Supplementary Information

J. Lanham[1]

[1]Department of Earth Sciences, University of Cambridge, UK

July 1, 2025

S1 Supplementary Information for Methods Section

S1.1 Basic OMP Machinery

We can write the series of linear mixing equations in vector form for n SWTs, where x_i is the relative fraction of SWT_i, $\mathbf{V}_i^{SWT}$ is the vector of SWT_i end members, and $\mathbf{V}^{Obs}$ is the vector of observed values with a residual vector $\mathbf{r}$:

$$\sum_{i=1}^{n} x_i * \mathbf{V}_i^{SWT} = \mathbf{V}^{Obs} + \mathbf{r} \tag{1}$$

The size of the vectors equals the number of variables used in the analysis. Both the end member vector and observed values vector are normalised by subtracting the mean and dividing by the standard deviation of the end members, such that:

$$\mathbf{Vnorm}_i^{SWT} = (\mathbf{V}_i^{SWT} - \overline{\mathbf{V}^{SWT}})/\mathbf{V}_\sigma^{SWT} \tag{2}$$

$$\mathbf{Vnorm}_i^{Obs} = (\mathbf{V}_i^{Obs} - \overline{\mathbf{V}^{Obs}})/\mathbf{V}_\sigma^{Obs} \tag{3}$$

Where $\overline{\mathbf{V}^{Obs}}$ and $\mathbf{V}_\sigma^{Obs}$ are the mean and standard deviation vectors across the whole data sample, respectively. The analysis also includes a series of user-defined weights, $\mathbf{W}$, to account for variation in the precision of measurement of each variable, or confidence in end-member definition. Our weighting matrix can be found in Table S1.2. The final cost function (i.e. the quantity to be minimised by the least-squares solver) can be written in terms of the residual vector $\mathbf{r}$:

$$\mathbf{r} = \mathbf{W} * \left(\sum_{i=1}^{n} x_i * \mathbf{Vnorm}_i^{SWT} - \mathbf{Vnorm}^{Obs} \right) \tag{4}$$

This is then solved in a least-squares optimisation to minimise $\mathbf{r}$, and only solutions in which $x_i \geq 0$ are accepted so that the analysis is 'non-negative'. We can write the water mass solution vector $\mathbf{x}$, in terms of the cost function in Equation 4, where $argmin$ specifies the value of $\mathbf{x}$ that minimises the quantity in square brackets:

$$\mathbf{x} = argmin \left[\mathbf{W} * \left(\sum_{i=1}^{n} x_i * \mathbf{Vnorm}_i^{SWT} - \mathbf{Vnorm}^{Obs} \right) \right] \tag{5}$$

As we state in Section 5.3, we opt to model the semi-conservative parameters NO and PO. These variables can be written as:

$$PO = O_2 + rt_{PO_4}^{O_2} PO_4 \tag{6}$$

$$NO = O_2 + \frac{rt^{O_2}_{PO_4}}{rt^{NO_3}_{PO_4}} NO_3 \tag{7}$$

Where rt^a_b denotes a user-specified exchange ratio of parameter a with respect to b. In this application, rt makes use of the Redfield ratios: we specify the $rt^{O_2}_{PO_4}$ value as -170 and the $rt^{NO_3}_{PO_4}$ value as 16. These are the default values in the original Karstensen & Tomczak [1] implementation.

S1.2 Additional Information on the Selection of End Members and Weighting Practices

We start with the original Temperature-Salinity end members from the Southern Ocean OMP from Pardo *et al.* [2], and make some minor tweaks to better fit the parameter space of our GLODAP data. These are shown in Figure 1. We use a K-D tree to find the nearest 1000 data points to the temperature and salinity end member pair, selecting the median value from this sample as the end member for each tracer. The full set of end members used is shown in Table S1.2. We elect to use the parameter weights shown in Table S1.2. These are the default values in the documentation from Shrikumar *et al.* [3], and relate to the uncertainties surrounding the measurement of each variable. In practice, these weights determine the strength of the constraint on each variable: i.e. the extent to which residuals in that variable will be tolerated in the solving of the least-squares equations. Where measurement uncertainty is relatively low, then the weight is increased to force lower residuals in that equation at the expense of variables with higher degrees of uncertainty. We define weights for every variable, for which the solver uses a hard mass constraint (which prioritises returning a residual of 0 in the mass equation). Further details can be found in Shrikumar *et al.* [3].

As described in section 5.5, we use a spatial mask to calculate Δ values, which we define as the change in the properties of the grid cells with the highest concentration of each water mass in the time-mean 20-year Argo dataset. We judge 1000 grid points to be a sufficient sample to describe the mean properties of the source water, but it is possible that this approach somewhat obscures any large changes in the distribution of the core of the source water. By definition, the spatial mask as defined in the 20-year mean encompasses both early and late distributions of each water mass; it is merely an approximation of the 3D distribution of the source water from a time-mean perspective. It is worth also noting that the definition of AABW Δ values is confined to the top 2000 m (the maximum depth range of Argo floats), such that it may not capture changes in the core of the water mass. However, as we describe in Sections 4 and S3.1.1, the main findings of the study are robust to significant perturbations in these values.

In Table S1.2 we show the full set of Δ values discussed in Section 5.5. We also show how zonal-mean temperature and salinity have changed in the Southern Ocean during the 20-year Argo period via the composite anomalies shown in Figure S22.

S1.3 The Importance of Tracers in the OMP Solution

To demonstrate the importance of tracers in the OMP-based classification of water masses, we show a sample of an OMP classification in the GLODAP climatology using only temperature and salinity (see Figure S16). This shows that the analysis struggles to differentiate between end members which are similar in temperature-salinity space. Specifically and most importantly for the focus of this study, the classification cannot distinguish between CDW, NADW, and AABW with any real degree of skill. As we state in Section 5.6, our method enables skillful classification in the absence of these tracers (provided that we have some tracer data to conduct an initial classification with which to train on). Figure S16 demonstrates that we cannot assess changes in water masses in the gridded Argo data from an OMP framework, without at least giving it a prior. We acknowledge that similar levels of skill may be achievable without machine learning, if we were to introduce a 'a-priori' weighting matrix into the OMP linear mixing equations.

Source Water End Members						
Source Water Type	Temp. (°C)	Salinity (Psu)	Oxygen (μmol kg^{-1})	Nitrate (μmol kg^{-1})	Phosphate (μmol kg^{-1})	Total Alkalinity (μmol kg^{-1})
STCW	13.50 ±2.33	35.30 ±0.17	214.49 ±11.8	8.63 ±1.96	0.72 ±0.15	2320.72 ±10.82
SAMW1	8.75 ±2.81	34.58 ±0.16	242.78 ±26.05	15.75 ±4.15	1.15 ±0.28	2284.11 ±12.49
SAMW2	5.00 ±1.18	34.14 ±0.13	310.21 ±13.88	20.00 ±3.75	1.43 ±0.23	2284.86 ±13.49
DSW	-1.91 ±1.13	34.82 ±0.14	282.86 ±25.32	30.42 ±2.24	2.15 ±0.16	2344.32 ±8.66
AASW	-1.85 ±0.51	33.82 ±0.16	326.26 ±16.18	26.13 ±2.69	1.84 ±0.21	2276.65 ±17.58
AAIW	3.13 ±0.39	34.14 ±0.03	291.17 ±12.28	26.34 ±1.32	1.84 ±0.09	2284.95 ±3.373
NADW	2.50 ±0.24	34.83 ±0.04	217.01 ±9.39	26.48 ±1.29	1.79 ±0.12	2341.39 ±4.00
CDW	1.00±0.08	34.74 ±0.01	206.12 ±10.02	31.06 ±0.73	2.16 ±0.07	2356.31 ±4.12
AABW	-0.50 ±0.14	34.64 ±0.01	252.93 ±6.81	32.41 ±0.57	2.24 ±0.04	2355.21 ±1.52

Table S1: End member values used for each of the 9 source waters used in the water mass classification. Uncertainty values are calculated retroactively as one standard deviation within the 1000 grid points of highest relative water mass fraction, taken to approximate the properties of the source water. These values are used in the sensitivity testing in Section S1.4.

Parameter Weights					
Temp.	Salinity	Oxygen	Nitrate	Phosphate	Alkalinity
24	24	7	2	2	2

Table S2: Relative weights used in the water mass classification for each of the 6 parameters considered.

Source Water Δ Values		
Source Water Type	Δ Temp. (°C)	Δ Salinity (Psu)
STCW	0.024	0.0166
SAMW1	0.181	0.0217
SAMW2	0.047	0.0092
DSW	0.127	-0.0001
AASW	0.091	0.0039
AAIW	0.096	-0.0059
AABW	0.077	0.0069

Table S3: End member Δ values used for each of the 9 source waters used in our water mass classification.

S1.4 Sensitivity Testing of the OMP Solution

Figures S17, S18, S19, and S20 show the results of end-member sensitivity testing in the climatological OMP solution at 0, 90, 180, and 240°E. Firstly, for each water mass, we define a source water mask as the 1000 data points with the highest fraction according to the original classification. We compute the standard deviation for each variable within this source water mask and use this as a perturbation to the original end members. The output shown in Figures S17, S18, S19, and S20 is the result of a simultaneous perturbation of 1 standard deviation in each variable. These values are shown in Table S1.2.

In almost all cases, the median impact of this perturbation is extremely close to 0. In the most sensitive cases, the inter-quartile range shows a change in water mass fractions on the order of 5-10%. This can be seen in the CDW/AABW perturbations at 0°E and in the SAMW1/SAMW2 perturbations at 90, 180, and 240°E. Unsurprisingly, perturbations to a given end member typically cause the greatest impact in fractions on itself and the end member closest to it in parameter space. For example, the change in CDW fractions that is induced by perturbing the CDW end member leads to an almost equal and opposite change in AABW (Figure S17). Likewise, when the SAMW1 end members are perturbed, the change in SAMWs fractions often induces a similar magnitude change of opposite sign in AAIW.

Figure S21 shows the impact of perturbing the weighting matrix on the OMP solution. Here, we reduce the relative weighting of each variable sequentially by 50%, and compare the change in the solution. Compared with perturbing the end members, the solution with perturbed weights is relatively insensitive to variability in the weights; the highest $\Delta WaterMassFraction$ values are only on the order of $\pm$ 0.6%. We show a sample of the output for just 0°E, but the impact is of similar order of magnitude elsewhere.

S1.5 Supplementary Information for Validation of the ML Model

Figure S23 shows a sample comparison between the output of the original GLODAP and algorithm-derived Argo water mass distribution. In Figure S24, we show the output of the uncertainty calculations used in the application of the RF model to the RG Argo gridded data. The uncertainty is calculated on a grid-point basis as the variance across the 5-fold ensemble members, normalised by the 99th percentile. This is shown for a variety of water masses at a number of depth levels.

Exclusion of water masses from R^2 values in RF model verification In Table 1, we show the results of the out-of-distribution testing of the RF via the exclusion of repeat sections. As we state in Section 5.6, there are some cases in which we remove the R^2 contribution from a particular water mass from the final R^2 value shown in the table. These are cases in which the relative fraction of a particular water mass is effectively 0 and the model exhibits no meaningful skill in prediction. In the Drake Passage (line SR01), we remove the low R^2 values associated with both NADW and STCW. Here, relative fractions of each are very close to 0 at all observed data points. In line P06, which is a zonal section across the Pacific at 30°S, we exclude the low R^2 values associated with both NADW and DSW. We do the same for the meridional section in the Pacific (line P18). In the Indian Ocean, the R^2 values associated with DSW and AASW are excluded from the zonal cross section in line IO5, whilst we exclude NADW from the meridional section in line IO7. In all such cases, the relative fraction of the water mass in all data points along that section is negligible. Finally, we also exclude the R^2 contribution from STCW and NADW in the zonal section across west Antarctic to the Ross Sea (line S04P), with the same reasoning.

S2 Supplementary Information for Results Sections

S2.1 Additional Samples of Water Mass Classification Output

Figures S1 and S3 show samples of DSW, AASW, STCW, SAMWs and NADW from the climatology / 'mean-state' GLODAP classification at a range of depths. Figure 4 shows the fractions

Expocodes used in Analysis and Validation		
Location	Line(s)	Expocode(s)
Weddell Sea	A12, A21	06AQ20060825 06AQ20080210 35MF20080207 740H20081226 740H20090307 29HE20190406
South of Tasmania	SR03	09AR20071216 09AR20080322 096U20180111
Drake Passage	SR01	74JC20151217 740H20090203 74JC20181103
West Pacific Ocean	P15	09SS20090203 096U20160426 49NZ20071122
Pacific Ocean (zonal)	P06	318M20091121 320620170703 320620170820 49NZ20170208
East Pacific Ocean	P18, P17	33RO20071215 33RO20161119 49NZ20170208
East Indian Ocean	I08	33RR20070204 33RR20080204 33RR20160208
Central Indian Ocean	I07	33RO20180423 49NZ20191229
West Indian Ocean	I06	325020190403
Indian Ocean (zonal)	I05	33RR20090320
Pacific Ocean (zonal)	S04P	320620180309

Table S4: List of all expocodes used in each of the location/line groupings using in the validation of the machine learning model in Section 5.6.

of AASW at the surface. These compliment the outputs in Figure 2. We also provide maps of the DSW and DSW + AABW distribution in the upper 1000 m of the mean-state classification in Figure S2. Our analysis successfully identifies the principal known locations of DSW production (Weddell and Ross Seas, Adelie Coast, Prydz Bay).

We show additional sample output from the water mass classification along repeat GLODAP lines. Figure S5 shows the classification along the I06S line in 2008, Figure S6 shows the classification along the I08S line in 2007, and Figure S7 shows the classification along the P18 line in 2007. The panels are plotted as scatter plots of point-wise observations.

We also provide maps of the absolute layer thicknesses, derived from the ML model applied to the Argo data. These are shown to complement the thickness change discussion in Section 3.2. Figure S9 shows the averaged 20-year integrated layer thickness for CDW, AABW+DSW, SAMWs and AAIW.

S2.2 Layer Thickness Calculations

In Section 3.2, we compute the layer thicknesses from the water mass classification output. The layer thickness at each grid point can be described as H, where D is depth, WMF is the water mass fraction and p is the total number of vertical levels:

$$H_{i,j} = \sum_{z=1}^{p-1} \left(\left(\frac{D_{z+1} - D_z}{2} + \frac{D_z - D_{z-1}}{2} \right) \times WMF_{i,j,z} \right) \tag{8}$$

We can then approximate the volume of each water mass at any given grid point, using an approximation for the surface area of each grid square:

$$v_{i,j} = H_{i,j} \times R^2 \times \Delta\lambda \times \Delta\Phi \times cos(\Phi) \tag{9}$$

Where R is the radius of the earth, λ is longitude and Φ is latitude. The total water mass volume in any given domain can therefore be estimated by integrating in both directions, such that $V_{WM} = \sum_{i=0}^{n} \sum_{j=0}^{m} v_{i,j}$, where n and m are the total number of meridional and zonal indices, respectively.

S3 Supplementary Information for Discussion and Conclusions Section

S3.1 Sensitivity Testing of the Poleward Expansion of CDW

S3.1.1 Varying Δ Values

Extreme AABW Warming In this section, we test the sensitivity of our conclusions to variations in the delta values that we pick to constrain the changes in end members over the 20-year Argo period. In particular, it is important to ensure that the major trend of poleward migration in the upwelling CDW water mass is not primarily a product of a substantial warming in the AABW end member, which is not captured by a AABW ΔT value that is too low. Our methodology for isolating the 20-year source water change (discussed in Section 5.5) gives an AABW ΔT value of 0.077. As we discuss, this is consistent with the range of observed rates of AABW warming. Here we repeat the analysis, but this time assume a very high AABW ΔT value of 0.16°C - at the upper-extreme of the observed warming rate range. We re-classify the water mass distribution, using an otherwise identical method to Section 5.3. We then re-train the ML model on this output, and re-calculate the layer thickness changes from Figure 5. These are shown in Figure S11. We also re-compute the EOF analysis and show this in Figure S12.

Perturbing Δ Values Additionally, we perform some perturbation experiments on the chosen Δ values to determine to what extent poleward migration of CDW is sensitive to the choice of these values. We compute the standard deviation of the Δ values (as described in Section 5.5), and simultaneously perturb temperature and salinity for every water mass by 1 standard deviation in each direction. We then repeat the method in exactly the same way, training the RF model on the output of the classification with the perturbed Δ values. Figure S13 shows that poleward migration of CDW remains the first EOF of the CDW layer thickness timeseries from Argo data in both perturbation scenarios.

S3.1.2 Training the ML Model on the Climatology Dataset

We also train a ML model on the output of the water mass classification in the GLODAP climatology dataset. This model is identical to that described in the main manuscript, except that it is trained on the GLODAP climatology OMP output described in Section 5.1. There is therefore no consideration of early and late periods, Δ values, or how end members may have changed during the 20-year Argo period that the data is applied to. The output of this model provides a baseline with which to compare our variable-end member model.

We show the layer thickness change derived from the climatology model applied to Argo data in Figure S14. Comparing with equivalent Figure 5 in the main manuscript reveals that there is very little change in the dominant trends between the two models. Likewise, we repeat the EOF analysis on the CDW layer thickness time series in Figure S15, and find that the poleward migration in CDW is still the dominant mode of non-seasonal variability.

This has a variety of important implications. Firstly, the high degree of agreement between the two models suggests that we don't need much spatial coverage to recover most of the water mass distribution and trends; the model in the main manuscript is trained on individual GLODAP lines (shown in Figure 1), whilst the one we show here is trained on the full 1°by 1°climatology grid with 58 vertical levels. Secondly, it implies that the principal trends in water mass distribution that we document in the main manuscript have a low degree of sensitivity to changing end members and the addition of Δ values. This is a useful indicator that the trends which we observe are not in any way the product of the use of Δ values.

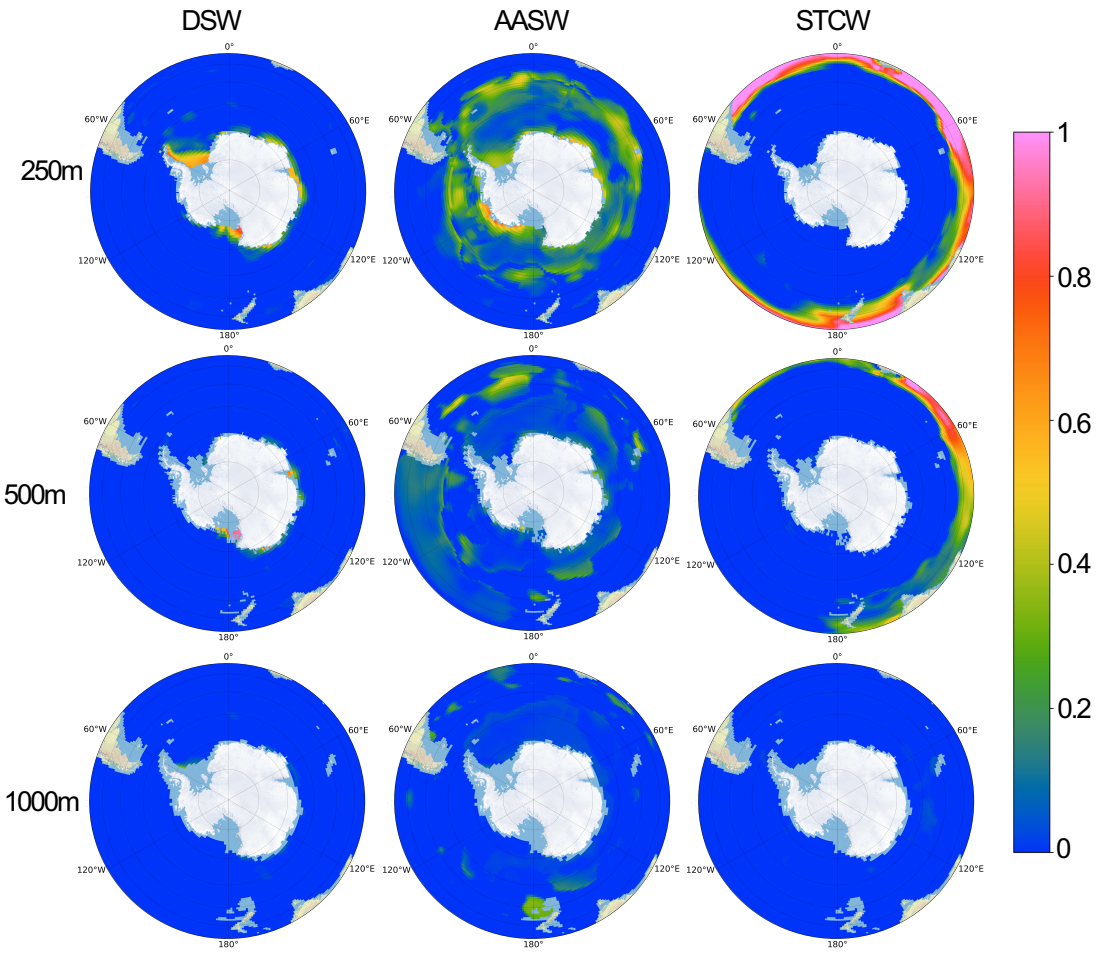

Figure S1: Additional samples of water mass classification output in the GLODAP climatology ('mean-state'), for DSW, AASW, and STCW at a range of depths.

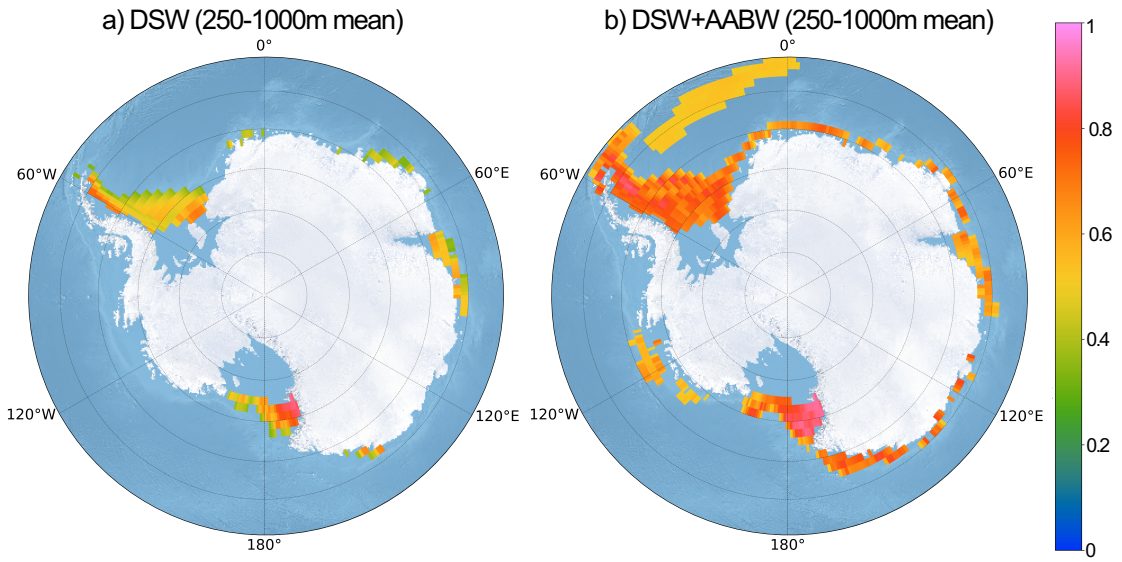

Figure S2: Additional samples of DSW and DSW+AABW output from the GLODAP climatology ('mean-state') water mass classification, showing the principal sites of DSW formation on the continental shelf. The plots show the mean concentration within the depth range 250-1000 m. Low concentrations are removed from the visualisation.

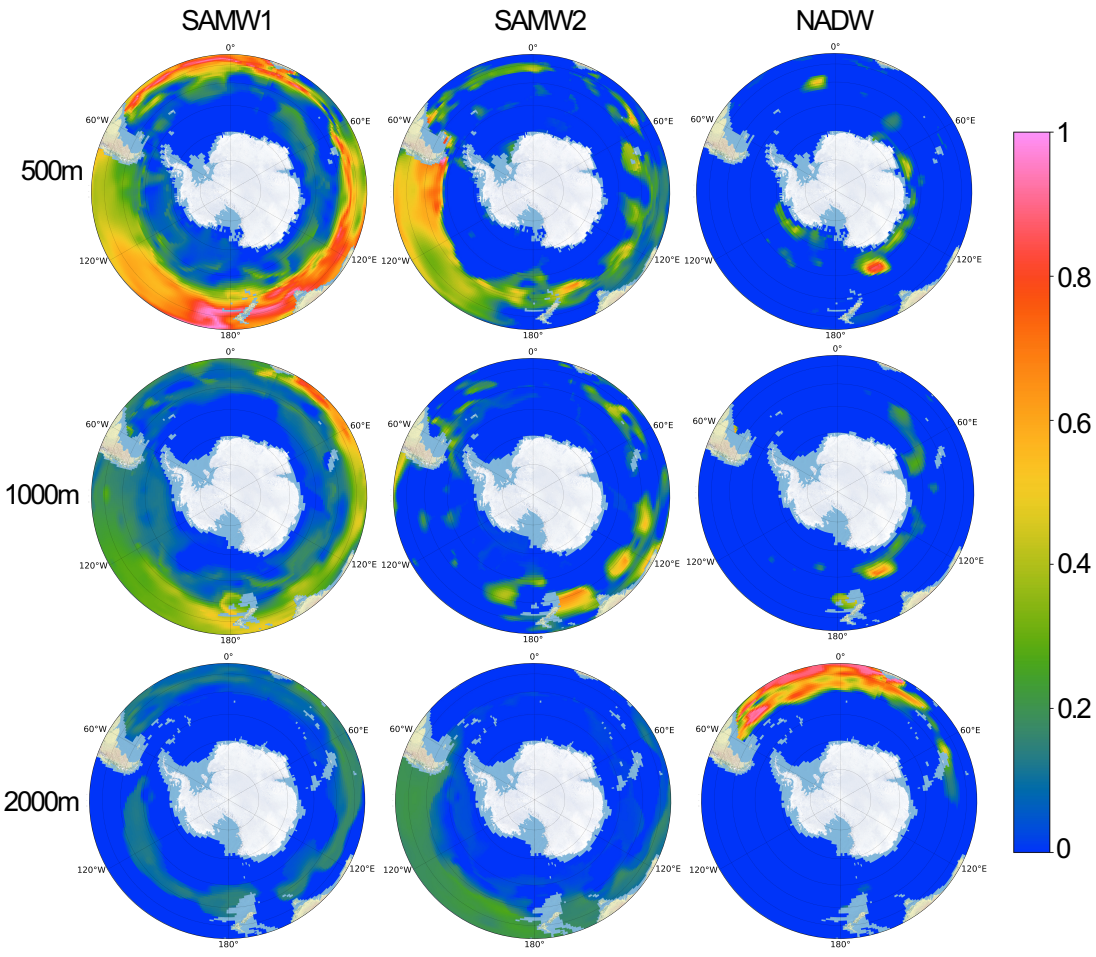

Figure S3: Additional samples of water mass classification output in the GLODAP climatology ('mean-state'), for SAMW1, SAMW2, and NADW at a range of depths.

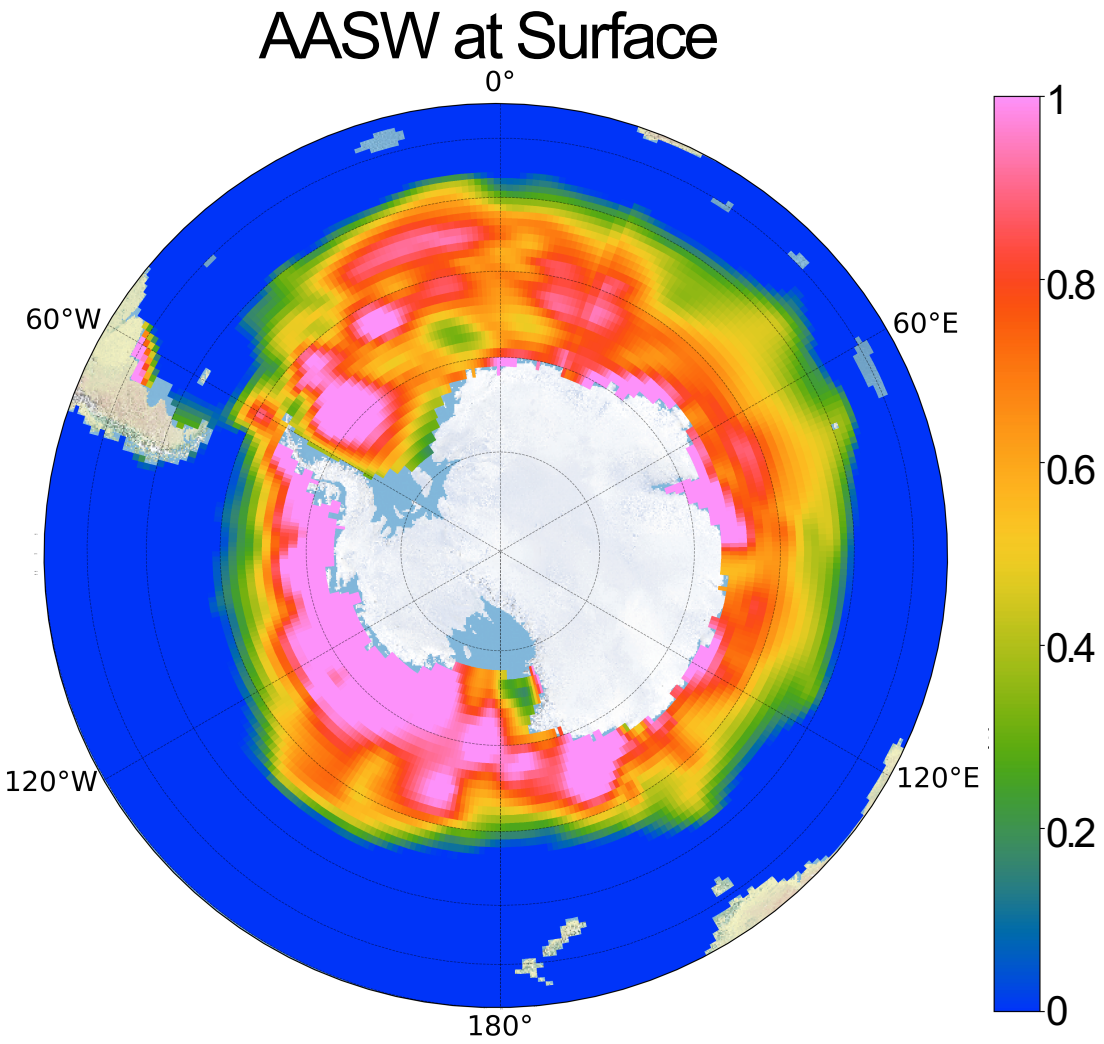

Figure S4: Concentrations of AASW at the surface in the GLODAP climatology ('mean-state') water mass classification.

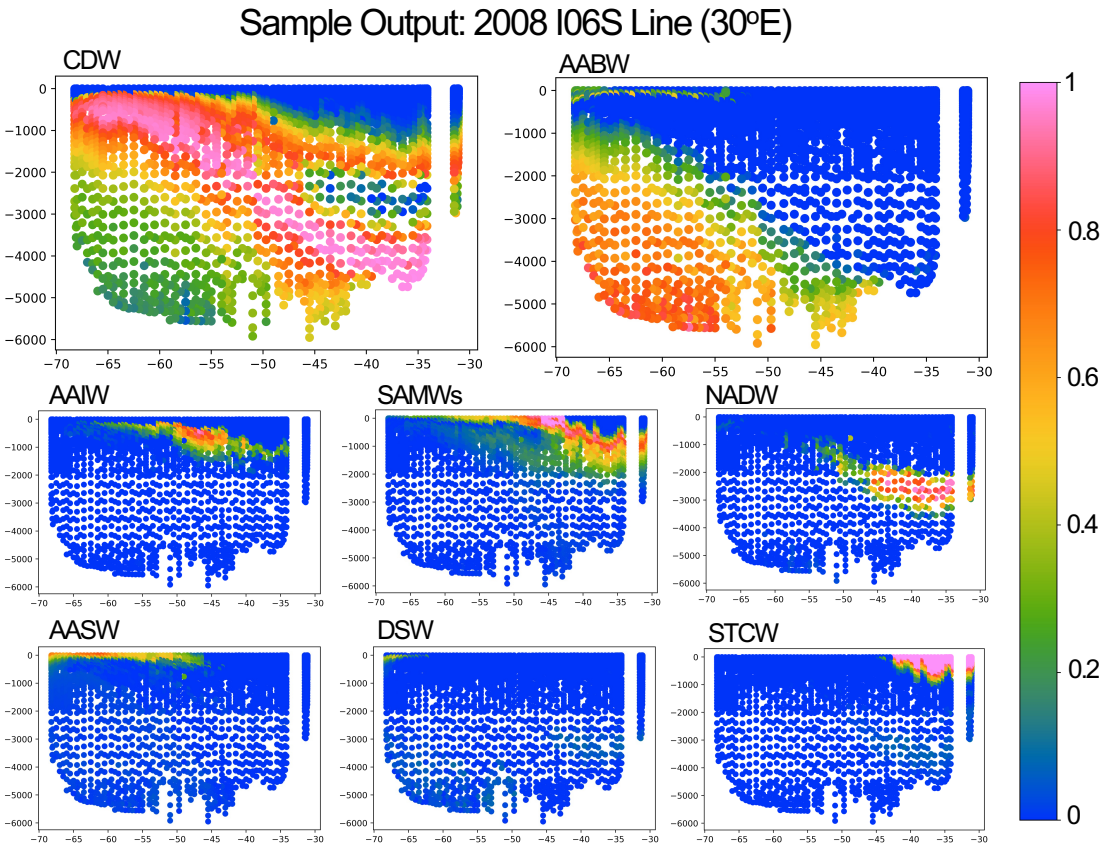

Figure S5: Additional sample of water mass classification along the I06S line.

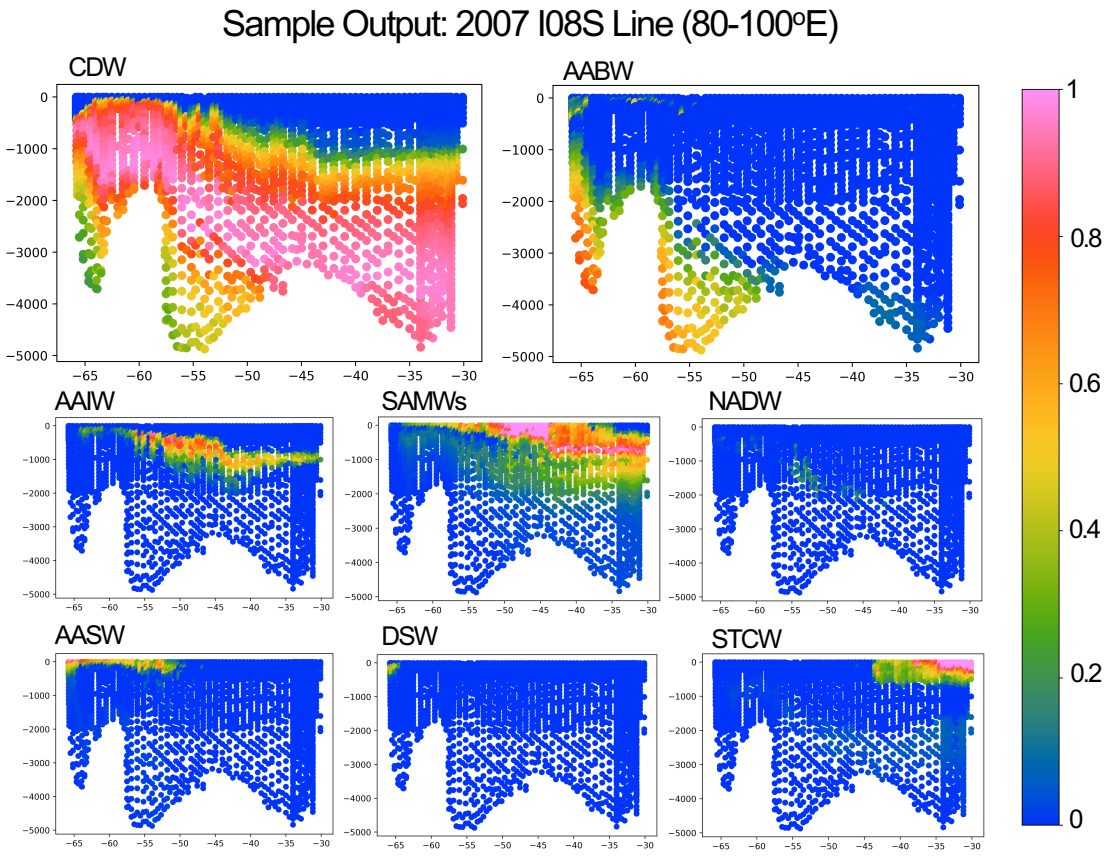

Figure S6: Additional sample of water mass classification along the I08S line.

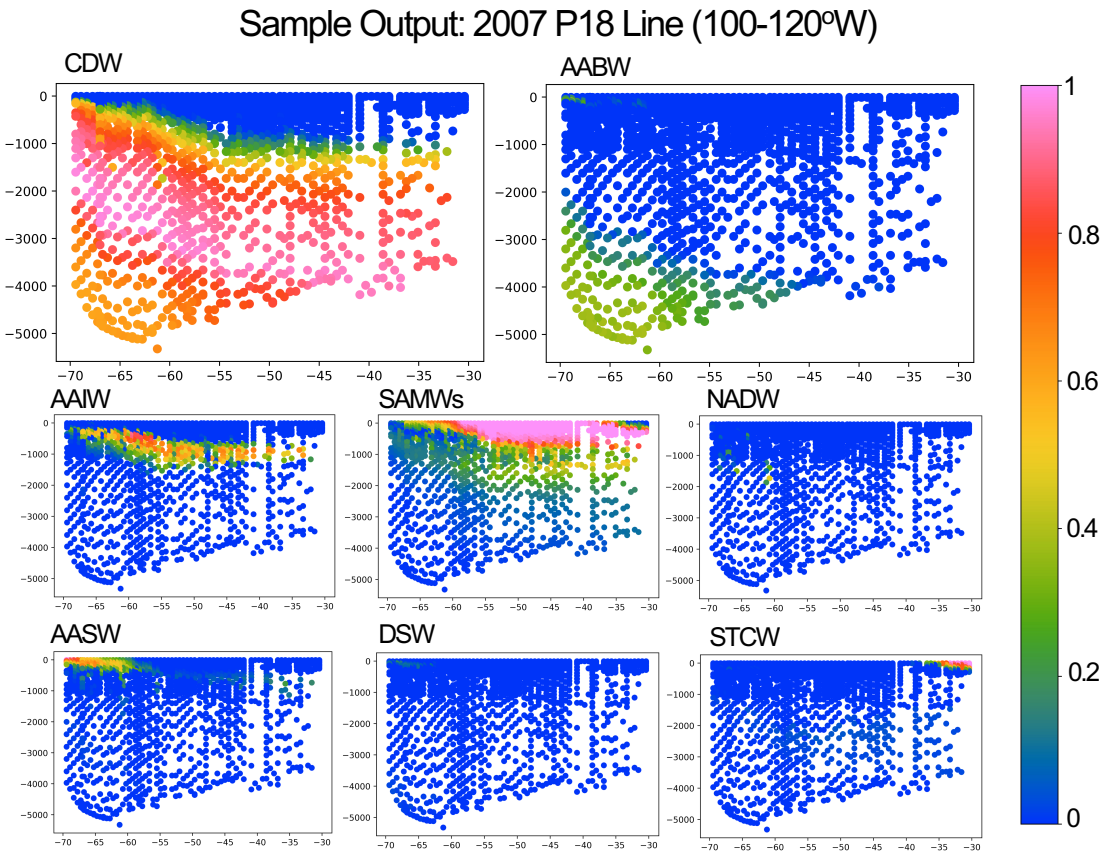

Figure S7: Additional sample of water mass classification along the P18 line.

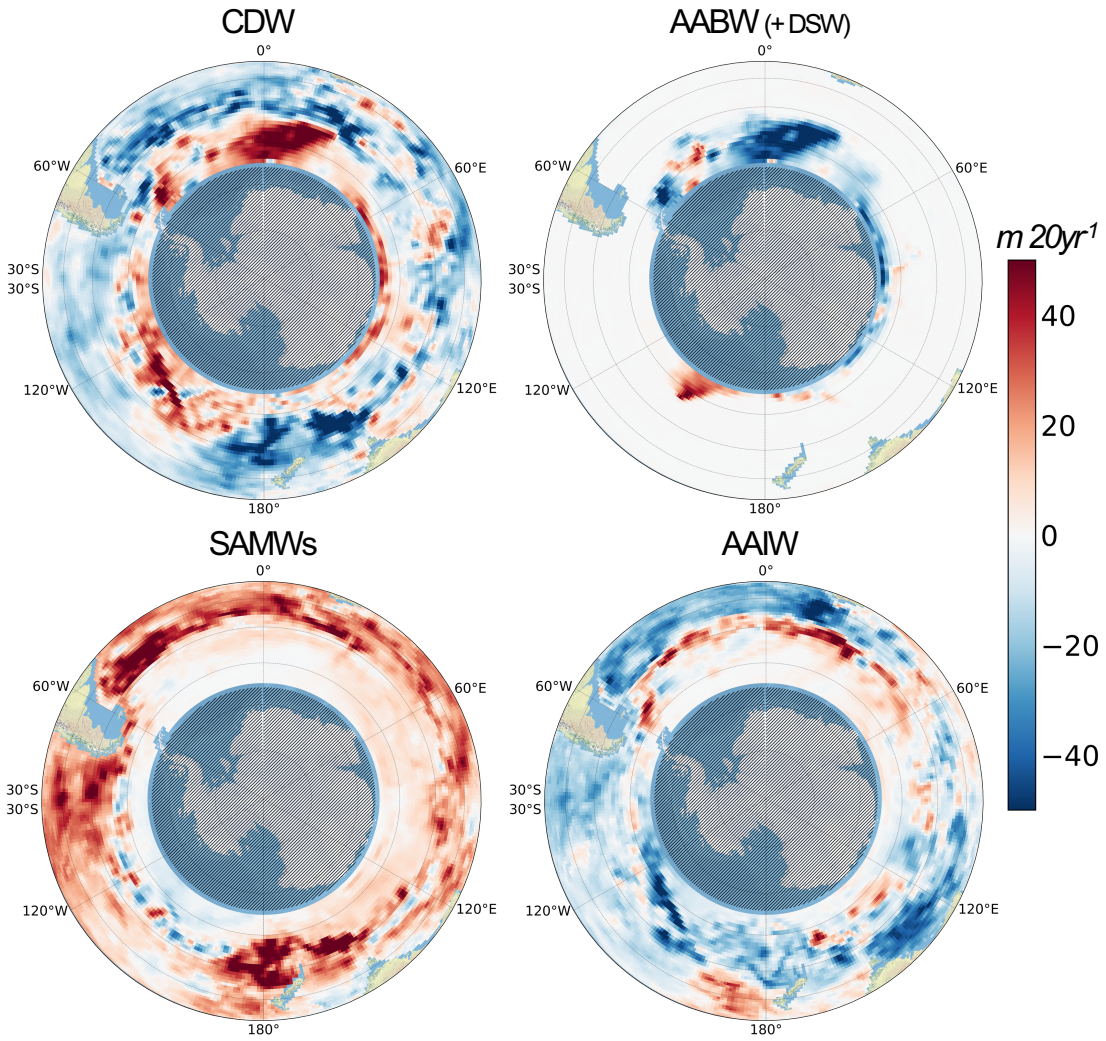

Figure S8: As for Figure 5, but showing the 20-year linear trend in integrated layer thickness.

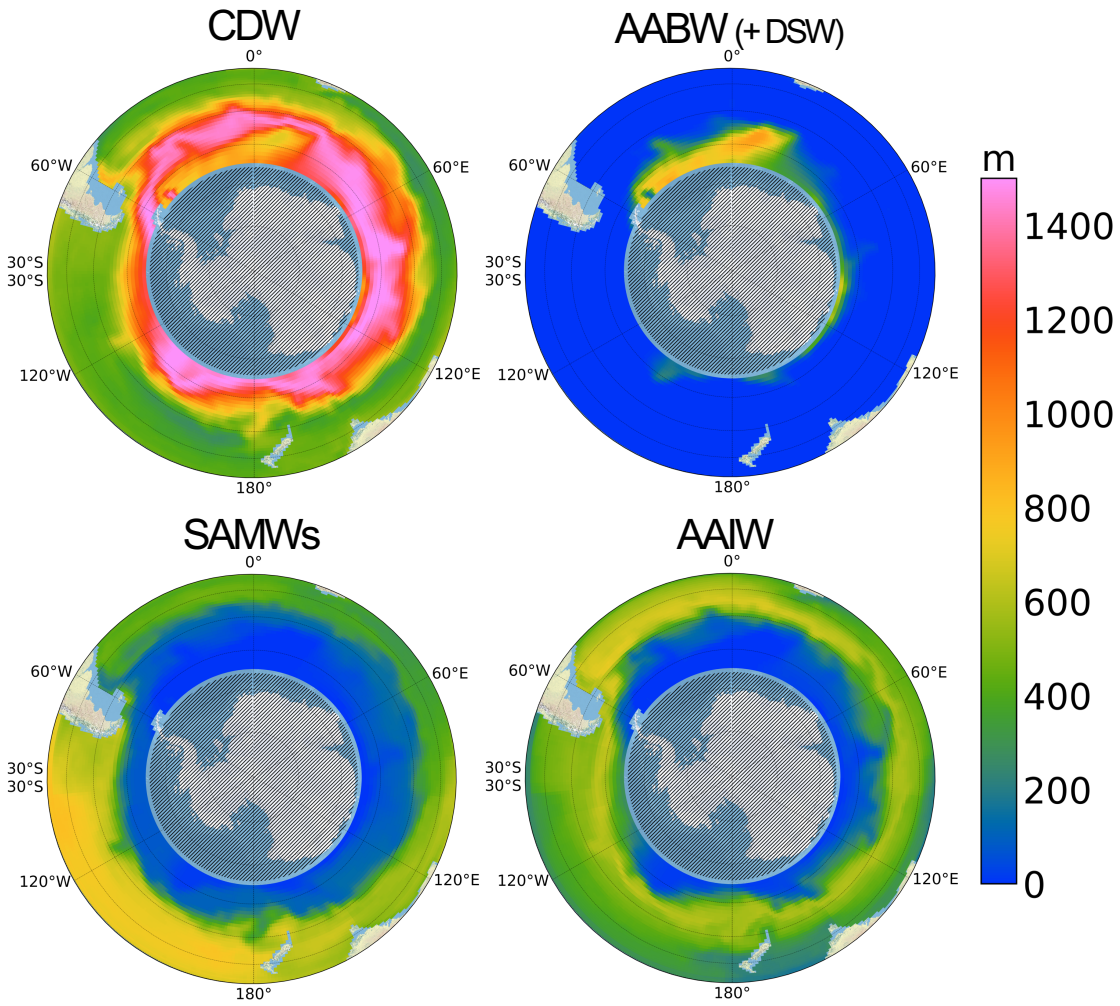

Figure S9: Absolute layer thicknesses in Argo data to accompany the changes shown in Figure 5, averaged over the 20-year period.

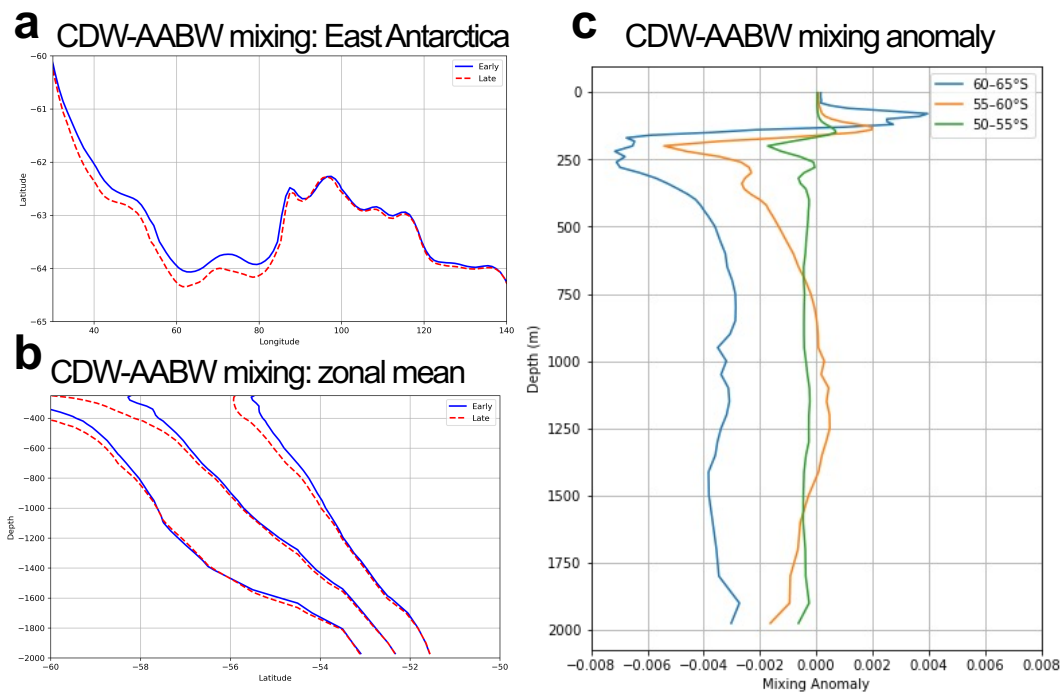

Figure S10: Mixing between CDW and AABW, calculated as the product of the OMP-derived CDW and AABW water mass fractions. Panel a) shows the longitudinal distribution across East Antarctica of the 0.08 mixing threshold contour for the early period (2004–2007 mean) and the late period (2021–2024 mean), illustrating the spatial extent of regions with relatively intense mixing. Panel b) shows the zonal mean mixing contours in the East Antarctic and Weddell Sea sectors for threshold values of 0.04, 0.10, and 0.15, highlighting the vertical and meridional structure of mixing intensity. Panel c) shows depth profiles of mean CDW–AABW mixing for discrete latitude bands, showing the vertical change in mixing strength. Together, these results indicate a modest southward migration of the primary mixing region over the study period, with no significant enhancement in mixing strength; the zonal mean and depth profiles instead generally suggest a slight decline in maximum mixing values (with the exception of a small increase in the 55-60°S band at mid-depths).

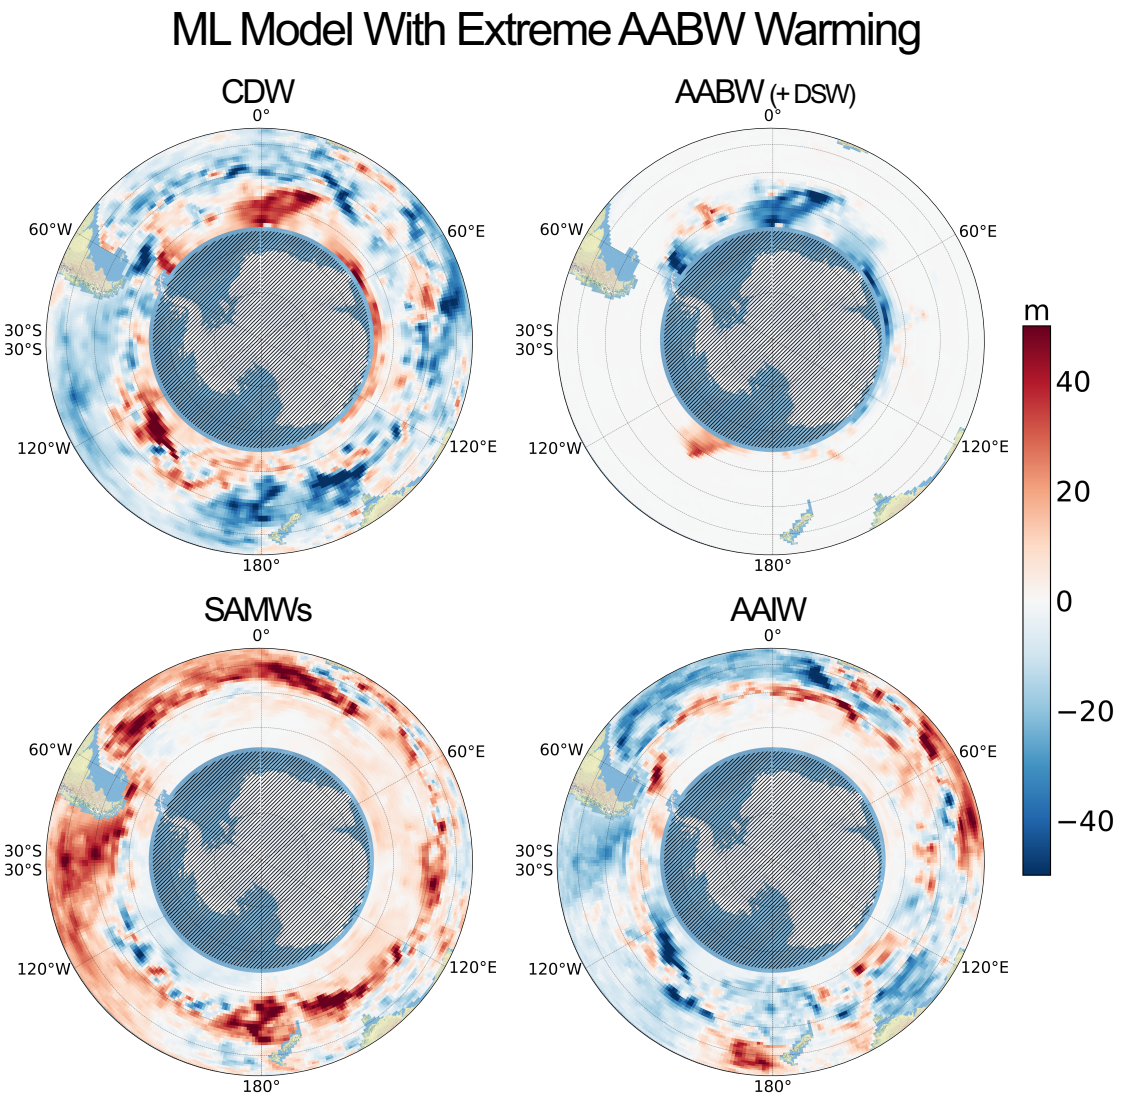

Figure S11: As for Figure 5, but using an RF model trained on a classification dataset in which an AABW ΔT value of 0.16°C is used.

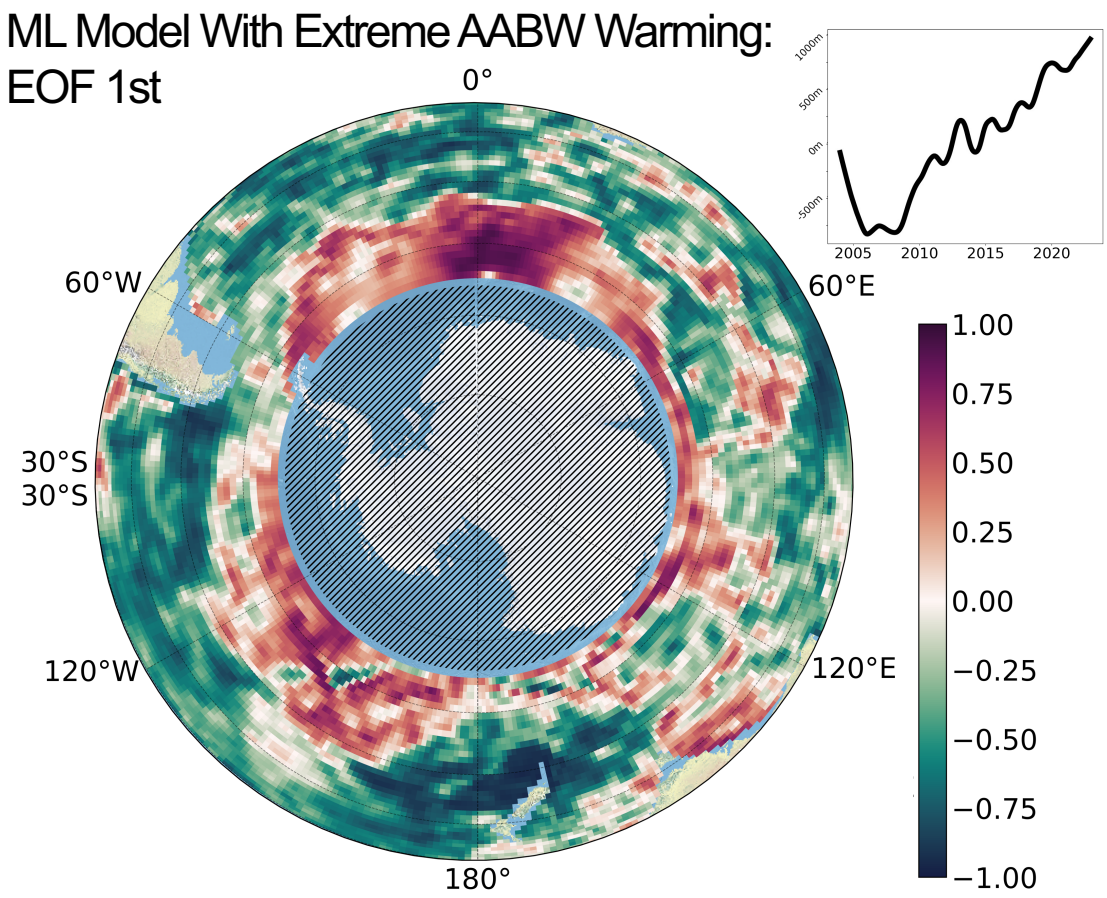

Figure S12: As for Figure 7, but using an RF model trained on a classification dataset in which an AABW ΔT value of 0.16°C is used. Note that only the first EOF mode and principal component are shown.

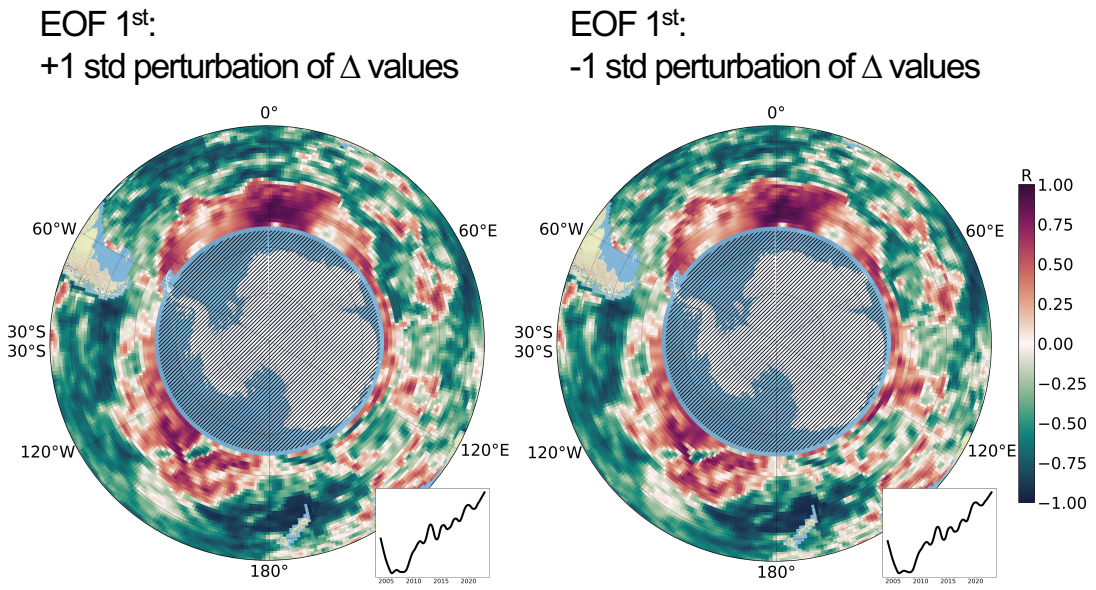

Figure S13: As for Figure 7, but using an RF model trained on a classification dataset in which ΔT and ΔS values are perturbed by 1 standard deviation in both directions. Note that only the first EOF mode and principal component are shown.

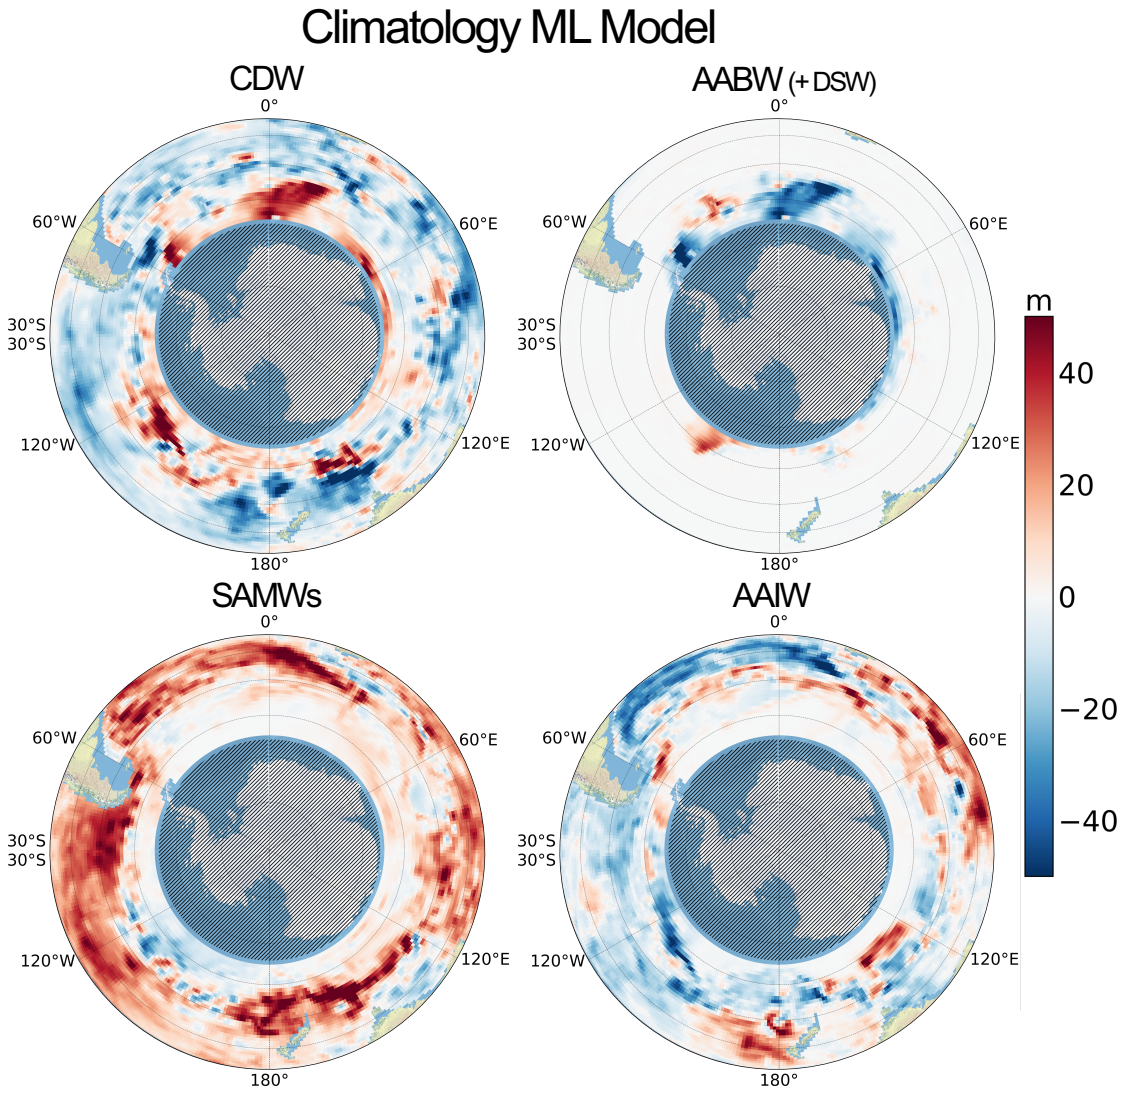

Figure S14: As for Figure 5, but using an RF model trained on the water mass classification from the GLODAP climatology ('mean-state').

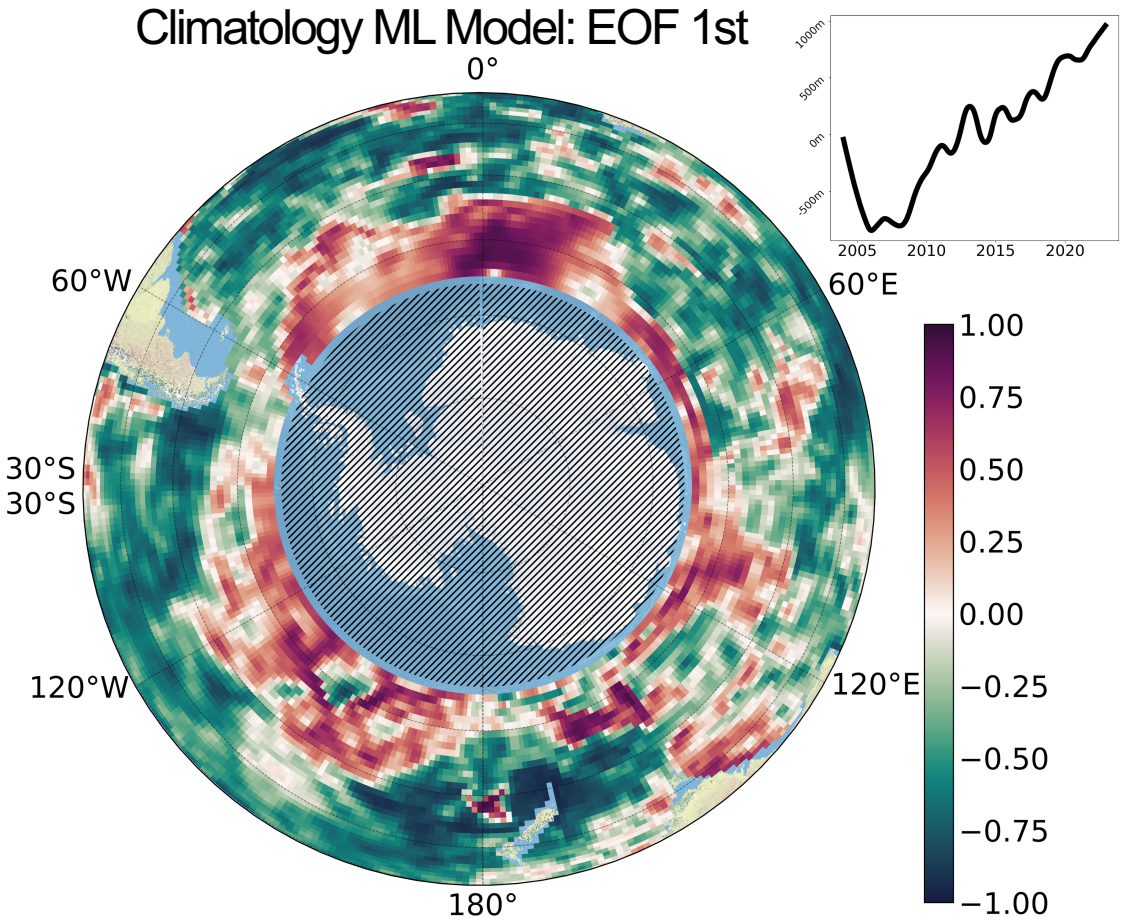

Figure S15: As for Figure 7, but using an RF model trained on the water mass classification from the GLODAP climatology ('mean-state'). Note that only the first EOF mode and principal component are shown.

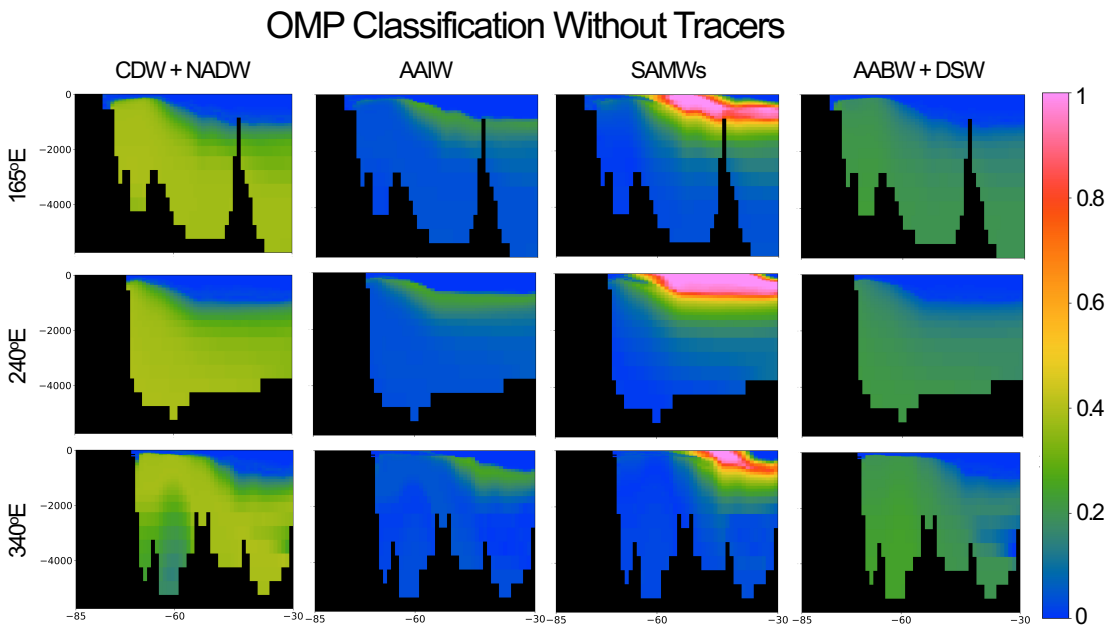

Figure S16: As for Figure 3, but using an OMP water mass classification in the GLODAP climatology which uses just temperature and salinity.

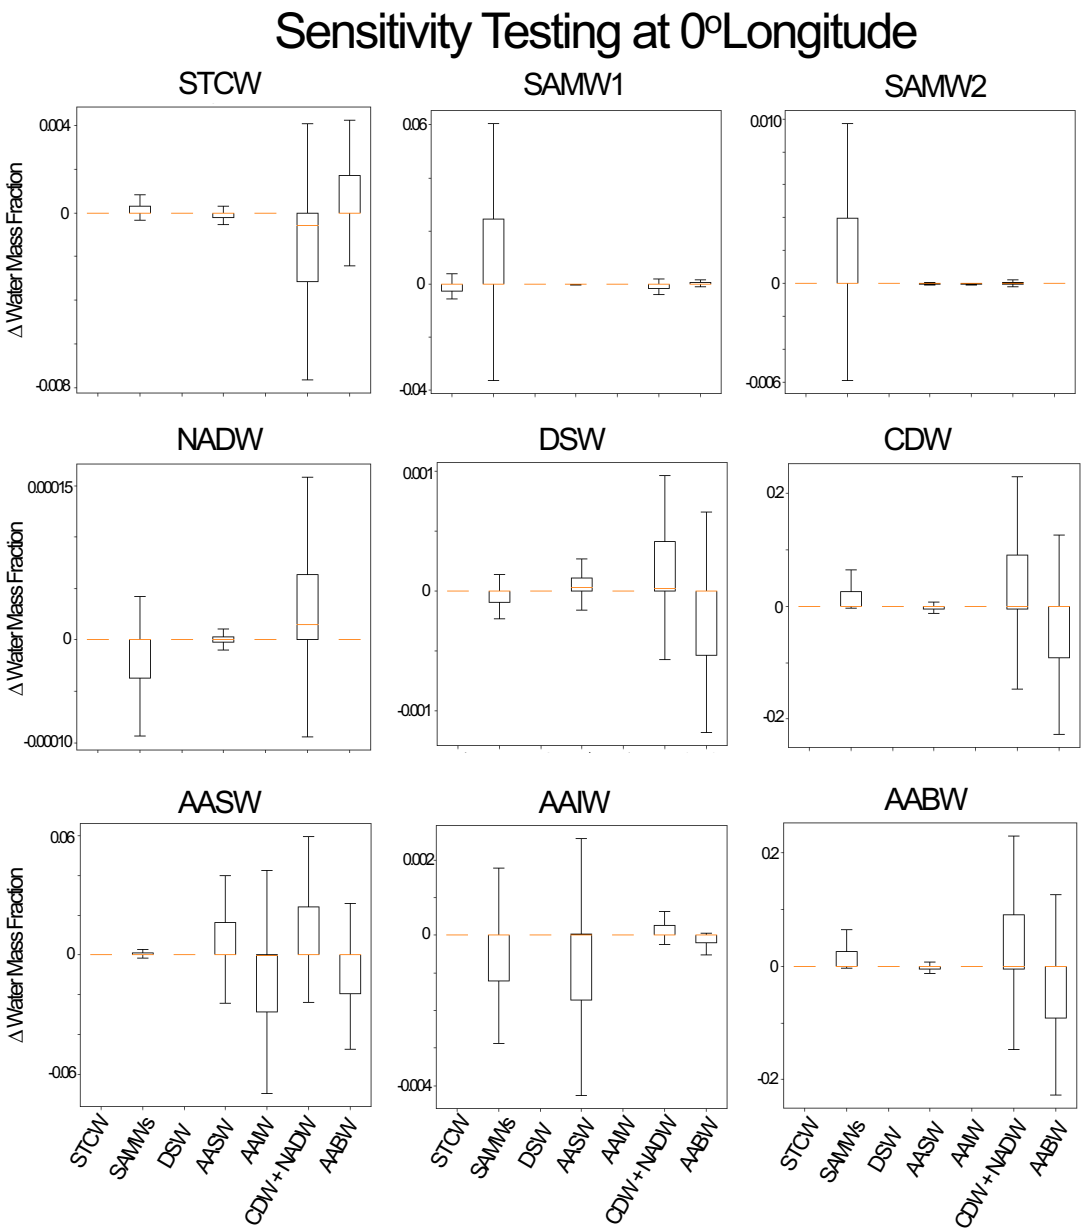

Figure S17: Sensitivity testing of the climatological OMP solution at 0°E. For each water mass, the properties of the source water are approximated by taking the 1000 grid points of the highest relative fraction in the original climatological OMP solution. Within this 'source water' mask, the standard deviation of each variable is calculated and used as a perturbation to investigate the sensitivity of the solution to end member definitions. The results shown here are the result of the simultaneous perturbation of the solution by 1 standard deviation in all variables (i.e. the STCW panel shows just the impact of varying the STCW end-member definition of temperature, salinity, oxygen, nitrate, phosphate, and alkalinity). The boxplots show the median, iter-quartile range, and 1.5* the inter-quartile range of the distribution of the point-wise difference between the non-perturbed and perturbed solutions.

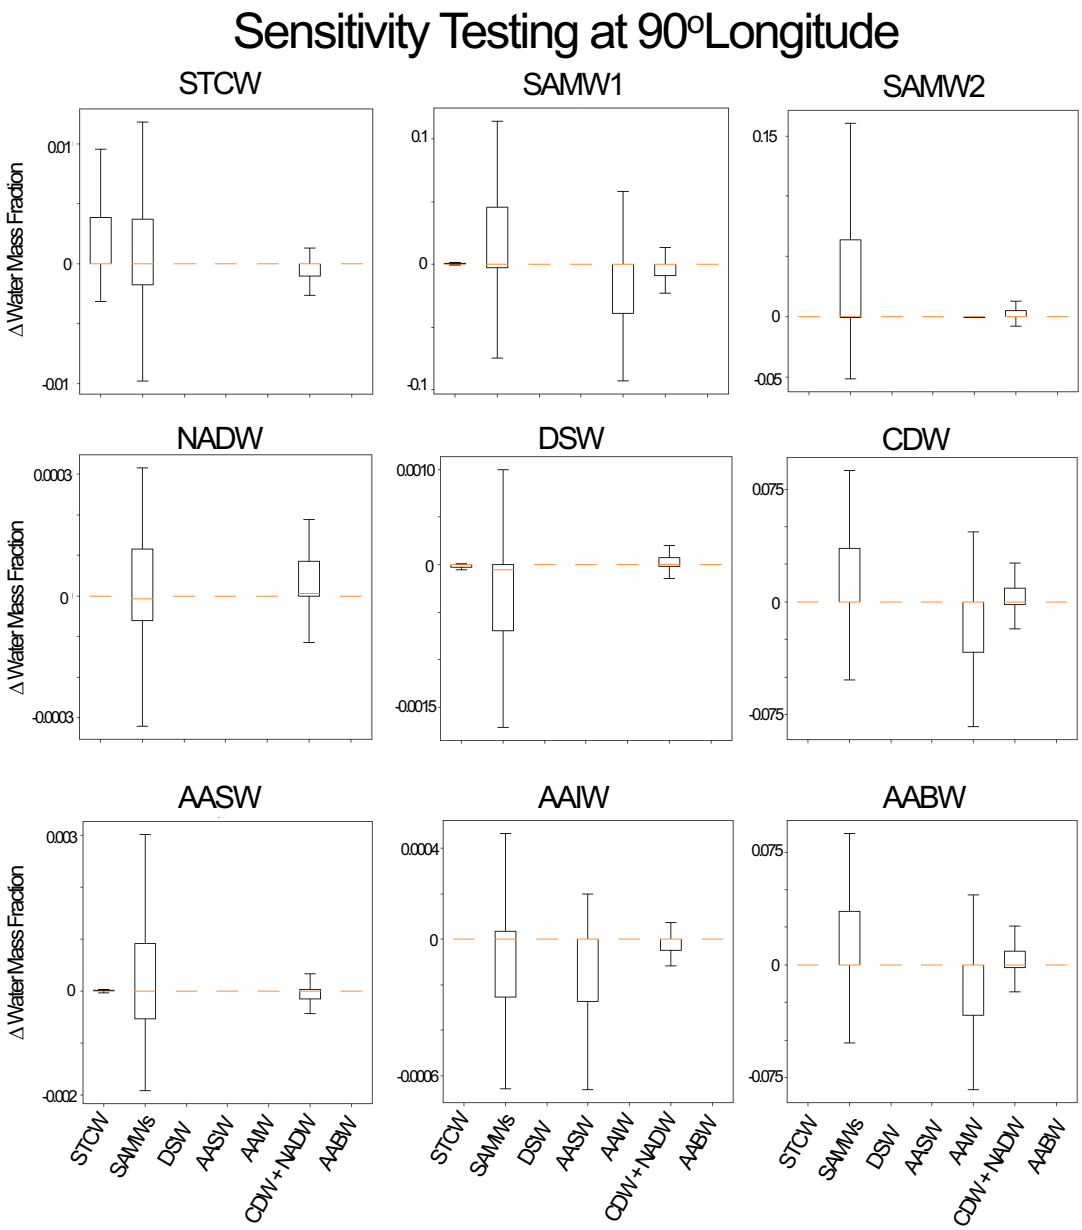

Figure S18: As for Figure S17, but for 90°E.

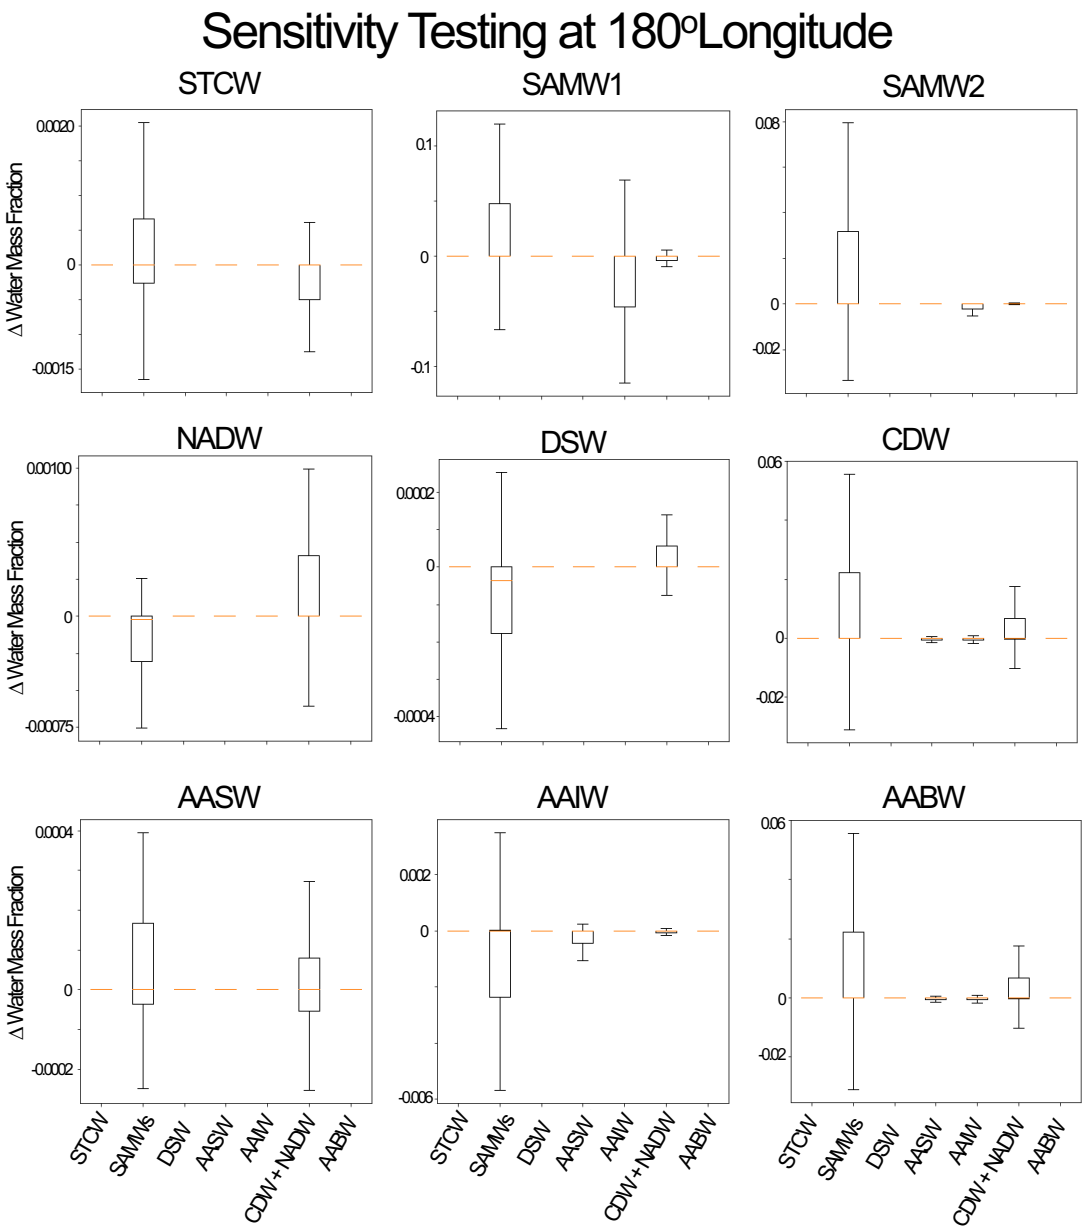

Figure S19: As for Figure S17, but for 180°E.

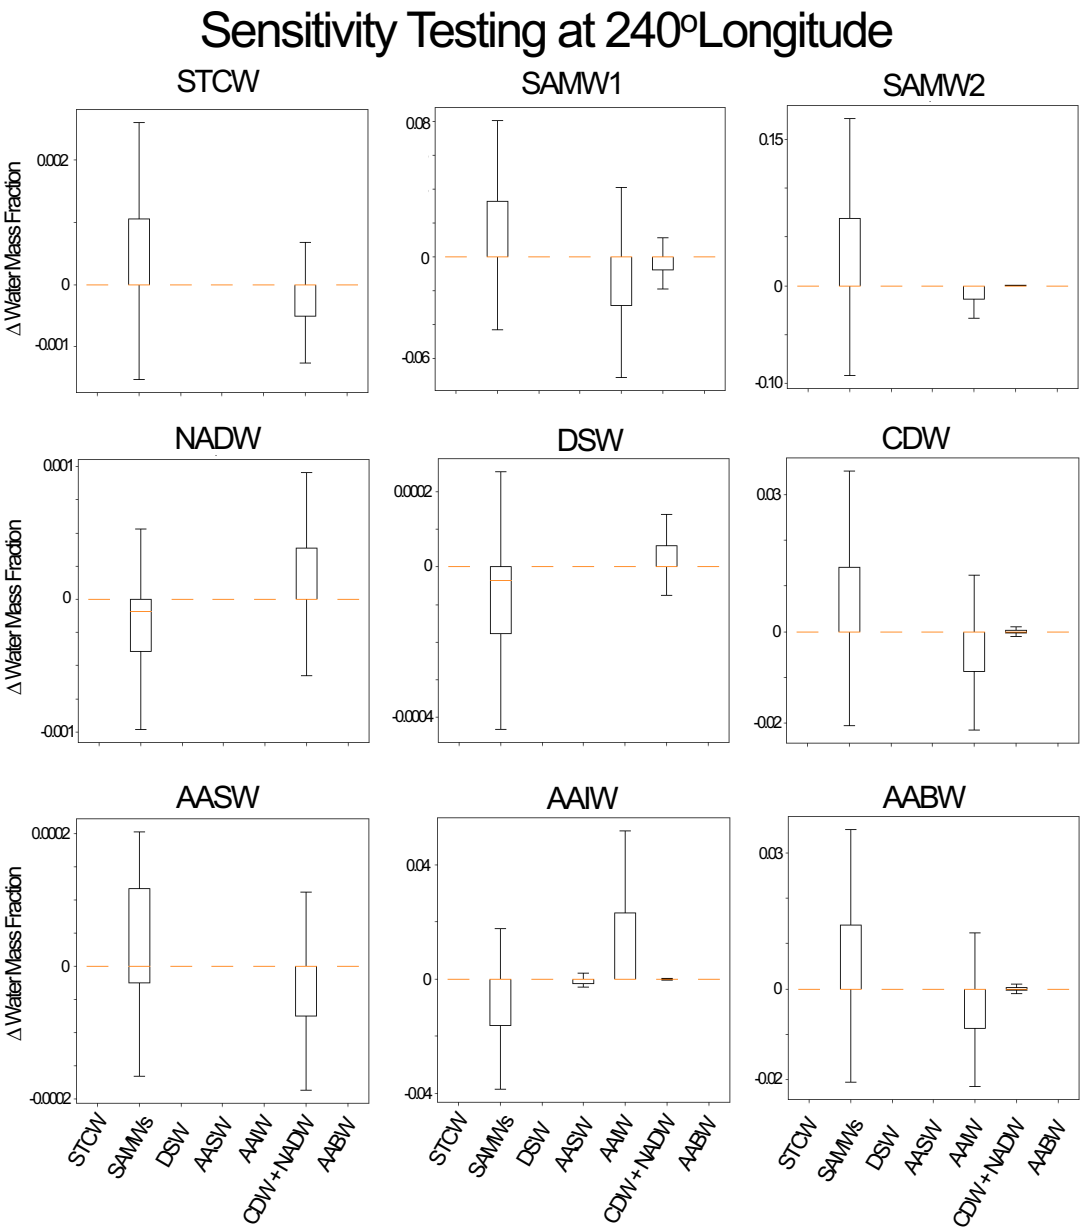

Figure S20: As for Figure S17, but for 240°E.

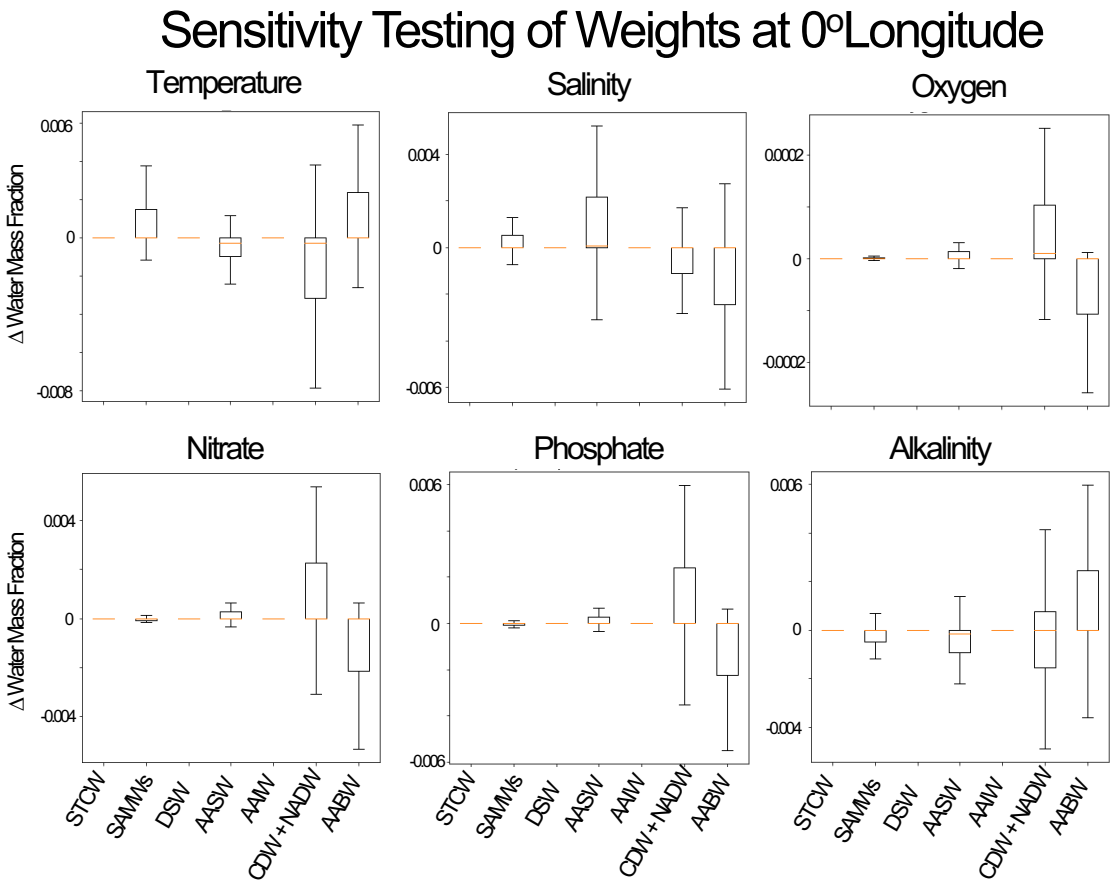

Figure S21: As for Figure S17, but testing the sensitivity of the OMP solution to perturbations in the weighting matrix. Each panel shows the impact on the solution of reducing the relative weight of each variable by 50%.

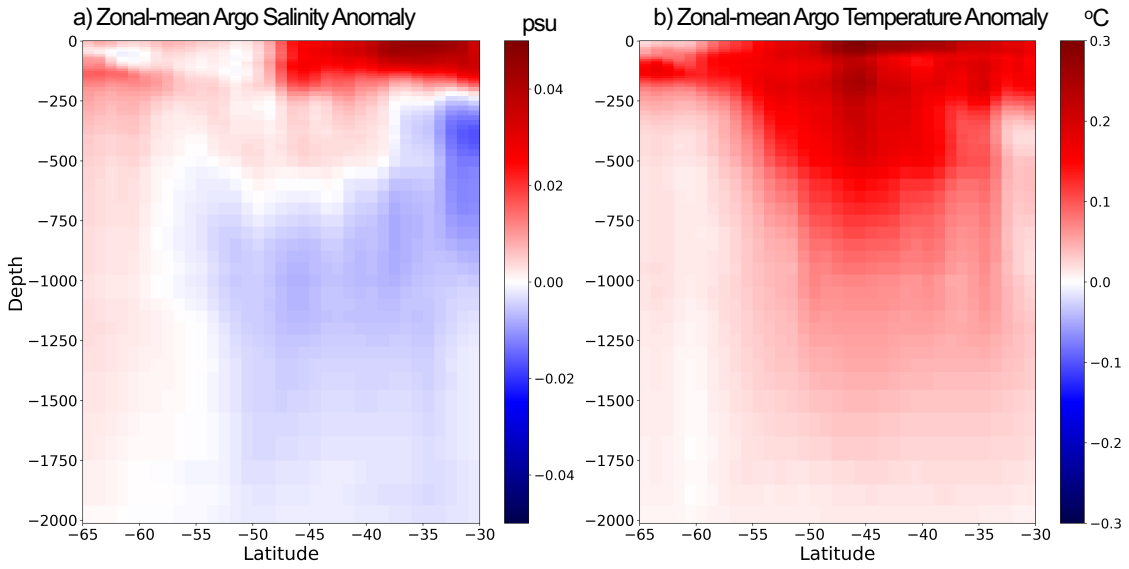

Figure S22: Zonal mean salinity and temperature changes in the Southern Ocean over the last 20 years, from the RG gridded Argo dataset. Panels a) and b) show a composite anomaly for salinity and temperature change respectively, calculated as the difference in means between the first and last 3 years of the 20-year Argo dataset.

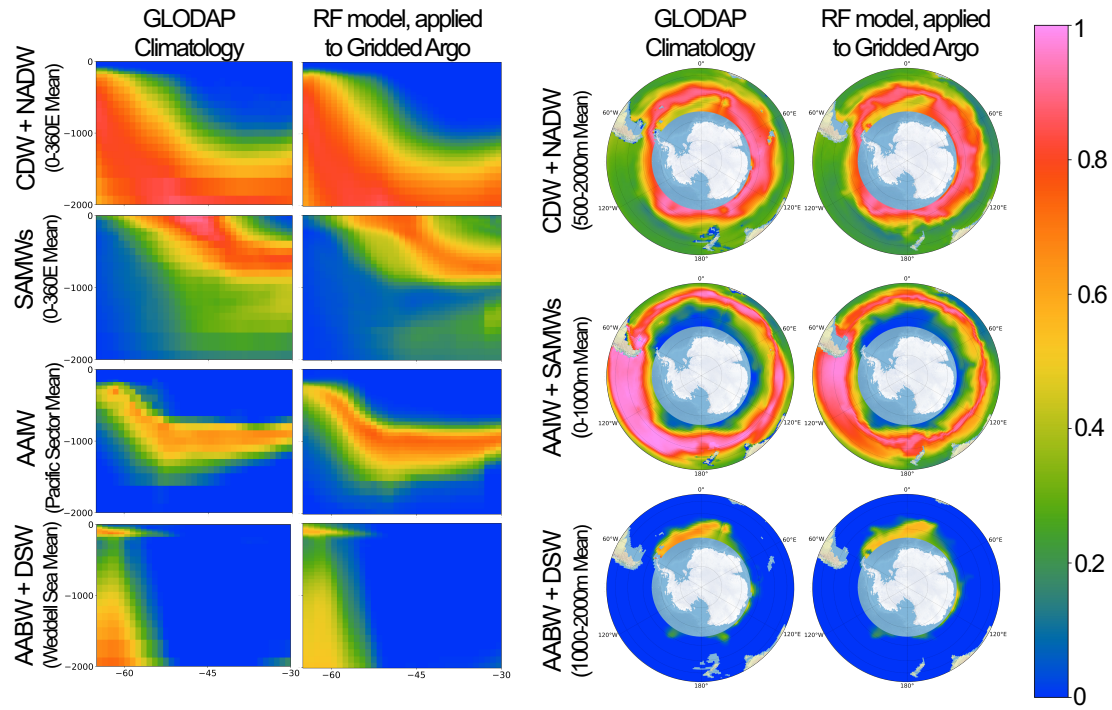

Figure S23: Comparison between water mass classification with the GLODAP climatology vs. from the machine learning model applied to Argo data.

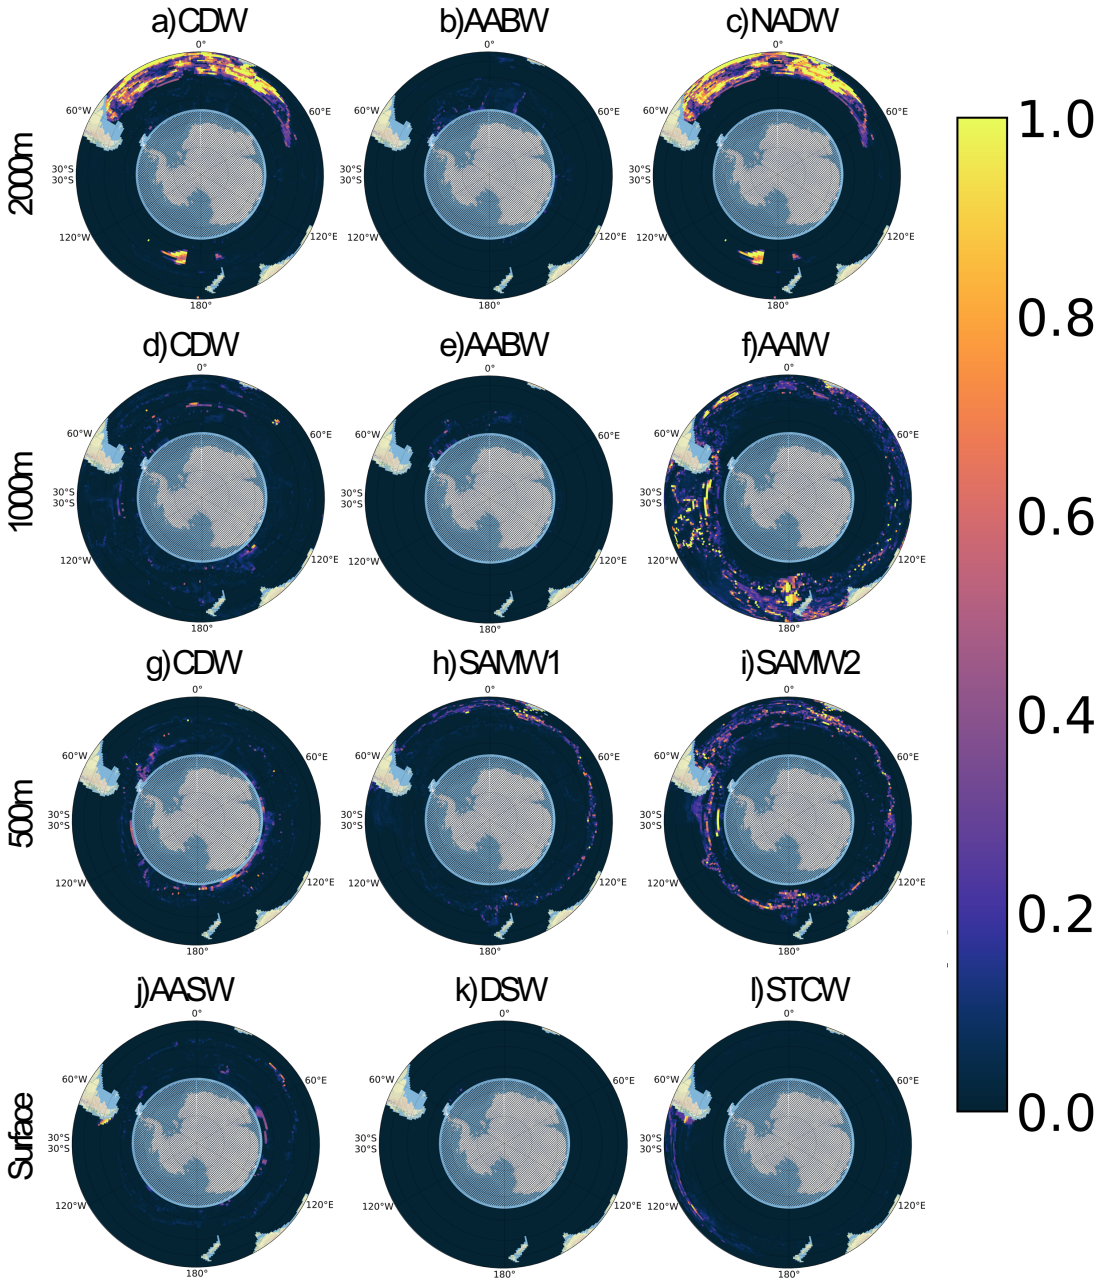

Figure S24: Uncertainty estimation in the application of the RF model to Argo data. Panels show the normalised variance across the 5 fold models at each grid point (see Section 5.6), for a range of water masses at a range of depths.