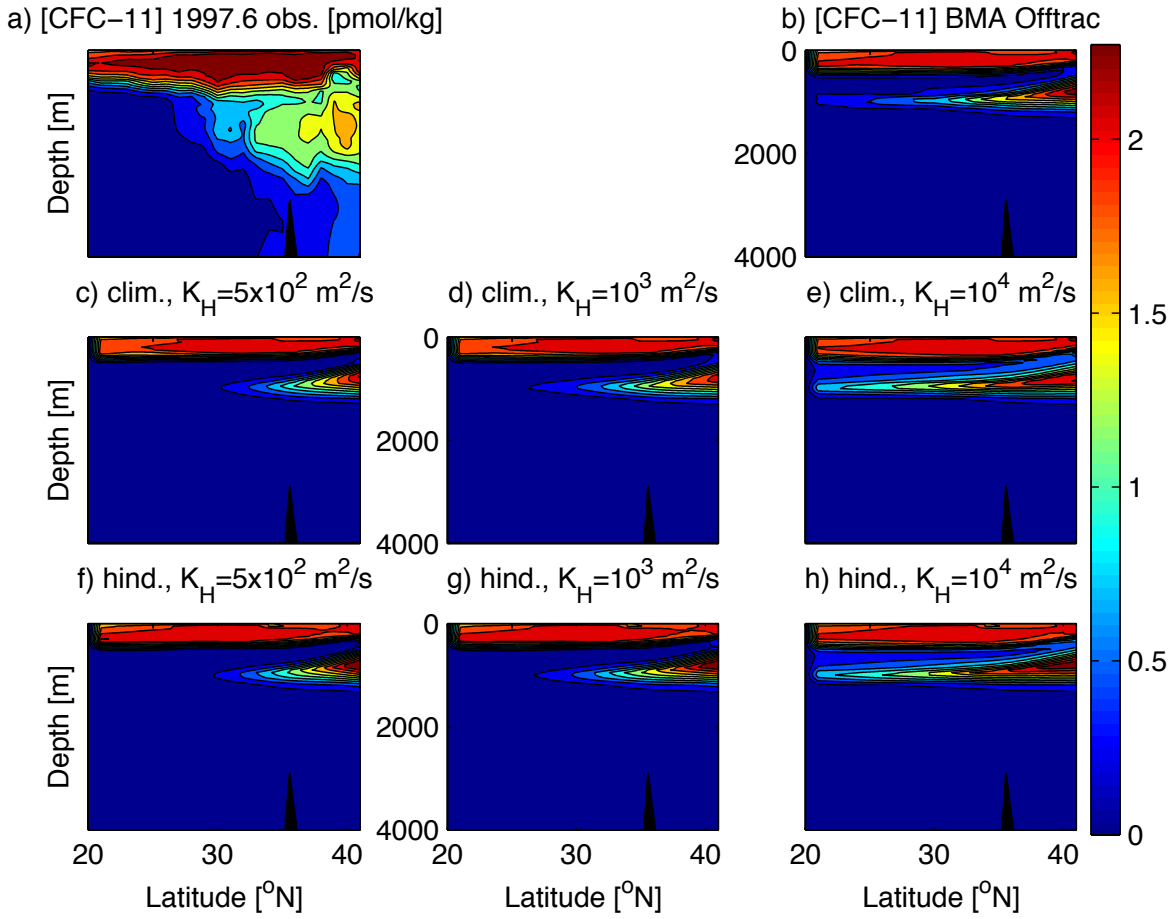

Figure A.5: Comparison plots of [CFC-11] in pmol/kg along A20 in 1997 from objectively mapped observations (a) and Offtrac (c-h) from all six configurations of model runs. The BMA estimates at each grid point from Offtrac output based on simulated and observed [CFC-11] is also shown (b).

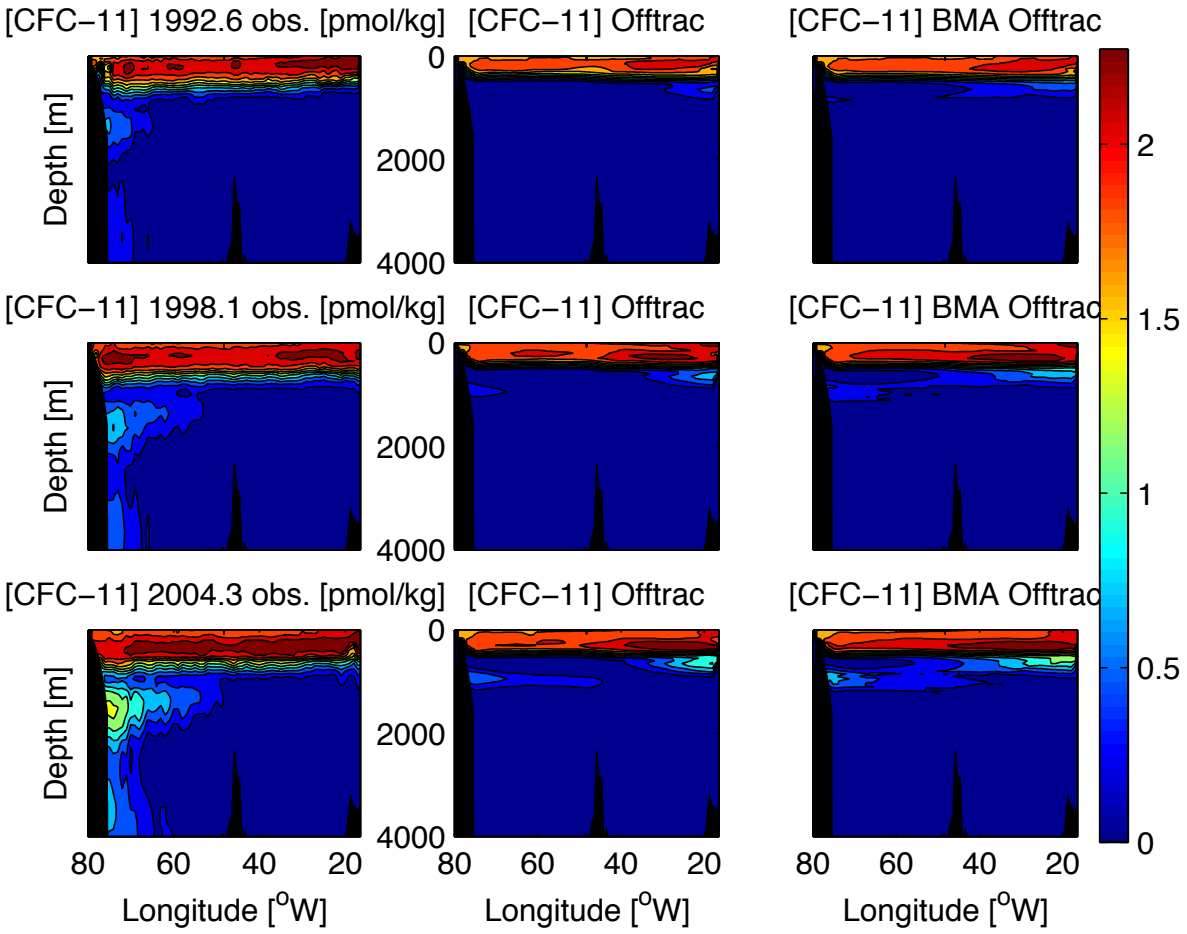

Figure A.6: Comparison plots of [CFC-11] in pmol/kg along A05 from objectively mapped observations (left), the climatological run with $\kappa_H = 10^3$ m^2 s^{-1} simulated by Offtrac (middle), and the BMA estimates at each grid point from Offtrac output based on simulated and observed [CFC-11] in 1992 (top), 1998 (middle), and 2004 (bottom).

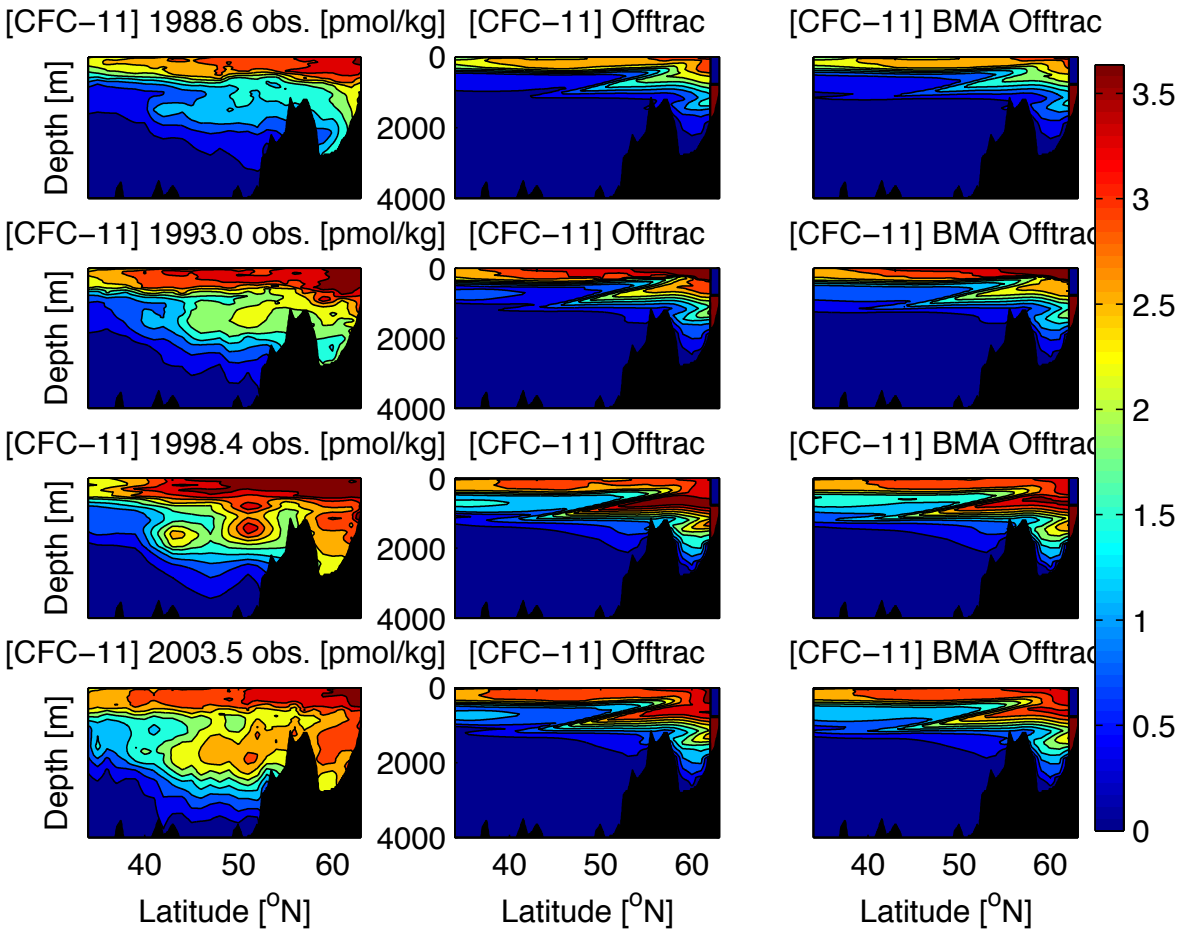

Figure A.7: Comparison plots of [CFC-11] in pmol/kg along A16 from objectively mapped observations (left), the climatological run with $\kappa_H = 10^3$ m^2 s^{-1} simulated by Offtrac (middle), and the BMA estimates at each grid point from Offtrac output based on simulated and observed [CFC-11] in 1988 (top), 1993 (second-to-top), 1998 (second-to-bottom), and 2003 (bottom).

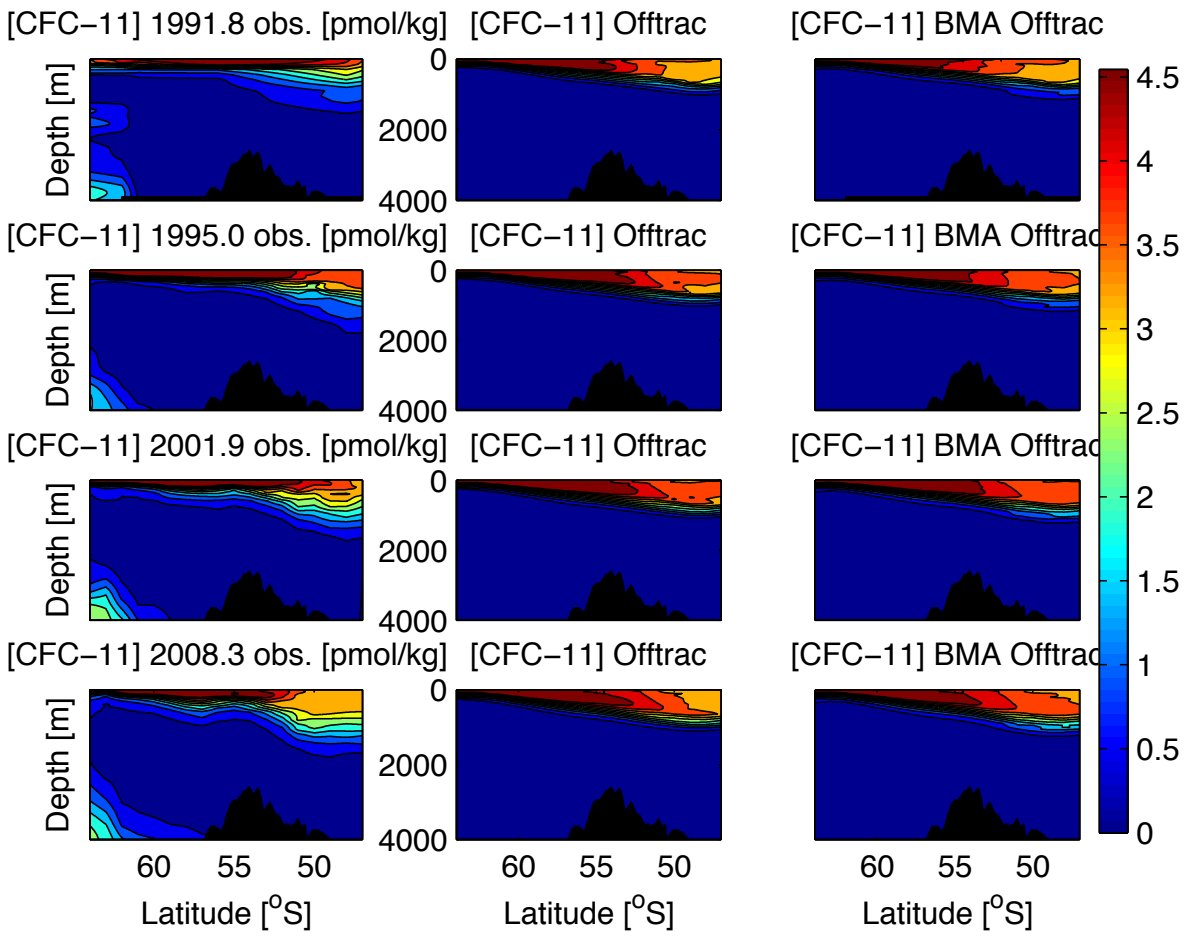

Figure A.8: Comparison plots of [CFC-11] in pmol/kg along SR03 from objectively mapped observations (left), the climatological run with $\kappa_H = 10^3$ m^2 s^{-1} simulated by Offtrac (middle), and the BMA estimates at each grid point from Offtrac output based on simulated and observed [CFC-11] in 1991 (top), 1995 (second-to-top), 2001 (second-to-bottom), and 2008 (bottom).

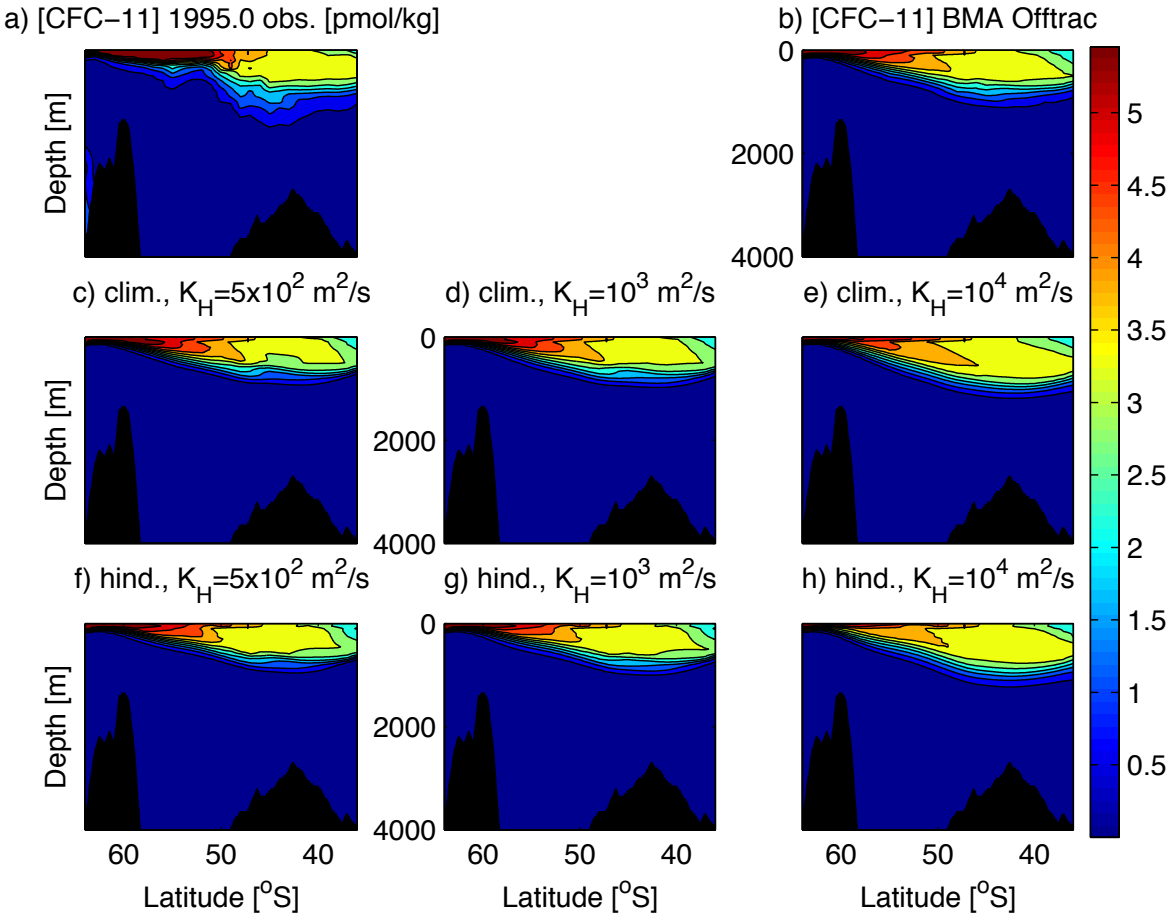

Figure A.9: Comparison plots of [CFC-11] in pmol/kg along I09 from objectively mapped observations (left) and Offtrac (c-h) from all six configurations of model runs. The BMA estimates at each grid point from Offtrac output based on simulated and observed [CFC-11] is also shown (b).