University of Southampton

Faculty of Environmental and Life Sciences

Ocean and Earth Science

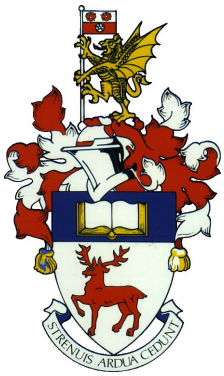

A Reconstruction of the Atlantic Meridional Overturning Circulation at 26°N

by

Emma Worthington

ORCiD: 0000-0002-6444-6461

Doctor of Philosophy

December 2021

University of Southampton

Abstract

Faculty of Environmental and Life Sciences

Ocean and Earth Science

Doctor of Philosophy

A Reconstruction of the Atlantic Meridional Overturning Circulation at 26°N

by Emma Worthington

A decline in Atlantic meridional overturning circulation (AMOC) strength has been observed between 2004 and 2012 by the RAPID-MOCHA-WBTS (RAPID –Meridional Overturning Circulation and Heatflux Array – Western Boundary Time Series, hereafter RAPID) array with this weakened state of the AMOC persisting until 2017. Climate model and paleo-oceanographic research suggests that the AMOC may have been declining for decades or even centuries before this; however direct observations are sparse prior to 2004, giving only "snapshots" of the overturning circulation. Proxy reconstructions extend the AMOC back in time, but tend to be either implicitly single-layer reconstructions of 20–30 years that do not capture all the observed variability, or longer reconstructions with high uncertainty and for the most part, not based on hydrography.

To bridge the gap between observations and proxy reconstructions, we developed a higher-fidelity statistical model of AMOC variability based on RAPID data and associated physically with changes in thickness of the persistent upper, intermediate, and deep water masses at 26°N and associated transports. We applied historical hydrographic data to the empirical model to create an AMOC time series extending from 1981 to 2016, and constrained the associated uncertainty by evaluating the statistical model within a high resolution ocean simulation.

Increasing the resolution of the observed AMOC to approximately annual shows multi-annual variability in agreement with RAPID observations, AMOC indices based on sea-surface temperatures, and palaeoclimate reconstructions. It captures a recovery of the AMOC during the 1990s and shows that the downturn between 2008 and 2012 was the weakest AMOC since the mid–1980s. This work creates the longest continuous AMOC reconstruction based on both shallow and deep sub-surface hydrographic data, and has created an independent AMOC index. The length of our AMOC time series is greater than reconstructions based on modern observations, and approaches the length of detectable climate-change related decline, but shows no overall AMOC decline as indicated by high-resolution climate models. Our results reinforce that adequately capturing changes to the deep circulation is key to detecting any anthropogenic climate-change-related AMOC decline.

Contents

List of Tables

List of Figures

Declaration of Authorship

I, Emma Worthington, declare that this thesis and the work presented in it are my own and have been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree of any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. When the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:
   (a) Worthington, E.L., Moat, B.I., Smeed, D.A., Mecking, J.V., Marsh, R., McCarthy, G.D., 2021. A 30-year reconstruction of the Atlantic meridional overturning circulation shows no decline. Ocean Sci. 17, 285–299. https://doi.org/10.5194/os-17-285-2021 (Worthington et al., 2021)

Acknowledgments

Firstly, I would like to gratefully thank my supervisors. Firstly, Gerard McCarthy, without whom this thesis would absolutely not exist. I could not have asked for more patience, support, belief, and encouragement during this PhD, and I'm very glad I didn't change projects when you moved to Maynooth University! Secondly, Jenny Mecking, thank you for your help and advice throughout, and especially for helping me get to grips with NEMO and JASMIN. Lastly, thanks to Bob Marsh, not only for dealing with administrative duties, but for your experience and ideas. Thank you all.

I would also like to thank my PhD panel chair Penny Holliday, who has been a fantastic source of support and excellent advice throughout. I would also like to thank Eleanor Frajka-Williams for encouraging me to start this PhD journey and co-authoring/walking me through publishing my MSc dissertation during the early part of my PhD. The RAPID project team have given me invaluable help, thanks especially to Ben Moat for his help with the AMOC code, and Ben and David Smeed for their contributions as co-authors. Thanks also to Harry Bryden, who helped me find data, answered and asked questions, and is a supportive presence for many NOC PhD students.

This PhD research was funded by the Natural Environment Research Council's SPITFIRE Doctoral Training Partnership, and I would like to thank them for the opportunity, and the SPITFIRE DTP and Graduate School administrative staff. This project uses observation data from the RAPID-MOCHA-WBTS array, where RAPID was funded by the U.K. Natural Environment Research Council, MOCHA was funded by the National Science Foundation, and WBTS was funded by the National Oceanic and Atmospheric Administration. I'd also like to thank Chris Roberts of the Met Office for his open source RapidMoc code, and for helping with its use and configuration.

Thanks to my PhD office mates and friends, especially Sam Hallam, Jess Cartwright and Liz Harris, for the mutual cheerleading, sympathy, and good advice. I've been fortunate to be included as part of the Maynooth A4 research group, and it's been helpful during the pandemic to connect with them as both people and researchers. Thanks also to Conor Murphy and Niamh Cahill of ICARUS for their statistical advice.

Thank you to my wonderful friends and family for patiently listening to me trying to explain what I've been doing for four years. And most of all, thanks to Pete, for being the most supportive and understanding husband I could wish for, especially during this last very difficult year. I couldn't have achieved any of this without you.

Finally, I'd like to express my gratitude to those no longer with us. To my Dad, Frank Wheeler, who always encouraged my interest in maths and science. To Pat and Harry Worthington, who were a great source of academic encouragement as well as wonderful in-laws. And to our beautiful sweet idiot dog Ludo, thanks for always making me laugh.

Abbreviations

AABW Antarctic Bottom Water

AAIW Antarctic Intermediate Water

AIC Akaike's Information Criterion

AMO Atlantic Multidecadal Oscillation

AMOC Atlantic meridional overturning circulation

AMV Atlantic Multidecadal Variability

AR Auto-regressive

BIC Bayes Information Criterion

BLUE Best linear unbiased estimator

CCHDO CLIVAR and Carbon Hydrographic Data Office

CLIVAR Climate and Ocean: Variability, Predictability and Change

CMIP Coupled Model Intercomparison Project

CTD conductivity-temperature-depth

DPSS Discrete Prolate Spheroidal Sequences

DWBC Deep Western Boundary Current

EOF Empirical Orthogonal Function

GCM General Circulation Model

GLS Generalised Least Squares

GLSAR Generalised Least Squares with Autoregression

GSW Gibbs Seawater

IPCC Intergovernmental Panel on Climate Change

LNADW Lower North Atlantic Deep Water

LSW Labrador Sea Water

MAR Mid-Atlantic Ridge

MOC Meridional overturning circulation

NAC North Atlantic Current

NADW North Atlantic Deep Water

NEMO Nucleus for European Modelling of the Ocean

OGCM Ocean General Circulation Model

OLS Ordinary Least Squares

ORCA Oceanic Remote Chemical/Optical Analyzer

Q-Q Quantile-Quantile

RAPID RAPID–Meridional Overturning Circulation Heat Array–Western Boundary Time Series

RMS Root Mean Square

SE Standard Error

SSH Sea-surface height

SSR Sum of Squared Residuals

SST Sea-surface temperature

TEOS-10 Thermodynamic Equation of Seawater-2010

TSS Total Sum of Squares

UMO Upper Mid-Ocean

UNADW Upper North Atlantic Deep Water

VIF Variance Inflation Factor

WOCE World Ocean Circulation Experiment

WOD World Ocean Database

Chapter 1

Introduction

1.1 The Atlantic meridional overturning circulation

1.1.1 Overview

Ocean circulation consists of three-dimensional currents that move properties such as heat, carbon, nutrients, and freshwater. Heat generally moves from low to high latitudes as water warmed by solar radiation in the tropics is carried polewards by surface currents. Surface currents also carry cooler and fresher water from high to low latitudes. In addition to these wind-driven surface currents, there is a density-driven thermohaline circulation. At deep water formation sites at high latitudes and through a variety of processes water becomes sufficiently dense to sink and flow initially equatorward at depth (Fig. 1.1).

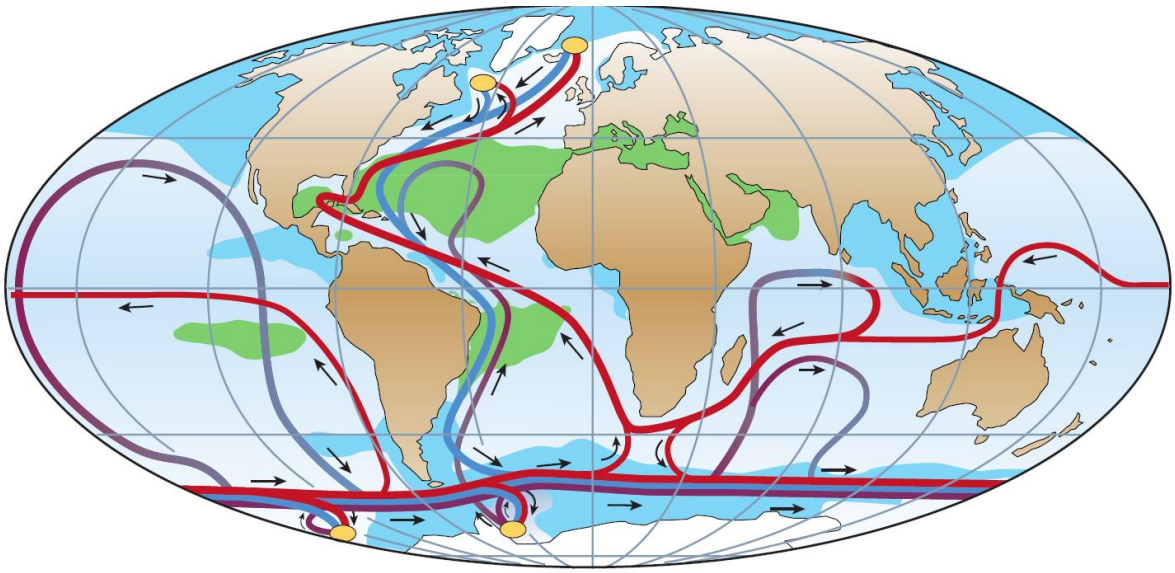

Figure 1.1: Schematic of the global thermohaline circulation. Surface waters are shown by red flowlines, deep waters in light blue and bottom waters in purple. Deep water formation sites are indicated by yellow circles. Green areas are high salinity (> 36 ‰) regions and blue areas are low salinity (< 34 ‰) regions. From Rahmstorf (2002), adapted from Broeker (1991).

In the present-day Atlantic, this thermohaline circulation is observed as the Atlantic meridional overturning circulation, or AMOC. The AMOC is characterised by the flow of warm, saline upper ocean water northwards from south of the equator, and a deep cold return flow. Fig. 1.2 shows how some of this warm, saline water comes from the Indian Ocean, carried by Agulhas rings (Biastoch et al., 2008; Gordon, 1986; Guerra et al., 2018), which move westwards across the South Atlantic and join the North Brazil Current to flow into the Gulf of Mexico. The water that leaves the Gulf of Mexico via the Florida Straits becomes the Florida Current, and this comprises the majority of the Gulf Stream. However the Antilles Current, a highly variable boundary current to the east of the Antilles and Bahamas, also makes a small contribution. Fig. 1.2 also shows how, of the northwards-flowing upper waters, a proportion recirculates back within the subtropical gyre, with the remainder reaching the sub-polar gyre via the North Atlantic Current, which is formed by the Gulf Stream combining with the Labrador Current. On reaching these high latitudes, the water releases heat to the overlying atmosphere, which helps to maintain north-western Europe's relatively mild climate for its latitude. The resulting buoyancy loss, combined with mixing, is sufficient for the water to sink to deeper than 3000 m to form North Atlantic Deep Water (NADW), which then returns equatorwards as the deep Western Boundary Current, forming the lower limb of the AMOC. The conventional view was that NADW was formed approximately equally by open-ocean convection in the Labrador Sea, and by Nordic Sea waters overflowing the Greenland-Scotland Ridge and entraining deeper water (Lozier, 2011; Pickart et al., 2003). However, recent studies using data from the OSNAP (Overturning in the Sub-Polar North Atlantic Programme, Lozier et al. (2017)) array have shown that the deep limb of the AMOC mainly comprises deep waters formed in the Irminger and Iceland basins and Nordic Sea (Lozier et al., 2019). Buoyancy forcing drives dense water formation in the eastern sub-polar region, and variability in its formation dominates variability in sub-polar overturning (Petit et al., 2020).

1.1.2 Transport of heat, freshwater, and carbon by the AMOC

The AMOC carries as much as 90% of all heat transported poleward by the ocean from the sub-tropics, and Atlantic Ocean meridional heat transport is disproportionately large compared to its relative volume. It is also uniquely northwards at all latitudes, unlike the Indo-Pacific, where the heat flux is poleward in both northern and southern

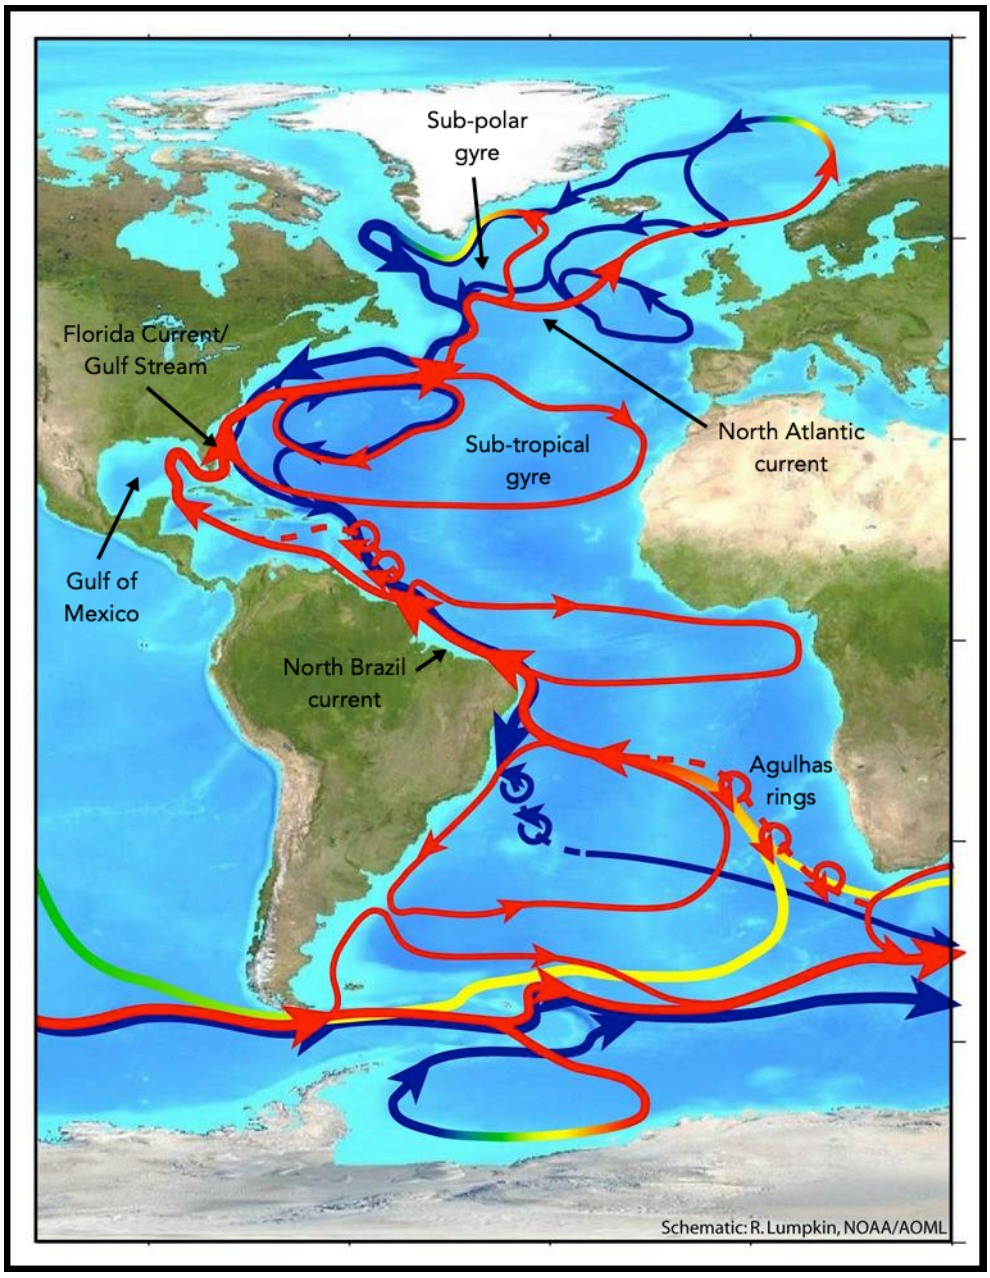

Figure 1.2: Schematic of the Atlantic overturning circulation, showing the surface waters as red flowlines, deep waters as blue and the transformation between them in yellow and green. Adapted from image by R. Lumpkin, NOAA/AOML.

hemispheres (Buckley and Marshall, 2016). The Atlantic Ocean meridional heat transport has a 10-day estimated mean of 1.3 ± 0.4 PW (1 PW = 10^{15} W) at a latitude of 26.5°N (Johns et al., 2011), which is close to the latitude of maximum heat flux. The AMOC also transports freshwater towards the equator (McDonagh et al., 2015), and the deep water formation at high latitudes moves carbon and heat into the deep ocean, helping to mitigate climate change (Kostov et al., 2014; Winton et al., 2013). The North Atlantic stores a disproportionately high 23% of global oceanic anthropogenic carbon dioxide, when its area is only 15% of the global ocean (Sabine et al., 2004). Between 1994 and 2007 the mean anthropogenic carbon dioxide storage at Atlantic mid-latitudes was double the global mean (Gruber et al., 2019).

1.1.3 The influence of the AMOC on climate

Meridional ocean heat transport by the AMOC impacts the climate of the North Atlantic region, moderating temperatures in north-west Europe when compared to the same latitude in north-west North America (Rhines et al., 2008; Seager et al., 2002). There is potentially a link between the AMOC and Atlantic Multidecadal Variability (AMV, also called Atlantic Multidecadal Oscillation or AMO), an index of low frequency North Atlantic sea-surface temperature anomalies, where warmer SST anomalies correspond to a positive AMV index, and colder anomalies correspond to a negative AMV index (Zhang et al., 2019). Although AMV has been linked to climate events as varied as decadal drought in the U.S. Mid-West (Patricola et al., 2015), multi-decadal variability in Indian and Sahel summer rainfall, and Atlantic hurricane activity (Zhang and Delworth, 2006), recent coupled-climate model experiments suggest there is no such internal multidecadal oscillation, and that these are signals forced by volcanic activity (Mann et al., 2021). The AMOC has however been shown to have a direct influence on the intensity of tropical hurricanes, with a weakened AMOC driving positive sea-surface temperature anomalies within their main development region, leading to more intense and thus damaging hurricanes (Hallam et al., 2019).

1.1.4 The AMOC in the past

Our knowledge of the AMOC prior to the 19th century is limited to palaeoclimate proxy reconstructions and climate model simulations. Examples of palaeoclimate proxy data are: oxygen isotopes in various organisms; relative abundance of plankton species; the ratios of trace metals in corals; and the grain size of sediments, related to deep boundary current velocities. Palaeoclimate proxy reconstructions show that over the past 120,000 years, the overturning circulation in the Atlantic has had three distinct modes: a stadial or 'cold' mode; an interstadial or 'warm' mode; and a Heinrich or 'off' mode (Rahmstorf, 2002). Our present circulation is typical of an interstadial mode: NADW is formed in the Nordic Seas, and the AMOC return limb is deep. The stadial mode had a shallow overturning profile, with the region of deep convection shifted to south of Iceland, a phenomenon that has been recreated in some coupled-climate models (Ganopolski and Rahmstorf, 2001). The Heinrich mode is so called because of its connection with Heinrich events, where vast amounts of freshwater were added to the North Atlantic by icebergs that had broken away from glaciers. In this mode, NADW production was minimal and deep Atlantic waters primarily originated in the Antarctic. The shifts between these modes were relatively abrupt. One example, the Younger Dryas, occurred approximately 12,000 years ago, and European pollen records and other proxies indicate there was a sudden change from a warm to a cold climate, which lasted around 800 years. The lack of an equivalent pollen record in North America (Broecker et al., 1985) and earth system model experiments (Meissner, 2007) suggest that it was related to a complete shut-down of the AMOC. Coupled ocean-atmosphere model simulations of an AMOC collapse show that impacts may include cooling temperatures and reduced precipitation over the northern hemisphere, and a southward shift of the Intertropical Convergence Zone (Vellinga and Wood, 2002).

Moving closer to the present day, Caesar et al. (2021) compared a number of independent reconstructions covering a wide range of time periods. Some of these used the palaeoclimate proxies mentioned above, but also included were reconstructions from sea surface and sub-surface temperature 'fingerprints', which are linked to meridional heat transport, and northern hemisphere temperatures using tree-ring and ice-core data. These reconstructions suggest there have been two AMOC declines since 1800: an initial

decline during the 19th century; and a faster decline from the 1950s. A slight recovery occurs during the 1990s before another decline after 2010, with the current AMOC state appearing to be the weakest for over a thousand years, and, with increased certainty, the weakest over the last 150 years. The post-1950s decline and recent short-lived recovery are obvious in Fig. 1.3 which shows the AMOC anomaly estimated from two different methods of SST fingerprinting. Caesar et al. (2018) created an AMOC index from average Nov-May sea-surface temperatures within a defined sub-polar gyre region, and subtracted the global mean SST. They used HadISST sea-surface temperature observations between 1870 to 2016. Rahmstorf et al. (2015) used annual HadCRUT4 sea-surface temperatures over a differently-defined sub-polar gyre region, and subtracted the northern hemisphere mean. Both used climate models to create a conversion factor between the AMOC index and AMOC anomaly. The decline and brief recovery seen in the mid- and late-twentieth century reconstructions are not observed in most historical simulations of coupled-climate model ensembles (e.g., Menary et al., 2020), although they are captured by a 120-year AMOC reconstruction at 50°N created using a Bernoulli inverse method and temperature and salinity observations (Fraser and Cunningham, 2021). However, between 1900 and 1930, this reconstruction shows a strengthening AMOC, which disagrees with the proxy reconstructions of Caesar et al. (2018); Rahmstorf et al. (2015); Thornalley et al. (2018), although the sparseness of early observations leads to high uncertainty in AMOC strength estimates before 1950.

Direct observations of the AMOC around 26°N go back to only 1957/58, the International Geophysical Year, when in October 1957 the Discovery II occupied the 24°N section (Fuglister, 1960). This section was not repeated until August–September 1981. AMOC strength has been estimated from these and later transatlantic hydrographic sections in July–August 1992, January–February 1998, April–May 2004 and January–February 2010 (Atkinson et al., 2012; Bryden et al., 2005). These 'snapshot' estimates of the AMOC appear to show a decline between 1957 and 2004 and then a re-strengthening in 2010. However the sparse resolution suggests that aliasing of the AMOC seasonal signal that has been observed by the RAPID array is extremely likely (Chidichimo et al., 2010; Kanzow et al., 2010). In addition to these hydrographic sections, there are extensive hydrographic data around 26°N going back to the early 1980s, particularly near to the western boundary, but there are insufficient simultaneous data at

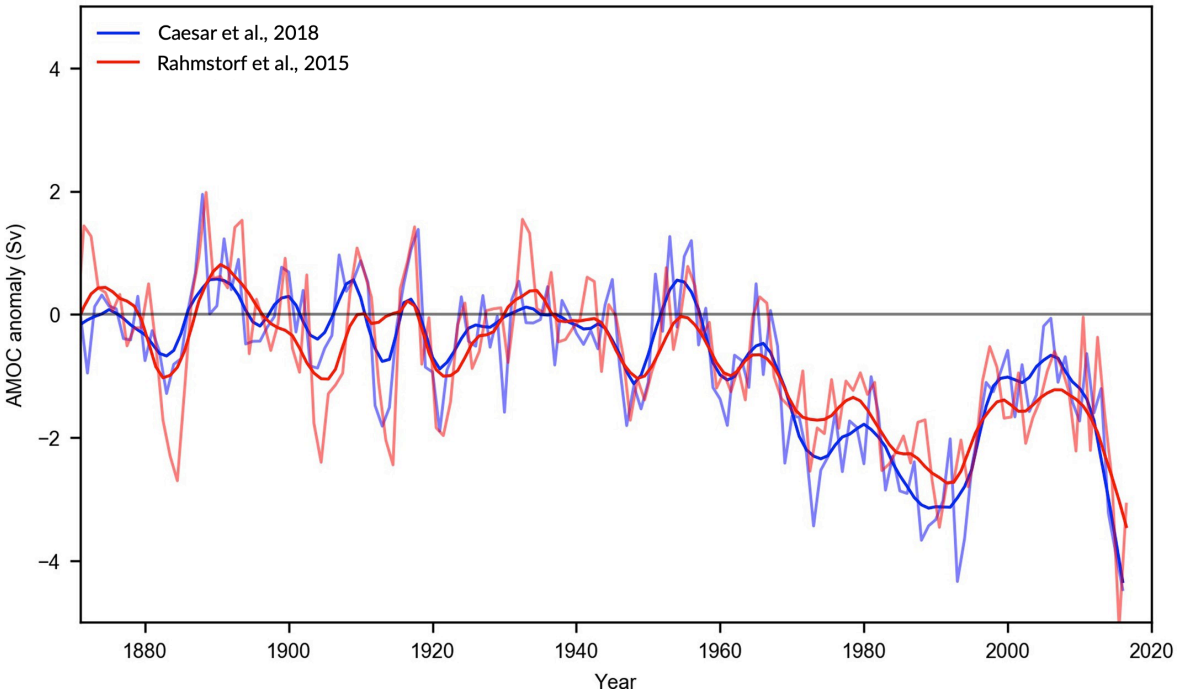

Figure 1.3: AMOC anomaly time series for two different methods of SST fingerprinting to create an AMOC index, adapted from Extended Data Fig. 6 of Caesar et al. (2018). The red line is the AMOC anomaly from Rahmstorf et al. (2015), which is based on HadCRUT4 data, and the blue line is the AMOC anomaly from Caesar et al. (2018), which is based on HadISST data.

both eastern and western boundaries to estimate the AMOC conventionally (Longworth et al., 2011). In an approach that was the starting point for this research, Longworth et al. (2011) used a mid-layer temperature anomaly at the western boundary as a proxy for geostrophic transport within the upper 800 m, or thermocline, finding that variability in the temperature anomaly explained 53% of the variance in the thermocline layer transport.

The era of satellite altimetry observations has led to several studies using sea-surface height anomalies as an effective proxy for the AMOC at mid-latitudes on interannual timescales. Zhang (2008) used satellite altimetry to identify a weakening of the sub-polar gyre during the 1990s that she linked to a strengthening of the AMOC. Willis (2010) also found a slight increase in the AMOC of 2.6 Sv since 1993 by combining satellite altimetry and data from Argo sub-surface floats at 41°N. Another proxy reconstruction of the recent AMOC was by Frajka-Williams (2015) who used satellite altimetry measured between 1993 and 2014 to estimate trans-basin baroclinic transport at 26°N. She used a simple linear regression between sea level anomaly and upper mid-ocean (UMO) transport. UMO transport is defined in Section 2.1.1, but is essentially the meridional transport in the top ~1100 m of the basin at 26°N that excludes the Florida Current and Ekman transport.

This method recovered over 80% of interannual UMO transport variability, however a decadal mean reduction in AMOC transport (comparing 1993–2003 and 2004–2014) was not reflected in the UMO decadal mean. It should be noted that most of the reconstruction methods described here, in particular those using sea-surface height anomalies and simple linear regressions, are implicitly single-layer, although they are attempting to reconstruct the AMOC, a multi-layer dynamic system.

The palaeoclimatic evidence shows that the AMOC has experienced large, abrupt changes, and more recent reconstructions indicate a general decline over the last 150 years (Caesar et al., 2018; Rahmstorf et al., 2015). Although the Intergovernmental Panel on Climate Change say that it is unlikely that the AMOC will stop this century, they state with medium confidence that a slowdown by 2050 due to anthropogenic climate change is very likely (Stocker et al., 2013). Concern over a possible AMOC collapse was one of the drivers for the creation of the RAPID trans-basin mooring array at 26°N in 2004 to provide a continuous observation record.

1.2 The RAPID mooring array

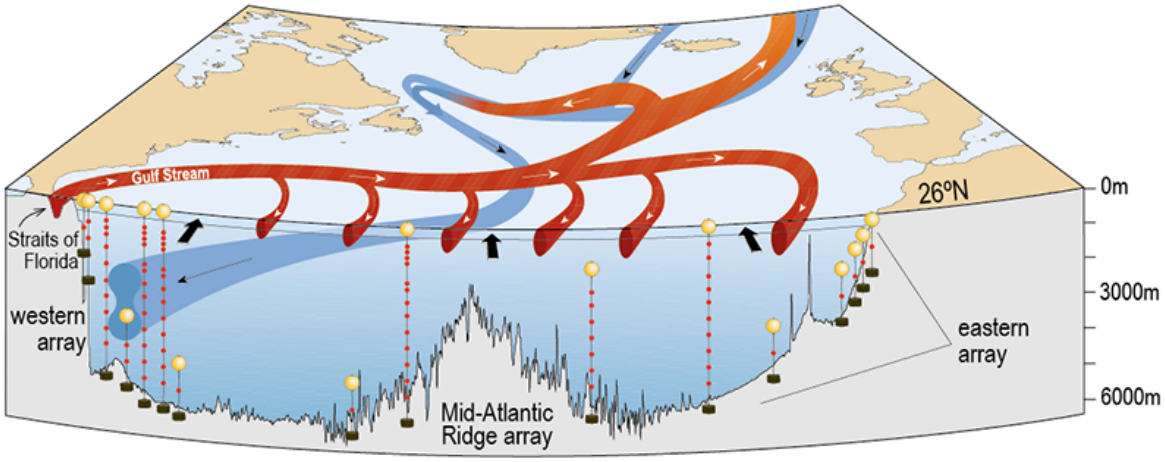

Figure 1.4: Schematic of the AMOC and the RAPID array at 26°N. The red flowlines show near-surface warm currents, including the Gulf Stream and how it divides between the upper limb of the overturning circulation and recirculation within the sub-tropical gyre. The blue flowlines show the deeper cold currents, with the Deep Western Boundary Current at 26°N divided into two lobes to indicate Upper and Lower North Atlantic Deep Water. The black arrows show the direction of wind-driven Ekman transport. The Florida current is measured by submarine cable within the Straits of Florida, and the array moorings measure temperature, salinity and in some cases current velocity and bottom pressure. Image from https://www.rapid.ac.uk.

The RAPID-MOCHA-WBTS (RAPID – Meridional Overturning Circulation and Heatflux Array – Western Boundary Time Series, hereafter RAPID) permanent mooring array at 26°N has observed the AMOC since April 2004 with subarrays of moorings deployed on the western and eastern boundaries and on either side of the Mid-Atlantic Ridge. These, along with the time-mean AMOC circulation are shown in the schematic in Fig. 1.4. Whilst the circulation as illustrated is simplified, it does show that at 26°N the Gulf Stream is mainly confined to the Florida Straits, with the exception of the Antilles Current to the east of the Bahamas. The Antilles Current has an estimated mean northward transport of 5-6 Sv (Johns et al., 2008; Lee et al., 1996), compared to the mean Florida Straits transport of 32.2 Sv between 1982 and 1998 (Baringer and Larsen, 2001). The partition of upper northwards flowing waters into those that reach high latitudes and those that recirculate within the sub-tropical gyre is also illustrated; as is the division of the deep return flow into Upper and Lower North Atlantic Deep Water, indicated by two distinct lobes in the schematic.

Details of the RAPID AMOC calculation are given in Section 2.1.1, but the resulting AMOC time series in Fig. 1.5 has given important insights into the variability that exists on a range of timescales as it has increased in length, from sub-annual and seasonal to interannual (e.g., Cunningham et al., 2007; Kanzow et al., 2010; McCarthy et al., 2012). Between 2004 and 2018, the annual mean AMOC has varied between 13.5 ± 4.4 Sv in 2009/10 to 20.9 ± 4.0 Sv in 2005/06, however all except three years have a mean between 17.2 and 18.9 Sv (Moat et al., 2020), compared to a time series mean of 17.0 ± 4.7 Sv from 2004 to 2018 (Table 1.1). Table 1.1 also shows the mean, standard deviation and integral time scale for the components of the RAPID AMOC, and the AMOC layer transports by depth. The integral timescale was calculated by integrating to the first zero crossing of the autocorrelation function (Thomson and Emery, 2014).

On interannual timescales, numerical models suggest that changes in boundary density dominate the AMOC (Hirschi and Lynch-Stieglitz, 2006), although McCarthy et al. (2012) found that both the geostrophic upper mid-ocean transport and the wind-driven Ekman transport contributed towards the exceptional downturn in the AMOC between December 2009 and March 2010, with these changes also reflected in the deep geostrophic transports. There is also a seasonal cycle in the AMOC with a peak-to-trough amplitude of 6.7 Sv,

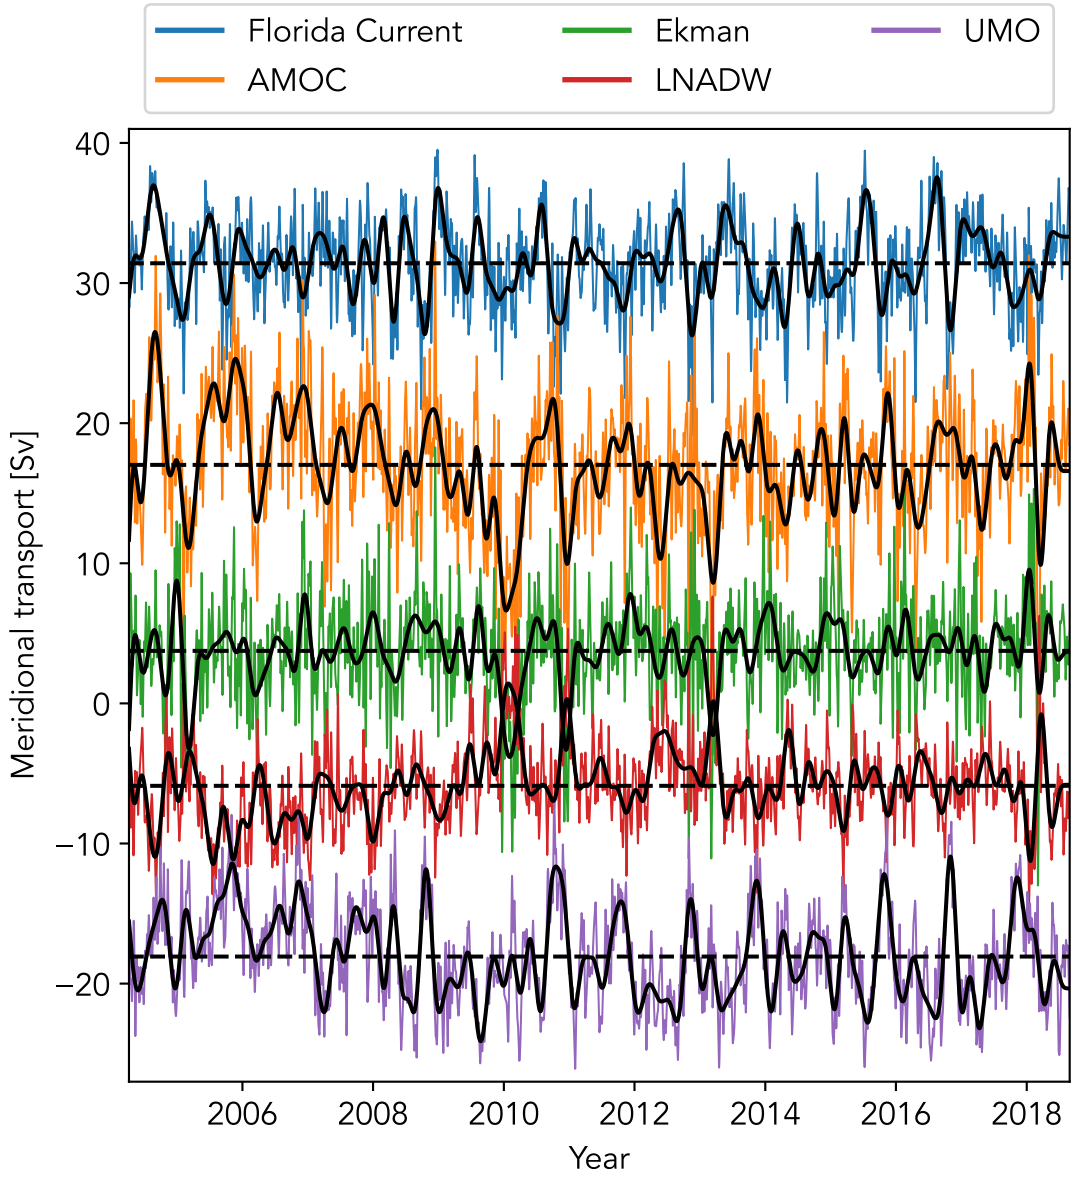

Figure 1.5: Ten-day (colours) and 90-day low-pass (black) time series of the Florida Current (blue), AMOC (orange), Ekman (green), LNADW (purple), and upper mid-ocean (red) transports for the period 7 April 2004 to 29 August 2018. Horizontal dashed black lines are the time series means. Positive transports correspond to northward flow. Time series data were downloaded from www.rapid.ac.uk and plotted as described in McCarthy et al. (2012).

Volume transport	**Mean [Sv]**	**STD [Sv]**	**Integral timescale [days]**
AMOC	17.0	4.7	24.9
UMO	-18.1	3.5	26.0
Florida Current	31.4	3.1	16.6
Ekman	3.7	3.5	10.9
0-800m	-18.8	3.4	26.3
800-1100m	0.4	0.7	12.8
1100-3000m	-11.9	2.6	14.7
3000-5000m	-5.9	3.0	36.8

Table 1.1: Mean, standard deviation and integral timescales of 10-day filtered RAPID transport time series between 7 April 2004 and 29 August 2018 shown in Fig. 1.5.

driven primarily by wind stress curl variability at the eastern boundary forcing changes in density (Chidichimo et al., 2010; Kanzow et al., 2010).

The observational error of annual mean AMOC transports is 0.9 Sv, and of ten day transports is 1.5 Sv (McCarthy et al., 2015), however a high resolution ocean model study suggests that due to the design of the RAPID array, it may underestimate the mean AMOC by ~1.5 Sv, although variability is captured very accurately (Sinha et al., 2018). Three main contributing factors to this bias are: the reference level assumed when calculating geostrophic transports; deep regions that are not sampled by the array; and ageostrophic transport.

The 2009/10 downturn highlights the impact of the AMOC on climate and sea level on shorter timescales. During this period, the AMOC strength declined by around 30%, to an annual mean of only 12.8 Sv (McCarthy et al., 2012). There was an associated reduction in northwards ocean heat transport of 0.4 PW, which caused anomalously cold sea-surface temperatures to the north of 25°N, and anomalously warm sea-surface temperatures to the south, driving a record-low negative North Atlantic Oscillation (Bryden et al., 2014). Subsurface cooling of the sub-tropical North Atlantic down to 2000 m was also observed between 2010 and 2012, primarily driven by AMOC weakening (Cunningham et al., 2013). The very cold winter of 2010-2011 over the UK and north-west Europe, and the intense tropical hurricane season of 2010 (Hallam et al., 2019) have both been linked to this AMOC downturn.

The RAPID time series shows a reduction in AMOC strength between 2004 and 2012 (Smeed et al., 2014), with that reduced state persisting until 2017 (Smeed et al., 2018). This decline is thought to be due to natural decadal variability rather than anthropogenic global warming (Roberts et al., 2014; Robson et al., 2014), and the most recent RAPID data suggest that the AMOC may be recovering (Moat et al., 2020), although the time series is still not sufficiently long to give a definitive answer. How long a time series needs to be to detect such a low frequency trend is addressed in the next section.

1.3 Detecting changes in the AMOC

Assessing the weakening trend in the AMOC between 2004 and 2012 (Smeed et al., 2014), Roberts and Palmer (2012) found that to detect whether a trend of the same rate of decline (-0.53 Sv/year) was significant, i.e., different from internal decadal variability, would require longer than a decade. To detect the climate change-related AMOC weakening more commonly predicted by CMIP5 models of around -0.1 Sv/year would require more than 40 years. Using a different numerical model (ECHAM5/MPI-OM), Baehr et al. (2008) showed that by using AMOC strength at 26°N alone, it would take approximately 60 years to distinguish a climate change-related weakening from the natural variability. Vellinga and Wood (2004) reduced detection time of climate change-related weakening of the AMOC within the HadCM3 general circulation model to around a decade by combining it with multiple other observations to improve the signal-to-noise ratio. This technique, called optimal fingerprinting, relies on observations such as the temperature at a region of deep water formation. This observation is effective as it is from deeper water, so is not directly influenced by the global warming signal, it is dynamically linked to the signal, and it has a good signal-to-noise ratio.

This fingerprinting technique was adapted by Baehr et al. (2007) by using the density difference between the western and eastern Atlantic boundaries at 26°N as both the climate signal and the fingerprint observation. Changes in density reflect changes in both temperature and salinity, and as AMOC strength is partly determined from the trans-basin density difference, they are strongly linked dynamically. Using a simulated RAPID array within a ECHAM5/MPI-OM general circulation model, they created a fixed spatial pattern of natural variability from a regression between the control run zonal

density differences and AMOC strength. The control run density field is then projected onto this fixed spatial pattern to give the upper and lower bounds of natural variability. The detection time is defined as when the climate vector (zonal density differences from a forced simulation, similarly projected onto the spatial pattern), moves outside these bounds. Fig. 1.6 shows the control (black) and forced (red) detection variables obtained using the trans-basin density difference over four different depth ranges. Only when the deep density difference (below 1700 m) was used was the detection time reduced by half, to around 30 years. Using actual RAPID observations rather than model data showed the detection variable had a similar magnitude of variability (Baehr, 2011). These results suggests that this adapted fingerprint method is feasible if the RAPID time series reaches a length of 3 decades. The detection methods described here are all based on numerical models, and McCarthy et al. (2017) found that some CMIP5 models, specifically HadGEM2-ES and HadGEM3-GC2, underestimated the cross-basin density difference, having averages of half and one-tenth of the RAPID observed values respectively. Both the detection techniques described and the limitations of the models emphasise the need for continuous long-term observations, such as those made by the RAPID and OSNAP arrays, to include deep density-related measurements.

1.4 The dynamics of the AMOC at 26°N

1.4.1 Water masses and transport at 26°N

Fig. 1.7 shows a meridional section of salinity in the Atlantic along the A16 transect, where the major water masses can be identified from their characteristic salinity (and temperature, not shown here). Arrows show the approximate pathways of the different water masses moving in opposite directions, with southwards-flowing North Atlantic Deep Water (NADW) between northwards-flowing Antarctic Intermediate Water (AAIW) and the deep Antarctic Bottom Water (AABW). In the upper layers, the water is also moving in wind-driven surface currents that form the subtropical and subpolar gyres.

Fig. 1.8 shows the zonal section of neutral density at 24°N and the changes in cross-basin density gradients with depth. Above around 1000 m, isopycnals deepen towards the west, indicative of southward flow in the thermocline; below 1000 m, isopycnals deepen

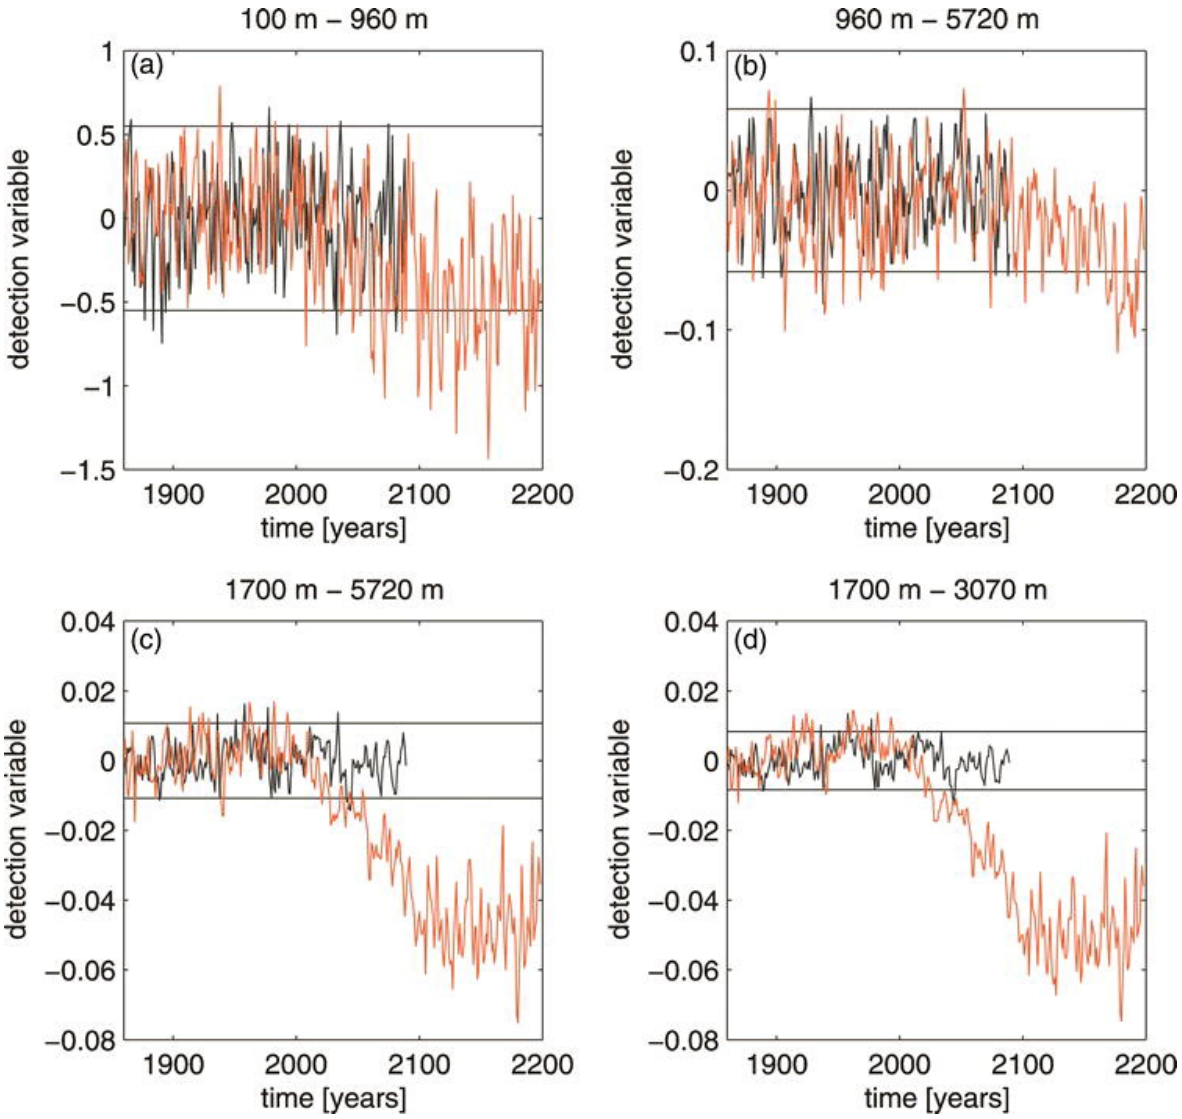

Figure 1.6: Detection variables for the control (black) and forced (red) simulation over depth ranges of (a) 100-960 m, (b) 6=960-5720 m (bottom), (c) 1700-5720 m, and (d) 1700-3700 m. The horizontal black lines are the bounds of the control detection variable (± 1.96 std). From Baehr et al. (2007).

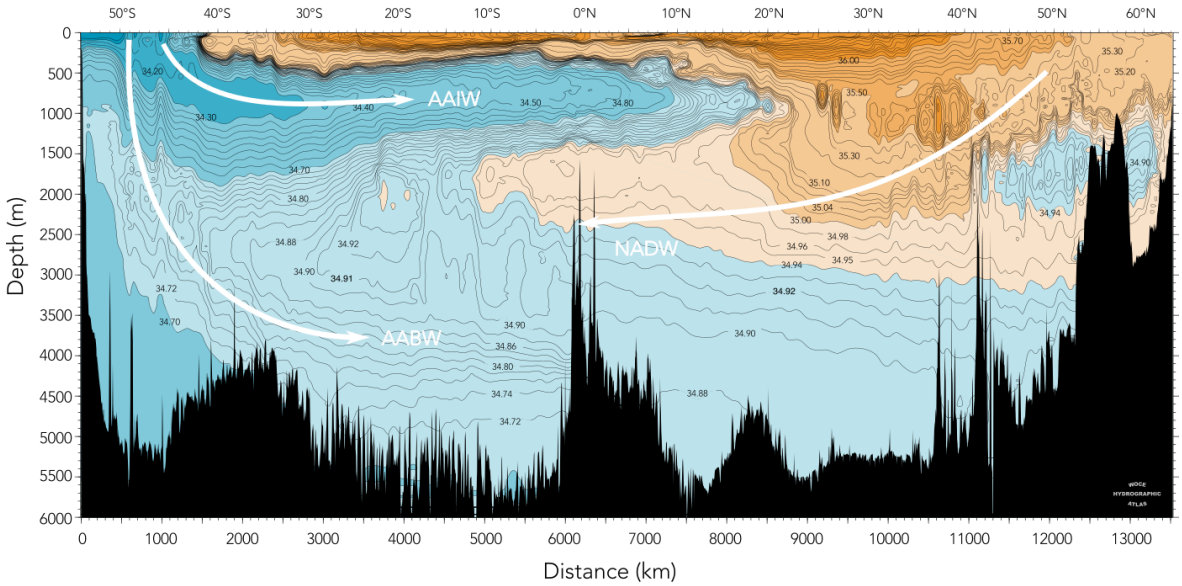

Figure 1.7: Distribution of salinity along the A16 meridional Atlantic Ocean section, from Iceland to the Southern Ocean, with the major water masses labelled and their approximate pathways shown by white arrows. The water masses shown are AAIW (Antarctic Intermediate Water), NADW (North Atlantic Deep Water), and AABW (Antarctic Bottom Water). Adapted from the WOCE Hydrographic Atlas Volume 3: Atlantic Ocean (Koltermann et al., 2011).

towards the east. The profile of the meridional mid-ocean transport driven by these zonal density gradients shows a small northwards AAIW transport around 1000 m, sandwiched between the overall southwards flows in the thermocline and NADW layers. The profile also shows that the NADW has two distinct layers, above and below approximately 3000 m. These are the upper (UNADW) and lower (LNADW) North Atlantic Deep Waters, and changes observed in one are not necessarily observed in the other: the reduction in AMOC strength between 2004 and 2012 (Smeed et al., 2014) was observed in LNADW but not UNADW, and Bryden et al. (2005) found southwards LNADW transport estimated from transatlantic hydrographic sections at 25°N decreased from -15 Sv in 1957 to less than -7 Sv in 1998 and 2004, while the UNADW transport remained between -9 and -12 Sv. Below the NADW layers, there is a small northwards transport below 5000 m, AABW, that flows along the western side of the Mid-Atlantic Ridge. The relative strength of the shallow and deep southwards water mass transports defines the strength of the overturning circulation: a weak AMOC is associated with greater recirculation within the upper layers of the thermocline and weaker deep return flow; whereas a stronger AMOC is associated with weaker recirculation within the thermocline and stronger deep NADW transport. For an empirical model to more fully represent AMOC dynamics, in particular lower frequency

changes, it must represent these deeper layers.

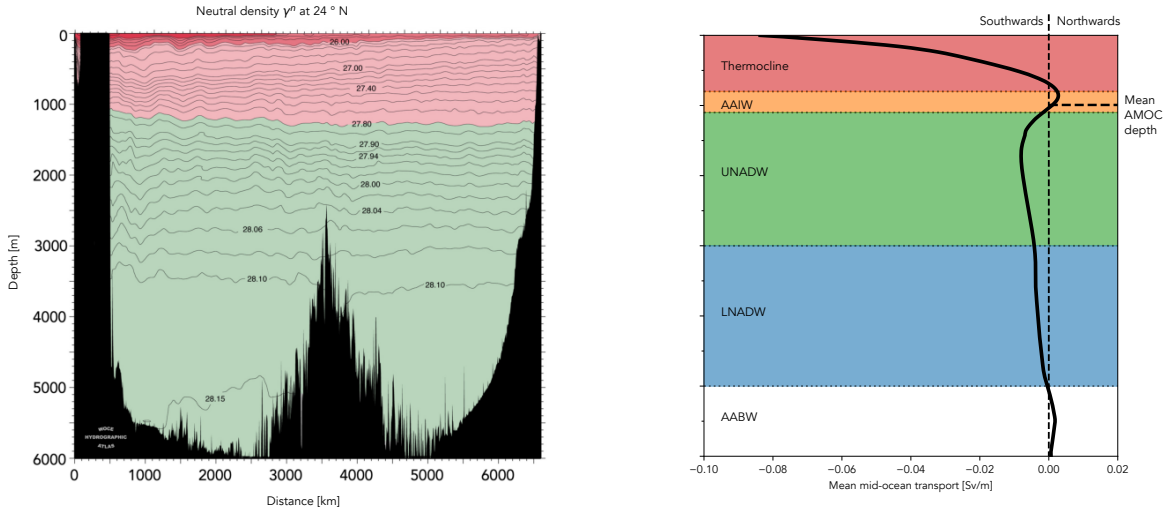

Figure 1.8: (Left) World Ocean Circulation Experiment (WOCE) North Atlantic A05 section of neutral density (γ^n) at 24°N. From the WOCE Atlantic Ocean Atlas Vol. 3. (Right) Profile of RAPID-estimated mean mid-ocean transport with depth and the RAPID-defined northwards and southwards layer transports. Mean AMOC depth is around 1100 m.

1.4.2 Meridional coherence of the AMOC

General circulation model studies have found meridional coherence of the AMOC on longer timescales where buoyancy forcing is the primary cause, for example, Polo et al. (2014) found that in experiments that isolated buoyancy forcing in a NEMO 1° ocean general circulation model, changes to deep water formation in the Labrador Sea could be related to AMOC anomalies at 26°N. In the same experiments that included wind-forcing, this relationship was obscured. In a related study, Polo et al. (2020) found that deep density anomalies propagated southwards from the sub-polar gyre along the western boundary at an average speed of 0.3 m/s, but that these buoyancy-forced density signals would only be distinguishable at 26°N for a time series approaching 30 years and filtered to remove wind-forced signals.

Bingham et al. (2007) investigated the meridional coherence of the AMOC within three numerical climate models, and found that higher frequency processes south of 40°N overpowered the meridionally-coherent mode. Their findings showed that interpreting intermittent observations with an assumption of meridional coherence is risky, which has

relevance to our investigations in Chapter 5.

Frajka-Williams et al. (2018) looked at coherence between observations made by the RAPID array at 26°N, and by the MOVE (Meridional Overturning Variability Experiment) array at 16°N. The MOVE array is not a basin-width array like RAPID, but covers the western side of the basin from the western boundary near Guadaloupe to the Mid-Atlantic Ridge. The array was deployed in early 2000 and measures temperature, salinity and pressure. Although the AMOC calculated by the two arrays show opposing trends, likely due to a different in methodology, the density anomalies below 1200 dbar show coherence over interannual timescales.

1.5 Motivation

The direct estimates and proxy reconstructions of the AMOC described previously inform this research. RAPID observations are accurate to within 1.5 Sv (0.9 Sv for annual means) (McCarthy et al., 2015) and encompass variability on timescales of weekly to decadal, but only exist from 2004. Direct estimates from transatlantic hydrographic sections have temporal resolutions of 5-10 years at best (Atkinson et al., 2012; Bryden et al., 2005). Proxy reconstructions extend the AMOC backward in time to 1993 using satellite altimetry (Frajka-Williams, 2015), to 1870 using sea-surface temperature fingerprinting (Caesar et al., 2018), and by up to a thousand years for a variety of palaeographic proxies. A thorough comparison is given in Caesar et al. (2021), but palaeoclimate proxies can be broadly categorised into those that use changes in sediment grain size as a proxy for changes in deep boundary currents, and those that use evidence for changes in nutrients and ocean productivity, such as isotope changes in corals and foraminifera, to reflect changes in water mass properties. Lastly, and closest to this research, Longworth et al. (2011) used a simple regression between temperature and volume transport anomalies, but only represented a single dynamic layer at 26°N, when the observed dynamics are more complex.

In summary, we have direct observations that are high resolution, low uncertainty, but short duration; direct observations that are longer duration but very low resolution; implicitly 'single-layer' reconstructions of 20-30 years that do not capture all the observed variability; and longer reconstructions with high uncertainty and for the most part, not based on hydrography. This research asks whether the AMOC can be more effectively

reconstructed using direct hydrographic observations, so avoiding the uncertainties inherent in proxy reconstructions. The hydrographic observations are also sub-surface, so minimise the effect of localised atmospheric changes. Finally, we ask whether a reconstruction that includes both shallow and deep observations, representing both the upper and lower limb of the overturning circulation, can better reconstruct the AMOC; and if, by including deep observations, we increase the likelihood of detecting low frequency changes in the AMOC (Baehr et al., 2007; McCarthy et al., 2017).

1.6 Thesis outline

The data and methods used to train and test empirical statistical models, and to reconstruct the AMOC using these models, are outlined in Chapter 2. Chapter 3 presents the development and comparison of the empirical regression models, then Chapter 4 describes how the best models are applied to hydrographic data to reconstruct UMO and LNADW transport anomalies, and hence the AMOC. Chapter 5 presents further investigation of the uncertainties, scope and limits of the regression models within a high resolution ocean simulation. The conclusions of this research are presented in Chapter 6.

Chapter 2

Data and Methods

This research uses observational data from the RAPID array and historical conductivity, temperature, and depth (CTD) profiles, and data from a high-resolution ocean simulation. The RAPID mooring array timeseries of in situ measurements and the resulting transport estimates are used to train and validate linear regression models. The most successful models are then applied to historical CTD profiles from the western boundary region around 26°N to reconstruct meridional volume transports, including the AMOC. Finally, the scope and limits of this reconstructive method are investigated within a hindcast run of the NEMO ORCA 1/12° ocean simulation.

2.1 Data

2.1.1 RAPID observations

The observational data from the RAPID array, outlined in Section 1.2, are a continuous time series of the AMOC and associated components and transports from April 2004 to the present, estimated from measurements collected from sensors deployed on the RAPID mooring array at 26°N in the North Atlantic.

The RAPID mooring array comprises four sub-arrays: the western boundary just east of The Bahamas; the eastern boundary just south of The Canaries; and either side of the Mid-Atlantic Ridge (Fig. 2.1). CTD sensors are sparsely deployed on these moorings, which can be as tall as 5000 m, reaching from the sea floor to within 100-200 m of the surface (Fig. 1.4). The vertical resolution of the sensors is depth-dependent, with sensors below 2000 dbar separated by approximately 500 dbar, those between 2000 and 500 dbar separated by around 200 dbar, and the shallowest instruments above 500 dbar are separated by approximately 100 dbar. The sensors measure in situ temperature, practical

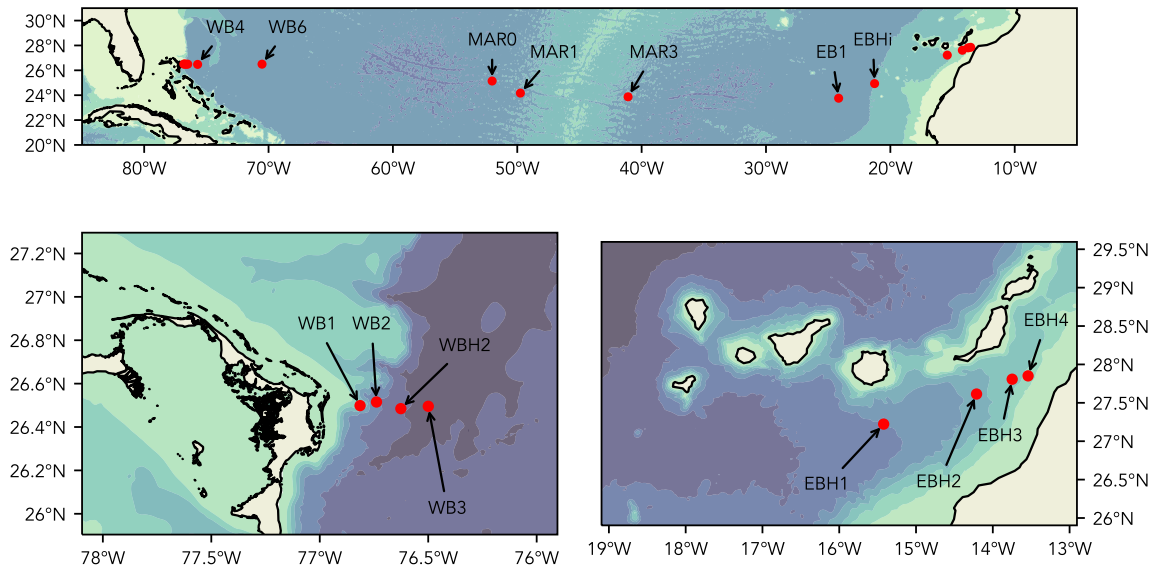

Figure 2.1: Locations of tall moorings in the RAPID array at 26°N (upper), with those close to the Bahamas (lower left) and Canaries (lower right) shown in more detail.

salinity from conductivity, and pressure every 15 minutes, which are binned to 12-hourly averages. CTD data is collected by research ship expeditions approximately every 18 months, processed by the RAPID project team (McCarthy et al., 2015; Rayner et al., 2011) and made available at www.rapid.ac.uk. McCarthy et al. (2015) describes the AMOC calculation in detail, but it is described here as we use interim data products as well as the resulting transport data.

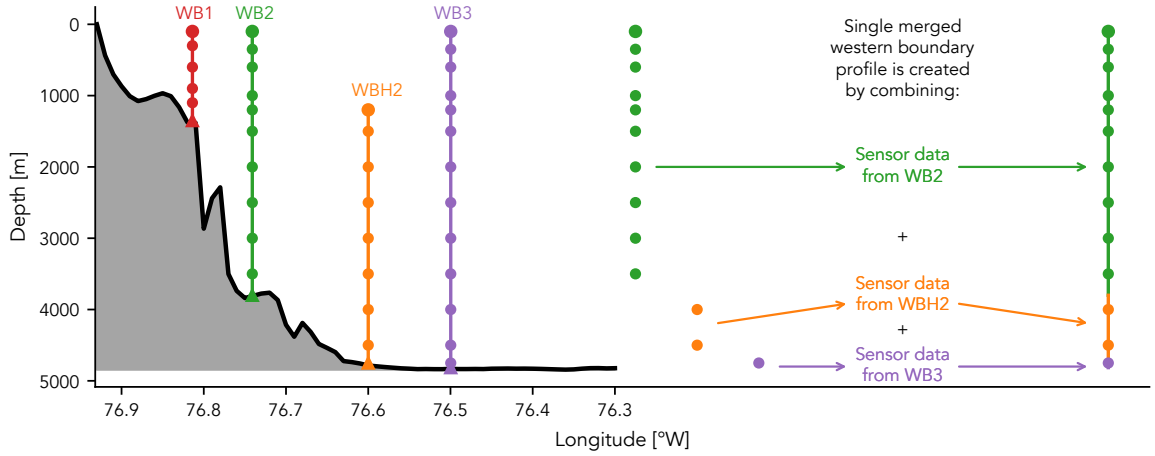

Figure 2.2: Schematic showing how sensor data from the primary RAPID array western boundary mooring, WB2 (green), is merged with data from deeper WBH2 (orange) and WB3 (purple) moorings to create a single full depth profile. The black line shows the approximate bathymetry, and coloured vertical lines represent the moorings with circles representing sensors and triangles as the mooring anchors.

During processing, a single full-height hydrographic profile is created for each sub-array by merging mooring data. For example, at the western boundary, most data comes from instruments deployed on the WB2 mooring, but to cover the full depth, data from the deeper, more eastwards WBH2 and WB3 moorings are added (Fig. 2.2). This merging creates four sparsely-resolved profiles, one for each sub-array. The profiles comprise 12-hourly in situ temperature and practical salinity data, that are then vertically interpolated over a pressure grid with 20 dbar resolution using a weighted monthly Hydrobase climatology for each location. In situ temperature is converted to conservative temperature, and practical salinity to absolute salinity, then density is calculated from conservative temperature and absolute salinity using the Gibbs SeaWater (GSW) implementation of the Thermodynamic Equation of Seawater-2010 (TEOS-10) equation of state. Dynamic height referenced to 4820 dbar (the deepest common measurement depth) is calculated from the merged and interpolated conservative temperature and absolute salinity profiles. Eq. 2.1 shows that dynamic height anomaly Φ is the integral of the specific volume anomaly δ with respect to pressure,

$$\Phi = \int \delta \, dp, \tag{2.1}$$

where the specific volume anomaly α is calculated from the in situ specific volume and the specific volume at the same pressure p with salinity $S = 35$ and temperature $T = 0 \circ CC$ (Eq. 2.2),

$$\delta = \alpha \, (S, T, p) - \alpha \, (35, 0, p), \tag{2.2}$$

and specific volume α is the inverse of density ρ (Eq. 2.3),

$$\alpha = \frac{1}{\rho}. \tag{2.3}$$

The dynamic height profiles are then extrapolated to the surface using a monthly Argo-based climatology, with the geostrophic shear captured by a model combining linear, quadratic and cubic terms. Eq. 2.4 shows how internal transport T_{int}, relative to a

reference level, is calculated between two dynamic height profiles at the eastern and western ends of a zonal section:

$$T_{int}(z) = \frac{1}{f}[\Phi_e(z) - \Phi_w(z)]. \tag{2.4}$$

The internal geostrophic transport relative to the reference depth of no motion at 4820 dbar is calculated between the four dynamic height profiles, one for each sub-array at the western boundary (WB), eastern boundary (EB), and on the western and eastern flanks of the Mid-Atlantic Ridge (MARWEST and MAREAST). Including the Mid-Atlantic Ridge profiles allow the contribution from each of the western and eastern basins and across the Mid-Atlantic Ridge at depths shallower than 3700 m to be calculated separately and then summed to give the basin-wide internal transport, as the schematic in Fig. 2.3 shows (Rayner et al., 2011).

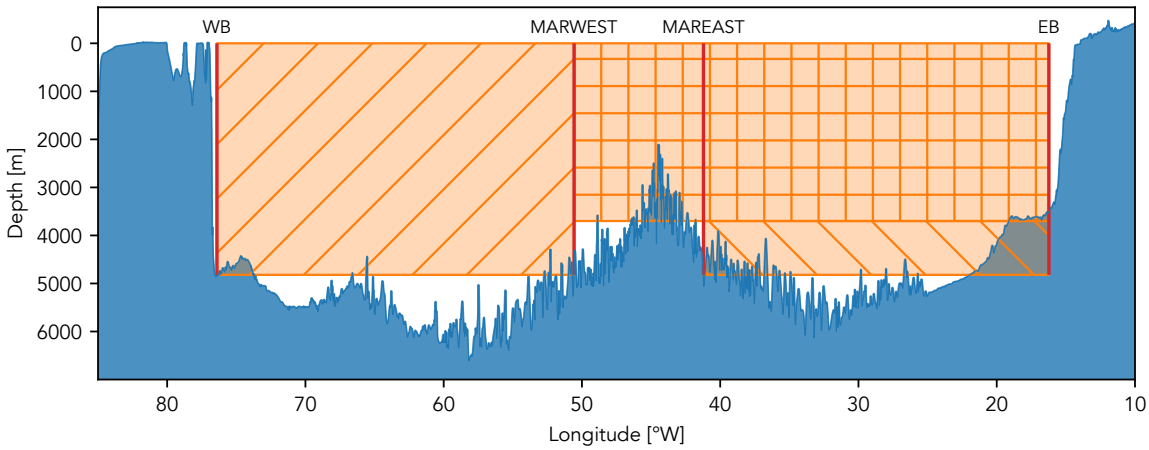

Figure 2.3: Schematic with red vertical lines showing the dynamic height profiles at the western (WB) and eastern (EB) boundaries, and either side of the Mid-Atlantic Ridge (MARWEST and MAREAST). The orange boxes represent the sections where internal transport is calculated between the dynamic height profiles for the western basin (diagonal hatching), for the deep eastern basin (back hatching) and for the shallow eastern basin (crossed hatching).

Eq. 2.5 shows how this internal transport (T_{int}) is used to estimate AMOC transport at 26°N (T_{amoc}) by summing it with three other directly observed components: Gulf Stream transport within the Florida Straits (T_{flo}), which is measured by submarine cables and calibrated by regular hydrographic sections (Baringer and Larsen, 2001; Meinen et al., 2010); Ekman transport (T_{ek}), which here is calculated from ERA-Interim reanalysis wind

fields; and Western Boundary Wedge transport (T_{wbw}), which is obtained from direct current measurements over the continental slope between the Bahamas and the WB2 mooring at 76.75°W. To the sum of these four components, an external transport (T_{ext}) is added to ensure that mass is conserved and that there is zero net flow across the section:

$$T_{amoc}(z) = T_{int}(z) + T_{flo}(z) + T_{ek}(z) + T_{wbw}(z) + T_{ext}(z). \tag{2.5}$$

Although the external transport has a uniform velocity, it is weighted by the width of the basin, so decreases with increasing depth as the basin narrows. The assumption of zero net flow holds on timescales longer than 10 days (Bryden et al., 2009).

The strength of the AMOC is then the maximum of T_{amoc} integrated over depth from the surface, i.e., the maximum of the transport streamfunction Ψ (Eq. 2.6);

$$\Psi(t, z) = \int T_{amoc}(t, z)dp, \tag{2.6}$$

and the AMOC depth (z_{amoc}) is the depth of that maximum at each time step, usually around 1100 m.

The upper mid-ocean (UMO) transport is defined by RAPID as the mid-ocean transport (T_{mo}) integrated between the surface and the AMOC depth, where the mid-ocean transport is the sum of the internal, external and Western Boundary Wedge (WBW) transports (Eq. 2.7):

$$T_{umo}(t) = \int^{z_{amoc}} T_{mo}(t, z)\, dz = \int^{z_{amoc}} \left[T_{int}(t, z) + T_{wbw}(t, z) + T_{ext}(t, z)\right]\, dz. \tag{2.7}$$

This net southwards UMO transport, which includes both the southwards gyre recirculation and the northwards Antilles Current, is the main contributor to changes in the AMOC (Smeed et al., 2014). Although the UMO transport is the sum of internal, external and WBW transports, internal geostrophic transports contribute most to its variability on interannual time scales (Frajka-Williams, 2015). The estimated uncertainties of the RAPID

array CTD measurements are 1 dbar for pressure, 0.002°C for temperature and 0.003 for salinity, leading to an estimated error of 0.8 Sv in the AMOC calculation (McCarthy et al., 2015).

During this study the RAPID project code that was current in October 2017 was converted from Matlab to Python. This was done to understand how the AMOC is calculated, and to allow the code to be edited more easily. Some differences exist between the now current RAPID methodology and ours, for example, as of June 2019 RAPID uses ERA5 wind stress at 25.5°N rather than ERA-Interim wind stress to calculate the Ekman transport. The primary change to the Python RAPID code was to calculate in situ density from the merged gridded temperature and salinity profiles using a Python version of the GSW toolbox for TEOS-10.

2.1.2 Hydrographic data

Several sources of ship-based hydrographic data were used within this study. Repeat occupations of the World Ocean Circulation Experiment (WOCE) A05 transatlantic line at 24.5°N were made during October 1957 (Fuglister, 1960), August–September 1981 (Roemmich and Wunsch, 1985), July–August 1992 (Parrilla et al., 1994), January–February 1998 (Baringer and Molinari, 1999), April–May 2004 (Cunningham, 2005), January–February 2010 (Atkinson et al., 2012), and December 2015–January 2016 (McDonagh et al., 2015). Apart from 1957 and 1981, all occupation datasets are available from CCHDO (CLIVAR and Carbon Hydrographic Data Office). The westernmost CTD profiles taken during the 1957, 1981 and 1992 cruises were taken at 24.5°N, but from 1998 onwards were taken at 26.5°N, to allow the use of Florida Current cable measurements when calculating section transports.

Historical CTD profiles were extracted from the World Ocean Database (WOD) 2018 (Boyer et al., 2018), from a region defined between 24 and 27°N, and 75 and 77.5°W. These were complemented by a dataset created by Longworth et al. (2011) from a variety of sources: the WOD 2001, Hydrobase 2003, and NOAA/AOML, some of which are no longer available. Any duplicate profiles were removed with the WOD 2018 profiles being preferred.

Each continuous CTD profile has practical salinity (PSS-78), in situ temperature and oxygen measured every 2 dbar. Using the Python-GSW package, practical salinity was converted to absolute salinity, and in situ temperature was converted to conservative temperature. Density was then calculated from absolute salinity and conservative temperature, and interpolated over a coarser 20 dbar grid to match the vertical resolution of the processed RAPID gridded profiles.

2.1.3 NEMO-ORCA high resolution ocean model

We used a high resolution (1/12°) version of the eddy-resolving NEMO ORCA012 (Nucleus for European Modelling of the Ocean, Oceanic Remote Chemical/Optical Analyzer) global ocean circulation model, hereafter NEMO. The NEMO family of circulation models are based on the Océan Parallélisé (OPA) primitive equation, free surface ocean model, with sea ice represented by version 2 of the Louvain-la-Neuve ice model (LIM2) (Madec and and the NEMO team, 2008). The global ORCA12 grid is tripolar, having two poles over land in Canada and Siberia in the northern hemisphere and one in the southern hemisphere. This is to avoid a pole over the Arctic Ocean, which causes numeric instability. Although the resolution is nominally 1/12°, the linear resolution varies from 9.27 km at the equator to 1.75 km in the Canadian Arctic (Hughes et al., 2018). At 26°N, the horizontal resolution is approximately 8.3 km. There are 75 depth levels, with vertical resolution decreasing from 1 m at the surface to 250 m at 5500 m. Bottom topography is derived from ETOPO2 (Center, 2006) and represented as partial steps. The Drakkar Surface Forcing dataset version 5.2 (DFS5.2) forces the surface air temperature and heat fluxes, precipitation, humidity and winds. Sea-surface salinity is relaxed towards climatology with a piston velocity of 33.33 mm/day/psu to prevent drifts in global salinity caused by issues with fresh water forcing. The output data are given as 5-day means, starting at the beginning of each year and ignoring the last day of leap years. Further details of this model are given in Hughes et al. (2018) and Moat et al. (2016). The NEMO data used here are from the historical N006 run between 1958 and 2015 from the Marine Systems Modelling group at the National Oceanography Centre Southampton, which is a single integration of the NEMO v3.6 circulation model (Hughes et al., 2018).

2.2 Methods

This research uses statistical linear regression models to relate the changes in meridional volume transport to changes in density within the water masses found at different depths at 26°N. The concept of parsimony within regression models means that the simplest model with the fewest variables that achieve a desired level of prediction or explanation should be selected. Following this concept, we start with the simplest linear regression model, an ordinary least squares regression with a single explanatory variable representing the density change within a single water mass. Additional variables representing deeper water masses, or layers, are added to try to improve the efficacy of the regression. In addition, the model must satisfy the assumptions of ordinary least squares regressions. Since linear regressions trained on time series data tend to violate at least one of these assumptions (Chatterjee and Simonoff, 2013), a different regression method, typically generalised least squares, may be required.

In addition to regression models, we use empirical orthogonal functions (EOFs) to analyse how the depths of the density anomalies chosen by algorithm for use in the regression models compare to the modes of variability of those density anomalies.

2.2.1 Linear regression models

This research uses both ordinary least squares and generalised least squares linear regression models. A description of each, the assumptions they rely on, the difference between them, and the justification for choosing them are given below. All linear regression models run during this research were part of the Python `statsmodels` package.

2.2.1.1 Ordinary Least Squares linear regression

A simple ordinary least squares regression model has the form:

$$y_i = \alpha + \beta \cdot x_i + e_i, \tag{2.8}$$

where y is the dependent (or response) variable, x is the independent (or explanatory) variable, α is the intercept of the regression line, β is the magnitude of the slope, or the

trend, and e is the residual or error, the difference between the regression and the observed values. The least squares method selects the 'line of best fit' that gives the minimum sum of residuals squared, minimising the distance between the regression line and the observed values. Using the square of the residuals makes negative and positive residuals the same and is easy to calculate, although using the least squares method means that outliers have an exaggerated effect on the regression line.

2.2.1.2 Assessing the fit of linear regression models

Several statistics are used to assess how well a linear regression model fits the data. Two of the most common are the R^2 value, sometimes called the coefficient of determination; and the standard error of the regression. For comparing models with different numbers of independent variables, statistics such as the adjusted R^2 and Akaike's Information Criterion and Bayes Information Criterion can be useful, so are also described.

The R^2 value is a number between zero and one (or 0% and 100%), and shows how much of the total variance in the response variable is described by the linear regression. An R^2 value of zero indicates that the model explains none of the variability of the observed data around the mean; and a value of one indicates that it explains all of the variability of the data around its mean. More formally, R^2 is defined by Eq. 2.9:

$$\mathrm{R}^2 = 1 - \frac{\mathrm{SSR}}{\mathrm{TSS}}, \tag{2.9}$$

where SSR is the sum of squared residuals, and is a measure of the difference between the observed values and those predicted by the regression. Eq. 2.10 shows the SSR for n observations y_i and predictions $\hat{y}_i$ is

$$\mathrm{SSR} = \sum_{i=1}^{n} (y_i - \hat{y}_i)^2, \tag{2.10}$$

and Eq. 2.11 shows the total sum of squares (TSS) is the total variance of the observations, a measure of the difference between the observed values and their arithmetic mean $\bar{y}$,

$$\text{TSS} = \sum_{i=1}^{n} (y_i - \bar{y}_i)^2. \tag{2.11}$$

Since the R^2 value can be increased just by adding explanatory variables, the adjusted R^2 value, which accounts for the model's degrees of freedom, is preferred for multiple linear regression models, as it allows direct comparison of models with differing numbers of explanatory variables. Eq. 2.12 shows how adjusted R^2 is calculated from the number of explanatory variables k, R^2 and the number of observations n:

$$\text{Adjusted } R^2 = 1 - \left[\frac{(1 - R^2)(n - 1)}{n - k - 1} \right]. \tag{2.12}$$

If the added explanatory variables are spurious, then it is likely that the adjusted R^2 value will decrease, whereas the R^2 value would increase. The R^2 and adjusted R^2 values are relative measures of fit and so are unitless, whereas the standard error (SE) of the regression is an absolute measure and has the same units as the model predictions. It is a measure of the average deviation of the model predictions from the observations, and for n observations for a model with p parameters, Eq. 2.13 shows the SE is given by:

$$\text{SE} = \sqrt{\frac{\text{SSR}}{n - p}}. \tag{2.13}$$

The standard error is an approximation of the 95% confidence interval, in that approximately 95% of observations should fall within ±2 SE. In addition to these statistics, two more that are reported by the `statsmodels` output are Akaike's Information Criterion and Bayes Information Criterion, and both are useful in directly comparing regression models. Akaike's Information Criterion (AIC) shows which model explains more of the variation using the fewest independent variables, and it penalises using more variables spuriously. A lower AIC is better, with AIC calculated from the number of independent variables K and log-likelihood estimate L, as shown in Eq. 2.14:

$$\text{AIC} = 2\ K - ln(L). \tag{2.14}$$

The log-likelihood estimate L is the likelihood that the model could have produced the dependent variable, and is a standard output statistic of the regression model. A model with an AIC that is lower by 2 than another is significantly better. The Bayes Information Criterion (BIC) is closely related to AIC, but also incorporates the number of data points n, as shown in Eq. 2.15:

$$\text{BIC} = ln(n)\ K - 2\ ln(L), \tag{2.15}$$

so more strongly penalises the spurious addition of explanatory variables. A lower BIC is also indicative of a better model.

2.2.1.3 Assumptions of linear regression models

The proof that an ordinary least squares model produces the best results is the Gauss-Markov theorem (Graybill, 1976), however there are several assumptions required for this proof to hold. These classical regression model assumptions are outlined below, together with the result of being violated, and common methods of diagnosis.

2.2.1.3.1 Linearity

The regression parameters, rather than the explanatory variables, e.g., α and β in Eq. 2.8, must be linear and the regression equation must be additive. If they are not, both the coefficients and standard errors are unreliable. Linearity can be detected by examining plots of explanatory against response variables.

2.2.1.3.2 Homoscedasticity

Homoscedasticity means that the variance in errors is independent of the explanatory variables. If there is heteroscedasticity, or non-constant error variance, then the standard error of the regression will be unreliable. Homoscedasticity is typically detected using a plot of the fitted values of the model against the residuals of those values. A diagnostic method such as the Breusch-Pagan test is a more formal option, and this is described in Section 2.2.1.4.

2.2.1.3.3 No perfect multicollinearity

In a multiple linear regression model, the explanatory variables must be truly independent, i.e., one must not depend upon another. If this assumption is violated, then both the coefficients and standard error of the regression are unreliable. Multicollinearity is detected by examining the correlation between explanatory variables, and by looking at variance inflation factors (VIF). VIFs use multiple linear regression to regress all but one of the explanatory variables onto the remaining explanatory variable, which becomes a response variable for the purposes of the test. This is done for each explanatory variable in the original multiple linear regression, for example, for a regression model with four explanatory variables (x_1, x_2, x_3, x_4), four multiple linear regressions are made. The first regresses x_2, x_3, x_4 on x_1; the second regresses x_1, x_3, x_4 on x_2; and so on. Each of these regressions produces an R^2 value, which is used to calculate the VIF for i = 1, 2, 3, 4 (Eq. 2.16):

$$\text{VIF}_i = \frac{1}{1 - R_i^2}. \tag{2.16}$$

When $R^2 = 0$, the three explanatory variables explain no variance in the fourth, so there is no multicollinearity. As R^2 increases, the denominator increases and so does the VIF. If the VIF for an explanatory variable is high, then the information it contains is already in the model in another explanatory variable. There is no defined threshold for VIF, but values of 4 or 5 are commonly used to indicate that multicollinearity is present.

2.2.1.3.4 Normality of error terms

The residuals of a linear regression model should be normally distributed and have a mean of zero. This can be assessed from a histogram plot of the residuals, or a quantile-quantile (Q-Q) plot, or by using several different diagnostics such as the Shapiro-Wilk, Anderson-Darling or Komolgorov-Smirnov tests. We use the Shapiro-Wilk test, which is described in Section 2.2.1.6. Violating this assumption can affect the standard error if the number of samples is small.

2.2.1.3.5 No autocorrelation

No autocorrelation means that each residual must be independent of the previous one. It is a common issue with regressions based on time series data, where a positive residual is more likely than not to be followed by another positive residual, and a negative residual is more likely than not to be followed by another negative residual. Methods of diagnosis are the Durbin-Watson test, to find first-order autocorrelation; and the Breusch-Godfrey test, which tests for autocorrelation beyond first order. Partial autocorrelation coefficients can indicate the order of autocorrelation. If autocorrelation is present, then the coefficients of the model are unaffected, so the model is still 'correct'. However both the standard error and the R^2 value of the model are affected and cannot be relied upon.

An ordinary least squares regression model that violates one or more of these assumptions is no longer BLUE (best linear unbiased estimator), and an alternative approach or regression model may be needed. Autocorrelation of residuals is a known problem with linear regression models trained on time series data, and here we use the Durbin-Watson test, described in Section 2.2.1.5, to detect first-order autocorrelation. Ordinary least squares regression models that show autocorrelation of residuals should be replaced with an alternate model, for example generalised least-squares with auto-regression (GLSAR), which models the residuals as an auto-correlative process with a specified lag (McKinney et al., 2019).

2.2.1.4 The Breusch-Pagan test for heteroscedasticity

The Breusch-Pagan test was developed by Breusch and Pagan (1979) to test for heteroscedasticity in linear regression models. The implementation of the Breusch-Pagan test in Python, `statsmodels.stats.diagnostic.het_breuschpagan` uses an adaption by Koenker (1981) that assumes equal and independent, rather than normal, distribution of residuals. The hypotheses are then:

H_0 (null hypothesis): The residuals are distributed with equal variance, so homoscedasticity is present in the model;

H_A (alternative hypothesis): The residuals are not distributed with equal variance but are dependent on the explanatory variable values. Heteroscedasticity is present in the

model.

The test first obtains the residuals from the regression, so if the following regression is fitted to a set of y values,

$$y_i = \alpha + \beta \cdot x_i + e_i,$$

a set of values for the residuals, $\hat{e}$ is obtained. These residuals are squared, and then used as the response variables in a new regression model:

$$\hat{e}^2 = \gamma + \zeta \cdot x_i + u_i. \tag{2.17}$$

Eq. 2.18 shows that how the Chi-Square test statistic χ^2 is calculated from the number of observations (n) and R^2_{new}, the R^2 value of the new regression (Eq. 2.17):

$$\chi^2 = n \cdot R^2_{new}. \tag{2.18}$$

If the p-value corresponding to χ^2 with degrees of freedom equal to the number of predictors in the regression is less than a significance level of 0.05, then we reject the null hypothesis and conclude that heteroscedasticity is present in the model.

2.2.1.5 The Durbin-Watson statistic

The Durbin-Watson statistic (Durbin and Watson, 1950) is used to test whether the residuals of a time series show autocorrelation of lag 1, i.e., whether there is correlation between consecutive residuals. It uses the following hypotheses:

H_0 (null hypothesis): The residuals are independent and show no correlation;

H_A (alternative hypothesis): The residuals are not independent and show first-order correlation.

The Durbin-Watson statistic d is calculated using Eq. 2.19, where T is the total number of observations, and e_t is the residual at timestep t:

$$d = \sum_{t=2}^{T}(e_t - e_{t-1})^2 \Big/ \sum_{t=1}^{T} e_t^2. \tag{2.19}$$

The test statistic d is always between 0 and 4, where d=2 indicates no autocorrelation of residuals; d<2 indicates positive serial correlation; and d>2 indicates negative serial correlation. The Durbin-Watson statistic can be compared against a table of critical values for a given sample size, number of independent observations and alpha level (Savin and White, 1977). If d is outside the critical values (upper or lower limits) given in the table, the null hypothesis can be rejected and either positive or negative autocorrelation is present. For positive autocorrelation, a generalised least squares regression model with autoregression of errors (GLSAR) is one recommended alternative, and these models are discussed next.

2.2.1.6 The Shapiro-Wilk test for normality

The Shapiro-Wilk test was devised by Shapiro and Wilk (1965) as a test for normality of distribution, and is widely used in statistics. It uses the following hypotheses:

H_0 (null hypothesis): The distribution of the sample is not significantly different from a normal distribution;

H_A (alternative hypothesis): The distribution of the sample is significantly different from a normal distribution.

Eq. 2.20 shows how the Shapiro-Wilk test statistic, W, is defined as:

$$\mathrm{W} = \frac{(\sum_{i=1}^{n} a_i \, x_{(i)})^2}{\sum_{i=1}^{n}(x_i - \bar{x})^2}, \tag{2.20}$$

where $x_{(i)}$ are the ordered sample data, i.e., $x_{(1)}$ is the first order statistic, or smallest; and a_i are constants generated from the means, variances and covariances of a normal distribution with n samples.

The `scipy.stats.shapiro` implementation of the test returns the W statistic and a p-value. If the p-value is not less than 0.05, the null hypothesis cannot rejected, and there is insufficient evidence to say that the sample is not from a normal distribution.

2.2.1.7 Generalised Least Squares linear regression

When an ordinary least squares regression such as Eq. 2.8 is based on a time series (Eq. 2.21):

$$y_t = \alpha + \beta \cdot x_t + e_t, \tag{2.21}$$

it commonly has serially-correlated errors. This violates the Gauss-Markov assumptions, and the standard error of the regression are unreliable.

The generalised least squares (GLS) method of linear regression is adapted from the ordinary least squares model (Eq. 2.8), and is used to model data with non-spherical errors, i.e., errors with homoscedasticity and/or autocorrelation (Aitken, 1936). The autocorrelated errors can be modelled by an auto-regressive (AR) process. Eq. 2.22 shows an AR(1) process, where the correlation is significant for lag 1

$$e_t = \rho \cdot e_{t-1} + \epsilon_t, \tag{2.22}$$

where ϵ_t is a normally-distributed random variable with a mean of zero and a variance of σ^2.

Then for time $t-1$, Eq. 2.21 becomes

$$y_{t-1} = \alpha + \beta \cdot x_{t-1} + e_{t-1}. \tag{2.23}$$

Multiplying Eq. 2.23 by ρ and subtracting it from Eq. 2.21 gives

$$y_t - \rho \cdot y_{t-1} = \alpha(1-\rho) + \beta(x_t - \rho \cdot x_{t-1}) + e_t - \rho \cdot e_{t-1}. \tag{2.24}$$

The final two terms of Eq. 2.24 are, from Eq. 2.22, equivalent to ϵ_t, so Eq. 2.24 becomes

$$y_t - \rho \cdot y_{t-1} = \alpha(1-\rho) + \beta(x_t - \rho \cdot x_{t-1}) + \epsilon_t. \tag{2.25}$$

Eq. 2.25 is now a generalised regression equation with normally-distributed random errors. One observation must be discarded, as when time $t = 0$, there is no y_{t-1}.

The parameter ρ can be determined by several methods, such as Cochrane-Orcutt (Cochrane and Orcutt, 1949) or Yule-Walker (Walker, 1931; Yule, 1927). The `statsmodels` package GLS model uses a two-stage iterative fit, alternating between a generalised least squares fit and the Yule-Walker method to estimate ρ to within a relative tolerance (McKinney et al., 2019).

2.2.2 Empirical orthogonal functions

Empirical orthogonal functions, or EOFs, were first used in meteorology by Lorenz (1956). Empirical orthogonal function analysis is similar to principal component analysis, and is used to identify the 'modes', or dominant patterns of variability in spatiotemporal data. EOFs allow observations which vary in both time and space to be expressed as the product of a time-independent basis function or pattern, and a time-dependent coefficient. The mode describes the spatial pattern of the data, and the principal component (PC) associated with it is a time series that describes how that pattern varies in time. The leading modes will contain a large percentage of the observed variance, and allow a data set containing many different variables to be reduced to one with many fewer new variables. It also removes noise from data, as random noise will have no pattern. The new variables are linear combinations of the originals, with the linear combinations selected to minimise the difference between the empirical orthogonal functions and the observations using a least-squares technique. The EOFs and principal components are chosen to minimise the sum of the errors, and this happens when the EOFs are orthogonal in space, and the principal components are orthogonal in time. Eq. 2.26 shows that if we have pressure p measured at m different depths at times $t_1, t_2, t_3, ...t_N$, then $p_m(t_i)$ can be expressed as

$$p_m(t) = \sum_{k=1}^{M} \alpha_k(t)\ \phi_{km}, \tag{2.26}$$

where ϕ_{km} are the time-independent basis functions, or EOFs, and $\alpha_k(t)$ are the time-dependent coefficients, or principal components. The two sides of the equation are only truly equal when $k{=}m$, however, as k increases, the variance of the coefficients

decreases quickly, so p_m can be approximated reasonably well when k is much less than m, with Eq. 2.27 showing the approximation for $k{=}K$, and the associated error:

$$p_m^K(t) = \sum_{k=1}^{K} \alpha_k(t)\ \phi_{km} + r_m^K(t). \tag{2.27}$$

In this study, which utilises the Python `eofs` package (Dawson, 2016), we also use the EOFs expressed as the correlation between the associated principal component and the input data time series.

2.2.3 Coherence analysis

Coherence is a measure of how similarly two time series vary at each frequency over a range of frequencies, and whether they have any statistically-significant shared signals. Eq. 2.28 shows how coherence $\hat{\gamma}^2$ between two stationary time series $x(t)$ and $y(t)$ is estimated from the auto-spectral density functions of each of these time series, $\hat{P}_{xx}$ and $\hat{P}_{yy}$, and the cross-spectral density function between them, $\hat{P}_{xy}$ (Bendat and Piersol, 2010):

$$\hat{\gamma}^2(f) = \frac{|\ \hat{P}_{xy}(f)\ |^2}{\hat{P}_{xx}(f)\ \hat{P}_{yy}(f)}. \tag{2.28}$$

Power spectra show the variance in the data projected onto frequency space. To estimate the power and cross spectra we use the multitaper method that was developed in the early 1980s by Thomson (1982). This method applies different tapers to the same data, allowing different sections of data to be preserved or attenuated. Slepian tapers, also known as discrete prolate spheroidal sequences (DPSS), provide good protection against spectral leakage, and, as each taper is orthogonal to all other tapers, the filtered signals are independent. The schematic in Fig. 2.4 shows how the tapers are applied to the data, and then a fast Fourier transform (FFT) is applied to each filtered signal. The average over the individual spectra that relate to each taper is taken to give the multitaper spectrum. The advantages of a multitaper spectrum are that it reduces both bias (the spectral estimate is not representative of the true spectra due to leakage) and variance (noise in the spectral estimate), compared to a single taper spectrum that reduces bias but not variance. A thorough explanation of the multitaper method, its advantages and method of use is given

in Prerau et al. (2017), and although they analysed electro-encephalogram output, the method has been widely applied in the analysis of geophysical signals, including ocean, atmospheric and seismological data.

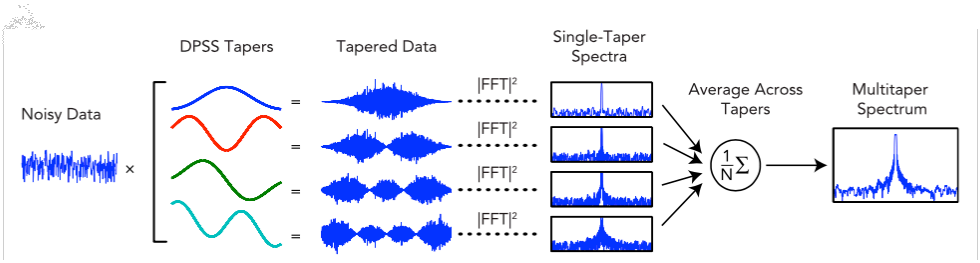

Figure 2.4: A schematic of the multi-taper method of spectral estimation, showing four orthogonal DPSS or Slepian tapers applied to noisy data, a FFT applied to each tapered data to obtain power spectra, and the averaging of these power spectra into a single multitaper spectrum. Taken from Fig. B of Prerau et al. (2017).

Multitaper spectral estimates require several parameters to be defined, which set both the number of Slepian tapers used and their properties. These parameters are M, the data segment length; NW, the time-half-bandwidth product; and K, the number of tapers. The data segment length should be the maximum length of time for which the data is stationary. The time-half-bandwidth product NW depends on both M and the spectral resolution (Δf), which is the width of the main lobe in the spectral estimate, and sets the minimum resolution distance (Eq. 2.29):

$$NW = \frac{M\Delta f}{2}. \tag{2.29}$$

Smaller spectral resolution allows higher resolution and variance; larger values will produce smoother spectra with lower resolution and variance, so the value of Δf chosen will depend on what spectral features we wish to resolve. Once a value for NW is chosen, the ideal number of tapers is then $K = 2NW - 1$.

Fig. 2.5 shows the first four Slepian tapers produced by the `dpss` function of the `scipy.signal.windows` Python package for M = 512 and NW values of 2.5 (Fig. 2.5 a)

and 4 (Fig. 2.5 b). The changing shape of each taper or window shows how each filters a different part of the data, and the increase in NW value shows the narrowing of the tapers, as for the same value of M, the spectral resolution (Δf) will be reduced.

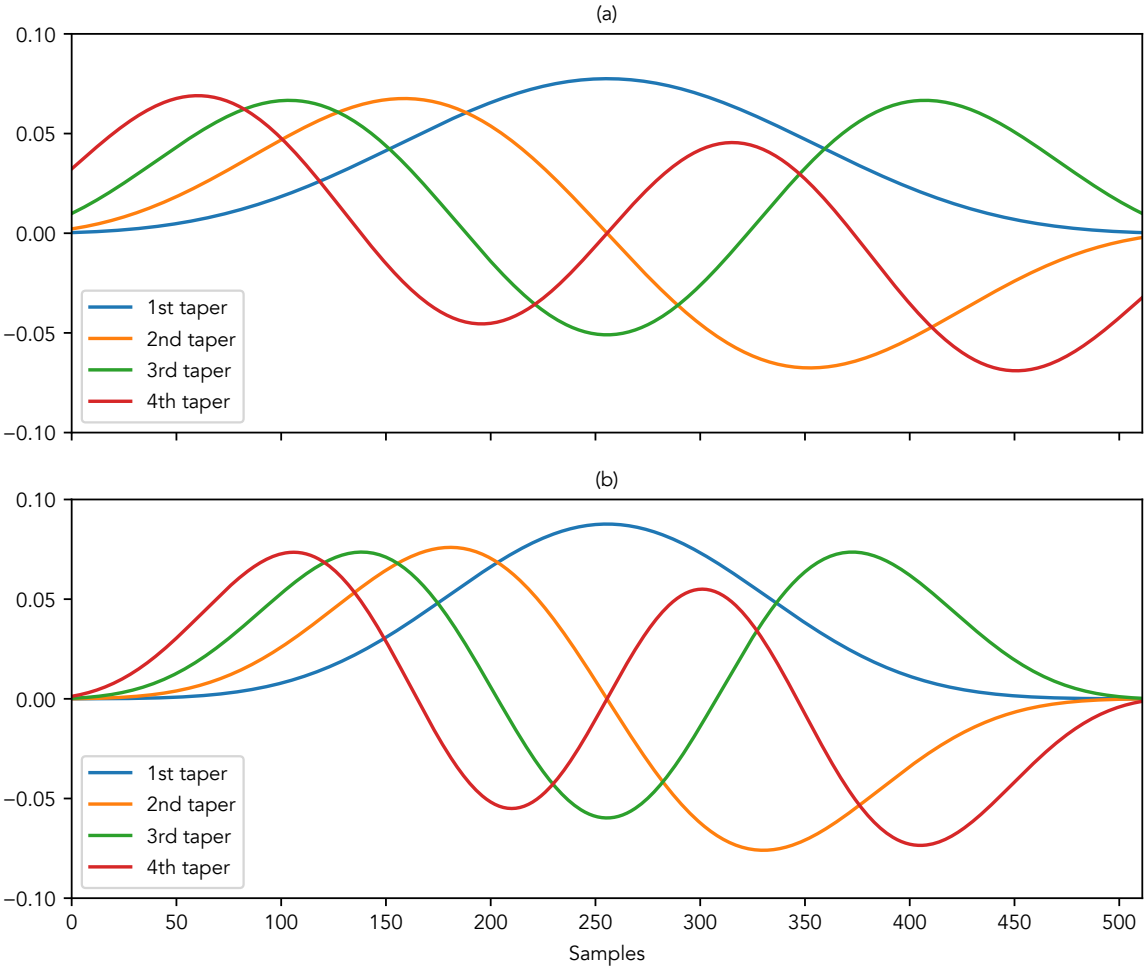

Figure 2.5: A comparison of the first four Slepian tapers produced by the `scipy.signal.windows.dpss` function for M = 512 and (a) NW = 2.5 and (b) NW = 4, showing how although the number of tapers is the same, their shape narrows as the spectral resolution value decreases. The tapers in (a) will result in smoother, lower resolution peaks than those in (b).

Finally, Eq. 2.30 shows how the 95% confidence interval is calculated from the number of tapers, K, for $\alpha = 0.05$:

$$CI_{\alpha} = 1 - \alpha^{\left[\frac{2}{2K-2}\right]}. \tag{2.30}$$

2.3 Summary

The RAPID observations and statistical methods and tests described in this chapter are used in the next chapter to develop empirical linear regressions to model UMO, LNADW and AMOC transport variability. The hydrographic data described in Section 2.1.2 are applied to these models in Chapter 4 to reconstruct historical transports. The regression models are then analysed within a high resolution ocean model (Section 2.1.3) in Chapter 5.

Chapter 3

An empirical model of the Atlantic meridional overturning circulation

3.1 Introduction

This chapter describes how we utilised 13 years of RAPID observations to produce an empirical statistical model of the AMOC. In doing so, we sought to balance the principle of parsimony in regression models – selecting the simplest model with the fewest explanatory variables that explains sufficient variation in the response variable – with one that also captures observed AMOC variability and is grounded in the dynamics at 26°N. Parsimony suggests that we begin with a simple ordinary least squares linear regression, and the starting point for this work was to revisit the statistical model from Longworth et al. (2011), which regresses the transport anomaly within the upper 800 m at 26°N onto the western boundary temperature anomaly at the mid-depth of this layer. The models were then improved by selecting different properties as explanatory variables, by choosing different transports as response variables, and by adding new explanatory variables to capture more of the variance in the selected transport anomaly. Using time series data such as these to train linear regression models can cause problems, so if any were found, an alternative regression model would be needed to address them.

3.2 Data

The empirical linear regression models developed here were trained on data created from the gridded RAPID data products described in Section 2.1.1. These data products comprise western and eastern boundary profiles of conservative temperature, absolute salinity and in situ density with a vertical resolution of 20 dbar and a temporal resolution of 12 hours. Additional RAPID data used were 10-day filtered AMOC transports,

components such as Florida Current and Ekman transport, and depth-defined layer transports such as those for the upper mid-ocean (UMO) and Upper and Lower North Atlantic Deep Water. Although RAPID data were available between 7 April 2004 and 21 February 2017 (Smeed et al., 2017), between 7 November 2005 and 26 May 2006 there was a collapse of the main western mooring, WB2. During this period, WB2 instrument measurements were replaced by those from the WB3 mooring located further to the east. We selected to use only data collected after the WB2 mooring collapse, between 27 May 2006 to 21 February 2017, to keep the training data consistent.

The AMOC and related transports that are calculated from these data (and that are used as response variables in the regression model) are filtered, so the same Butterworth 10-day, low pass filter was applied to the gridded data profiles. These filtered temperature/salinity/density gridded data and the transports were then resampled to monthly means. Anomalies of the monthly mean data were created by detrending for the period 27 May 2006 to 21 February 2017, also ensuring that each time series was stationary. These data were used to train each of the regression models described in the following sections.

3.3 Methods: developing the empirical models

3.3.1 Modelling UMO transport using temperature and salinity

Our initial approach was to revisit the regression in Longworth et al. (2011), an ordinary least-squares (OLS) regression of the transport anomaly within the upper 800 m or thermocline layer on the conservative temperature anomaly at 400 dbar at the western boundary (Eq. 3.1):

$$T_{thermo} = \alpha\, t_{wb}^{400} + \beta. \tag{3.1}$$

The data used by Longworth et al. (2011) were 39 repeat CTD profiles made between 1986 and 1998 during the Abaco current meter deployments (Lee et al., 1996). The temperature anomaly was created by subtracting the group mean of 17.60°C. The coefficient and intercept found for the regression in Eq. 3.1 were $\alpha = -4.75$ (± 1.50)

and $\beta = 0.00$ (± 0.42). Longworth et al. (2011) suggested that increasing southwards thermocline transport (a negative transport anomaly) causes more warm water to recirculate close to the western boundary, and so the temperature at any particular depth within the thermocline will increase (a positive temperature anomaly). They chose 400 dbar as both the mid-point of the thermocline layer, and a depth at which each of the 39 CTD profiles had a sensor deployed within 50 m. In contrast, the RAPID western boundary profile provided 132 conservative temperature anomalies at 400 dbar. For this and all subsequent regression models, monthly mean values described in Section 3.2 were used for all variables. The left-hand scatter plot in Fig. 3.1 shows that the same linear regression using the RAPID monthly mean temperature and thermocline transport anomalies explains only 22% of the variance in the observed thermocline transport anomaly, compared with 53% found by Longworth et al. (2011) using CTD profiles.

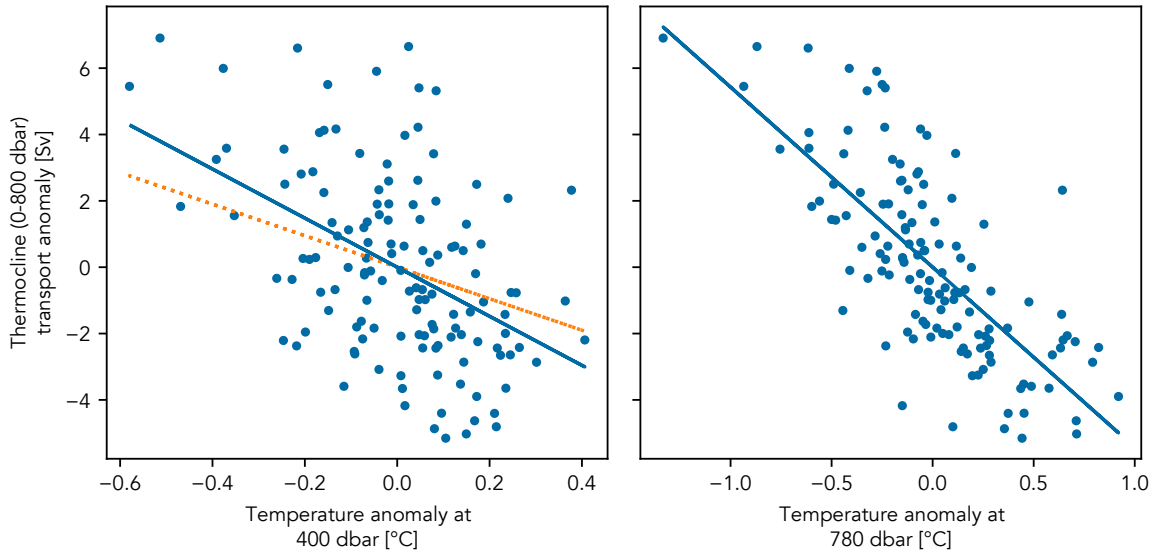

Figure 3.1: Linear regression of the monthly mean RAPID thermocline (0–800 m) transport anomaly on the monthly mean conservative temperature anomaly at (left) 400 dbar and (right) 780 dbar from the RAPID western boundary profiles. The blue line is the regression equation for these data; the orange dashed line in the left-hand figure shows the regression equation from Longworth et al. (2011) applied to the same data.

To investigate whether the temperature anomaly at a different depth within the thermocline layer would better model the thermocline transport, an algorithm was created that repeated the regression using the conservative temperature anomaly every 20 dbar between 220 and 800 dbar. The algorithm selected a depth of 780 dbar, at which the

regression, shown in the right-hand scatter plot of Fig. 3.1, gave the highest explained variance of 53% (equivalent to $R^2 = 0.53$).

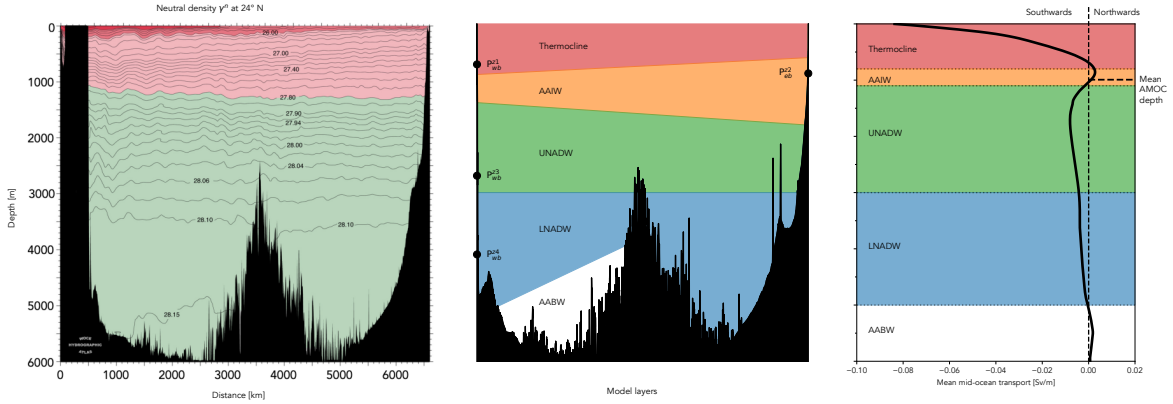

Figure 3.2: (Left) World Ocean Circulation Experiment (WOCE) North Atlantic A05 section of neutral density (γ^n) at 24°N. From the WOCE Atlantic Ocean Atlas Vol. 3. (Middle) Schematic of four dynamic layers to be represented within the regression models by a property P (black circles) at the western and eastern boundaries at depths $z1 - z4$ within each layer. (Right) Profile of RAPID-estimated mean mid-ocean transport with depth and the RAPID-defined northwards and southwards layer transports. Mean AMOC depth is around 1100 m.

Key to the research here is the idea that a single-layer proxy cannot effectively model the overturning circulation at 26°N, and that further explanatory or independent variables need to be added to the linear regression to represent the changes in water masses and circulation with depth. The mean mid-ocean transport profile in Fig. 3.2 shows the transport layers defined by RAPID that correspond to different water masses: the thermocline, which has already been described, between the surface and 800 dbar; Antarctic Intermediate Water (AAIW) between 800 and 1100 dbar; Upper North Atlantic Deep Water (UNADW) between 1100 and 3000 dbar; and finally Lower North Atlantic Deep Water (LNADW) between 3000 and 5000 dbar. RAPID also defines an even deeper, northwards-flowing water mass, Antarctic Bottom Water (AABW), but as this is not located at the western boundary, but leans against the western flank of the Mid-Atlantic Ridge, it is omitted from these regression models. As AMOC variability on most timescales is dominated by variability at the western boundary (Bryden et al., 2009; Elipot et al., 2014; Worthington et al., 2019), and density is dominated by temperature at this latitude, initially each water mass or layer was represented in the regression model by a western boundary temperature anomaly, a natural extension of Longworth et al. (2011) using the 400 dbar

temperature anomaly as a proxy for thermocline transport. The exception was the proxy for the intermediate depth layer of northwards-flowing AAIW. Seasonal variability in the AMOC is driven largely by eastern boundary density anomalies with a sub-surface maximum around 1000 dbar, and these density anomalies are strongly influenced by salinity (Chidichimo et al., 2010). The AAIW water mass, or layer, was thus represented by a salinity anomaly on the eastern boundary between 800 and 1100 dbar. For the transport anomaly chosen as the response variable, UMO transport was preferred to thermocline layer transport, since the UMO is the main contributor to the AMOC (Smeed et al., 2014). The central schematic in Fig. 3.2 shows the boundary properties that represent each of these water masses and the changes in their meridional volume transport within the regression models. For example, P_{wb}^{z1} is the property at the western boundary at depth $z1$ that is a proxy for transport within the thermocline layer. We constructed regression models using one, two, three and four of these proxy boundary properties as explanatory variables, with their depths ($z1$, $z2$, $z3$ and $z4$), and associated coefficients (α, β, γ and ζ) to be found by algorithm (Eq. 3.2):

$$T_{umo} = \alpha\, t_{wb}^{z1}, \tag{3.2a}$$

$$T_{umo} = \alpha\, t_{wb}^{z1} + \beta\, s_{eb}^{z2}, \tag{3.2b}$$

$$T_{umo} = \alpha\, t_{wb}^{z1} + \beta\, s_{eb}^{z2} + \gamma\, t_{wb}^{z3}, \tag{3.2c}$$

$$T_{umo} = \alpha\, t_{wb}^{z1} + \beta\, s_{eb}^{z2} + \gamma\, t_{wb}^{z3} + \zeta\, t_{wb}^{z4}. \tag{3.2d}$$

For linear regressions with one, two and three explanatory variables; combinations of 132 monthly mean conservative temperature or absolute salinity anomalies every 20 dbar were used. For example, for the second linear regression (Eq. 3.2b), each western boundary conservative temperature anomaly every 20 dbar between 220 and 780 dbar was combined in turn with each eastern boundary absolute salinity anomaly every 20 dbar between 800 and 1080 dbar, and regressed against the UMO transport anomaly. The two depths that resulted in the highest adjusted R^2 value for the regression were then selected as $z1$ and

$z2$. Due to the number of iterations required by adding a fourth explanatory variable, the western boundary anomalies at depths $z3$ and $z4$ which represent the UNADW and LNADW water masses were selected every 100 dbar rather than every 20. This also reflects the lower resolution of deeper-deployed RAPID instruments, justified by the decrease in both variability and vertical temperature gradients with depth.

3.3.2 Modelling UMO transport using density

Although the initial regression models built on the work done by Longworth et al. (2011) used temperature and salinity anomalies, the geostrophic UMO transport is estimated using dynamic height, which is directly related to density (Eq. 2.6). Frajka-Williams et al. (2016) showed strong positive correlation between UMO transport and isopycnal displacements on the western boundary around 820 m, and negative correlation between LNADW transport and isopycnal displacements on the western boundary between 1500 m and the bottom.

The seasonal cycle of the AMOC is driven largely by seasonality at the eastern boundary (Chidichimo et al., 2010; Pérez-Hernández et al., 2015). The annual maximum northward transport at the eastern boundary and of the AMOC occur around October (Vélez-Belchí et al., 2017) and are driven by changes in the circulation of the Canary Current (Casanova-Masjoan et al., 2020; Hernández-Guerra et al., 2017) and at intermediate depths (700–1400 dbar) by seasonal changes in the Intermediate Poleward Undercurrent (Hernández-Guerra et al., 2017; Vélez-Belchí et al., 2017). Eastern boundary density anomalies have maximum subsurface variability around 1000 dbar (Chidichimo et al., 2010).

These studies provided sufficient justification for replacing temperature and salinity anomalies with density anomalies as the proxy properties at both western and eastern boundaries, and repeating the process. The resulting regression equations (Eq. 3.3) again show the depths and coefficients for each explanatory variable to be found by algorithm using the same method described in Section 3.3.1:

$$T_{umo} = \alpha\ \rho_{wb}^{z1}, \tag{3.3a}$$

$$T_{umo} = \alpha\ \rho_{wb}^{z1} + \beta\ \rho_{eb}^{z2}, \tag{3.3b}$$

$$T_{umo} = \alpha\ \rho_{wb}^{z1} + \beta\ \rho_{eb}^{z2} + \gamma\ \rho_{wb}^{z3}, \tag{3.3c}$$

$$T_{umo} = \alpha\ \rho_{wb}^{z1} + \beta\ \rho_{eb}^{z2} + \gamma\ \rho_{wb}^{z3} + \zeta\ \rho_{wb}^{z4}. \tag{3.3d}$$

3.3.3 Modelling LNADW transport using deep density

The greater influence of the LNADW in the AMOC decline compared to the UNADW (Smeed et al., 2018), and the importance of deep density changes in detecting long-term AMOC decline (e.g., Baehr et al., 2007) suggested the utility of an additional linear regression model. As shown in Eq. 3.7, this would regress the LNADW transport anomaly onto a western boundary density anomaly at a depth $z4$ within the LNADW layer, and the Ekman transport anomaly:

$$T_{lnadw} = \alpha\ \rho_{wb}^{z4} + \beta\ T_{ekman}. \tag{3.4}$$

The Ekman transport is included in the model as Frajka-Williams et al. (2016) found that LNADW transport showed a deep baroclinic response to changes in Ekman transport. We applied a similar algorithm to the previous regression models, repeating the regression for the deep density anomaly every 20 dbar between 3000 dbar and 4820 dbar, and reporting the maximum explained variance as the adjusted R^2 found at depth $z4$ with coefficients α and β.

Finally, as stated in Section 2.2.1.3, there are several assumptions that must hold for an ordinary least squares regression (OLS) model to be fit for purpose. Each of the OLS regressions in Eq. 3.2, Eq. 3.3 and Eq. 3.7 were checked against these assumptions, and any necessary changes were made to the regression models used.

3.4 Results

3.4.1 Ordinary Least Squares Regression Models

3.4.1.1 Modelling UMO transport using temperature and salinity

The coefficients and depths of the ordinary least squares regression models (Eq. 3.2) that were selected by algorithm as having the highest adjusted R^2 values are shown in Eq. 3.5:

$$T_{umo} = -5.4\ t_{wb}^{780}, \tag{3.5a}$$

$$T_{umo} = -5.1 t_{wb}^{780} - 30.5\ s_{eb}^{1080}, \tag{3.5b}$$

$$T_{umo} = -4.8\ t_{wb}^{780} - 29.8\ s_{eb}^{1080} - 3.4\ t_{wb}^{1200}, \tag{3.5c}$$

$$T_{umo} = -4.8\ t_{wb}^{780} - 31.2\ s_{eb}^{1080} - 3.5\ t_{wb}^{1200} - 4.8\ t_{wb}^{3000}, \tag{3.5d}$$

where t is a conservative temperature anomaly in °C, and s is an absolute salinity anomaly, with the superscript showing the depth of each anomaly, and the subscript whether the anomaly is from the western (wb) or eastern (eb) boundary profile. For each regression, the independently-run algorithm selected the same depth, for example the thermocline temperature anomaly was selected at 780 dbar each time. The coefficients vary slightly for comparable anomalies, but are the same to first-order approximation.

The UMO transport anomaly predicted by each regression model is compared against the RAPID-observed UMO transport anomaly in Fig. 3.3, and shown with the adjusted R^2 of each regression. The figures shows each additional explanatory variable increases the adjusted R^2, from 0.49 for the single temperature anomaly at 780 dbar, to 0.58 by adding the eastern boundary salinity anomaly at 1080 dbar, to 0.59 by adding the deeper western boundary temperature anomaly at 1240 dbar. Almost no improvement (~0.004) in adjusted R^2 is made by adding the deepest western boundary density anomaly at 3000 dbar. The shaded orange area around each predicted UMO transport anomaly is the 95%

Linear regression model	**Adj. R^2**	**SE [Sv]**	**Mean error [Sv]**	**σ error [Sv]**	**RMS error [Sv]**	**AIC**	**BIC**
Thermocline	0.491	2.05	1.63	1.22	2.04	566.2	572.0
Therm. + AAIW	0.577	1.87	1.44	1.17	1.85	542.9	551.5
Therm. + AAIW + UNADW	0.586	1.85	1.44	1.12	1.82	541.0	552.6
Therm. + AAIW + UNADW + LNADW	0.590	1.84	1.41	1.13	1.81	540.5	554.9

Table 3.1: Comparison of model statistics from the ordinary least squares linear regressions in Eq. 3.5. Statistics shown are the adjusted R^2 ; standard error of the regression (SE); the mean, standard deviation (σ) and root-mean-square (RMS) of the error, where the error is the difference between the actual and model-predicted values; Akaike's Information Criterion (AIC); and Bayes Information Criterion (BIC).

prediction interval. This interval should contain 95% of future observable realisations with the same explanatory variable values. The observed UMO transport anomaly is within the prediction interval over 80% of the time for each regression model, however the mean prediction interval is large, 4.1 Sv for the simple regression, and reducing to 3.7 Sv for the multiple linear regressions.

As a measure of the actual difference between the 132 predicted and observed monthly mean UMO values, the standard error of the regression (see Section 2.2.1.2) is improved by the addition of each explanatory variable. Table 3.1 shows it reduces from 2.05 Sv to 1.87, 1.85 and 1.84 Sv with the addition of each variable to the regression. It also shows that adding variables also reduces the mean model error, or the difference between the values predicted by the model and the observed values it was trained on, from 1.63 to 1.41 Sv. The root-mean-square (RMS) error also decreases from 2.04 to 1.81 Sv, although this is expected as the RMS error is a very similar statistic to the standard error of the regression, dividing the squared sum of the residuals by the number of variables n, rather than by $n-p$, where p is the number of parameters (Eq. 2.13). The standard deviation (σ) of the error also decreases with each additional variable, from 1.22 to 1.17 then 1.12 Sv, however adding the fourth variable increases the standard deviation by 0.01 Sv. Similar improvement is seen in the Akaike's Information Criterion (AIC) metric, which is reduced from 566.2 to 542.9, 541.0 and 540.5. The exception is the Bayes Information Criterion metric, which after reducing from 572.0 to 551.5, then increases slightly with the addition of the third and fourth explanatory variables to 552.6 and 554.6. This suggests that the improvement by the last two regression models is marginal when using temperature and salinity anomalies.

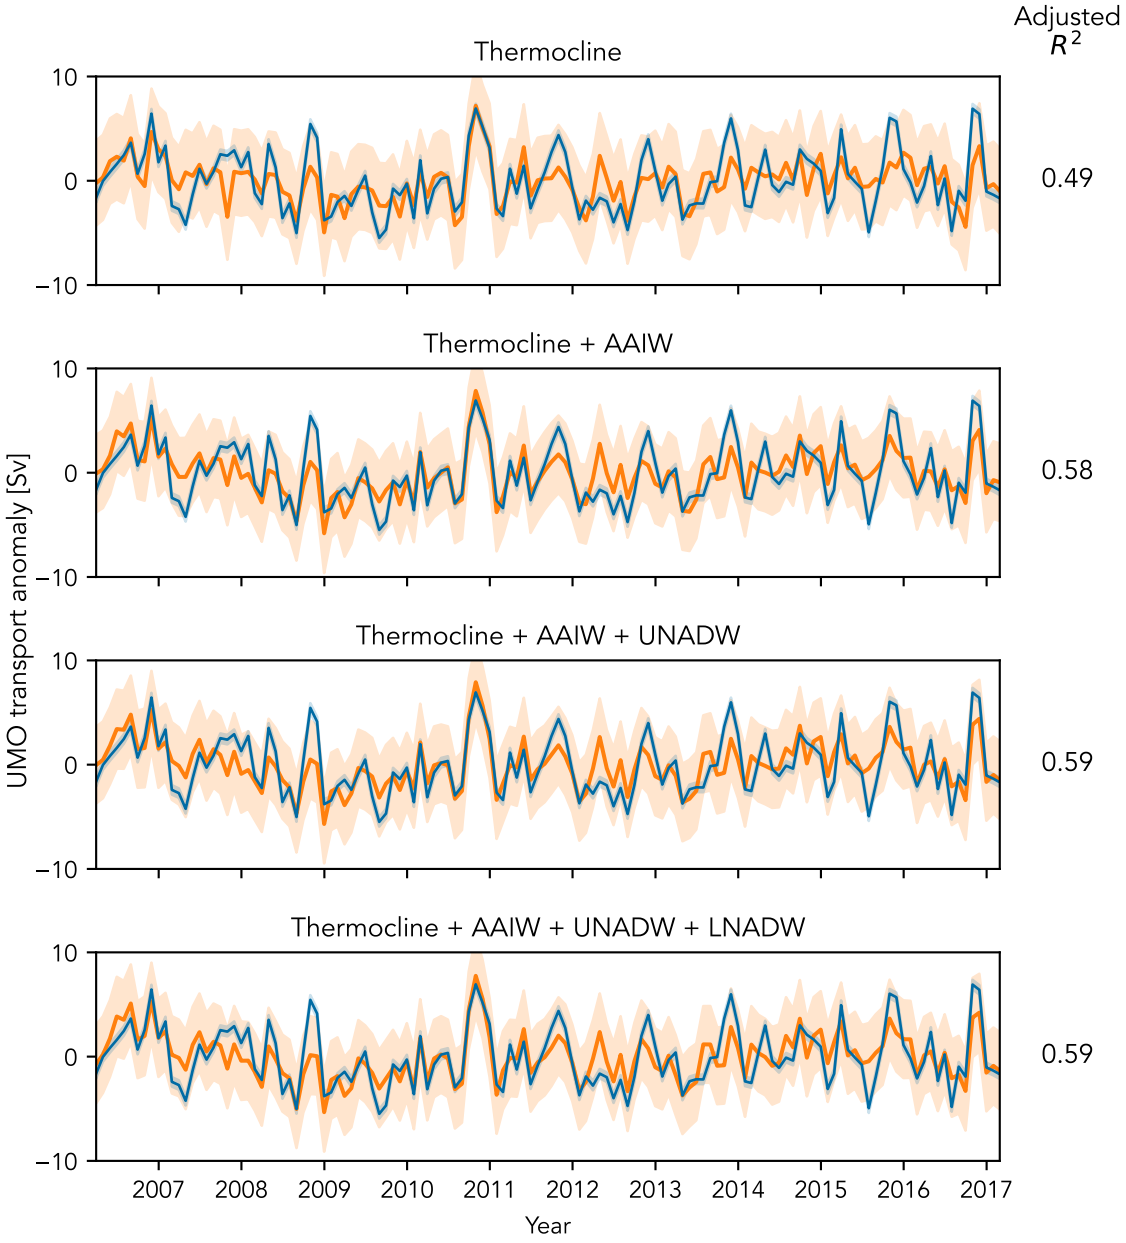

Figure 3.3: Comparison of monthly mean UMO transport anomalies predicted by OLS linear regression models (orange line), and observed by RAPID (blue line). The models have one, two, three and four explanatory variables, representing different water masses at 26°N, which are named above each plot. The explanatory variables are temperature (thermocline, UNADW and LNADW) and salinity (AAIW) anomalies. Adjusted R^2 values for each model are shown to the right of the related plot. The orange shaded area around each model prediction line shows the 95% prediction interval of the model.

3.4.1.2 Modelling UMO transport using density

Eq. 3.6 show the depths and coefficient obtained by algorithm for the regression equations in Eq. 3.3 when the boundary temperature and salinity anomalies are replaced by density anomalies:

$$T_{umo} = 52.0\ \rho_{wb}^{780}, \tag{3.6a}$$

$$T_{umo} = 45.9\ \rho_{wb}^{760} - 107.4\ \rho_{eb}^{920}, \tag{3.6b}$$

$$T_{umo} = 45.2\ \rho_{wb}^{760} - 111.0\ \rho_{eb}^{920} + 47.4\ \rho_{wb}^{2740}, \tag{3.6c}$$

$$T_{umo} = 41.2\ \rho_{wb}^{740} - 110.2\ \rho_{eb}^{920} + 37.5\ \rho_{wb}^{1200} + 44.1\ \rho_{wb}^{3000}, \tag{3.6d}$$

where ρ is the in situ density anomaly (kg/m^3), and again the superscript and subscript indicate the anomaly depth and the boundary respectively.

Linear regression model	**Adj. R^2**	**SE [Sv]**	**Mean error [Sv]**	**σ error [Sv]**	**RMS error [Sv]**	**AIC**	**BIC**
Thermocline	0.501	2.03	1.62	1.20	2.02	63.7	569.4
Therm. + AAIW	0.781	1.34	1.10	0.75	1.33	455.8	464.4
Therm. + AAIW + UNADW	0.795	1.30	1.03	0.76	1.28	448.2	459.7
Therm. + AAIW + UNADW + LNADW	0.804	1.27	1.00	0.75	1.25	443.5	457.9

Table 3.2: Comparison of model statistics from the ordinary least squares linear regressions in Eq. 3.6. Statistics shown are the adjusted R^2 ; standard error of the regression (SE); the mean, standard deviation (σ) and root-mean-square (RMS) of the error, where the error is the difference between the actual and model-predicted values; Akaike's Information Criterion (AIC); and Bayes Information Criterion (BIC).

Using density anomalies as proxies for each layer transport makes a slight difference for the simple linear regression: the adjusted R^2 improves from 0.49 to 0.50; the standard error reduces from 2.05 to 2.03 Sv; the mean error reduces from 1.63 to 1.62 Sv; the standard deviation of the error from 1.22 to 1.20 Sv; the RMS error from 2.04 to 2.02

Sv; and the depth of the density anomaly (780 dbar) selected by algorithm remains the same. The AIC and BIC statistics both improve slightly, AIC decreasing from 566.2 to 563.7 and BIC from 572.0 to 569.4. The coefficient is obviously different, positive rather than negative, reflecting that a positive UMO transport anomaly, indicating weakened southward sub-tropical gyre recirculation, is related to less warm (and hence more dense) water recirculating at the western boundary.

The selected depths for the multiple linear regressions show more variation than when trained on temperature and salinity anomalies, but there is a maximum difference of 40 dbar between $z1$ values in Eq. 3.6 a-d. Comparing the depths selected between the temperature/salinity- and density-based models, the greatest differences are between the AAIW eastern boundary anomaly depths, which are 1080 dbar in Eq. 3.5 and 920 dbar in Eq. 3.6. The UNADW density anomaly is selected at 1200 dbar for each regression except Eq. 3.6c where it is 2740 dbar, and the deepest LNADW western boundary anomaly depth is consistently selected at 3000 dbar.

Adding a second density anomaly at the eastern boundary at 920 dbar (Eq. 3.6b) increases the adjusted R^2 significantly, from 0.50 to 0.78, and reduces the standard error from 2.03 to 1.34 Sv, and the RMS error similarly from 2.02 to 1.33 Sv. The mean error also decreases from 1.62 to 1.10 Sv, and the standard deviation of the residuals also decreases from 1.20 to 0.75 Sv. These statistics show a marked improvement when compared to the difference made by adding a second explanatory variable to the temperature/salinity-trained models. Adding the third and then fourth explanatory variables representing the upper and lower NADW water masses improves the adjusted R^2 to 0.80 (0.795 and 0.804 respectively), the standard error to 1.30 and 1.27 Sv and the RMS error to 1.28 and 1.25 Sv respectively (Table 3.2). The mean error also improves to 1.03 and 1.00 Sv, however the standard deviation remains between 0.75 and 0.76 Sv. The best model fit is found by combining western boundary density anomalies at 740, 1200 and 3000 dbar with an eastern boundary density anomaly at 920 dbar, however the improvement by adding the fourth variable is very small. Comparing predicted and observed UMO transport anomalies in Fig. 3.4 also shows the prediction interval reduces in width compared to the temperature/salinity-based regressions. The mean prediction interval for the simplest regression is 4.0 Sv, and reduces to 2.7, 2.6 and 2.6 Sv for each

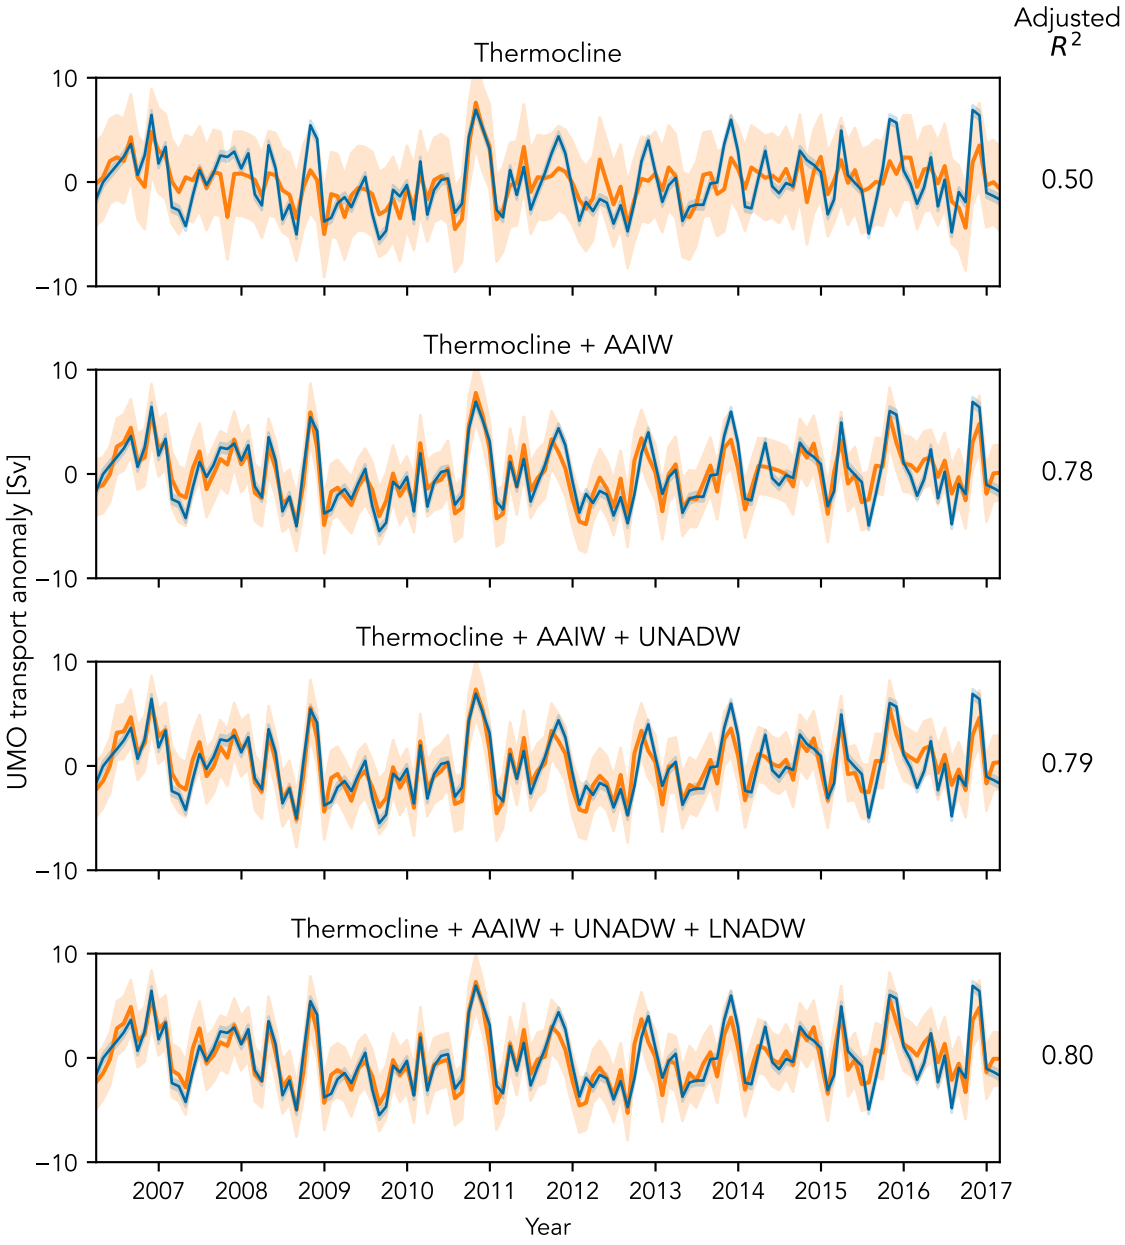

Figure 3.4: Comparison of monthly mean UMO transport anomalies predicted by OLS linear regression models (orange line), and observed by RAPID (blue line). The models have one, two, three and four explanatory variables, representing different water masses at 26°N, which are named above each plot. The explanatory variables here are each density anomalies. Adjusted R^2 values for each model are shown to the right of the related plot. The orange shaded area around each model prediction line shows the 95% prediction interval of the model.

additional explanatory variable. The AIC and BIC metrics follow a different pattern to the temperature/salinity-based regressions, decreasing by more than 2 with almost each variable added, which indicates a better model fit, with the exception an increase in BIC of only 1.8 between the 3- and 4-variable regressions.

Although the adjusted R^2 increase between the third and fourth models is only 0.09, there is a 0.03 Sv improvement in the standard error. Also, although BIC penalises additional parameters more harshly, it still shows a small (although not necessarily significant) improvement in model fit by adding the LNADW proxy variable. Given all the model fit statistics, the final regression using four boundary density anomalies was selected to model the UMO transport anomaly.

3.4.1.3 Modelling LNADW transport using deep density

The algorithm for the regression in Eq. 3.7 found that a depth of 3000 dbar for the deep density anomaly gave the highest R^2 value (adjusted $R^2 = 0.754$) and standard error of the regression, or average deviation of the model predictions from the observations, of 1.04 Sv (Table 3.3):

$$T_{lnadw} = -173.3\ \rho_{wb}^{3000} - 0.40\ T_{ekman}. \tag{3.7}$$

The coefficient of the density anomaly ρ_{wb}^{3000}, -173.3, shows that a decrease in density at 3000 dbar is associated with an increase in LNADW transport, and the magnitude of the coefficient shows that, for a constant Ekman transport, a LNADW transport change of 1 Sv requires a density anomaly change of less than 0.006 kg/m^3.

Linear regression model	**Adj. R^2**	**SE [Sv]**	**Mean error [Sv]**	**σ error [Sv]**	**RMS error [Sv]**	**AIC**	**BIC**
LNADW	0.754	1.04	0.81	0.62	1.03	387.2	395.8

Table 3.3: Comparison of model statistics from the ordinary least squares linear regression in Eq. 3.7. Statistics shown are the adjusted R^2 ; standard error of the regression (SE); the mean, standard deviation (σ) and root-mean-square (RMS) of the error, where the error is the difference between the actual and model-predicted values; Akaike's Information Criterion (AIC), and Bayes Information Criterion (BIC).

The time series in Fig. 3.5 show that the LNADW transport anomaly predicted by this regression model compares well with the RAPID observation, with the error having a mean of 0.81 Sv and a standard deviation of 0.62 Sv. The RMS error is again very close to the standard error of the regression, at 1.03 Sv. This regression model captures the anomaly variability well, effectively representing periods of increasing and decreasing LNADW transport, and the shift from one to the other. This, combined with the high model fit statistics, meant this model was also selected to reconstruct LNADW transport.

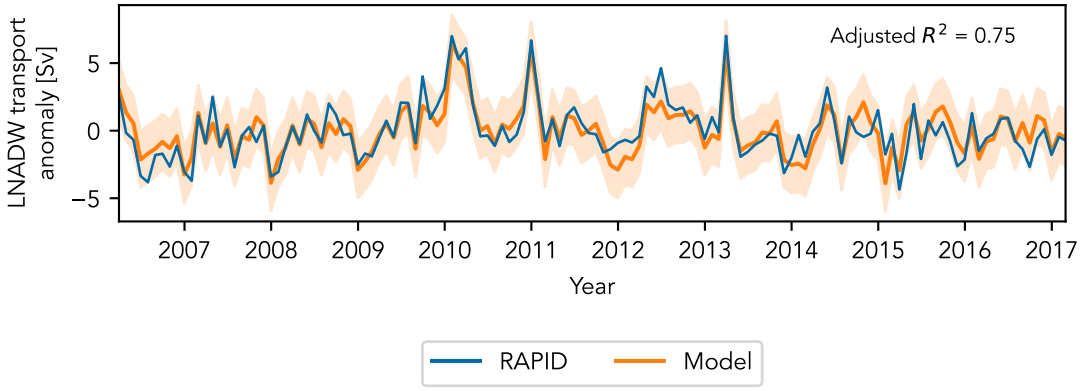

Figure 3.5: Comparison of monthly mean LNADW transport anomalies predicted by the OLS linear regression model in Eq. 3.7 (orange line), and observed by RAPID (blue line). The model R^2 value, adjusted for the degrees of freedom of each model is also shown. The orange shaded region shows the model 95% prediction interval.

3.4.2 Testing assumptions of OLS regression models

The results of the ordinary least squares regression models were promising, however the selected models (shown again in Eq. 3.8 and Eq. 3.9) needed to be assessed to ensure that they satisfied the assumptions of linear regression described in Section 2.2.1.3:

$$T_{umo} = 41.2\ \rho_{wb}^{740} - 110.2\ \rho_{eb}^{920} + 37.5\ \rho_{wb}^{1200} + 44.1\ \rho_{wb}^{3000}, \tag{3.8}$$

$$T_{lnadw} = -173.3\ \rho_{wb}^{3000} - 0.40\ T_{ekman}. \tag{3.9}$$

3.4.2.1 Linearity

The scatter plots in Fig. 3.6 and Fig. 3.7 of explanatory against response variables for the regression equations Eq. 3.8 and Eq. 3.9 show that the regression parameters are linear, although for the UMO model, that linearity becomes less clear with increasing depth of density anomaly. The scatter plots do not show any obvious non-linearity, for example a parabolic shape.

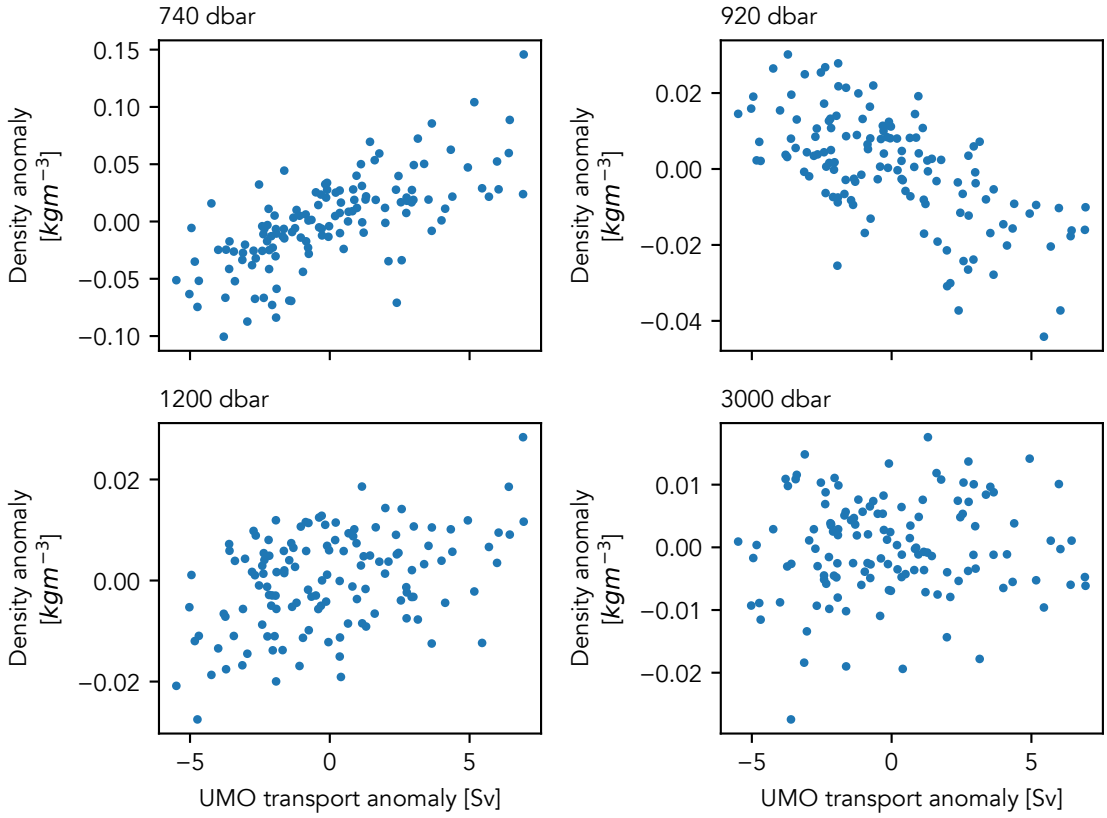

Figure 3.6: Explanatory variables (density anomalies at indicated depths) vs. response variable (UMO transport anomaly) for the ordinary least squares UMO regression model (Eq. 3.8) to check for linearity.

Fig. 3.8 shows the relationship between the two explanatory variables in the LNADW linear regression model. As both variables show relatively linear relationships with the LNADW transport anomaly, it is not surprising that they show some weak positive correlation. The 10-day filtered Ekman transport has an integral timescale of 10.9 days (Table 1.1) while the 10-day filtered western boundary density anomaly at 3000 dbar has an integral timescale of 18.3 days. In contrast, the 10-day filtered LNADW transport has an

integral timescale of 36.8 days. To create the model, the monthly mean values of all time series are taken, which will further smooth the Ekman transport and deep density anomaly time series and bring their timescales closer to that of the LNADW transport.

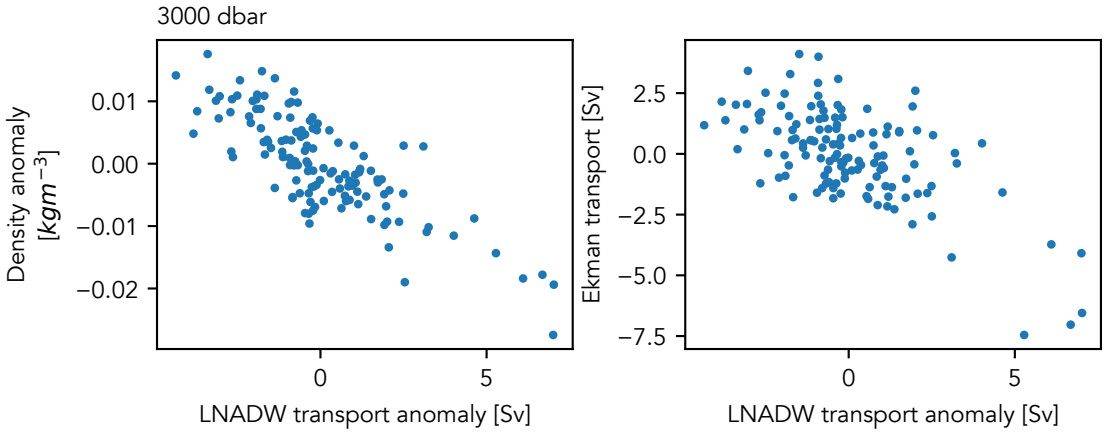

Figure 3.7: Explanatory variables (western boundary density anomaly at 3000 dbar and Ekman transport anomaly) vs. response variable (LNADW transport anomaly) for the ordinary least squares LNADW regression model (Eq. 3.9) to check for linearity.

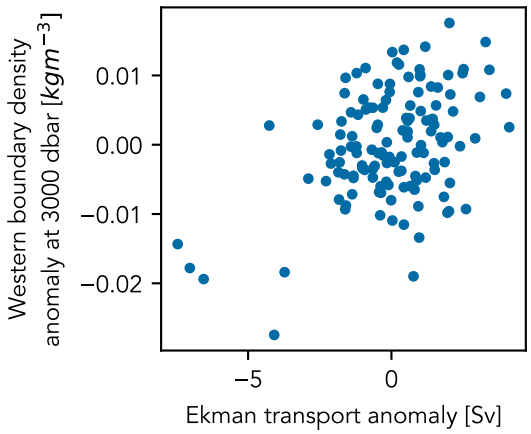

Figure 3.8: Scatter plot of monthly mean Ekman transport anomaly against the deep density anomaly at 3040 dbar used to create the LNADW regression model

3.4.2.2 Homoscedasticity

We tested the regression models for homoscedasticity both informally by plotting the residuals against the fitted values, and formally using the Breusch-Pagan test. The scatter plots in Fig. 3.9 show that variance in the residuals is not always constant, and tends to

be greater for fitted values closer to zero in each. Since there is some doubt, we use the Koenker version of the Breusch-Pagan test to check for homoscedasticity in each regression.

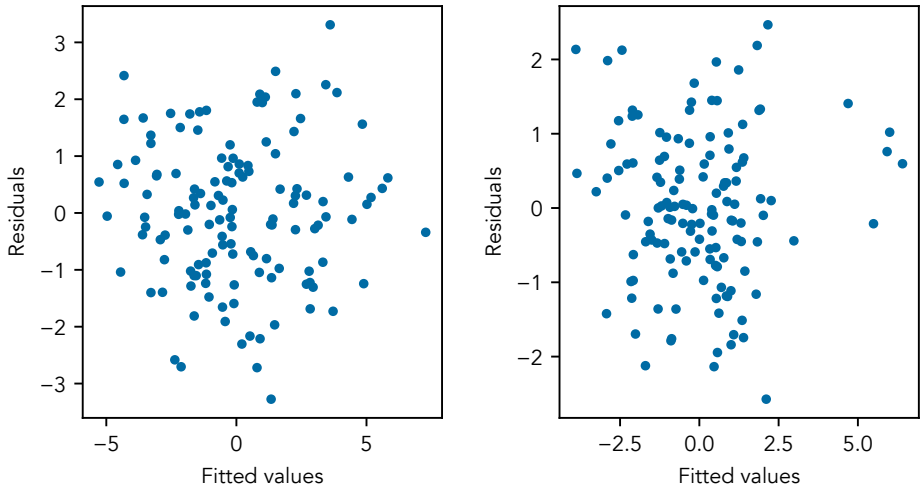

Figure 3.9: Scatter plot of residuals against fitted values for linear regression models of (left) UMO and (right) LNADW transport anomalies.

The Koenker version of the Breusch-Pagan test gives a p-value of 0.68 for the UMO model and 0.79 for the LNADW model. As these are both greater than 0.05, we cannot reject the null hypothesis in either case, and can assume that homoscedasticity is present in these regression models. We can conclude that the variations in residual variances observed in Fig. 3.9 were not significant enough to cause concern.

3.4.2.3 No perfect multicollinearity

We examined the variance inflation factors (VIF) to test for independence of our explanatory variables in the regression shown in Eq. 3.8. As Table 3.4a shows, each of the explanatory variables has a VIF of 1.14 or less, which is below any threshold VIF commonly used. Thus there is no multicollinearity for these explanatory variables, and they are independent of each other. The same is true for the regression in Eq. 3.9, with a VIF for Ekman transport of 1.23 (Table 3.4b) showing there is no multicollinearity with the deep western boundary density anomaly ρ_{wb}^{3000}.

Explanatory variable	VIF
ρ_{wb}^{780}	1.00
ρ_{eb}^{980}	1.13
ρ_{wb}^{1200}	1.07
ρ_{wb}^{3000}	1.14

(a) UMO model

Explanatory variable	VIF
ρ_{wb}^{3000}	1.00
T_{ekman}	1.23

(b) LNADW model

Table 3.4: Variance inflation factors (VIF) for each explanatory variable from the (left) UMO regression model (Eq. 3.8) and (right) LNADW regression model (Eq. 3.9), with density anomaly (ρ) superscripts showing the pressure in dbar and subscripts showing the boundary.

3.4.2.4 Normality of error terms

We used quantile-quantile (Q-Q) and histogram plots (Fig. 3.10 and Fig. 3.11) to examine the normality of error terms for our regression models. For a perfectly normal distribution, the residuals should lie exactly on the diagonal line in the Q-Q plot. The residuals from Eq. 3.8 are not perfectly on the line, but close enough to suggest an approximately normal distribution. The histogram plot in Fig. 3.10 also shows a distribution that approaches normality. Both the Q-Q plot and histogram in Fig. 3.11 indicate a residual distribution very close to normality.

More formally, we use the Shapiro-Wilk test implementation `scipy.stats.shapiro`. For our UMO regression model residuals, the Shapiro-Wilk test returns a p-value of 0.93, which is greater than 0.05, so we can reject the null hypothesis, and conclude that there is insufficient evidence that the residuals do not come from a normal distribution. The Shapiro-Wilk p-value for the LNADW regression model residuals is 0.88, again confirming the nearly-normal distribution.

3.4.2.5 Autocorrelation of residuals

Autocorrelation of residuals was seen in all of the ordinary least squares regressions, an example for the selected UMO regression model Eq. 3.8 can be seen in the autocorrelation function plot in Fig. 3.12, which shows significant autocorrelation at lags of one and four months, and indicates that autocorrelation of residuals is a problem for this ordinary least squares model. The partial autocorrelation function plot (Fig. 3.12) shows autocorrelation

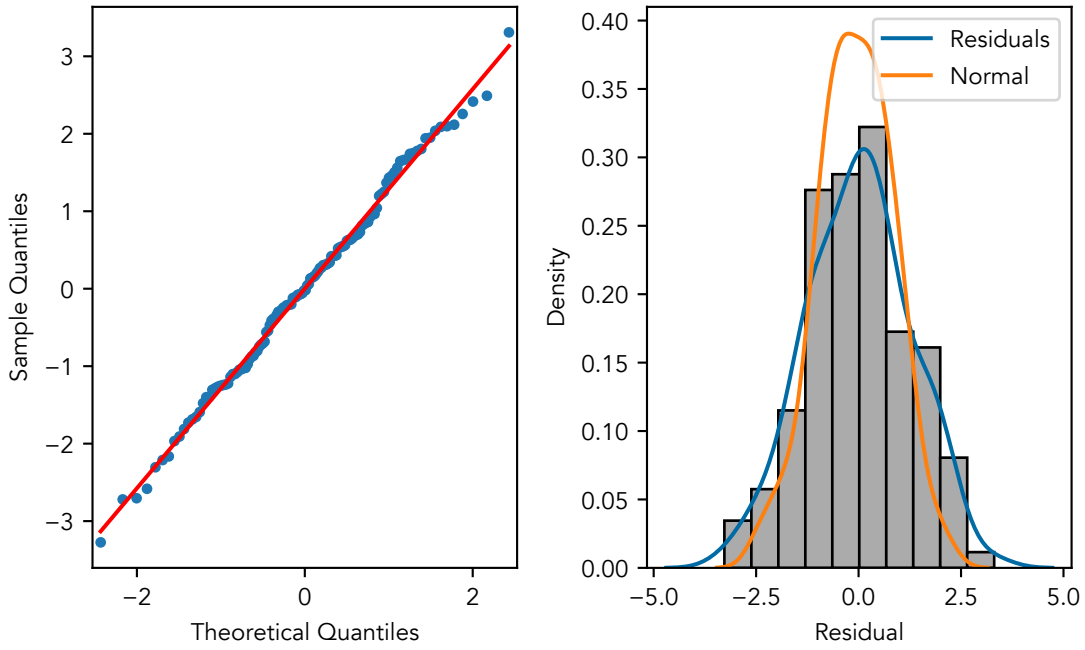

Figure 3.10: (Left) A quantile-quantile (Q-Q) plot and (Right) histogram plot of the residuals for the UMO regression model in Eq. 3.8. The lines show the probability density curve for the residuals (blue) and a random normal distribution of the same size (orange).

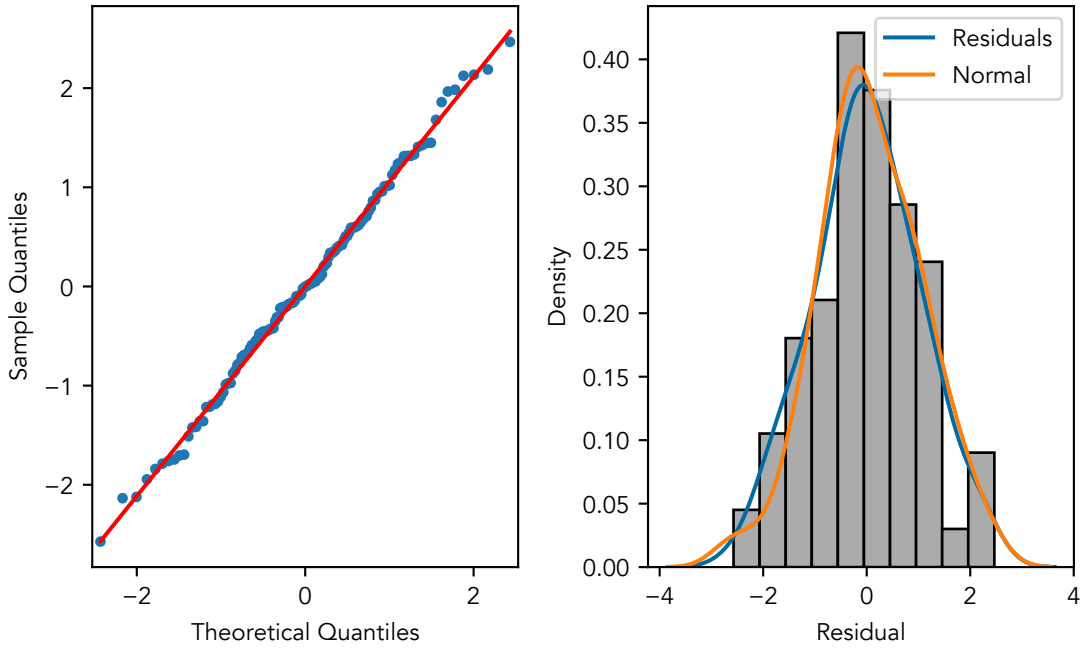

Figure 3.11: (Left) A quantile-quantile (Q-Q) plot and (Right) histogram plot of the residuals for the LNADW regression model in Eq. 3.9. The lines show the probability density curve for the residuals (blue) and a random normal distribution of the same size (orange).

between two residuals with the correlation due to any intervening residuals removed, and is used to find the lag for any autoregressive model. It shows partial autocorrelation is significant for a lag of one month, and this is the case for each of the ordinary least squares regressions. This is consistent with the decorrelation length scale associated with the AMOC of 40 days, shown by (Smeed et al., 2014).

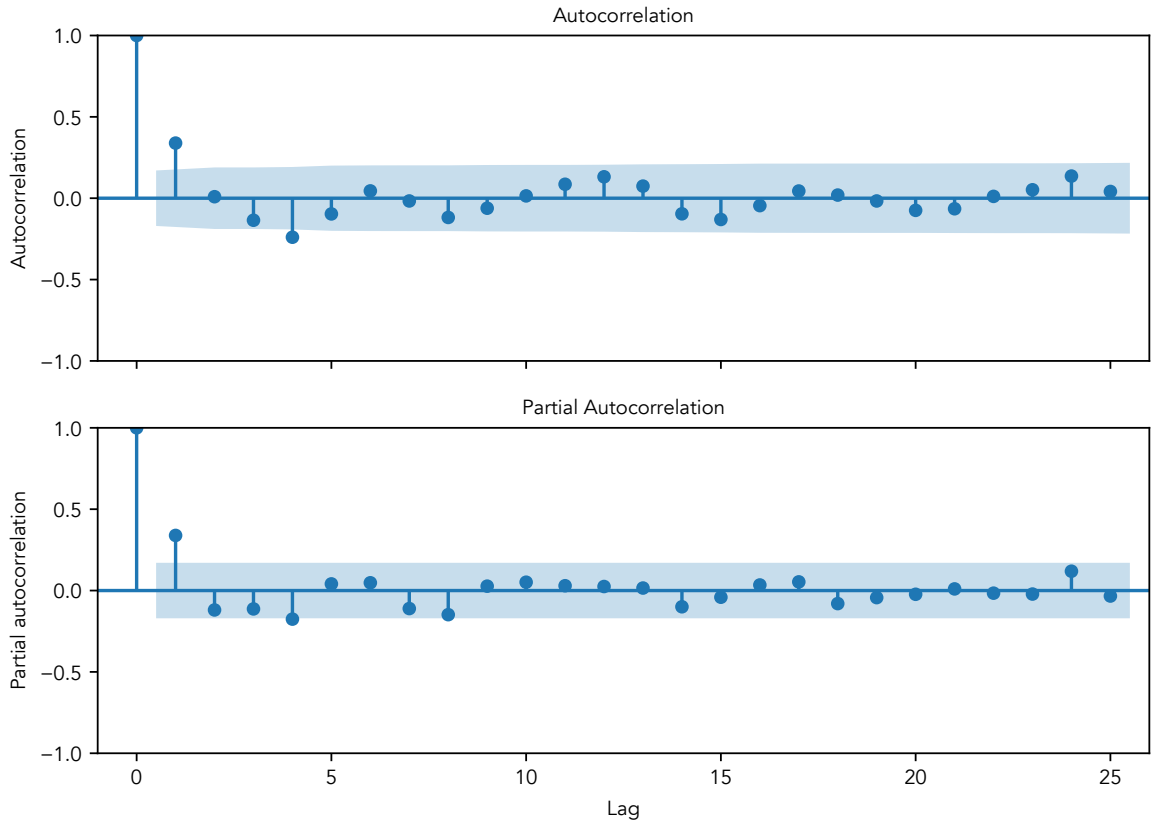

Figure 3.12: (Top) Autocorrelation and (Bottom) partial autocorrelation of residuals for the first 25 lags for the UMO regression model, with the significance level shown by the shaded region.

The Durbin-Watson statistic, used to detect first-order autocorrelation, is 1.31 for the residuals of Eq. 3.8, and 1.17 for the residuals of Eq. 3.9. As they are both less than 2, they indicate positive serial correlation within both ordinary least squares regressions. The Durbin-Watson statistic was checked against a lookup table giving 5% significance points, where upper and lower limits are given for k regressors and sample size n. We have $n = 132$ for both regressions, and $k = 4$ and $k = 2$ for the UMO and LNADW regressions respectively. As limits are given only for $n = 100$ and $n = 150$, we use the next lowest sample size of $n = 100$. For four regressors, the table gives a lower bound of 1.592 and an upper bound of 1.758, and for two regressors, it gives a lower bound of 1.634 and an upper

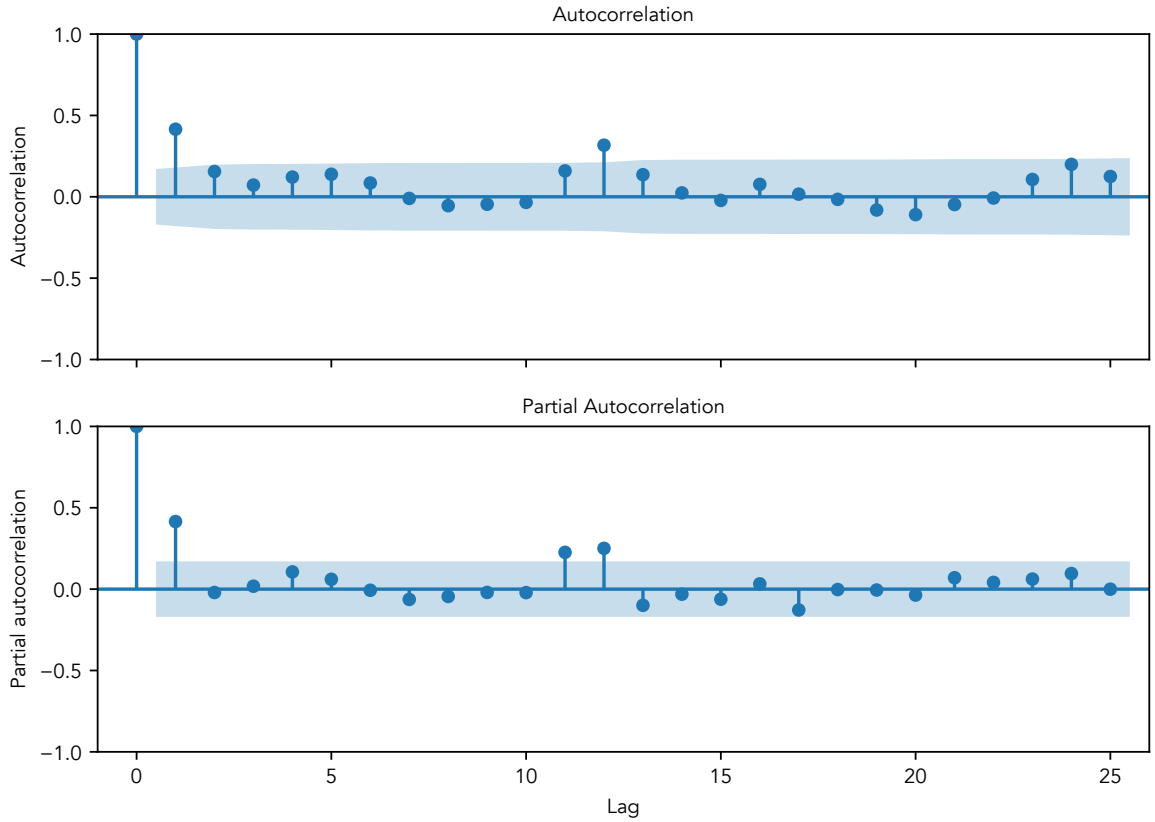

Figure 3.13: (Top) Autocorrelation and (Bottom) partial autocorrelation of residuals for the first 25 lags for the LNADW regression model, with the significance level shown by the shaded region.

bound of 1.715. As the Durbin-Watson statistic for each regression is less than the lower bound, in each case we reject the null hypothesis of non-autocorrelated errors in favour of the alternative hypothesis that there is positive first-order autocorrelation.

As described in McKinney et al. (2019), one option when autocorrelation of residuals occurs is to use a generalised least squares (GLS) model with an auto-regressive process that models the required lag. Fig. 3.12 and Fig. 3.13 showed significant partial autocorrelation for a lag of one month, so a first-order auto-regressive or AR(1) process is used. The `statsmodels` package provides a `GLSAR` function which uses a two-stage iterative fit to estimate the GLS model with AR(1) errors, and we used this to replace the ordinary least squares regression models described in the previous sections.

3.4.3 Generalised Least Squares Regression Models

3.4.3.1 Modelling UMO transport using density

Eq. 3.10 shows the depths and coefficients selected by algorithm for the regressions in Eq. 3.3 using a generalised least squares regression with an auto-regressive process modelling a lag of one month:

$$T_{umo} = 49.3\ \rho_{wb}^{780}, \tag{3.10a}$$

$$T_{umo} = 44.0\ \rho_{wb}^{780} - 101.5\ \rho_{eb}^{920}, \tag{3.10b}$$

$$T_{umo} = 43.7\ \rho_{wb}^{780} - 105.4\ \rho_{eb}^{920} + 46.7\ \rho_{wb}^{2720}, \tag{3.10c}$$

$$T_{umo} = 39.6\ \rho_{wb}^{740} - 105.4\ \rho_{eb}^{920} + 39.8\ \rho_{wb}^{1200} + 44.4\ \rho_{wb}^{3000}. \tag{3.10d}$$

The depths selected were similar to those of the ordinary least squares regression models (Eq. 3.6), differing by at most 20 dbar.

Linear regression model	**Adj. R^2**	**SE [Sv]**	**Mean error [Sv]**	**σ error [Sv]**	**RMS error [Sv]**	**AIC**	**BIC**
Thermocline	0.488	1.84	1.62	1.20	2.02	533.0	535.9
Therm. + AAIW	0.761	1.26	1.08	0.78	1.34	435.3	441.1
Therm. + AAIW + UNADW	0.775	1.23	1.03	0.78	1.29	428.9	437.5
Therm. + AAIW + UNADW + LNADW	0.787	1.19	1.00	0.76	1.25	421.9	433.4

Table 3.5: Comparison of model statistics from the generalised least squares linear regressions in Eq. 3.10. Statistics shown are the adjusted R^2 ; standard error of the regression (SE); the mean, standard deviation (σ) and root-mean-square (RMS) of the error, where the error is the difference between the actual and model-predicted values; Akaike's Information Criterion (AIC); and Bayes Information Criterion (BIC).

Table 3.5 shows that the simplest model selected by algorithm, regressing UMO transport on the western boundary density anomaly at 780 dbar, gives an adjusted R^2 value of

only 0.49 and a standard error of 1.85 Sv. This resulting time series (Fig. 3.14, top) also shows relatively large model prediction intervals (mean = 3.7 Sv), indicated by the shaded region around the orange line of the predicted UMO anomaly. However, as Fig. 3.14 shows, adding the eastern boundary density anomaly at 920 dbar markedly increases the maximum adjusted R^2 value to 0.76, reduces the standard error to 1.26 Sv and decreases the prediction interval range. Adding the ρ_{wb}^{z3} and then the ρ_{wb}^{z4} western boundary density anomalies increases the adjusted R^2 value of the model to 0.78 and 0.79 respectively. Although increasing the number of explanatory variables from two, to three, to four doesn't increase the adjusted R^2 greatly, it does reduce the standard error of the regression from 1.84 Sv for the single variable model, to 1.26, 1.23 and 1.19 for the two, three and four explanatory variable models respectively (Table 3.5). The mean prediction intervals also reduce from 3.7 Sv, to 2.5, 2.5 and 2.4 Sv. The decrease in the mean and standard deviation of the error, or residuals, is similar to that seen in the ordinary least squares models. The mean error reduces from 1.62 to 1.08, 1.06 and 1.00 Sv with each additional explanatory variable, while the standard deviation reduces from 1.20 Sv to 0.78, 0.78 and 0.76 Sv. The RMS error reduces from 2.02 Sv for the simple linear regression, to 1.34, 1.29 and 1.25 Sv with each iteration of the model. The AIC and BIC metrics also improve significantly, increasing by more than 2 with each additional explanatory variable.

Although multicollinearity tests showed that our explanatory variables were sufficiently independent, and the regression statistics all showed improvement with each additional regressor, the slight improvement in fit from adding the third and fourth regressors, when compared to the large improvement after adding a second variable, brought into question whether the improvements were indistinguishable from noise. To test this, we created two randomly generated anomalies with the same standard deviation as the third and fourth explanatory variables to replace them. The resulting model had an adjusted R^2 of 0.757 with a standard deviation of 0.004, showing that any improvement is above statistical uncertainty.

These results provided good evidence that despite only increasing adjusted R^2, by less than 3%, the addition of the two deep density anomalies to the model was justified. As using four explanatory variables represents all four water masses above 4800 dbar at 26°N, and also explains the greatest variance in the UMO transport anomaly and has the lowest

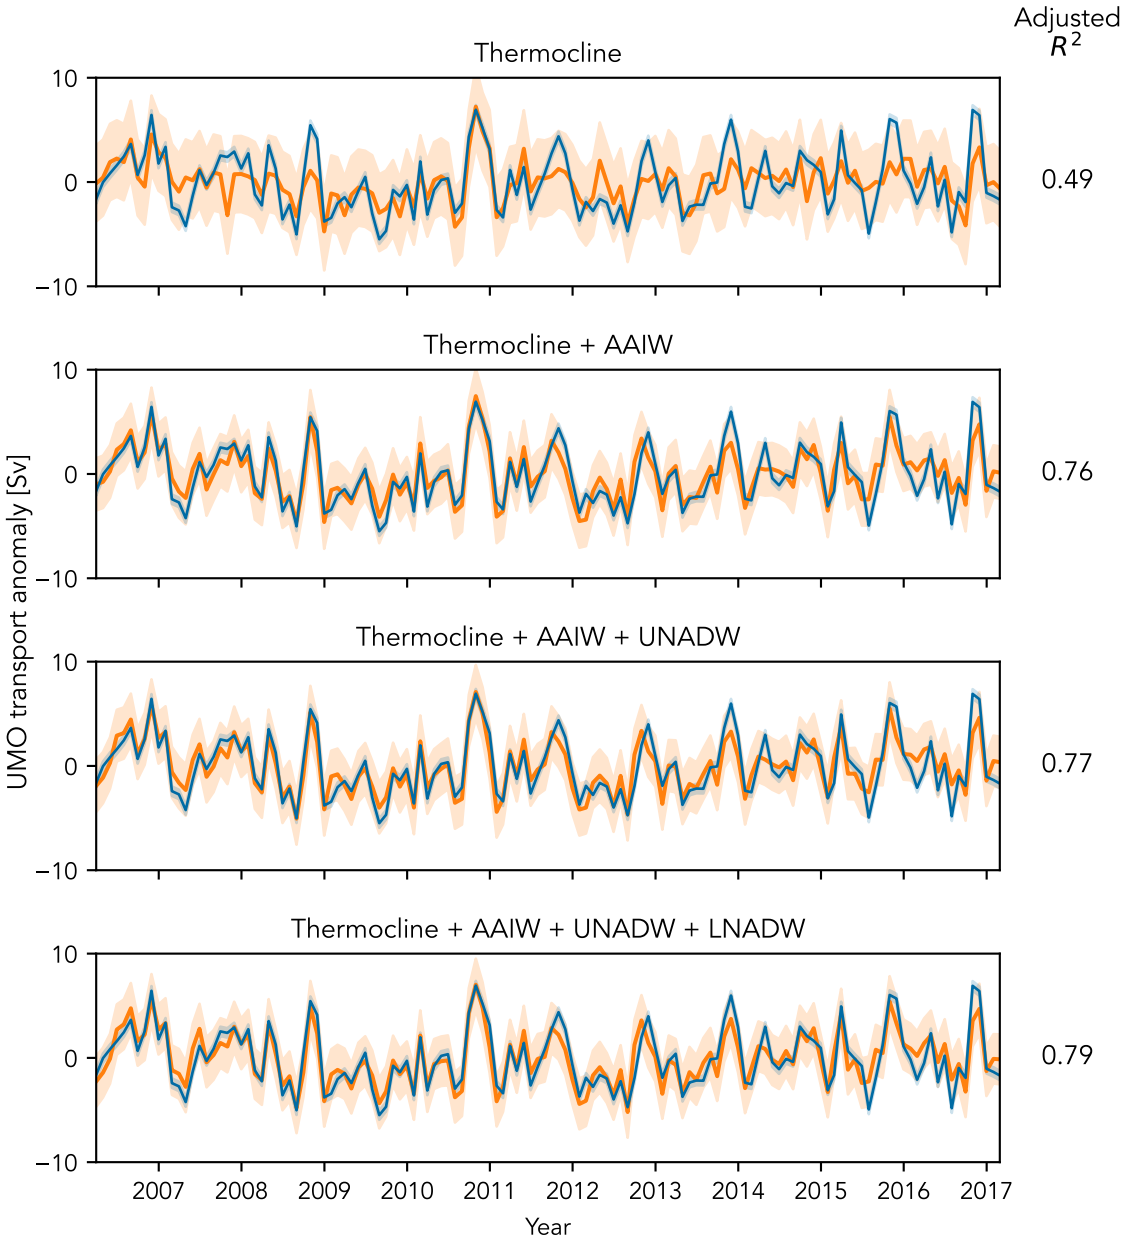

Figure 3.14: Comparison of monthly mean UMO transport anomalies predicted by the GLSAR(1) linear regression models in Eq. 3.10 (orange line), and observed by RAPID (blue line). The models have one, two, three and four explanatory variables – western and eastern boundary density anomalies – representing different water masses at 26°N, which are named above each plot. Adjusted R^2 values for each model are shown to the right of the related plot. The orange shaded area around each model prediction line shows the 95% prediction interval of the model.

standard error, Eq. 3.10d was selected to apply to the historical hydrographic data to reconstruct the UMO and AMOC transports.

3.4.3.2 Modelling LNADW transport using deep density

For the LNADW linear regression, the algorithm selected the western boundary density anomaly at 3040 dbar, close to the RAPID-defined boundary between the upper and lower North Atlantic Deep Water layers at 3000 dbar. The resulting linear regression (Eq. 3.11)

$$T_{\text{lnadw}} = -174.5\rho_{wb}^{3040} - 0.4\, T_{ekman} \tag{3.11}$$

Linear regression model	**Adj. R^2**	**SE [Sv]**	**Mean error [Sv]**	**σ error [Sv]**	**RMS error [Sv]**	**AIC**	**BIC**
LNADW	0.744	0.94	0.82	0.63	1.03	357.8	363.5

Table 3.6: Comparison of model statistics from the generalized least squares linear regression in Eq. 3.11. Statistics shown are the adjusted R^2 ; standard error of the regression (SE); the mean, standard deviation (σ) and root-mean-square (RMS) of the error, where the error is the difference between the actual and model-predicted values; Akaike's Information Criterion (AIC), and Bayes Information Criterion (BIC).

has an adjusted R^2 value of 0.74, a standard error of 0.94 Sv, and a mean prediction interval of 1.9 Sv. The differences between the modelled and observed time series have a mean of 0.82 Sv and standard deviation of 0.63 Sv, while the RMS of the differences is slightly larger than the standard error of the regression at 1.03 Sv. The coefficients show that a positive anomaly in LNADW transport is associated with both a negative density anomaly and negative Ekman transport anomaly. This means a reduction in the deep southwards LNADW flow is linked to lower density water at 3040 dbar, which we would expect with a reduction in overturning. The inverse relation between LNADW and Ekman transports reflects the statistically significant inverse correlation (r = −0.58) found by Frajka-Williams et al. (2016) between the same transports.

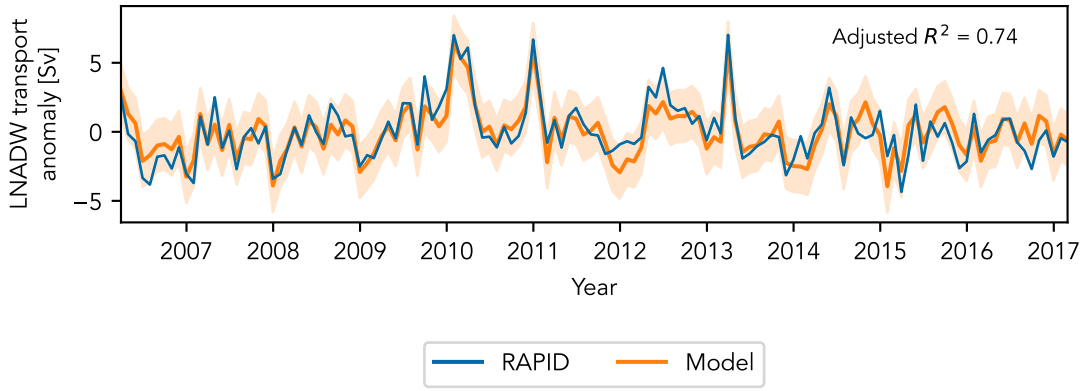

Figure 3.15: Comparison of monthly mean LNADW transport anomalies predicted by a the GLSAR(1) linear regression model in Eq. 3.11 (orange line), and observed by RAPID (blue line). The model regresses LNADW transport anomaly on a deep western boundary density anomaly at 3040 dbar and the Ekman transport anomaly. The model adjusted R^2 value is shown on the figure. The orange shaded area around the model-predicted LNADW shows the 95% prediction interval.

3.4.4 Testing selected models

The two GLSAR regression models were tested against RAPID data that were not used to train them, which were made available following the most recent expedition and cover the period from 5 April 2017 to 29 August 2018 (Smeed et al., 2019). These data have a 12-hour temporal resolution and were smoothed with a 10-day Butterworth filter. The model-predicted UMO transport anomaly shown in Fig. 3.16 shows that it reproduces the trends and variability, although not always the magnitude, of the observed values well, having a strong positive Pearson's correlation coefficient between the two time series (r = 0.75). The mean error, or difference between the observed and modelled values, was 1.75 Sv, and had a standard deviation of 1.26 Sv (Table 3.7). The RMS error was 2.16 Sv, which compares well to the model's mean prediction interval of 2.4 Sv (Section 3.4.3.1), and is encouraging given the changed timescale from monthly mean to 10-day filtered values.

Predicted LNADW transport anomaly from the same RAPID data and the LNADW-Ekman regression model is shown in Fig. 3.17, and again the trends and variability of the observed LNADW transport anomaly are captured by the model predictions, with strong positive correlation (r = 0.73) between the two time series. The model tends to underestimate the magnitude of the LNADW transport anomaly, however.

The empirical model was trained on monthly mean density anomalies, but to be useful,

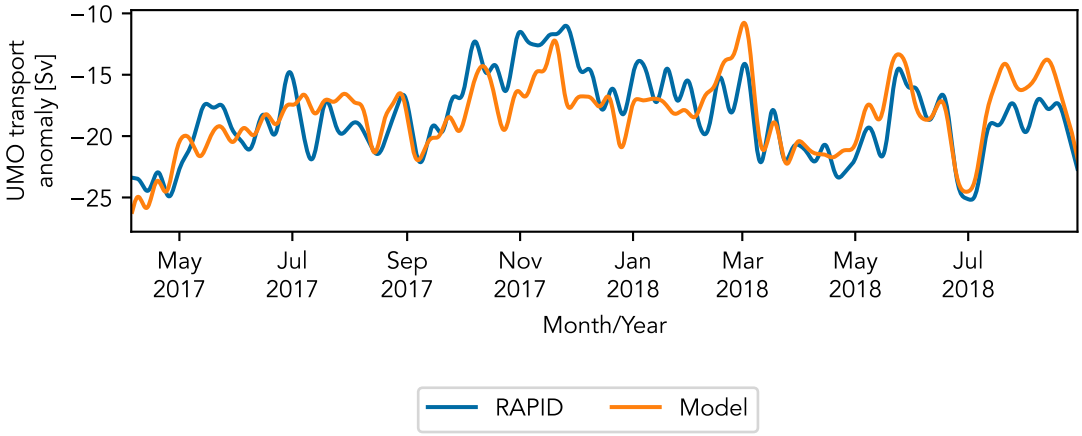

Figure 3.16: UMO transport anomaly estimated from the most recent RAPID observations (blue) compared against that predicted by the four-layer linear regression model (orange) using 12-hourly, 10-day filtered density anomalies from the same RAPID data.

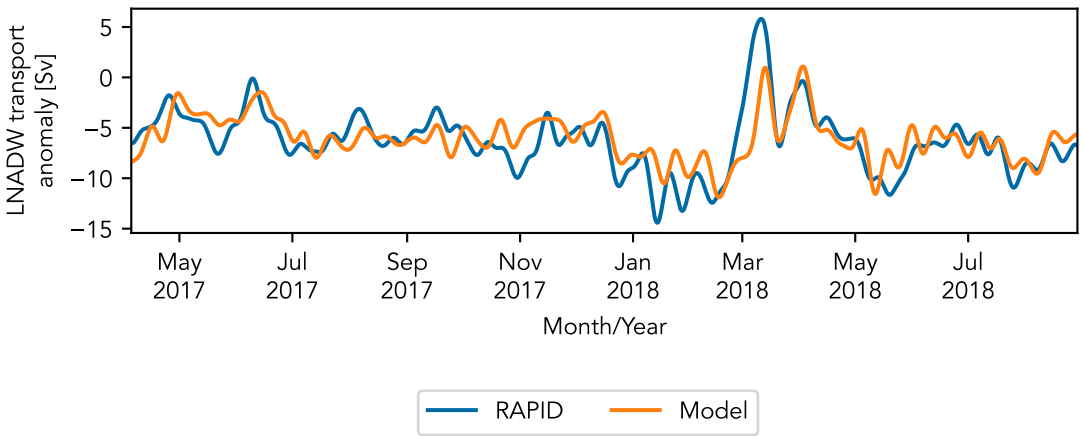

Figure 3.17: LNADW transport anomaly estimated from the most recent RAPID observations (blue) compared against the LNADW transport anomaly predicted by the linear regression model that combines the western boundary density anomaly at 3040 dbar derived from the same RAPID data, and the Ekman transport (orange). The data were 12-hourly and 10-day filtered.

Regression model	**Mean error [Sv]**	**σ error [Sv]**	**RMS error [Sv]**
UMO	1.75	1.26	2.16
UMO inc. EB climatology†	1.86	1.17	2.20
LNADW	1.56	1.47	2.14

Table 3.7: Mean, standard deviation (σ) and root-mean-square (RMS) of the error between RAPID-observed and model-predicted transport anomalies in Fig. 3.16 and Fig. 3.17

would need to be applied to data from much shorter timescales, for example single CTD profiles. To evaluate how well the model would perform, 'snapshot' data were simulated by randomly selecting 20 single points from the 7691 available 12-hourly, 10-day filtered RAPID data values. These were applied to the model, and the predicted UMO compared to the observed monthly mean UMO for the same time. Bootstrapping showed that around 65% of the observed UMO values were within the prediction interval of the corresponding model UMO. The distribution of the model error, or the difference between the observed monthly mean and predicted 'snapshot' UMO transport anomalies in Fig. 3.18 shows that the standard deviation of this error was 2.7 Sv, and the root mean square (RMS) error was also 2.7 Sv, which again compares well with the mean prediction interval of 2.4 Sv. Repeating this simulation but replacing the eastern boundary density anomaly with a climatology, described in Section 4.2.1, reduced the RMS error to 2.5 Sv.

3.4.5 Coherence analysis

To evaluate the co-variability of each explanatory variable density anomaly and the observed UMO transport anomaly, we determined the coherence between them using a multi-taper spectrum following Percival and Walden (1998). As described in Section 2.2.3, this technique reduces spectral leakage while minimising the data loss associated with other tapers. The number of tapers used was $K = 2NW - 1$, where $NW = 4$, with the 95% confidence level given by $1 - 0.05^{\frac{1}{(K-1)}}$. Higher values of NW result in narrower tapers, reducing spectral resolution. As we were interested in lower frequency features, a value of $NW = 4$, smoothed the spectra by removing higher frequency noise, but still allowed features with seasonal and longer periods to be seen.

The coherence and phase relationships in Fig. 3.19 show that the western boundary density anomaly at 740 dbar, which represents thermocline transport, exhibits significant

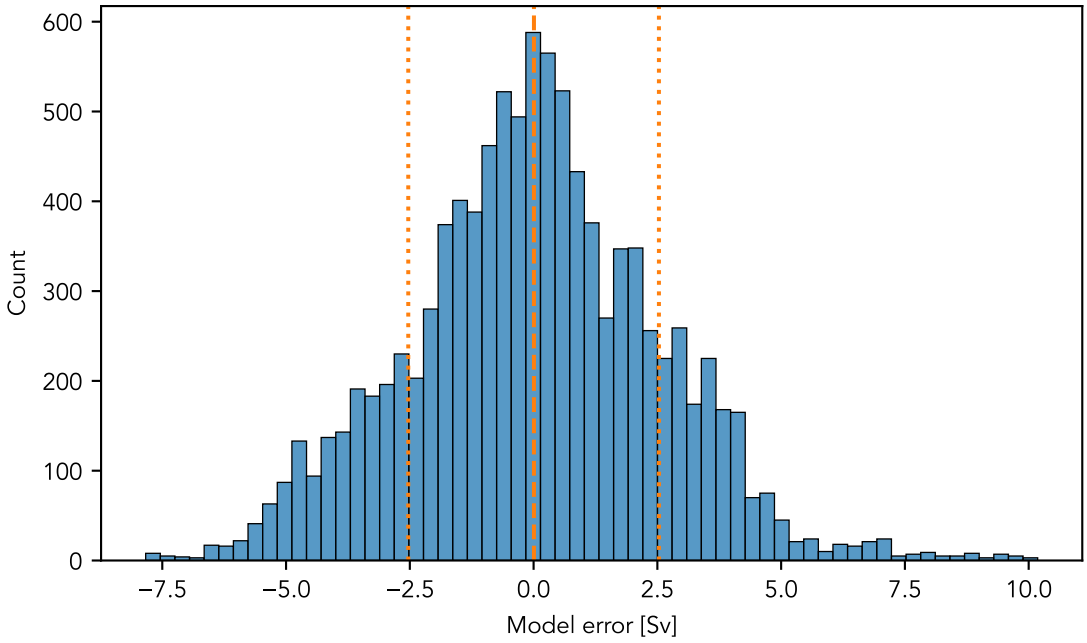

Figure 3.18: Distribution of model error (difference between predicted and actual UMO transport anomaly) for 10,000 randomly selected data points from the RAPID 12-hourly data for 2017-2018. The vertical orange lines are the mean (dashed) and one standard deviation (dotted).

coherence with the UMO transport anomaly at almost all periods, with the exception of 77-90 and 104-113 days. The highest significance occurs for periods of around 65 days, 140 days and from 2 years onwards, and the coherence is in phase at all periods. The eastern boundary density anomaly at 920 dbar, representing AAIW transport, shows high significant coherence for periods of around 110 days, between 150 and 200 days, and between 265 days and 1.5 years, indicating strong coherence at semi-annual and annual frequencies. The coherence for this variable is out of phase for most periods of significant coherence which is consistent with its negative coefficient. The western boundary density anomaly at 1200 dbar, representing UNADW transport, shows strong coherence for periods around 70 days, and slightly significant coherence for periods of 90 and just over 180 days. The significant coherence is approximately in phase, with the observed UMO transport anomaly lagging the 1200 dbar density anomaly at 70 and 90 days . Finally, the western boundary density anomaly at 3000 dbar, representing LNADW transport, shows significant coherence only for periods of 136 days and over 1 year. For this latter period, the co-variability is also approximately in phase, with the 3000 dbar density anomaly leading the UMO transport anomaly slightly.

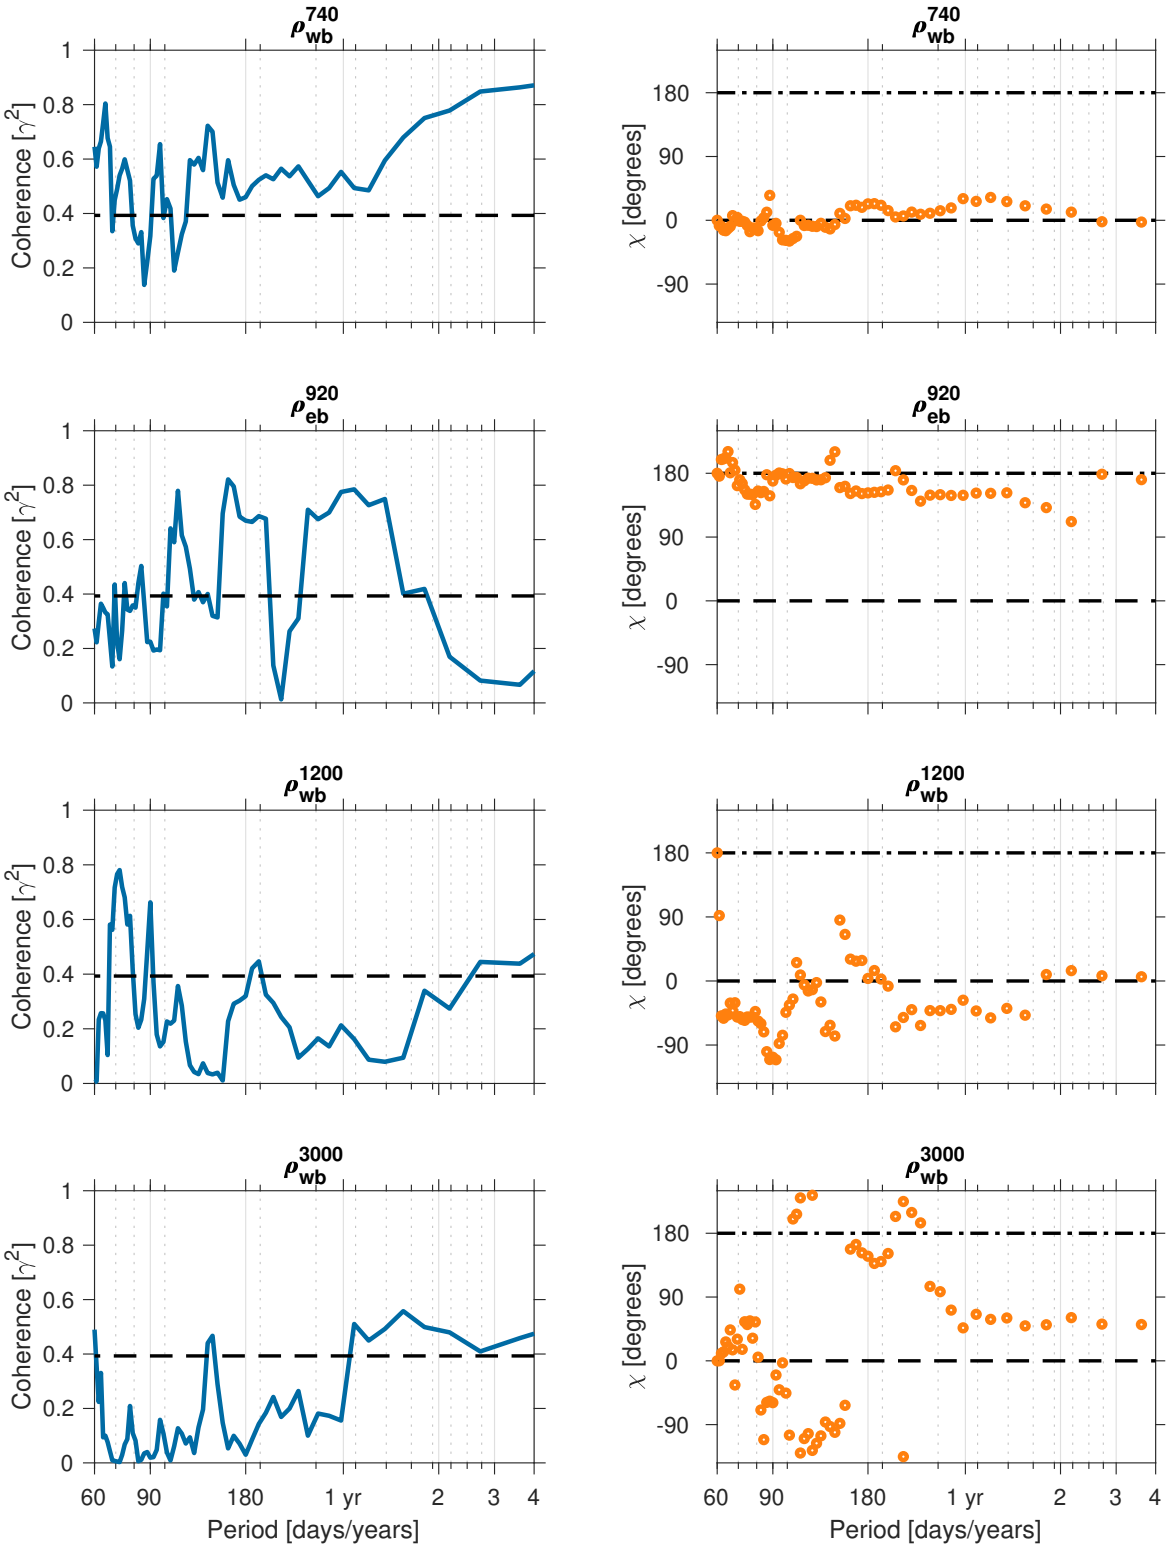

Figure 3.19: Multi-taper spectrum coherence (left) and phase relationship (right) between the UMO transport anomaly observed by RAPID, and the density anomaly that is each explanatory variable in the linear regression model. Significance of coherence (95% confidence) is indicated by the black dashed horizontal line (left). The horizontal black lines in the right-hand panels show where the time series are exactly in phase (dash) and exactly 180° out of phase (dot-dash).

3.4.6 EOF analysis

Fig. 3.20 shows how the standard deviation of the regression model RAPID training data – monthly mean density anomalies for the western and eastern boundaries – varies with depth, and shows the high level of noise in the upper few hundred dbar of the profiles, particularly on the western boundary.

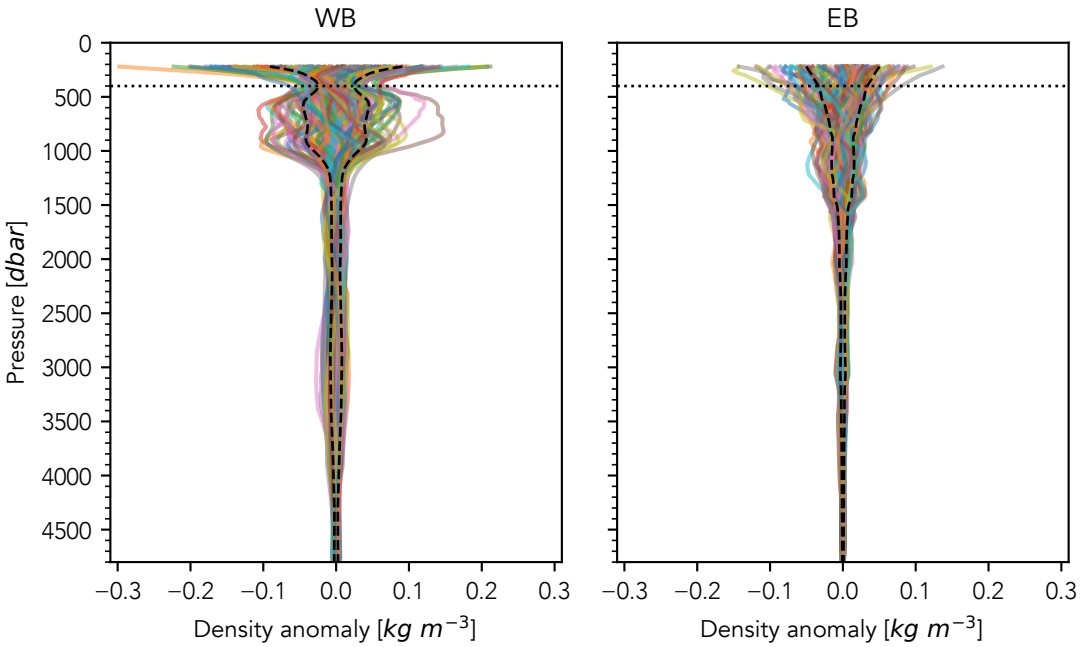

Figure 3.20: Variability of RAPID density anomalies over the full depth of the western (left) and eastern (right) profiles used to train the regression models. The standard deviation is shown by black dashed lines. The horizontal dotted line shows the upper depth limit of EOF analysis.

The high near-surface variability meant that the shallowest density anomalies were excluded from the empirical orthogonal functions found for the regression model training data at the western and eastern boundaries, to keep the focus on patterns of lower-frequency variability. At the western boundary, the narrowing of the standard deviation at 400 dbar made that an obvious truncation depth. On the eastern boundary it was less obvious, so the same depth was chosen to limit a strong surface signal in the leading mode. The vertical structure associated with the first three EOF modes for each boundary are shown in Fig. 3.21, together with the variance of each mode. The selected depths of the model regressors are shown by horizontal black dashed lines. The first three modes account for 92.7% of the total variance on the western boundary, and for 90.4% on the eastern

boundary. On the western boundary, the first EOF mode shows a strong signal between 540-760 dbar, the second between 460-580 dbar, and again between 840-1080 dbar. The third EOF mode has a much broader signal between 1100-4800 dbar with a maximum negative signal at 3000 dbar. The first and third modes align closely with the depths selected for the thermocline and LNADW regressors at 740 and 3000 dbar, however no EOF mode coincides with the UNADW regressor at 1200 dbar. On the eastern boundary, the EOF modes are less clear, all three show a strong signal at 300 dbar. Below this, the second mode shows a negative signal between 1340-1440 dbar, while the third mode shows two maxima: a negative signal below 1000 dbar and a broader positive signal between 1400-2100 dbar.

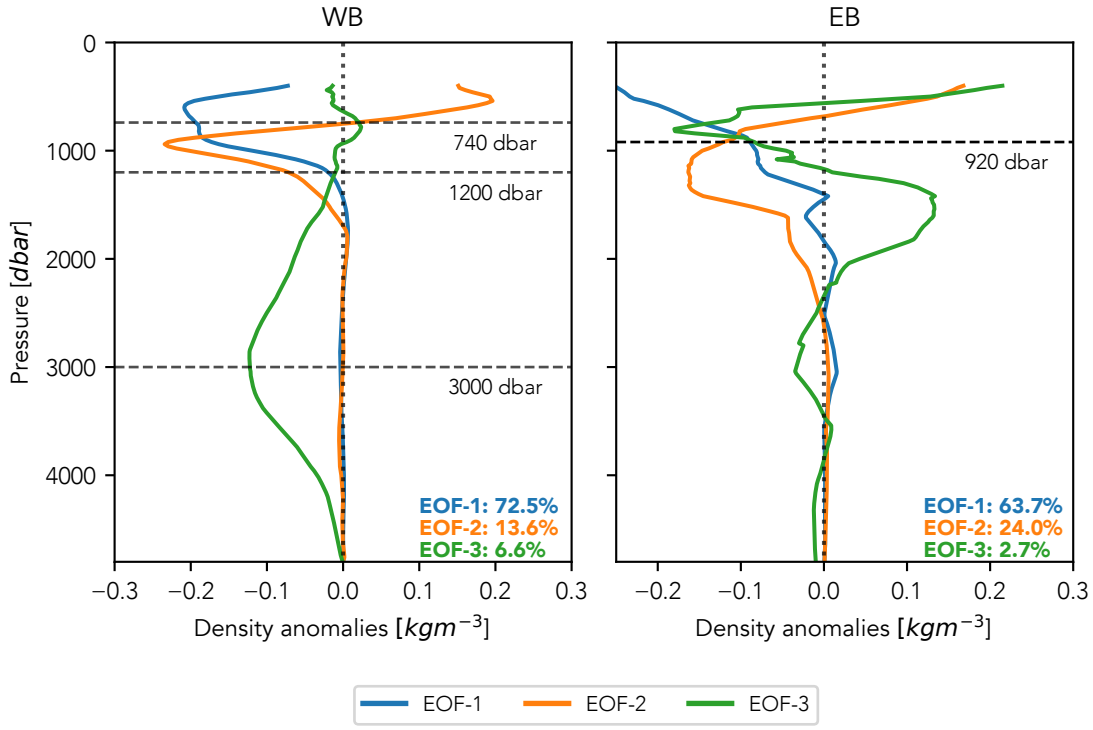

Figure 3.21: Profiles of RAPID density anomalies between 400 and 4800 dbar associated with the first three leading EOF modes of density variability for the western (left) and eastern (right) boundary profiles, and the variance explained by each mode. Horizontal dashed lines show the depths of density anomalies used in the UMO regression model.

In addition, we calculated the correlation between the principal component time series related to each of the modes shown in Fig. 3.21, and the density anomaly time series at each pressure interval (Fig. 3.22). On the western boundary, the depths of the maximum negative correlation associated with the first, second, and third leading modes are very close

to the depths of the western boundary density anomalies representing the thermocline, UNADW and LNADW respectively. On the eastern boundary, the depth of the AAIW density anomaly is between the maximum negative correlations for both the first and second leading modes.

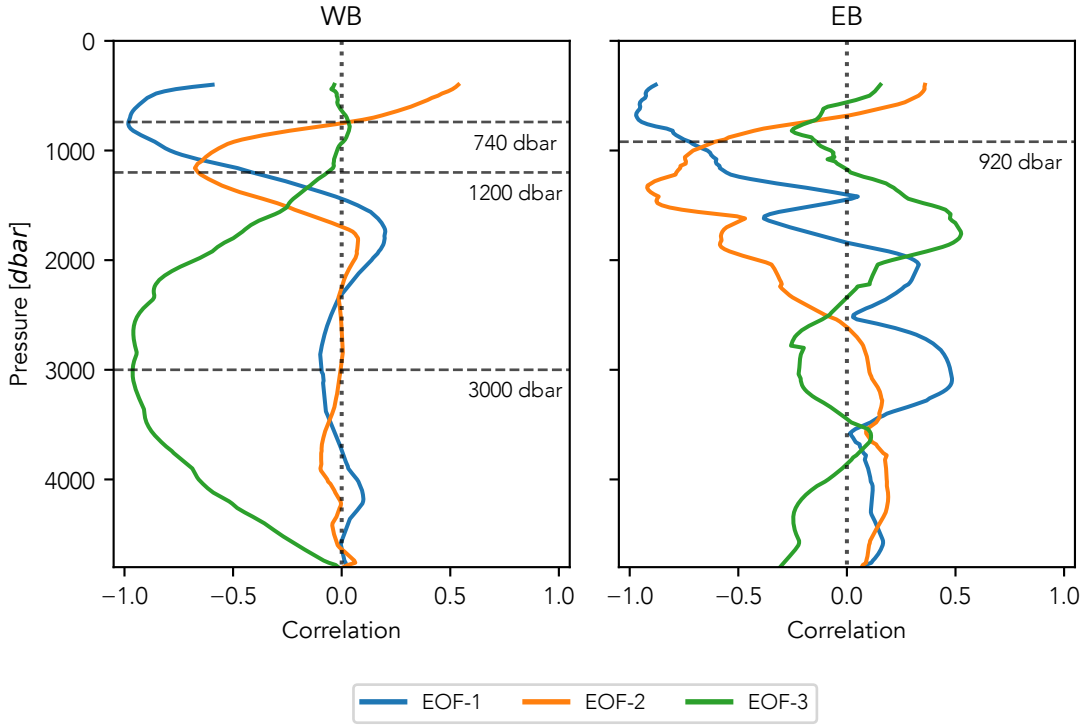

Figure 3.22: Correlation between the principal components associated with each of the three leading EOF modes, and the RAPID density anomalies at each depth for the western (left) and eastern (right) boundary profiles.

3.5 Conclusion

In this chapter we have shown how a simple linear regression model is insufficient to explain more than 53% of the variance of the upper mid-ocean geostrophic meridional transport at 26°N, and that multiple explanatory variables are needed for a regression model to explain over 70%. Our initial regressions using ordinary least squares models showed promising results, and suggested that density should be preferred to temperature and salinity as the boundary anomaly properties. However, the ordinary least squares models showed first-order autocorrelation of residuals, common in regression models trained on time series data, so the regressions were changed to generalised least squares

with a first-order auto-regressive process to model the autocorrelation.

We also showed that although the bulk of improvement in our regression models was obtained by adding a second explanatory variable, in this case an eastern boundary density anomaly at intermediate depth, adding a third and then a fourth variable improved the model fit statistics even further, including explained variance and standard error. Variance inflation factors showed that the regressors were sufficiently independent of each other, and we also showed that improvement in the model was significantly greater than would be explained by adding two random regressors. Coherence analysis of the relationship between UMO transport and the individual explanatory variables suggested that including deep density anomalies to represent changes in the upper and lower North Atlantic Deep Water transport is necessary to capture the full range of variability shared with UMO transport. This is further supported by EOF analysis that showed maximum negative correlation between the principal component time series and density anomalies at the depths selected by algorithm, suggesting that the second and third leading modes at the western boundary are represented by the two deepest density anomalies.

An additional model showed that a single deep density anomaly combined with the Ekman transport anomaly was successful in reconstructing LNADW transport, the deepest limb of the AMOC in the subtropical North Atlantic. It was not described here, but similar models replacing the single deep density anomaly with conservative temperature and then absolute salinity were also created. Compared to the adjusted R^2 of 0.75 for a 3000 dbar density anomaly, the algorithms selected a temperature anomaly at the same pressure, and the resulting model explained 71% of the variance, while the salinity anomaly was selected at 3240 dbar, with that model explaining only 53% of LNADW transport anomaly variance. These results suggest that temperature dominates deep density changes.

The selected linear regressions had RMS errors of 1.25 Sv for the UMO model, and 1.03 Sv for the LNADW model, however these show only the difference between the modelled and observed values for the time series used to create, or train, the regression models. RMS error is very similar to the standard error of the regression, in that both quantify how well the model fits the training data. For example, in a simple linear regression, this would be how well the regression line fits the data points. The model's prediction intervals give a measure of the expected uncertainty when it is applied to new explanatory variables. The

prediction intervals vary depending on the explanatory variables, but the mean prediction interval for the UMO model applied to the training data was 2.4 Sv, and for the LNADW model was 1.9 Sv. However, prediction intervals do not necessarily quantify the expected uncertainty when applied to data with time scales different to the training data's. Testing of both models against more recently available RAPID observations showed good agreement between model-predicted and observed transport anomalies, reproducing variability in particular well, and having an RMS error of 2.16 Sv, while the mean internal transport remained at 2.4 Sv. We also showed that when used to predict UMO transport anomalies from 'snapshot' data such as a single CTD profile, that the model uncertainty remains realistic, with a RMS error of 2.7 Sv. These 'snapshot' profiles were however created from 10-day filtered RAPID data, so the model error may again be different for predictions from individual CTD profiles. This will be addressed in the next chapter, which describes how these successful models were used to reconstruct UMO, LNADW and AMOC transport anomalies from historical hydrographic data.

Chapter 4

Reconstructing the Atlantic meridional overturning circulation from historical hydrographic data

4.1 Introduction

This chapter describes how the empirical regression models created and tested in the previous chapter were applied to historical CTD data from the western boundary of the subtropical North Atlantic, to reconstruct the UMO and LNADW transport anomalies at 26°N. We then reconstructed the AMOC between 1981 and 2016 by summing the UMO transport with observed Florida Current and Ekman transports.

4.2 Data

4.2.1 Replacing the eastern boundary with a climatology

In the available historical hydrographic datasets there are many more profiles close to the western boundary than to the eastern. For the pre-RAPID period of 1980-2005, there were less than a third of the number of profiles (including transatlantic sections) at the eastern boundary compared to the western (Longworth et al., 2011). Even if there were more eastern boundary data, each may not have an equivalent western boundary data point sufficiently close in time. The UMO regression model as described in Chapter 3 requires western and eastern hydrographic profiles that were sampled at approximately the same time, so, to allow us to make full use of the western boundary hydrographic data, a change to the regression model was made.

Since density at the eastern boundary shows strong seasonal variability, especially at intermediate depths (Chidichimo et al., 2010; Pérez-Hernández et al., 2015), we replaced the eastern boundary density anomaly in the UMO regression model with a monthly climatology created from the merged RAPID eastern boundary profile data. Fig. 4.1 shows that at 920 dbar (the selected regressor depth), the monthly climatology compares well to the observations, reproducing the strong negative anomalies well. The differences between the density anomaly and climatology have a mean of 0.008 kg/m^3 and standard deviation of 0.006 kg/m^3. The RMS difference is 0.010 kg/m^3. However strong positive anomalies are not captured as well by the climatology, most noticeably during Jan-Feb and to a lesser extent Apr-Jun for the years 2007-2008 and 2012-2015.

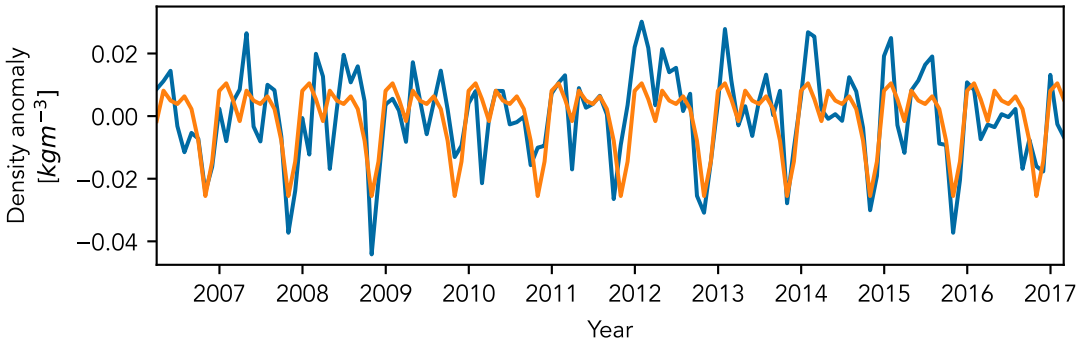

Figure 4.1: Time series of the eastern boundary density anomaly at 920 dbar (blue) and its monthly climatology (orange).

Fig. 4.2 reproduces the UMO transport anomaly predicted from RAPID observations between 5 April 2017 and 29 August 2018 made in Section 3.4.4, but also shows the effect of replacing the eastern boundary density anomaly at 920 dbar in the regression model with a monthly climatology. The climatology-based model result reproduces most of the variability of the original model, but doesn't entirely reproduce the magnitude of some of the peaks and troughs in the UMO transport.This may not be entirely disadvantageous, as it can avoid the over-estimation seen at some points in the standard model prediction, for example during Feb–Mar 2018. Table 4.1 shows that the Pearson's correlation coefficient between the climatology-model and the RAPID observations is only slightly reduced compared to the standard model (r = 0.73 compared to r = 0.75). The RMS error also increases only slightly, from 2.16 to 2.20 Sv, so using the climatology appears to be an acceptable alteration to the model selected in Chapter 3.

920 dbar EB explanatory variable	r	Mean error [Sv]	σ [Sv]	RMS error [Sv]
Density anomaly	0.75	1.75	1.26	2.16
Climatology	0.73	1.86	1.17	2.20

Table 4.1: Pearson's correlation coefficient (r), standard deviation (σ) and RMS error between RAPID-observed and model-predicted UMO time series in Fig. 4.2, where the model uses the 920 dbar density anomaly, and where this is replaced by the climatology at the same depth.

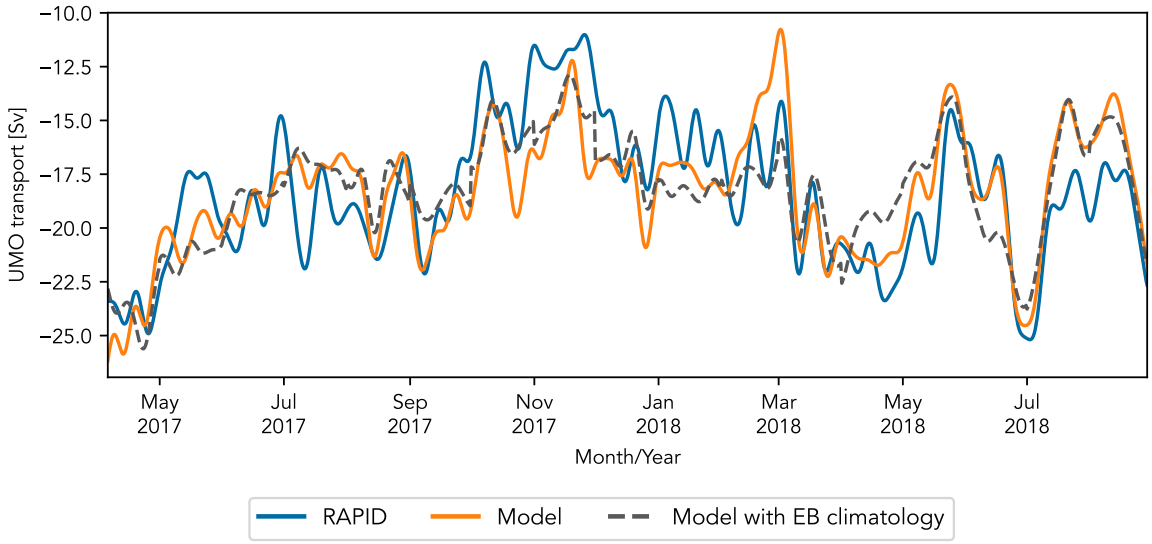

Figure 4.2: UMO transport anomaly estimated from the most recent RAPID observations (blue) compared with that predicted by the four-layer linear regression model using density anomalies from the same RAPID data (orange), and the same model with the eastern boundary density anomaly replaced by an eastern boundary climatology (grey dashed). The RAPID data were 12-hourly and 10-day filtered.

4.2.2 Hydrographic data from transatlantic sections

The regression models from Chapter 3 and Section 4.2.1 were first used to reconstruct transport anomalies from the approximately decadal transatlantic ship sections at 24.5°N and 26.5°N described in Section 2.1.2. Data for the sections from Jul-Aug 1992, Jan-Feb 1998, Apr-May 2004, Jan-Feb 2010, and Dec 2015-Jan 2016 were available from CCHDO. The section data from Aug-Sep 1981 were provided by Gerard McCarthy, but were missing time data, so these were kindly provided following a request to Woods Hole Oceanographic Institution. AMOC transports and related component and layer transports for these sections have been calculated and compared by Bryden et al. (2005) and Atkinson et al. (2012) among others, and values from these studies where available were included for

comparison with our estimates.

To be applied to the regression models, the ship-board CTD data were processed to have the same structure as the model training data: monthly mean in situ density anomalies over a 20 dbar pressure grid. Western and eastern profiles were then created for each section from the section CTDs. For each CTD, we determined the meridional distance between it and the boundary at the depth of the shallowest western boundary density anomaly, 740 dbar. The bathymetry data used to identify the boundary were the GEBCO 2014 Grid, version 20150318 (freely available at https://www.gebco.net). This distance was then compared to the equivalent distance between the boundary and WB2, the main RAPID western mooring, which was 12.95 km. The CTD with the closest matching distance was selected (Fig. 4.3). The CTD selected from each section had measurements extending to 3000 dbar, so no merging with other, deeper CTD data was required. A similar process was used for the eastern boundary, but with the distance to the boundary at 940 dbar determined and compared to that of the RAPID eastern mooring EBH3 (20.9 km). EBH3 was chosen as RAPID use measurements from this mooring between the surface and around 1400 m to create the eastern merged profile (Fig. 4.4).

4.2.3 Other hydrographic data

Historical hydrographic data were obtained from the World Ocean Database 2018 (WOD2018) (Boyer et al., 2018), and from the raw datasets processed by Longworth et al. (2011). If there were duplicate data, those provided by Longworth et al. (2011) were removed. Data were selected initially based on a region defined between 24°N and 27°N, and between 75°W and 77.5°W. We justified the meridional range as the AMOC shows meridional coherence of buoyancy anomalies within 100 km. CTD casts with WOD temperature or salinity quality flags were excluded, with the only exception data from the Malcolm Baldrige cruise of July 1986, where the salinities from all casts were flagged but the reason was not clear. CTD profiles were then grouped by date, with a group defined as separated by 3 days or more. This prevented treating very temporally and spatially similar CTD profiles, for example those taken on consecutive days during a hydrographic cruise, as separate data points. From each group, we again selected CTDs based on their having the most similar distance to the continental slope as the RAPID WB2 mooring, as

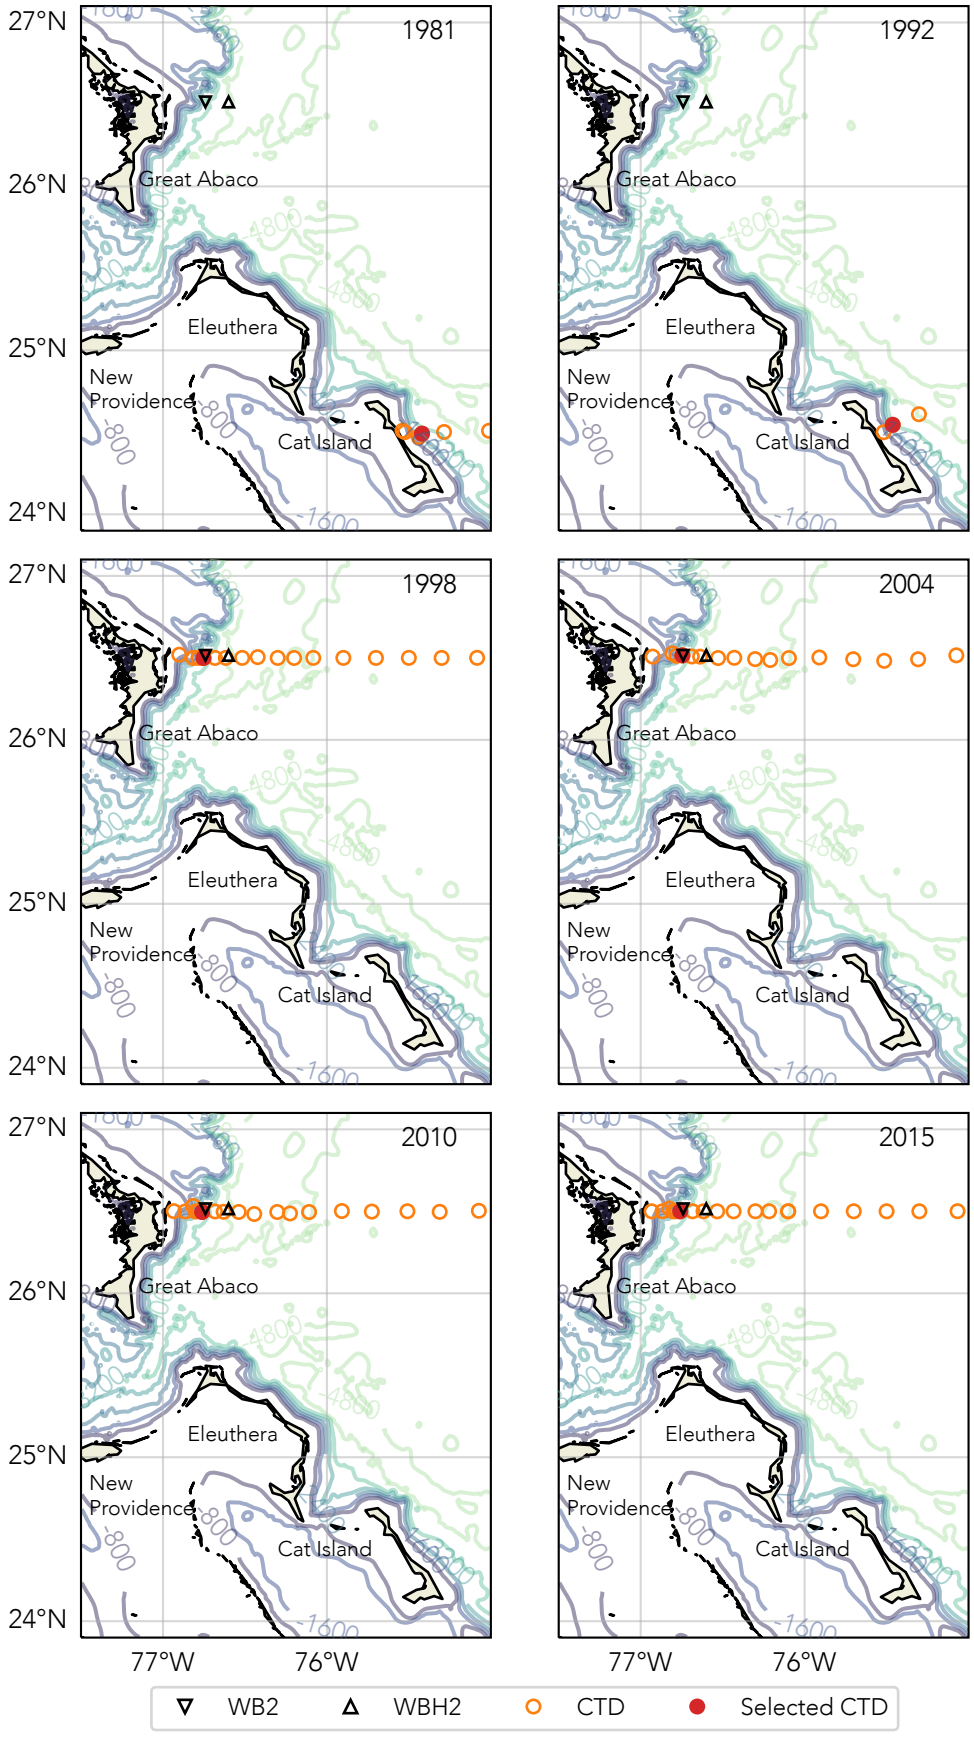

Figure 4.3: Location of selected CTD profiles (red circles) chosen from all western boundary region transatlantic section profiles (orange circles) between 1981 and 2015. The selection criterion was similarity of distance from the boundary at 720 dbar as the WB2 RAPID array mooring (black inverted triangle).

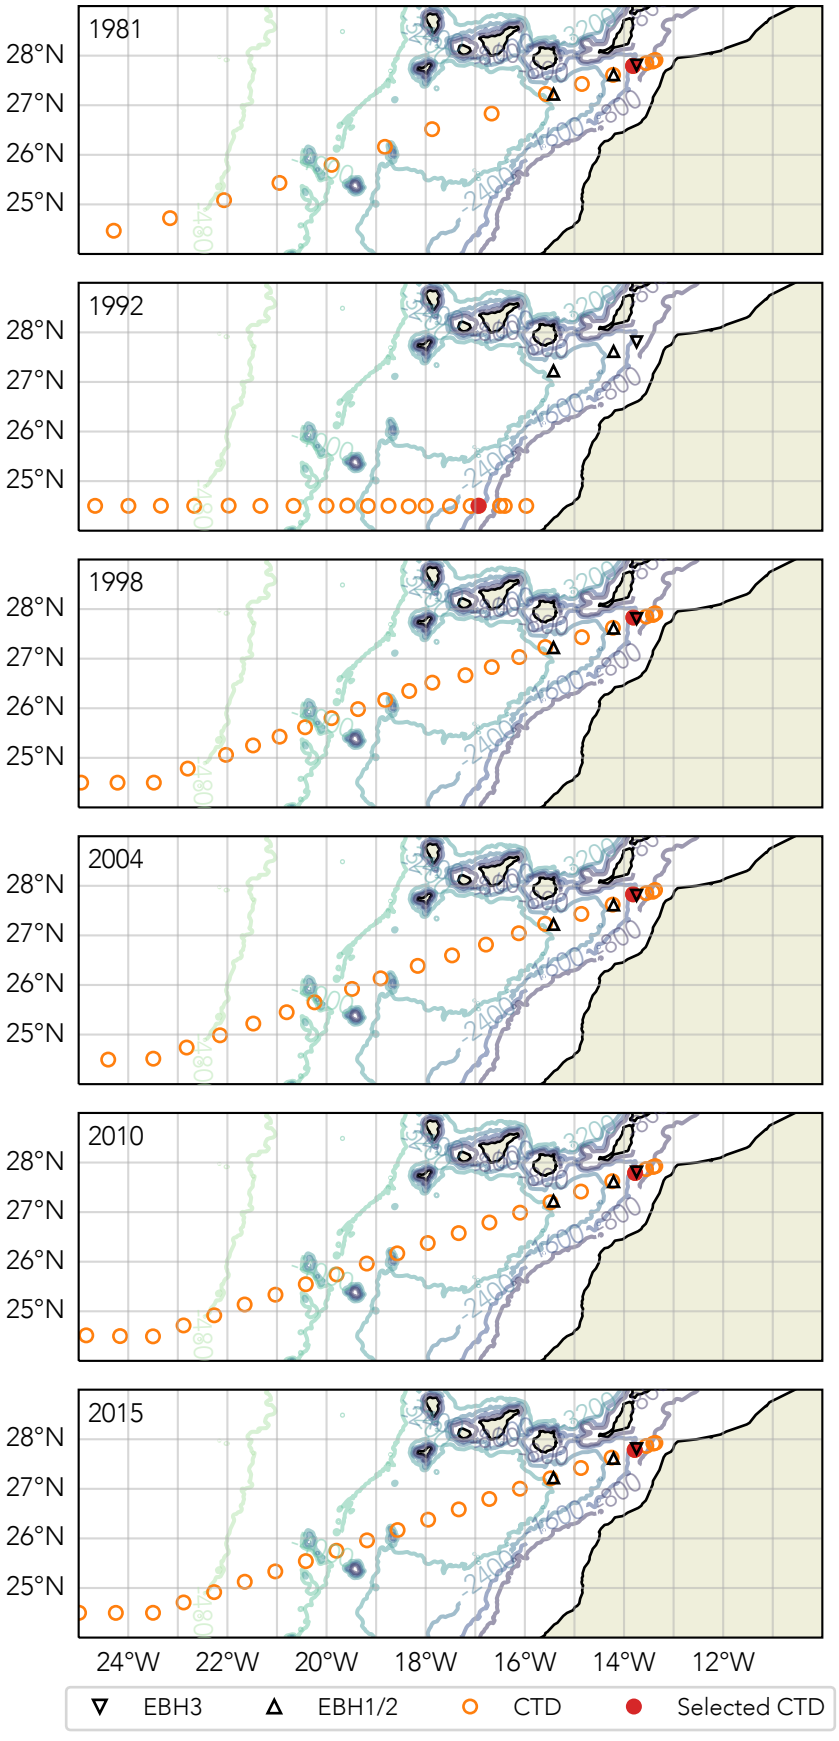

Figure 4.4: Location of selected CTD profiles (red circles) chosen from all eastern boundary region transatlantic section profiles (orange circles) between 1981 and 2015. The selection criterion was similarity of distance from the boundary at 940 dbar as the EBH3 RAPID array mooring (black inverted triangle).

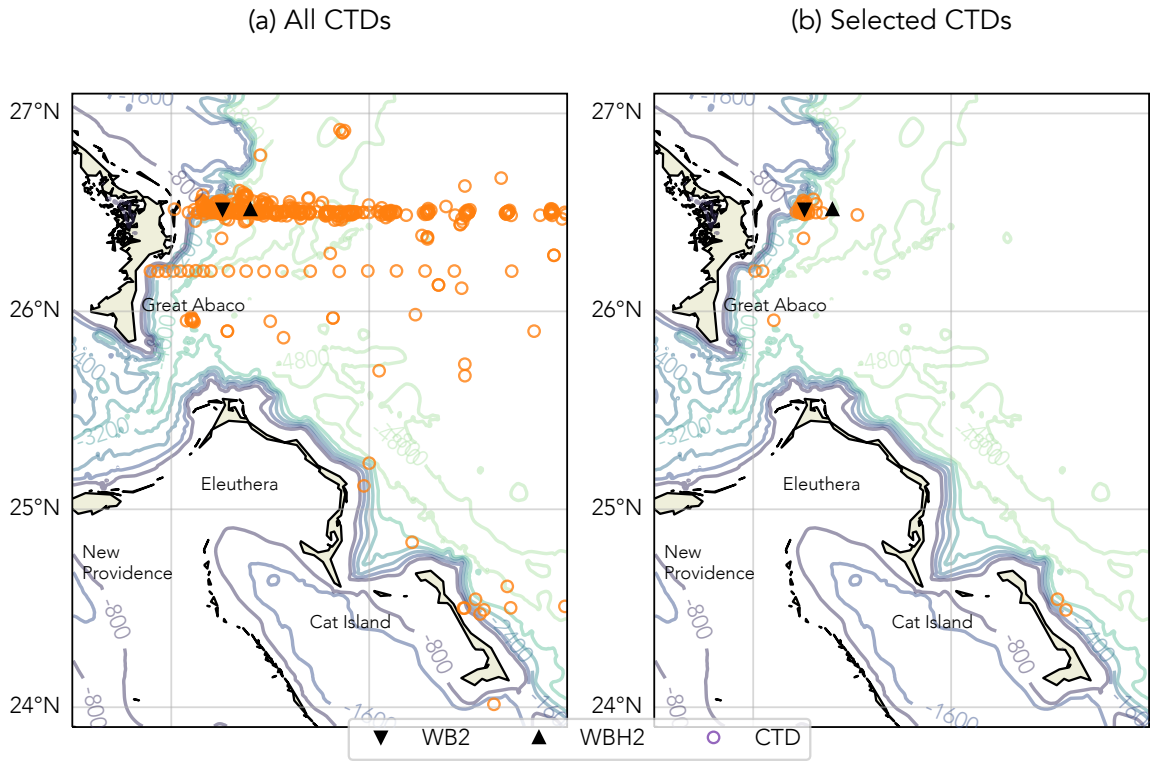

Figure 4.5: (a) All CTD profiles from the World Ocean Database 2018 in the region shown, compared with (b) those selected as being the most similar distance from the western boundary as the WB2 and WBH2 RAPID mooring

variability at the western boundary increases with distance from it (Kanzow et al., 2009). The RAPID western boundary profile that the models were trained on uses the WB2 mooring for measurements down to around 3850 dbar, so all model regressor depths are shallower than this. The distances between the western bathymetry and the WB2 mooring, and between the western bathymetry and all CTDs were determined at each of the western boundary regressor depths: 740, 1200 and 3000 dbar. Then, for the CTDs within each date group, the profile most similar in distance from the boundary to the WB2 mooring at 740 dbar was selected. For example, for the density anomaly ρ_{wb}^{740}, the WB2 mooring was 12.95 km from the boundary at 740 dbar. If there were 5 CTD profiles within the group, with distances at 740 dbar of 9.8 km, 12.4 km, 15.0 km, 18.9 km, and 27.5 km, then the profile that was 12.4 km from the boundary was selected. If this profile did not have data extending as deep as 1200 or 3000 dbar, then one or two additional profiles were selected using the same method, but by comparing the distance at each of these depths rather than at 740 dbar. Thus between one and three CTD profiles were selected from each group to use in the model. If more than one profile were selected, they were merged, with data from

the second and/or third profiles providing the deeper measurements. The two regional plans in Fig. 4.5, which compare all the available CTD profiles and those selected for use in the models, show that the majority were located close to the RAPID western boundary mooring array at 26.5°N, with a few around 26°N and 24.5°N. Table 4.2 shows the date, position and source of each of the 62 selected CTD profiles shown in Fig. 4.5.

4.2.4 Florida Current and Ekman transports

The Florida Current data used were the Western Boundary Time Series (WBTS) daily mean transport estimates from submarine cable voltage. The Florida Current cable and section data were freely available from the Atlantic Oceanographic and Meteorological Laboratory web page (https://www.aoml.noaa.gov/phod/floridacurrent/) and are funded by the DOC-NOAA Climate Program Office - Ocean Observing and Monitoring Division. These data have a gap between 22 October 1998 and 19 June 2000. Although there were no hydrographic profiles selected during this gap, it was filled with the time series mean to allow filtering. The Ekman data used were the same ERA-Interim reanalysis-derived transports as used by RAPID prior to June 2019. The ECMWF ERA-Interim reanalysis data were freely available at https://apps.ecmwf.int/datasets/. A 10-day Butterworth filter was applied to each of these to match the RAPID transport data processing, then the filtered Florida Current and Ekman transports were resampled to give monthly mean values.

4.3 Methods

The selected CTD profiles to be applied to the regression models had irregular sampling with approximately annual resolution, so the UMO and LNADW transports predicted by the regressions would also have irregular intervals. To allow comparison between these and the regular RAPID observations, we used two different approaches. Since Smeed et al. (2014) calculated RAPID transport means for two 4-year periods, 1 April to 31 March 2004–2008 and 2008–2012, we extended this approach by defining additional mean periods as 1 April–31 March for the years 1984-1988, 1988–1992, 1992–1996, 1996–2000, and 2012–2016. The single data point prior to 1 April 1984 was omitted from these mean calculations. To smooth the model results while allowing for the irregular sampling

Date	Lat.	Lon.	Source
04 Sep 1981	24.49	−75.42	wod2018
24 Apr 1985	26.56	−76.75	wod2018
12 May 1985	26.54	−76.68	wod2018
29 Aug 1985	26.53	−76.75	wod2018
19 Jan 1986	26.55	−76.73	wod2018
05 Feb 1986	26.55	−76.75	wod2018
23 Jul 1986 †	26.53	−76.74	wod2018
18 Nov 1986	26.51	−76.73	wod2018
19 Mar 1987	26.54	−76.76	hybas
27 Mar 1987	26.55	−76.75	wod2018
03 Sep 1987	26.55	−76.77	hybas
07 Oct 1988	26.53	−76.73	wod2018
13 Mar 1989	26.55	−76.75	wod2018
28 Feb 1990	26.55	−76.75	hybas
18 Jun 1990	26.49	−76.82	hybas
30 Jun 1990	26.49	−76.74	hybas
09 Jan 1991	26.51	−76.76	hybas
15 Jun 1991	26.53	−76.76	hybas
22 Sep 1991	26.52	−76.76	hybas
20 Jul 1992	24.55	−75.46	hybas
11 Aug 1992	26.50	−76.72	wod2018
14 Aug 1992	24.55	−75.46	wod2018
02 Jun 1993	26.50	−76.68	aomlw
13 Jul 1994	26.50	−76.68	aomlw
08 Mar 1996	26.50	−76.69	aomlw
10 Jul 1996	26.50	−76.75	hybas
31 Jul 1996	26.50	−76.75	aomlw
11 Jun 1997	26.51	−76.75	hybas
22 Feb 1998	26.50	−76.75	wod2018
28 Apr 2001	26.50	−76.73	wod2018
26 Jun 2002	26.50	−76.76	wod2018
11 Feb 2003	26.50	−76.72	wod2018
27 Mar 2004	26.51	−76.76	wod2018
08 Apr 2004	26.52	−76.75	wod2018
02 Oct 2004	26.50	−76.77	wod2018
08 May 2005	26.50	−76.76	wod2018
23 May 2005	26.50	−76.76	wod2018
20 Sep 2005	26.50	−76.77	wod2018
19 Mar 2006	26.50	−76.77	wod2018
28 Sep 2006	26.50	−76.74	wod2018
08 Oct 2006	26.50	−76.74	wod2018
18 Sep 2007	26.50	−76.74	wod2018
18 Apr 2009	26.50	−76.74	wod2018
01 May 2009	26.37	−76.75	wod2018
23 Nov 2009	26.50	−76.74	wod2018
04 Dec 2009	26.53	−76.77	wod2018
09 Jan 2010	26.50	−76.78	wod2018
31 Mar 2010	26.57	−76.70	wod2018
07 Apr 2010	26.50	−76.74	wod2018
11 Mar 2011	26.20	−76.99	wod2018
15 Apr 2011	26.50	−76.74	wod2018
28 Apr 2011	26.50	−76.65	wod2018
29 Feb 2012	26.50	−76.74	wod2018
26 Sep 2012	26.49	−76.75	wod2018
27 Feb 2013	26.50	−76.74	wod2018
17 Mar 2014	26.50	−76.75	wod2018
28 Mar 2014	26.50	−76.74	wod2018
23 Feb 2015	26.49	−76.74	wod2018
05 Oct 2015	26.49	−76.75	wod2018
16 Oct 2015	25.95	−76.90	wod2018
14 Dec 2015	26.50	−76.76	wod2018
24 Feb 2016	26.49	−76.74	wod2018

Table 4.2: Selected western boundary CTD profiles. For merged profiles, the dates, position and data source are for the shallowest profile. The data sources are: wod2018 - World Ocean Database 2018; aomlw - NOAA/AOML; and hybas - Hydrobase 2003. The latter two were data obtained by Longworth et al. (2011).
† Cruise flagged for salinity.

intervals, we calculated an annual Gaussian-weighted rolling mean using a four-year window, weighting the results according to a normalised Gaussian distribution with its maximum at mid-point and decreasing to zero at the window edges. Finally, we removed the annual cycle, taken as the monthly mean of the RAPID observations, from all observed and predicted transport anomalies, and calculated the annual Gaussian-weighted rolling four-year means of these deseasonalised data.

4.4 Results

4.4.1 Reconstruction of transport anomalies from transatlantic section data

Our initial reconstruction applied density anomalies derived from transatlantic sections at 24.5°N from 1981, 1992, 1998, 2004, 2010 and 2015 to the UMO and LNADW regression models. As each section allowed us to create western and eastern profiles from approximately the same time, it gave an additional opportunity to compare the UMO model predictions against fully reconstructed transports from all sections except 2015 (Atkinson et al., 2012; Bryden et al., 2005), and to further validate the use of an eastern boundary climatology. Although the same transatlantic section was occupied during 1957, the regression model uncertainty for the 1957 section model estimate was much larger than for later sections. Since there were also no suitable hydrographic data between 1957 and 1981 and the 1957 transport would have had too much influence over any trend with time, we chose to omit the 1957 section from our results.

The predicted UMO transport anomalies in Fig. 4.6 show generally very good agreement between the original, non-climatology model and Bryden et al. (2005) estimates, although the model's UMO transport is consistently weaker southwards than the observations. The error bars show the model's prediction intervals, which give the range of UMO transport anomalies that we have 95% confidence will occur for that combination of boundary density anomalies. The prediction intervals vary slightly, but are between 2.4 and 2.6 Sv, with a mean of 2.5 Sv. Only the 2004 estimate from Bryden et al. (2005) is not within the original model's prediction intervals, however the model estimate appears to be in better agreement with the magnitude of the RAPID observations that begin very shortly

afterwards. The model UMO estimates for 2010 and 2015 are also consistent with the RAPID UMO, for example reflecting the peak in 2015. The estimates from the older sections do not show greater disagreement, suggesting the regression model works for data outside the training data period, which is not always the case for linear models.

The difference between the model predictions and observations is slightly increased by the use of an eastern boundary climatology, with the Bryden et al. (2005) estimates for both 1998 and 2004 now outside the model uncertainty, however the agreement between the 2010 and 2015 model estimates and RAPID observations is arguably improved. The differences between the original and climatology-based regression model predictions are likely due to the increased underestimation noted in Section 4.2.1.

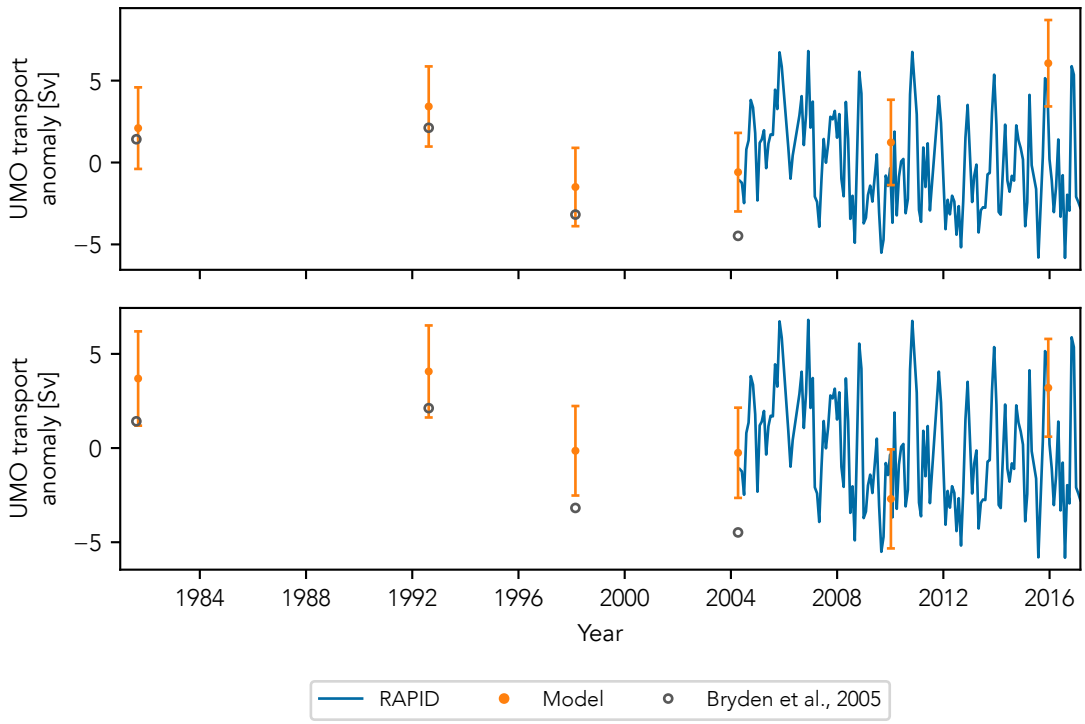

Figure 4.6: (Top) UMO transport anomaly estimated by regression model (orange) using density anomalies from six transatlantic hydrographic sections, compared to estimates from Bryden et al. (2005) (grey circles) and RAPID (blue). (Bottom) Same but replacing the eastern boundary regressor with a monthly climatology. The uncertainties shown for the model-derived values are the model's prediction intervals.

Fig. 4.7 shows similar plots for the AMOC transport anomaly. The AMOC model prediction is obtained by adding the predicted UMO transport anomaly to the monthly mean Florida Current and Ekman transport anomalies. This is in contrast to Bryden et

al. (2005), who used a constant northwards Ekman transport of 5.4 Sv, and a constant northwards Florida Current transport of 30.2 Sv for the 1981 and 1992 sections, and 32.2 Sv for the 1998 and 2004 sections. The differences in methodology may account for the decreased agreement between the model predictions and Bryden et al. (2005) estimates in 1992 and 1998 compared to the UMO estimates. The difference is reduced for the 1998 section when the eastern boundary climatology is used, again likely due to underestimation, so when the model predictions are compared to observations, a slightly weaker southwards UMO transport translates into a slightly stronger northwards AMOC transport. Atkinson et al. (2012) used annual mean Florida Current transports from the same Florida Straits cable measurements used in this study, while their annual mean Ekman transport were derived from the NOC v1.1 wind stress climatology (e.g., Josey et al. (2002)), and their estimated AMOC transports are all outside the regression prediction intervals with the exception of 1981, and show greater disagreement with the model-predicted AMOC than the Bryden et al. (2005) estimates. Despite this, the RAPID-observed AMOC transport anomaly is well within the expected uncertainty of both model predictions, and the models appear to capture the observed 2009-2010 downturn and 2015 peak well, although most of the downturn is likely due to changes in Ekman transport, as the model-predicted UMO transport anomaly is 1.3 Sv for 9 Jan 2010, compared to a mean RAPID value of −3.7 Sv for January 2010.

Lastly, we compare the estimated LNADW transport anomalies from the regression model, Bryden et al. (2005), Atkinson et al. (2012) and RAPID observations in Fig. 4.8. These model anomalies show a similar underestimation of the strength of the southwards deep return transport to the UMO model, although all but the 1992 Bryden et al. (2005) estimate are within the model's prediction intervals. The 2010 model estimate shows again a slightly weaker southwards LNADW anomaly than observed by RAPID, but it is closer to the peak than the 2010 Atkinson et al. (2012) estimate. The 2015 model estimate is markedly and unusually stronger than the RAPID LNADW transport.

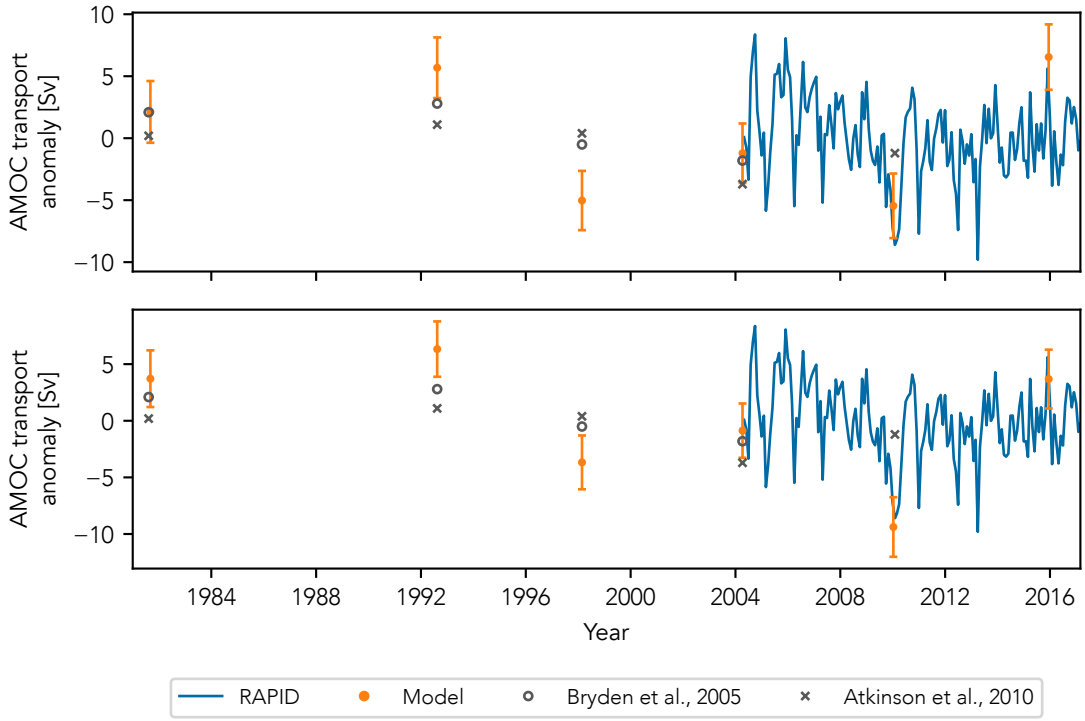

Figure 4.7: (Top) AMOC transport anomaly estimated by regression model (orange) using density anomalies from six transatlantic hydrographic sections, compared to estimates from Bryden et al. (2005) (grey circles), Atkinson et al. (2012) (grey crosses), and RAPID (blue). (Bottom) Same but replacing the eastern boundary regressor with a monthly climatology. The uncertainties shown for the model-derived values are the model's prediction intervals.

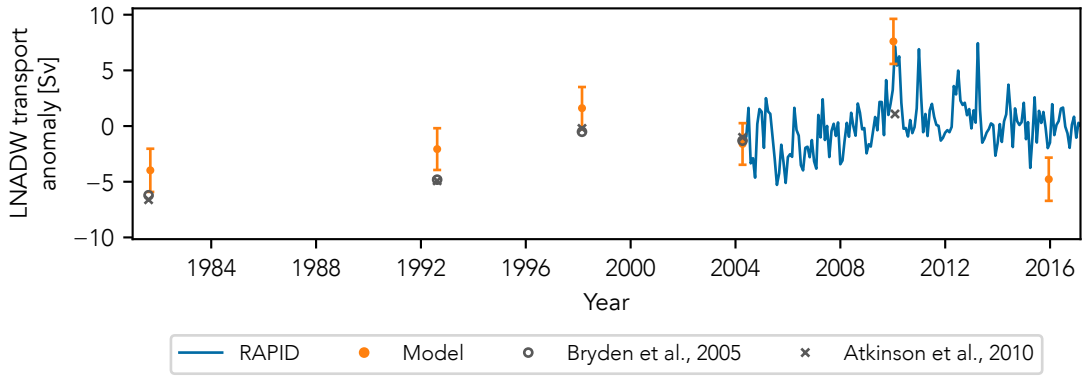

Figure 4.8: LNADW transport anomaly estimated by regression model (orange) using density anomalies from six transatlantic hydrographic sections, compared to estimates from Bryden et al. (2005) (grey circles), Atkinson et al. (2012) (grey crosses), and RAPID (blue). The uncertainties shown for the model-derived values are the model's prediction intervals.

4.4.2 Reconstruction of the AMOC from all available western boundary hydrographic data at 26°N

We applied both UMO and LNADW regression models to western boundary density anomalies from the hydrographic profiles selected in Section 4.2.3, with the UMO model including an eastern boundary climatology. The resolution of the resulting UMO time series in Fig. 4.9 is irregular and approximately annual, and is sufficient to show decadal and multi-annual variability. The standard deviation of the RAPID monthly mean anomalies is ±2.9 Sv, and the UMO transport anomaly estimated by the model is stronger southward than this in March 1989, June 1993, February 2003, and March 2004, with each anomaly lower than −3.3 Sv. This is less than any negative anomaly predicted by the model during the RAPID period, which have a minimum of −2.75 Sv.

The UMO is weaker southward than the RAPID standard deviation in September 1981, February 1986, July and August 1992, October 2004, and October and December 2015, with each being greater than 3.0 Sv, and with the exception of July 1992 and December 2015, greater than 3.5 Sv. When compared directly to the RAPID observations, the model predictions generally agree well with the overall trends, with one notable exception during the RAPID WB2 mooring failure of late 2005 to early 2006, when the model estimates a stronger southwards UMO than was observed by RAPID. The weighted rolling annual mean shows decadal variability of just under ±1.6 Sv, with 2007 having the weakest UMO anomaly of −16.9 Sv, and 2002 having the strongest of −20.1 Sv. We must consider however the lower resolution of hydrographic data from the late 1990s to early 2000s, making the rolling mean more susceptible to influence by more extreme values. Over the whole time series, a Mann-Kendall test, which is used to detect monotonic trends (Kendall, 1948; Mann, 1945), shows that there is no significant trend (95% significance).

The AMOC transport anomalies shown in Fig. 4.10 are again obtained by summing the UMO transport anomalies shown in Fig. 4.9 with the monthly mean Florida Current and Ekman transport anomalies, and including the monthly mean Ekman transport allows the model to capture the 2009–2010 downturn well, with anomalies of −9.4 and −6.5 Sv for CTD profiles on 9th January and 31 March 2010 respectively. Prior to 2004, an AMOC anomaly less than −5 Sv is only seen in March 1987 (−5.6 Sv). The strongest AMOC

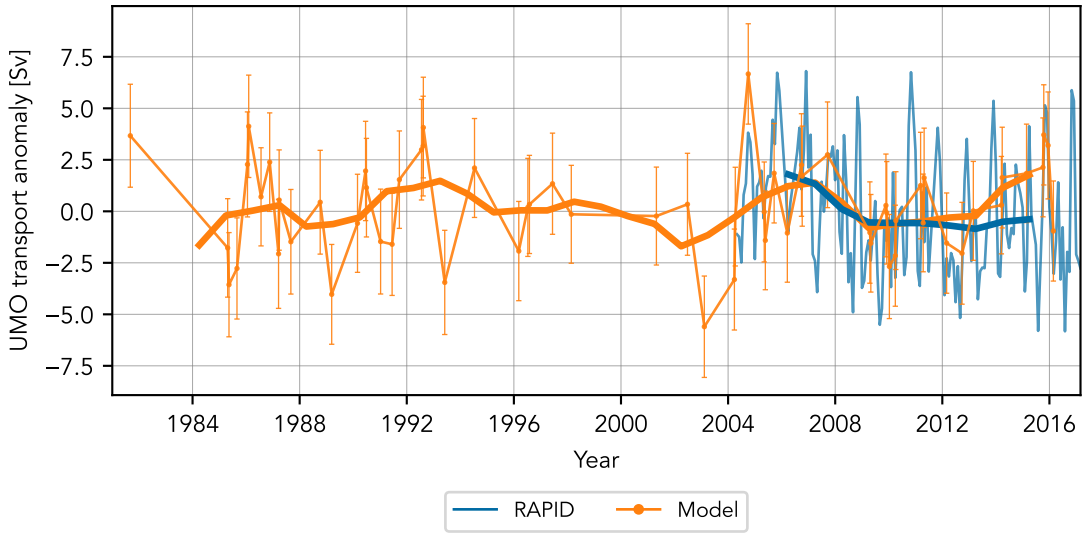

Figure 4.9: The thin lines are the UMO transport anomaly (orange) estimated by regression model using density anomalies from hydrographic CTD profiles, compared to the RAPID observations (blue). The 980 dbar eastern boundary density anomaly for the UMO model was replaced by RAPID monthly climatology. The thick lines are the annual Gaussian-weighted rolling mean using a four-year window.

anomalies (greater than +5 Sv) estimated by the regression model are for CTD profiles taken on 05 February 1986 (5.4 Sv), 18 November 1986 (7.4 Sv), 11 August 1992 (5.6 Sv), and 02 October 2004 (7.8 Sv). Given the nature of the regression models to underestimate the anomalies, it is likely that these positive anomalies are actually less than than the model predicts, while the negative anomalies are actually greater. The Mann-Kendall test for the model-estimated AMOC time series also shows no significant overall trend with 95% significance.

The standard deviation of the RAPID monthly mean LNADW anomalies is 2.2 Sv. Of the LNADW transport anomalies that are close to, or larger than twice the standard deviation, the strongest positive anomalies estimated from hydrographic profiles are on 19 March 1987 (+6.6 Sv), 9 January 2010 (+7.6 Sv) and 11 March 2011 (+4.5 Sv). Four of the strongest negative anomalies occur before 1992: on 4 September 1981 (−4.3 Sv), on 29 August 1985 (−4.2 Sv), on 19 January 1986 (−6.7 Sv), and on 9 January 1991 (−4.1 Sv). After 1992, on 18 September 2007 the minimum LNADW transport anomaly is −4.8 Sv. The last two hydrographic profiles, taken on 14 December 2015 and 24 February 2016, also show negative LNADW anomalies of −4.8 and −4.1 Sv respectively, although these are much lower than any RAPID observations for this period. The LNADW anomaly of

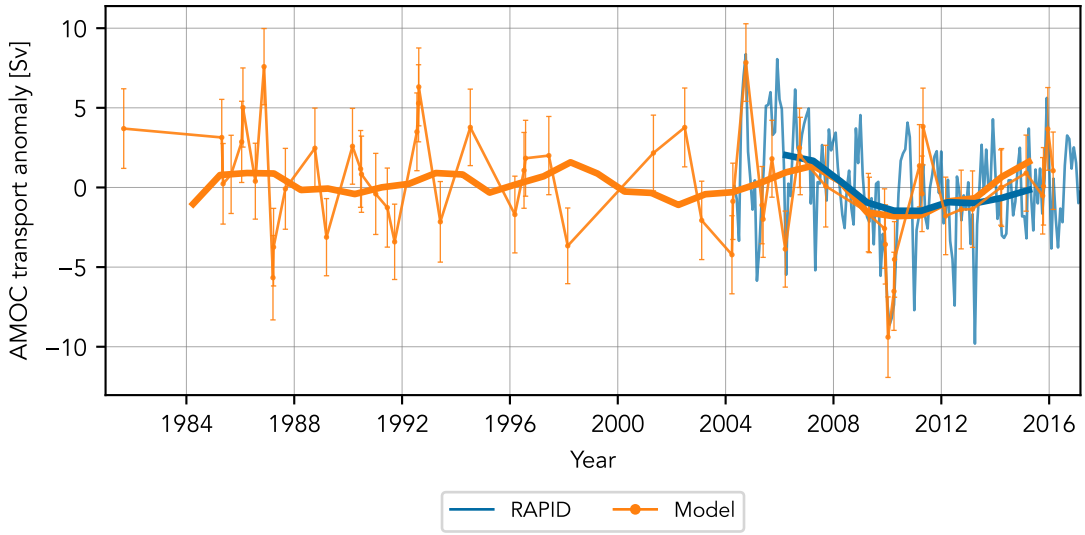

Figure 4.10: As for Fig. 4.9, but the orange lines show the AMOC transport anomaly as the sum of the UMO transport anomaly in Fig. 4.9 and the monthly mean observed Florida Current and Ekman transport anomalies.

−6.7 Sv on 19 January 1986 is the lowest estimate for the whole time series, and is lower than the minimum RAPID monthly mean LNADW anomaly of −5.8 Sv on 31 July 2005, if not lower than the minimum RAPID anomaly of −7.7 Sv. The Mann-Kendall test for the model-estimated LNADW time series also shows no significant overall trend.

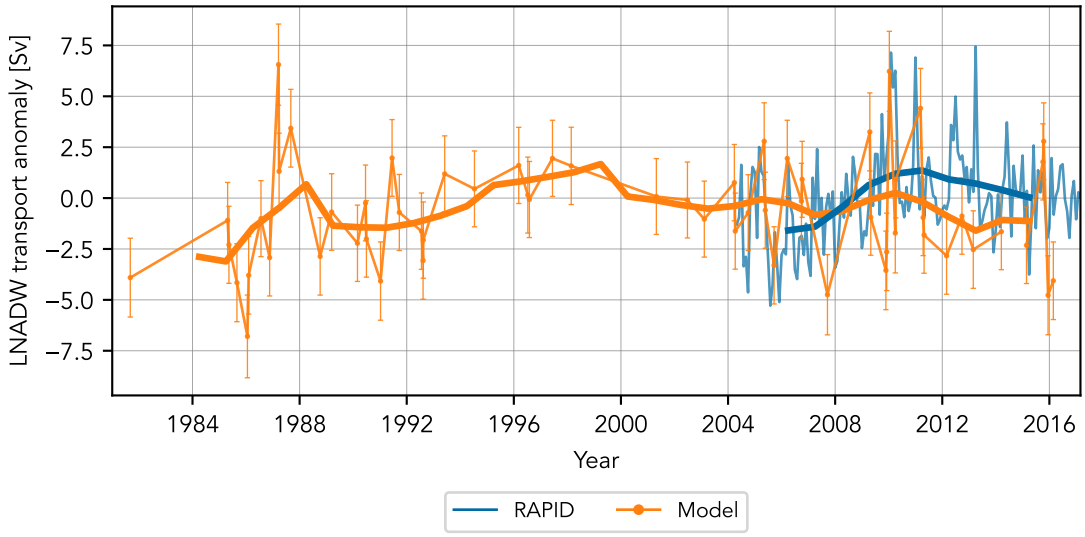

Figure 4.11: The thin lines are the LNADW transport anomaly (orange) estimated by regression model using density anomalies from hydrographic CTD profiles, compared to the RAPID observations (blue). The thick lines are the annual Gaussian-weighted rolling mean using a four-year window.

Fig. 4.12 shows the same estimates shown individually in Fig. 4.9, Fig. 4.10 and Fig. 4.11, now added to the mean transports that were subtracted to create the original RAPID training data anomalies. The plot also includes estimates from Bryden et al. (2005) and Atkinson et al. (2012), and the monthly mean Florida Current and Ekman transports to allow a direct comparison of all estimated transports.

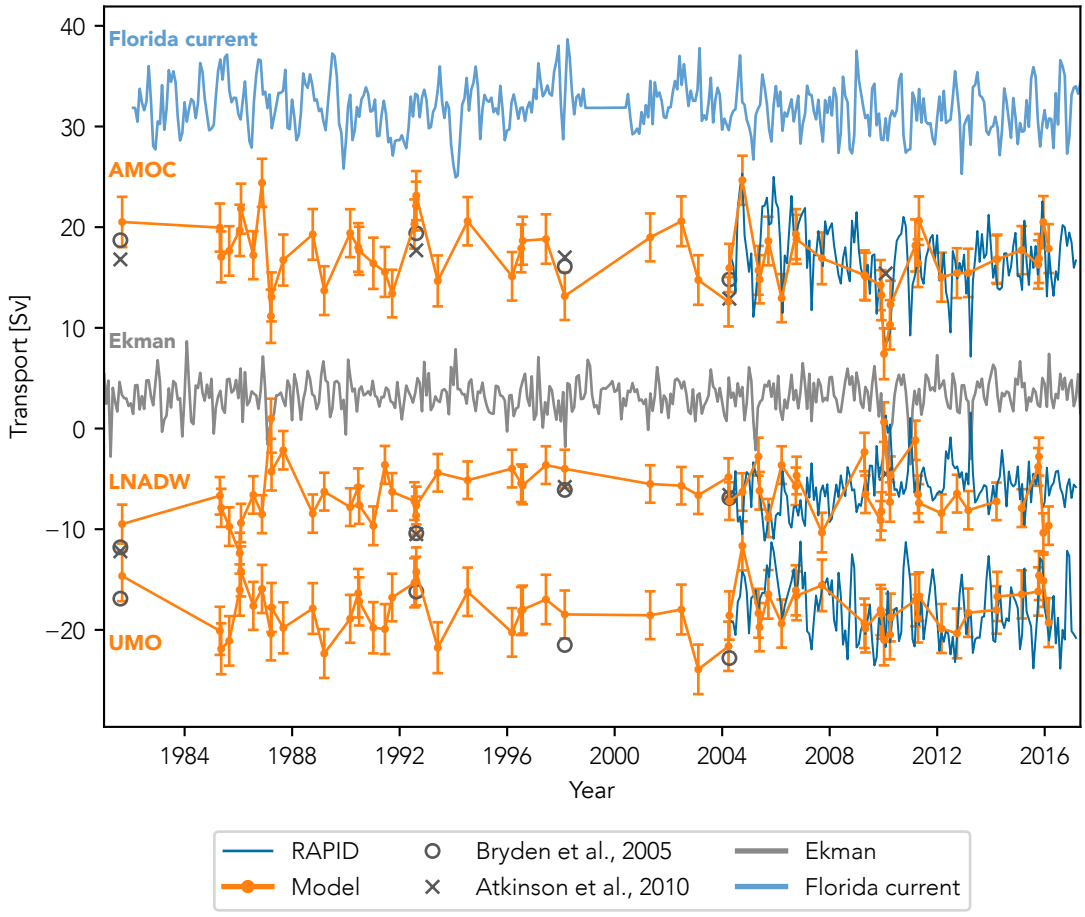

Figure 4.12: UMO and LNADW transports estimated by empirical models using density anomalies from hydrographic CTD profiles, compared to estimates from RAPID and Bryden et al. (2005). The 920 dbar eastern boundary density anomaly for the UMO model was replaced by RAPID monthly climatology. The monthly mean Florida Current and Ekman transports are also shown, and were added to the UMO model-estimated transports to give the estimated AMOC transport.

The time series in Fig. 4.13 are the annual Gaussian-weighted rolling means with a four-year window, applied to the time series in Fig. 4.12 with the RAPID-based annual cycle removed from each transport. They show that between 1984 and 2000, the UMO strength was within 1.9 Sv of the RAPID mean UMO of −18.0 Sv, taken for the period

used to create the model. Removing the annual mean from the model estimates increases the decadal variability of the rolling mean from ±1.6 to ±1.8 Sv. The rolling mean also suggests that the multi-year UMO variability was low during the 1990s, followed by a period of strengthened southward transport in the early 2000s. It also suggest a weakening of southwards UMO transport from 2012 not seen in the observations. The AMOC rolling mean largely reflects the UMO mean, however the increasing Florida Current transport between 1992 and 1999 result in a strengthening trend in the AMOC rolling mean from 1989 to 1998, reaching a time series maximum of 19.8 Sv in 1998. It then declines to to an overall minimum of 15.0 Sv in 2010, before recovering to back above the RAPID mean of 17.0 Sv. The rolling mean LNADW transport suggests a non-monotonic weakening trend in the southwards deep return flow between 1985 and 1999, from 8.8 Sv southwards in 1985, decreasing to 3.8 Sv southwards in 1999. In 2000, it strengthens slightly to around the mean RAPID transport of −5.9 Sv, and then varies by less than 0.51 Sv between 2000 and 2011, before strengthening again to a post-2000 maximum southwards transport of −6.7 Sv.

Table 4.3 shows simple four-year means estimated for the model-predicted and RAPID-observed transports, allowing comparison with the four-year mean RAPID transports from Smeed et al. (2018). The four-year simple means start from 1984 as there was only a single pre-1984 profile. The small number of data points means that the standard errors of each mean for the model-predicted transports are much higher than for the RAPID-observed transports. The period 2000–2004 has both the strongest southwards UMO and weakest northwards AMOC mean transport at −20.5 ±1.2 Sv and 14.9 ±1.7 Sv respectively. These are more extreme than during the RAPID reduced periods of 2008–2012 and 2012–2016. During these two periods, the model and RAPID four-year means compare well, the UMO transport differing by 0.4 Sv or less, and the AMOC by 0.7 Sv or less. The combined uncertainty for these periods are however ±0.9 and ±0.5 Sv for UMO transports, and ±1.4 and ±1.2 Sv for AMOC transports. In contrast, the model mean UMO for 2012–2017 is 1.5 ±0.7 Sv less than the RAPID mean, and the difference for the AMOC is 1.3 ±1.3 Sv. The four-year mean LNADW transports show the strongest southward values for 1988–1992 and 2012–2016, the latter reflecting the unusually low anomaly seen predicted from the 2015 transatlantic section data. The most recent period's mean is 1.7 ±0.9 Sv less than the RAPID mean, whereas the differences for 2004-2008 and

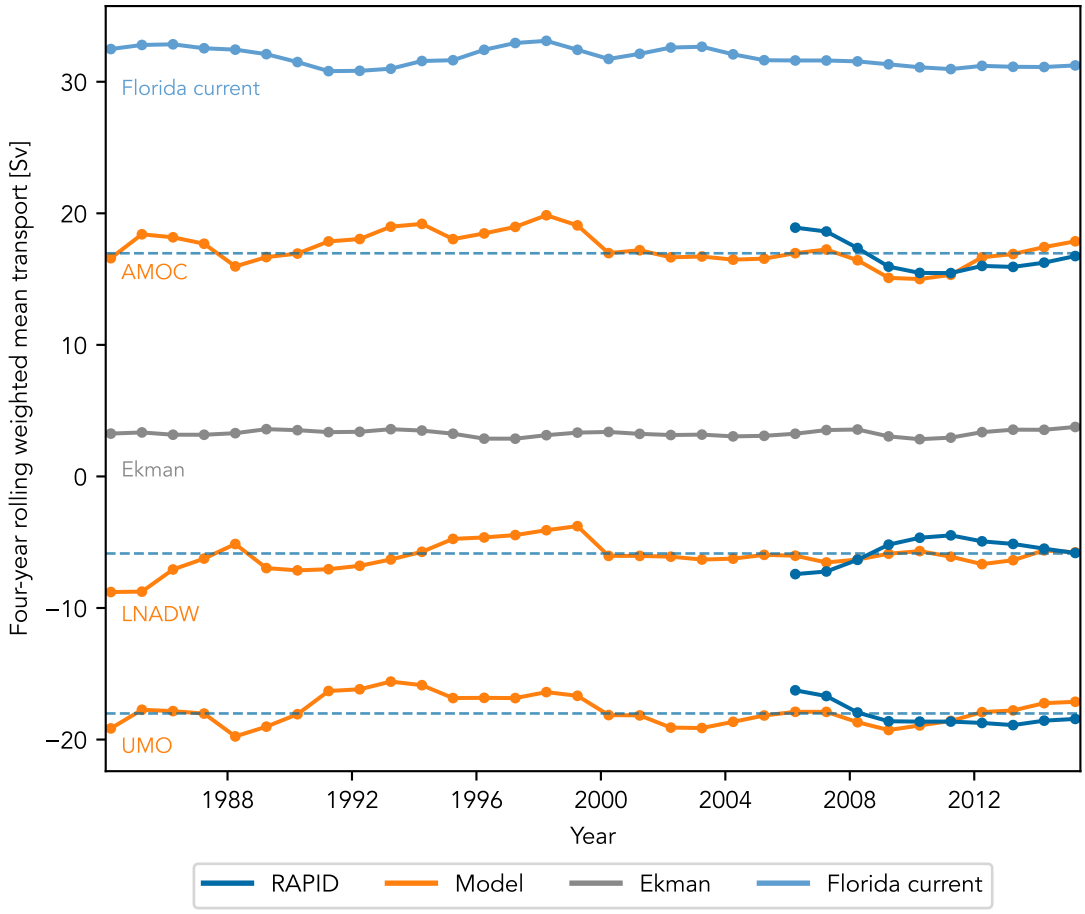

Figure 4.13: Gaussian-weighted 4-year rolling mean for AMOC, LNADW and UMO transports estimated by regression model (orange) and RAPID observations (dark blue). The equivalent rolling means for Florida Current (light blue) and Ekman (dark grey) transports are also shown. All transports had the annual cycle removed before the rolling means were calculated. Dashed horizontal blue lines show the means for the RAPID observed transports.

2008–2012 are 1.1 ±0.9 Sv and 0.4 ±1.0 Sv. These mean transports suggest LNADW flow was weakest during the late 1990s, but has since recovered to its strongest level, although the relatively large uncertainties should be taken into account.

4.4.3 Direct comparison with RAPID observations

The results from the previous section provide us with an additional method of estimating the uncertainty associated with using this regression model, trained on monthly mean anomalies, to predict transports from single data points. For the 28 CTD profiles that were taken during the period of RAPID mooring deployment, we made a direct comparison between each transport anomaly predicted by the regression model and the equivalent RAPID transport anomaly closest in time to that prediction.

Start year	End year	n	UMO		AMOC		LNADW	
			RAPID	Model	RAPID	Model	RAPID	Model
1984	1988	10		−18.5 (0.8)		17.3 (1.4)		−6.7 (1.2)
1988	1992	8		−18.6 (0.7)		16.2 (1.2)		−6.9 (0.6)
1992	1996	6		−17.2 (1.1)		17.8 (1.7)		−6.2 (0.7)
1996	2000	4		−17.9 (0.3)		18.1 (0.8)		−4.7 (0.4)
2000	2004	4		−20.5 (1.2)		14.9 (1.7)		−5.7 (0.3)
2004	2008	9	−16.5 (0.1)	−16.9 (0.8)	18.8 (0.1)	18.1 (1.3)	−7.3 (0.1)	−6.2 (0.8)
2008	2012	11	−18.6 (0.1)	−18.9 (0.4)	15.9 (0.1)	15.3 (1.1)	−5.0 (0.1)	−5.6 (0.9)
2012	2016	9	−18.7 (0.1)	−17.2 (0.6)	16.1 (0.1)	17.4 (1.2)	−5.3 (0.0)	−7.0 (0.9)

Table 4.3: Four-year mean UMO, AMOC and LNADW transports in Sverdrup estimated by RAPID and regression models. The four year period is defined between 31 March of the start year and 01 April of the end year. The standard error of the mean is shown in parentheses after the mean, and n is the number of values used to determine the mean for each period.

Transport	Error statistics [Sv]		
	Mean	σ	RMS
UMO	1.61	1.16	1.99
MOC	1.43	1.14	1.83
LNADW	2.82	1.82	3.36

Table 4.4: The mean, standard deviation (σ) and root-mean-square (RMS) of the error, where the error is the difference between the transport predicted by the regression model from the hydrographic data, and the nearest RAPID observed transport, for the period of the RAPID deployment.

Fig. 4.14 shows that for most of the UMO transport anomalies predicted by the model during the RAPID deployment period, the nearest RAPID observation is within the model prediction interval, shown by the error bars. The actual errors, taken as the residuals between the predicted and observed values, have a mean of 1.61 Sv (σ = 1.16) and a root-mean-square of 1.99 Sv. This is lower than the 2.7 Sv RMS error estimated in Section 3.4.4, however that was based on data entirely outside the duration of the time series used to train the regression, whereas the predictions in Fig. 4.14 occur mostly within the training data period of 27 May 2006 to 21 February 2017.

Fig. 4.15 shows that the errors are even smaller for the model-predicted AMOC transport anomalies, with a mean error of 1.43 Sv (σ = 1.14) and a RMS error of 1.83 Sv. The addition of either or both of the Florida Current and Ekman transport anomalies to

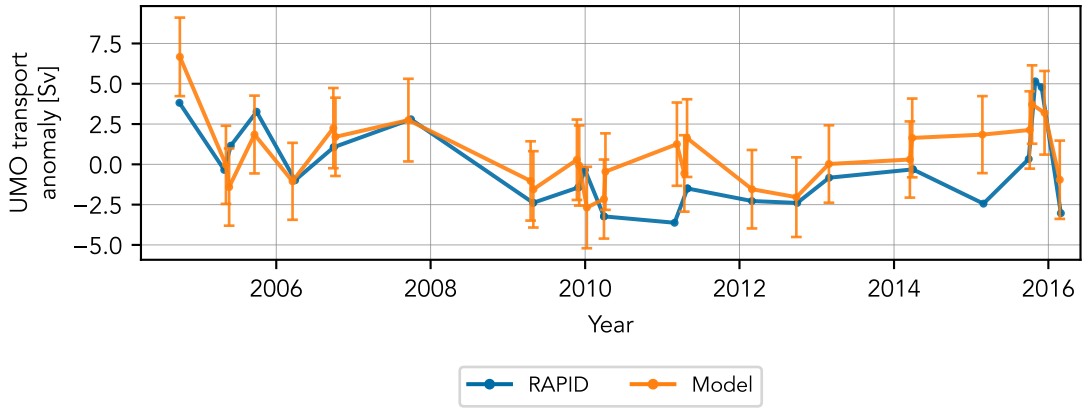

Figure 4.14: The UMO transport anomaly (orange) estimated by regression model using density anomalies from 28 hydrographic CTD profiles, compared to the RAPID observation closest to each in time (blue). The 980 dbar eastern boundary density anomaly for the UMO model was replaced by RAPID monthly climatology.

the predicted UMO anomalies reduced the overall error, and comparing Fig. 4.14 and Fig. 4.15 suggests that this had the most impact during 2011 and 2015.

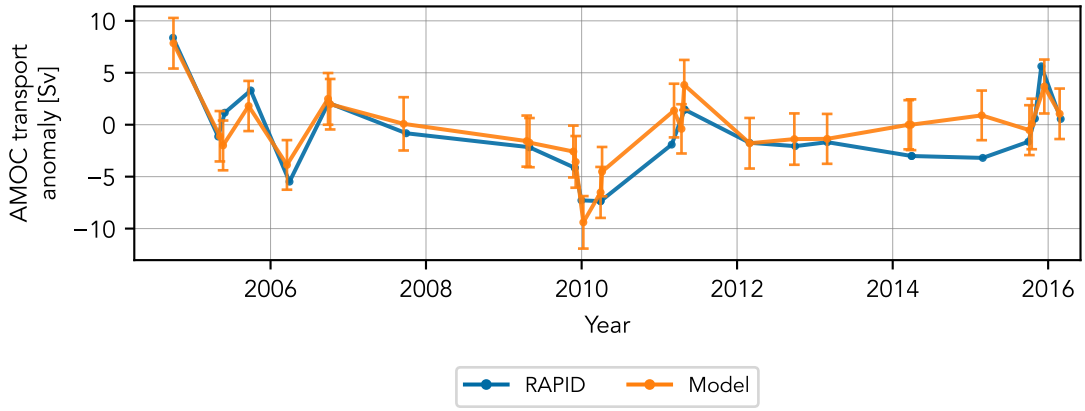

Figure 4.15: As for Fig. 4.14 but the orange lines show the AMOC transport anomaly as the sum of the UMO transport anomaly in Fig. 4.14 and the monthly mean observed Florida Current and Ekman transport anomalies.

Fig. 4.16 shows that the errors for the predicted LNADW transport are much larger, with a mean error of 2.82 Sv (σ = 1.82) and RMS error of 3.36 Sv.

4.4.4 Comparison with other proxy reconstructions

Here we compare the AMOC reconstruction from our empirical model with proxy reconstructions from Caesar et al. (2018), Thornalley et al. (2018) and Frajka-Williams

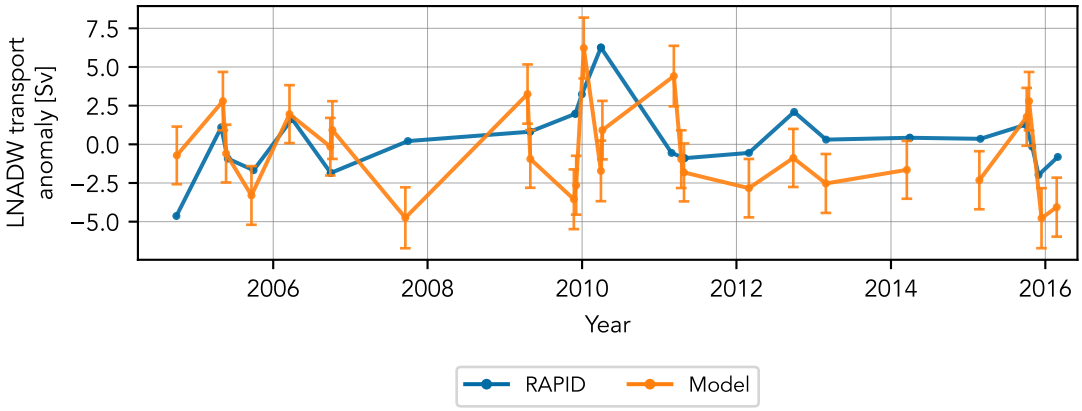

Figure 4.16: The LNADW transport anomaly (orange) estimated by regression model using density anomalies from hydrographic CTD profiles, compared to the RAPID observation closest to each in time (blue).

(2015). These reconstructions use independent proxies and the first two cover a much longer period, but we focus on the 1980-2018 period. The model AMOC time series is the Gaussian-weighted four-year rolling mean of the model predictions with the annual cycle removed. Determining the standard error of the weighted time series is beyond the scope of this study, but the root-mean-square of the differences between the model predictions and RAPID observations is 1.1 Sv. The simple four-year means (Table 4.4) have an average standard error of 1.3 Sv, so this seems a reasonable approximation of the uncertainty in the smoothed time series.

Thornalley et al. (2018) used sortable-silt mean grain size from two sediment cores taken from beneath the Deep Western Boundary Current near Cape Hatteras. The mean grain size is a proxy for current speed close to the seabed, with a stronger current associated with a stronger AMOC. They showed that this grain size proxy agrees well with other AMOC proxies such as deep Labrador Sea annual density, sub-polar gyre heat content, and several sub-surface temperature fingerprints obtained from climate models, ocean reanalyses and sediment cores. Caesar et al. (2018) used HadISST sea-surface temperatures (SST) to create an AMOC index between 1870 and 2016, defined as the SST anomaly of a defined sub-polar gyre region with respect to the global mean SST. The SST anomaly was converted from Kelvin to Sverdrup using a conversion factor of 3.8 Sv/K, which was calculated from CMIP5 ensemble data. Fig. 4.17 shows this annual AMOC index compares well to our RAPID-model AMOC reconstruction, after the SST-index is offset by 2.5 Sv and lagged by 6 years. We determined the 6 year lag from cross-correlation of

the two time series, and the magnitude of this lag is comparable to the 1 to 5 year lags found between a SST dipole index and AMOC at 30°N in 10 CMIP5 models by Roberts et al. (2013). Frajka-Williams (2015) regressed UMO transport onto monthly satellite sea surface height anomaly centred at 30°N, 70°W, creating a single-layer model that showed a significant correlation.

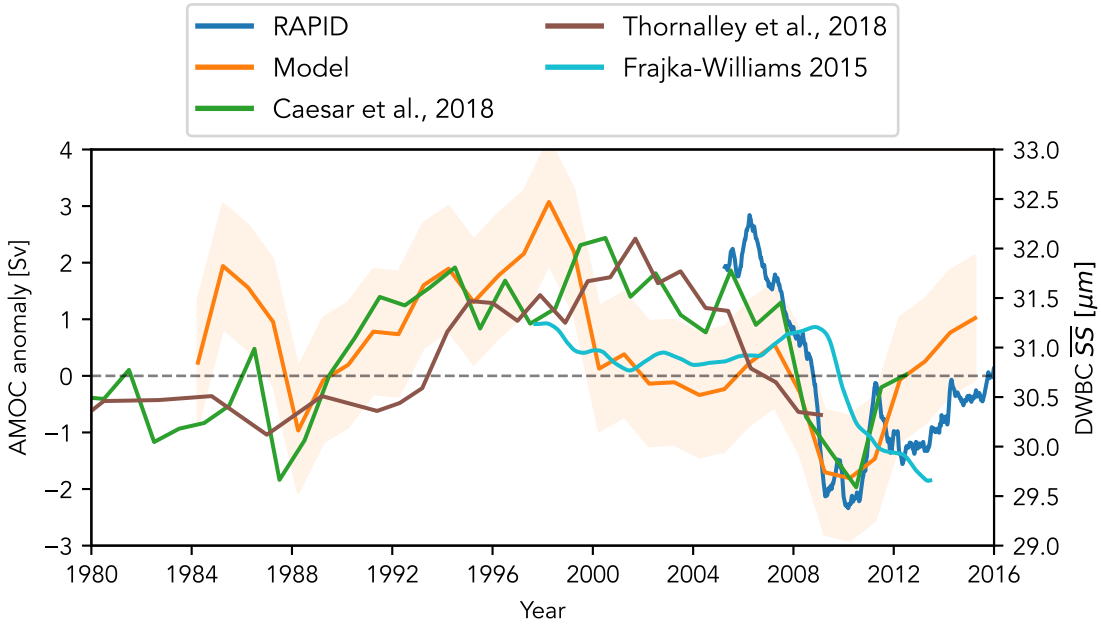

Figure 4.17: Comparison of AMOC anomalies from: RAPID (dark blue, smoothed with 4-year rolling mean); our regression model prediction (orange, smoothed with 4-year Gaussian-weighted rolling mean); AMOC anomaly index from Caesar et al. (2018) (green, offset by +2.5 Sv and −6 years); AMOC reconstruction from Frajka-Williams (2015) (light blue, smoothed with a 4-year rolling mean); and the mean grain size of sortable silt under the Deep Western Boundary Current from Thornalley et al. (2018) (brown). The orange shaded area shows a ±1.1 Sv uncertainty.

Fig. 4.17 shows that the annual mean of each proxy reconstruction show the same overall pattern: a strengthening AMOC from the late 1980s until the mid-to-late 1990s; then a steeper decline until around 2010, depending on the reconstruction. Our reconstruction and that of Caesar et al. (2018) then suggest a recovery from this minima. The RAPID observations also show the same sharp decline and recovery before and after the 2009-2010 period. The ±1.1 Sv estimated uncertainty of the smoothed reconstructions is sufficiently small compared to the 5.2 Sv range seen in the smoothed RAPID AMOC time series to give us confidence in comparing the model results with other reconstructions.

Fig. 4.18 compares UMO transport anomalies from two proxy reconstructions: our regression model and Frajka-Williams (2015)'s satellite altimetry study. Although both reconstructions reproduce the lower frequency variability observed in the RAPID transport, only our reconstruction captures the higher frequency variability and the magnitude of variability. This is shown in the decadal means; whereas there is less than a 0.2 Sv difference between the 1993-2003 and 2004-2014 altimetry proxy UMO mean transports, our reconstruction shows a 0.9 Sv difference.

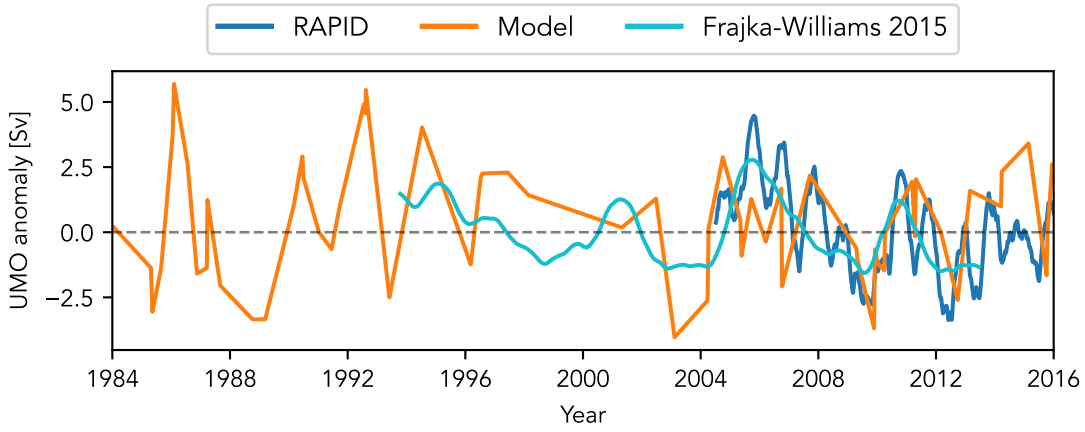

Figure 4.18: Comparison of UMO anomalies from RAPID (dark blue, smoothed with 1-year rolling mean); our regression model prediction (orange); and the AMOC proxy reconstruction from Frajka-Williams (2015) (light blue).

4.5 Conclusions

In this chapter, we have described the reconstruction of the UMO and LNADW transports from our empirical models, and the subsequent AMOC reconstruction obtained by summing the predicted UMO transport with observed Florida Current and Ekman transports. The 'best-fit' regression models from Chapter 3 were applied to western boundary density anomalies obtained first from transatlantic sections, then to historical hydrographic CTD profiles, with the eastern boundary density anomaly replaced by a monthly climatology for the latter. The model-predicted time series from 1981 to 2016 have irregular but approximately annual resolution, which is sufficient to show pentadal to decadal variability. In estimating uncertainty, where the hydrographic reconstructions overlapped the RAPID deployment period, the RMS error between 28 model predictions and the nearest RAPID observations was 2.0 Sv, less than the mean model prediction

interval of around 2.5 Sv.

There is no overall trend in any transport time series estimated by the model from hydrographic data, but four-year means, following Smeed et al. (2018), suggest that there were stronger southwards UMO and stronger northwards AMOC transports between 2000 and 2004 than at any time observed by RAPID. The sea-surface height model developed by Frajka-Williams (2015), which is implicitly single-layer, estimated the mean UMO for 1993-2003 and 2004-2014 to vary by less than 0.2 Sv. By contrast, our model estimated the same decadal means as −19.1 and −18.2 Sv, a difference of 0.9 Sv. The regression model 1982-1992 mean was also −18.2 Sv, showing that our four-layer model captures decadal mean UMO transport changes of nearly 1 Sv. The model mean UMO transport for 2004-2014 of −18.2 Sv also agrees well with the RAPID equivalent of −17.9 Sv. The importance of representing the deep layers can be seen in repeating the predictions using the two- and three-layer models described earlier. Adding a second or even a third density anomaly to the regression model gives a maximum decadal mean difference in UMO transport of only 0.3 Sv. In order to obtain the full 0.9 Sv difference between decadal mean UMO transport, all four density anomalies must be included in the regression.

In addition to the four-layer UMO regression model, we also created a similar model regressing LNADW transport on the deep western boundary density anomaly at 3040 dbar and Ekman transport. The rolling mean transports suggest that the AMOC and LNADW were both at their strongest in the late 1990s, however they also suggest that the southwards LNADW transport was at its strongest during the 1980s, and the weakening of LNADW associated with the 2004-2012 AMOC decline is not reflected by our reconstruction. The reconstructed UMO transport does show some decline, however the tendency of the regression model to underestimate transport magnitudes is clear. This underestimation then also suggests that the UMO may have been even weaker during the 1990s than the reconstruction suggests, with a consequently stronger AMOC.

Although this model increases the temporal resolution of AMOC estimates, the resolution is still coarse compared to RAPID and the time intervals between profiles are inconsistent. The longest period where no interval is greater than one year is October 1988 to July 1994. There are only two intervals longer than two years: September 1981 to April 1985; and February 1998 to April 2001. The longest interval is 1328 days, the mean is 210

days. Although this resolution is sufficient to show multi-year variability, as shown by the four-year means, the length of some of the sampling intervals and their inconsistency means the model cannot show interannual variability reliably.

Although we have reconstructed the AMOC time series from 1981 to 2015, the limited amount of data makes it difficult to assess how realistic our uncertainty estimates are. In the next chapter we address this further, and explore the scope and limits of these regression models, by evaluating them within a high resolution ocean simulation.

Chapter 5

Evaluating the empirical model in a high resolution ocean circulation model

5.1 Introduction

Evaluating the empirical models described in Chapter 3 within a numerical ocean model allowed us to move beyond the spatial and temporal restrictions of limited observational data. Using a high resolution ocean model – the NEMO ocean circulation model with a 1/12° resolution ORCA global tripolar grid (hereafter NEMO) –allowed us to further test the accuracy of the model predictions, as well as investigate how the models perform when the latitude and longitude of hydrographic data move away from the position of the RAPID western boundary mooring. The differences between regression models trained on RAPID and NEMO data could also provide insight into how well this numerical model represents the overturning circulation.

5.2 Data and Methods

5.2.1 Data

5.2.1.1 NEMO high resolution ocean model

The numerical model data used were 5-day means of potential temperature, salinity, meridional velocity and zonal velocity from the NEMO-ORCA0083 N006 historical run described in detail in Section 2.1.3. Initially, trans-Atlantic basin-wide sections or hyperslabs were extracted at a latitude of approximately 26.5°N, between longitudes of 12.7°W and 81.1°W, using the NCO `ncrcat` function to concatenate the 5-day data files into annual data files between 1958 and 2015. The domain data – bathymetry, bathymetry

levels and mesh_zgr – were also extracted for the same hyperslab. In situ densities were calculated from potential temperature and salinity data using the CDFTOOLS package `cdfsiginsitu` function. Since `cdfsiginsitu` produced in situ density values below the defined bathymetry, the bathymetry data were used to mask these values, and the masked NEMO density data at 26°N used to create RAPID-style single full-depth profiles at the western and eastern boundaries.

5.2.2 Methods

5.2.2.1 Replicating RAPID boundary profiles

Each of the methods for creating profiles for use in the regression model are similar in concept to the RAPID method for merging data from several moorings into a single profile (Section 2.1.1). Here we have the additional constraint that the gridded model data have zonal and meridional coordinates that do not exactly match latitude and longitude, due to the ORCA tripolar grid (see Section 2.1.3 for full details). However, the methodology remains the same: where a single profile is insufficiently deep, it is merged with the lower part of a deeper profile. On the western boundary, this will be a profile to the east of the initial profile; and on the eastern boundary, it will be to the west. This continues until the required depth of merged profile (4750 m) is reached.

To create NEMO training data for the regression models, we replicated the RAPID profiles as closely as possible. Data were selected from the grid points closest to the WB2, WBH2 and WB3 moorings on the western boundary, and to the EB1, EBH1, EBH2 and EBH3 moorings on the eastern boundary, and were merged as follows. On the western boundary, the data from the grid point closest to the WB2 mooring location was used as the main profile, with data extending from 0.5 m to 3513 m. From 3515 m down to 4688 m, the data for the grid point closest to the WBH2 was added, and from 4688 m to 4888 m, data from the grid point closest to WB3 was added to give a single profile from 0.5 m to 4888 m. A similar method was used on the eastern boundary, however the grid point nearest to EBH3 was a land point in NEMO. Thus the top part of the profile was comprised of data from the grid point closest to EBH2, from 0.5 m to 36 m. This was merged with data from the grid point closest to EBH1 between 36 m and 2262 m, and then with the data from the grid point closest to EB1 between 2262 m and 5089 m.

These NEMO profiles were then interpolated over the same depth and temporal grid as the RAPID profile data. Depth was converted to pressure using the Python `gsw` toolbox, and then the profiles were interpolated over a 20 dbar pressure grid. Both western and eastern boundary profiles were resampled to give monthly means, and then restricted to the same period as the RAPID model training data (March 2006 to February 2017). Finally monthly mean density anomalies were created by detrending the shortened time series. An eastern boundary monthly climatology was also created. Although an ocean model has simultaneous eastern and western boundary data, again this allowed us to better compare the RAPID- and NEMO-trained models. Fig. 5.1 shows that NEMO does not represent the strong seasonal signal of the AMOC, seen in observations of density anomalies at intermediate depths, particularly well. The much smaller range of the 1080 dbar density anomaly climatology reflects this, when compared to the RAPID-based equivalent in Fig. 4.1. The large negative density anomaly observed in October in the RAPID 920 dbar eastern boundary density anomaly and in the RAPID climatology is much smaller in NEMO, and there are some years (e.g., 2008 and 2010) when it is absent entirely. The NEMO climatology reflects this, showing much smaller seasonal extremes.

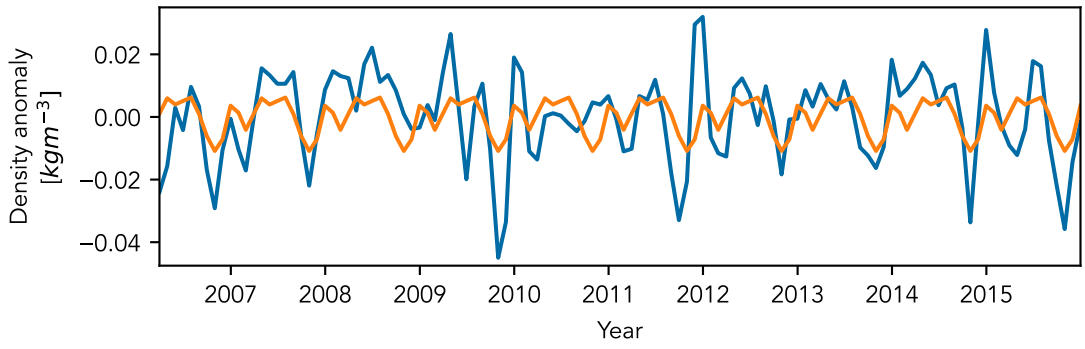

Figure 5.1: Time series of the NEMO eastern boundary density anomaly at 1080 dbar (blue) and its monthly climatology (orange).

5.2.2.2 Reconstructing UMO and AMOC transports from NEMO data

To calculate the UMO and AMOC transports from the NEMO data, we used the RapidMoc Python software package developed and made available by Chris Roberts (Roberts, 2017). The software calculates volume, freshwater and heat transports from ocean general circulation model data using a method analogous to the RAPID methodology.

The configuration requires setting longitudinal limits for the Florida Current and Western Boundary Wedge, and internal geostrophic transport. The model analogue to the Florida Current is defined between the minimum and maximum Florida Current longitudes (*fc_minlon* and *fc_maxlon* in Fig. 5.2). The model analogue to the Western Boundary Wedge containing the Antilles Current and Deep Western Boundary Current is defined between *fc_maxlon* and the eastern limit of the WBW region, *wbw_maxlon*. The geostrophic interior transport is calculated between *wbw_maxlon* and the eastern boundary limit, *int_maxlon*. The best regression model results were obtained when *wbw_maxlon* was set close to the longitude of the RAPID WB3 mooring, 76.5°W, and *int_maxlon* was set close to the longitude of the RAPID EBH3 mooring, 13.7°W. This was in accordance with the RapidMoc configuration guidelines (Roberts, 2017), where the model analogues to the northward flowing Antilles Current and southward flowing Deep Western Boundary Current should lie between *fc_maxlon* and *wbw_maxlon*.

Florida Current and Western Boundary Wedge transports are calculated from full-field model velocities, and meridional Ekman transport is calculated from the zonal wind stress component. Dynamic height is calculated from model temperature and salinity using the Jackett and McDougall Equation of State (Jackett and Mcdougall, 1995). Geostrophic 'interior' transport relative to a depth of no motion of 4750 m is then calculated from the dynamic height profiles between the *wbw_maxlon* and *int_maxlon* longitudes. Finally a RAPID-style mass compensation transport is applied as a barotropic velocity over the same region, to ensure there is zero net transport across the section. The uncertainties associated with using a level of no motion are discussed in Roberts et al. (2013).

Upper mid-ocean transports are calculated by the RapidMoc code as part of the AMOC calculations, but the RAPID depth-defined layer transports (thermocline, AAIW, UNADW and LNADW) are not. These were created using the same method as RAPID where the mid-ocean transport, which is the total compensated transport minus the Florida Current and Ekman transports, is integrated over the required depths: 0-800 (thermocline), 800-1100 (AAIW), 1100-3000 (UNADW) and 3000-5000 dbar (LNADW). We would not expect the RAPID and RapidMoc NEMO layer transports to match, but Fig. 5.3 compares their overall mean values and variability, and shows that the NEMO thermocline layer has a slightly stronger southwards transport (mean = −20.5 Sv) than that observed by RAPID

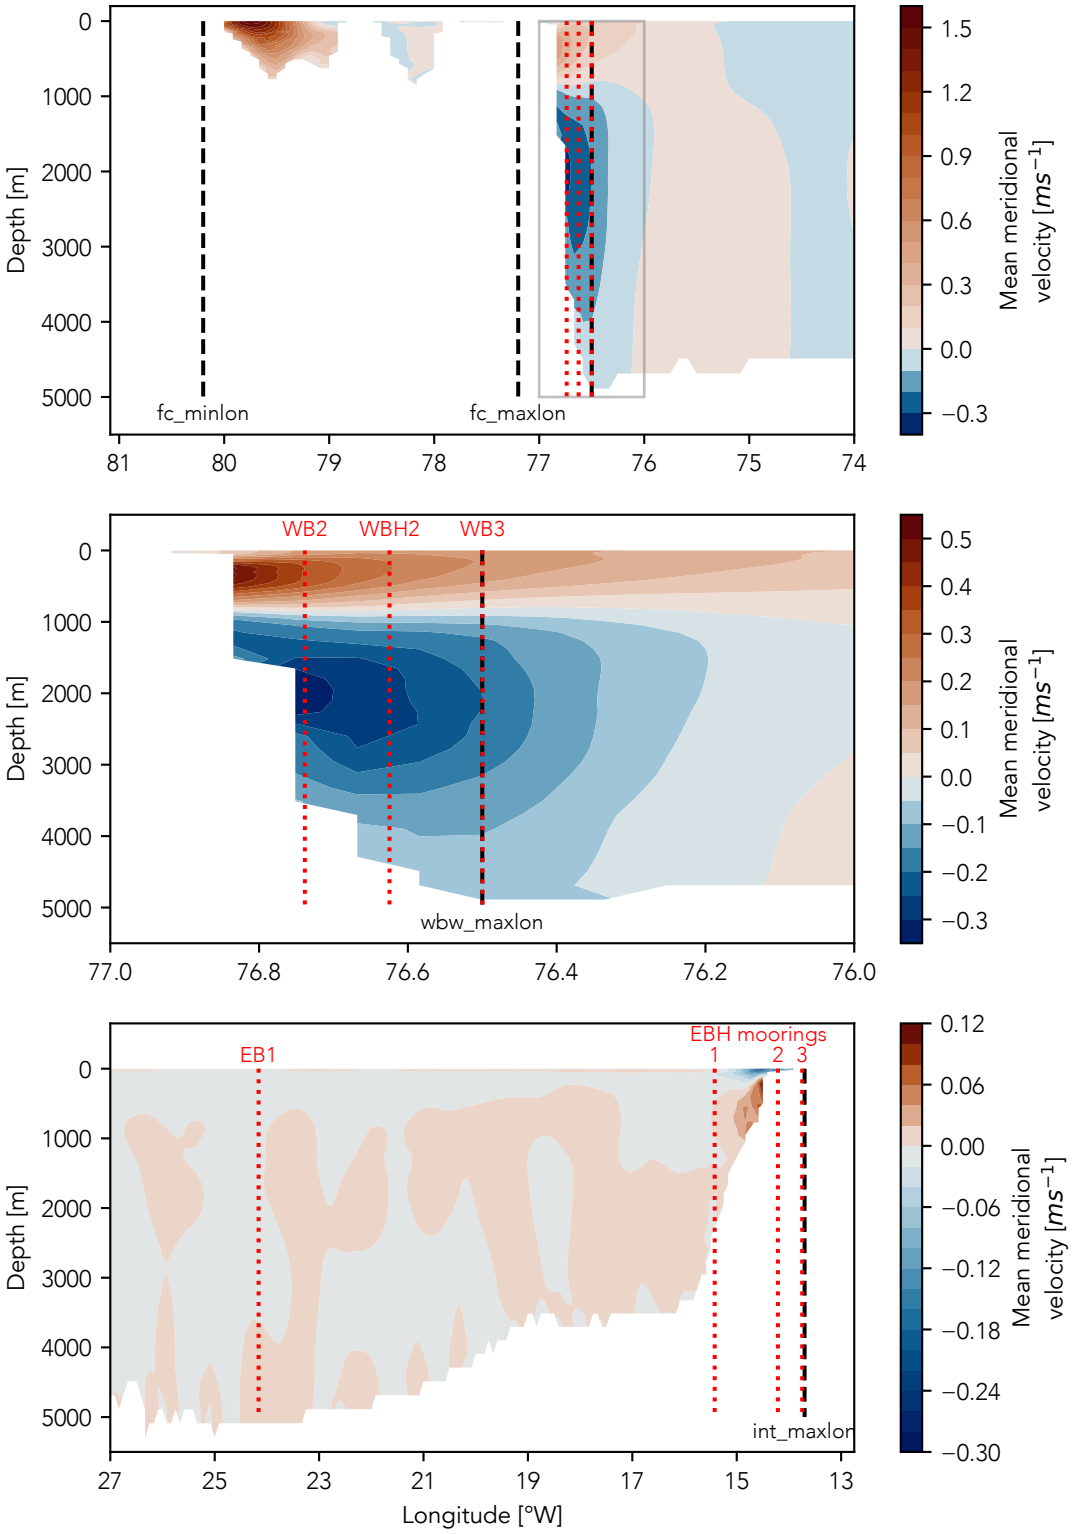

Figure 5.2: NEMO mean meridional velocity for 1959–2015 at approximately 26°N showing the longitudes of RAPID moorings (red dotted lines) and RapidMoc configuration region bounds (black dashed lines). (Top) Western boundary region showing the minimum and maximum longitudes defining the Florida Current model analogue. The grey box shows the western boundary current region, which is highlighted in: (Middle) Western boundary current region showing the RAPID western boundary moorings and the eastern limit of the model analogues of the Antilles and Deep Western Boundary Currents. (Bottom) Eastern boundary region showing the RAPID eastern boundary moorings and the eastern limit of the geostrophic interior region.

(mean = −18.7 Sv), whereas the NEMO LNADW layer has a generally weaker southwards transport (mean LNADW: NEMO = −3.4 Sv; RAPID = −5.9 Sv). The NEMO layers transports all show less variability than RAPID, particularly LNADW, which has standard deviations of 2.9 and 2.3 Sv for the RAPID and NEMO transports respectively.

Transport	**Mean [Sv]**		**σ [Sv]**	
	RAPID	**NEMO**	**RAPID**	**NEMO**
Thermocline (0-800m)	−18.7	−20.5	3.3	3.0
AAIW (800-1100m)	0.5	0.3	0.6	0.5
UNADW (1100-3000m)	−11.9	−12.8	2.5	2.4
LNADW (3000-5000m)	−5.9	−3.4	2.9	2.3

Table 5.1: Mean and standard deviation (σ) of NEMO RapidMoc layer transport time series shown in Fig. 5.3.

5.2.2.3 Creating an empirical model using NEMO data

Exactly the same approach as Chapter 2 was used, with algorithms selecting the particular combination of boundary density anomalies that gave the best-fitting linear regression models. The only difference was the density anomalies used as training data were created from NEMO 5-day mean temperature and salinity data rather than the RAPID mooring data. Fig. 5.4 shows the profiles of the mean mid-ocean transport from RAPID and calculated by RapidMoc from NEMO data for the same period. Although the NEMO profile shows slightly shallower thermocline and AAIW layers, the difference was not enough to justify changing the depths defined within the algorithms, i.e., the shallowest density anomaly was still found between 220 and 800 dbar.

5.2.2.4 Replicating historical hydrographic data

One advantage of using an ocean simulation such as NEMO is that they can be used to reproduce situations that in reality have only a limited number of observations, for example, the hydrographic data used to reconstruct the AMOC in Chapter 4. For the AMOC reconstruction, we used the regression model's prediction intervals as a measure of uncertainty, however the regression models are based on monthly mean anomalies, and the hydrographic data are effectively 'snapshots'. In Section 3.4.4 we tested how well

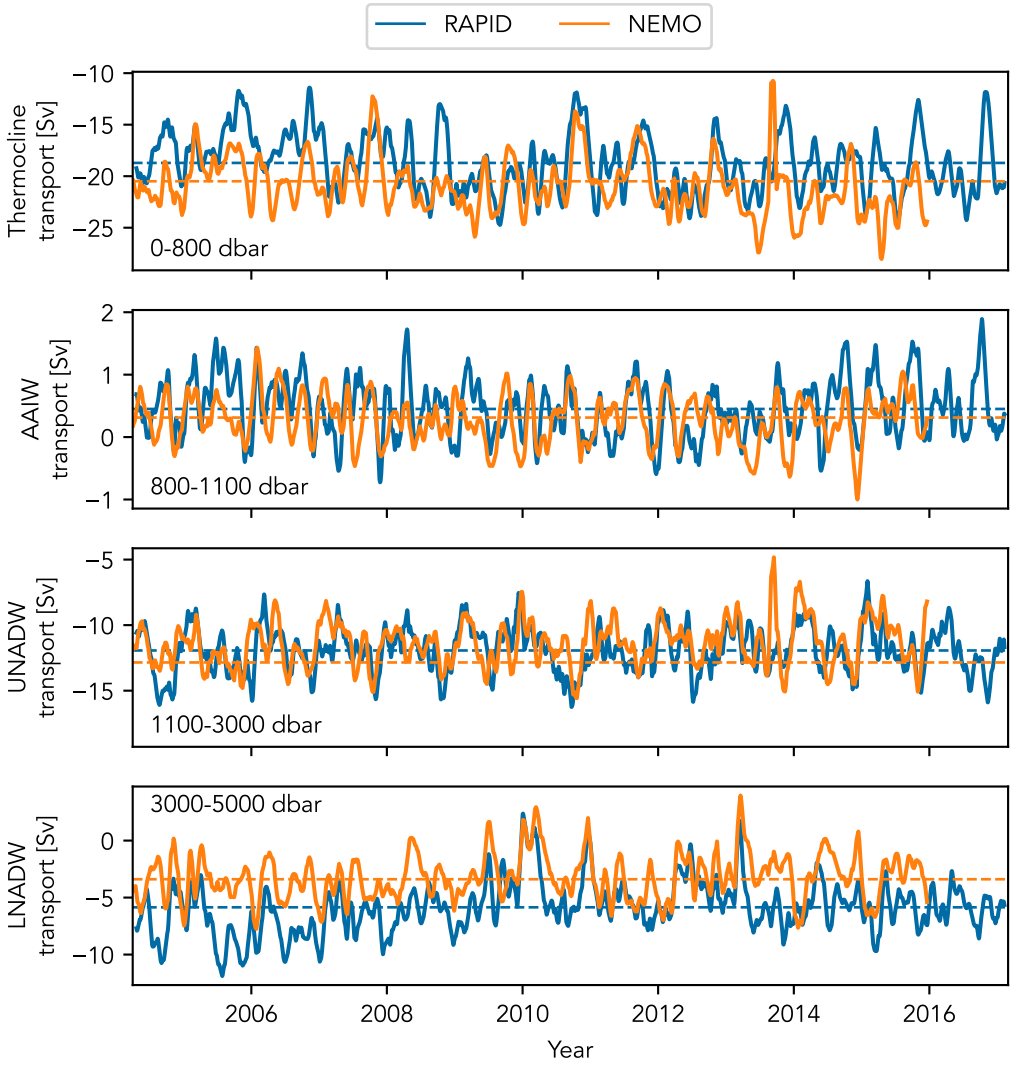

Figure 5.3: Time series of RAPID (blue) and RapidMoc NEMO (orange) transports between April 2004 and February 2017 for the thermocline (top), AAIW (upper middle), UNADW (lower middle) and LNADW (bottom) layers. The horizontal dashed lines show the time series mean for each layer transport. The NEMO data ends in December 2015.

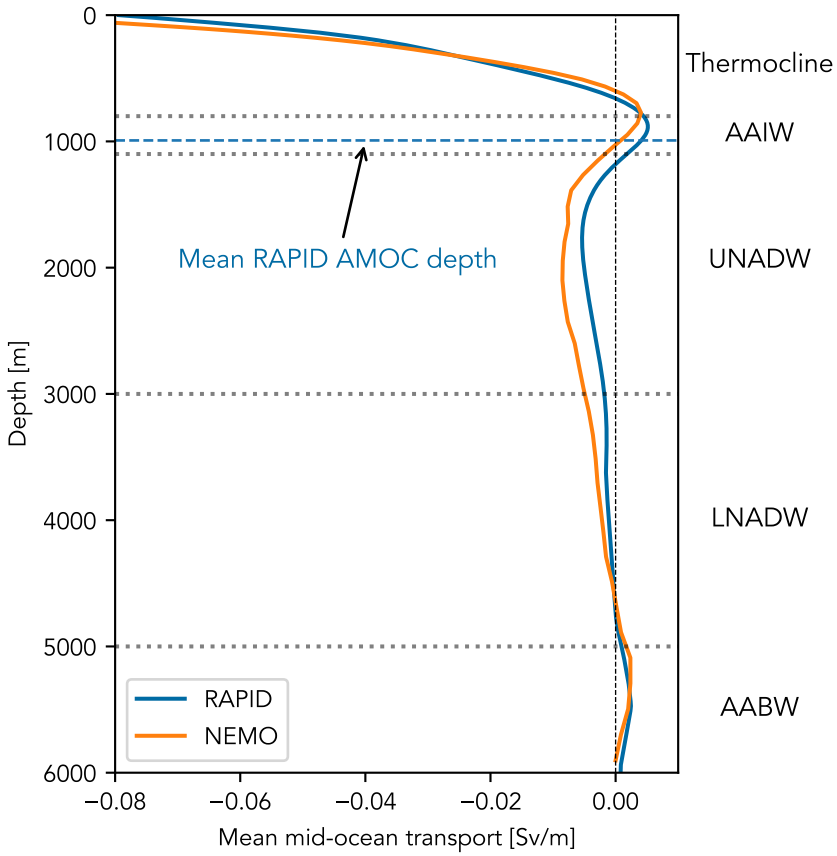

Figure 5.4: Mean mid-ocean transports between 2 April 2004 and 28 February 2017 from RAPID and calculated by the RapidMoc package from NEMO 5-day mean temperature and salinity data using methodology analogous to RAPID. The RAPID mean AMOC depth is indicated by the blue dashed line, and black dotted lines show the depth boundaries between the RAPID-defined layer transports named to the right.

these 'snapshot' anomalies, when applied to the regression model, reproduced the observed UMO. These data however were taken from RAPID mooring observations, so were very closely related to the regression model training data. Within NEMO we could reproduce the more varied sampling locations of the hydrographic data, and further constrain the regression model confidence intervals.

We replicated the CTDs by selecting model grid points as close as possible to the CTD locations given the 1/12° ORCA grid resolution, by using latitudes and longitudes from the NEMO temperature and salinity data. Fig. 5.5 shows that the majority of CTDs were located between five ORCA grid points between 26.5°N and 26.55°N. If the profile at a grid point was too shallow, then, as for the training data profiles, it was merged with the deeper section of the next grid point to the east. From the locations of the 62 replicated CTD profiles, we selected density anomalies with the same times as the original

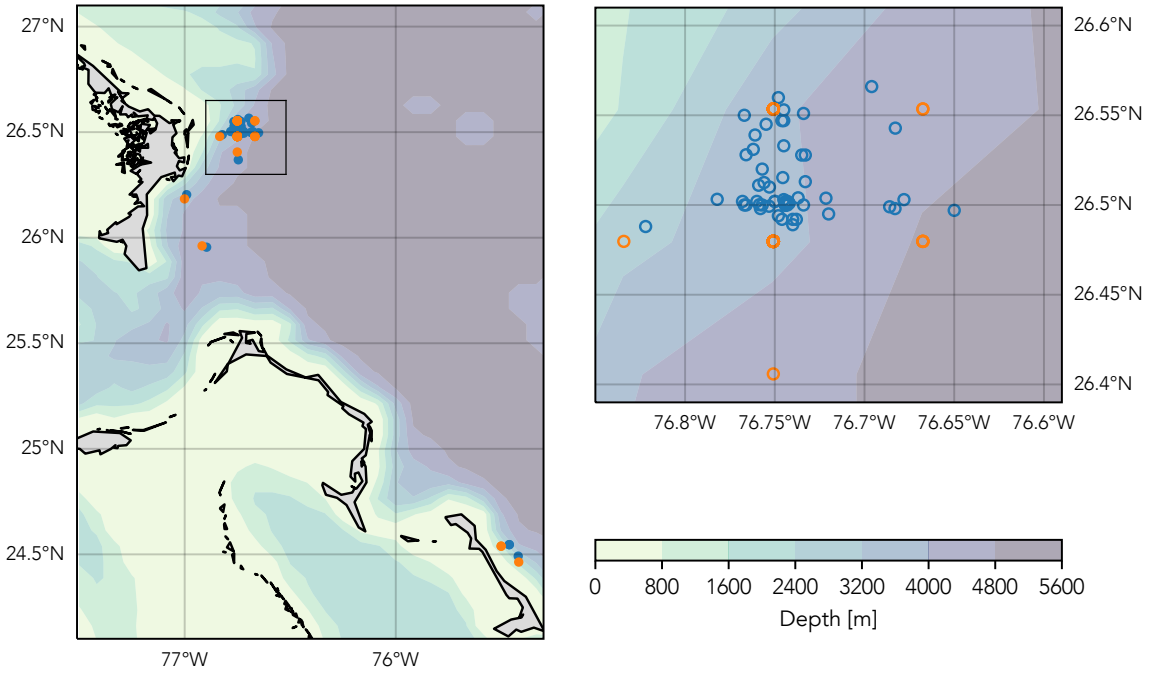

Figure 5.5: (Left) Western boundary region showing location of CTD profiles used to reconstruct the AMOC (blue circles), and the NEMO-ORCA model grid points selected to replicate them (orange circles). Contours show the ORCA version 3.3 *mesh_zgr* domain bathymetry. The black rectangle shows the area highlighted in the right-hand plot.

hydrographic profiles, and compared the UMO predicted from these with the RapidMoc UMO. Then, to estimate the uncertainty of these predictions, we selected density anomalies from the same locations but at 62 different times, randomly selected from the same time period of 1981 to 2016, and made the same regression model prediction. This was repeated 500 times to give a distribution of the errors, with the error being the difference between the predicted and RapidMoc UMO transport anomalies for each run.

5.2.2.5 Boundary profiles for other latitudes

The effectiveness of the RAPID AMOC calculations at 26°N is due in part to the steep gradient of the continental slope at the western boundary, which suppresses mesoscale eddy variability (Kanzow et al., 2009; McCarthy et al., 2015). Evans et al. (2022) found that the kinetic energy of eddies offshore of the Bahamas is dissipated by their interaction with the steep and rough topography. The last questions to be investigated using this NEMO N06 run were: firstly, whether our empirical model would still effectively reproduce the UMO transport when the density anomaly profiles were taken from latitudes other than 26°N; secondly, how far from the continental slope can density profiles be taken for the regression

model to remain effective? The first question impacts on our reconstruction of the AMOC from hydrographic data, as several of the CTDs were sampled around 24.5°N. The second question is also relevant to our CTD selection, as we assumed that density variability would increase with increasing distance from the boundary and adversely affect the model results, so chose CTDs a similar distance from the boundary as the RAPID WB2 mooring.

We created density anomaly profiles at each model grid point between approximately 41.8°N and 8.7°N, and between the western boundary (defined as where the depth became shallower than 150 m when moving westwards from the open ocean) and an eastern limit approximately 10 degrees west of this boundary. As described previously, grid points with insufficiently deep density data were merged with data at successively eastwards grid points to create full depth density anomaly profiles at each model grid point within the required region. There was no restriction made on the distance between different grid points that could be used to create a merged profile. East of The Bahamas, the steep bathymetry meant this distance was not great, but to the north, the gentler bathymetry meant that the separate profiles used to create a single merged profile had a much greater separation. The density anomaly profile at each of these model grid points was applied to the UMO regression model. The correlation between each predicted UMO transport anomaly and the RapidMoc UMO transport anomaly at 26°N was calculated using a Pearson's correlation coefficient. In addition, we calculated the correlation between the same time series smoothed with a 30-day rolling mean.

5.3 Results

5.3.1 Generalised least squares regression models

Eq. 5.1 shows the depths and coefficients selected by algorithm for the four regressions trained on NEMO monthly mean density anomalies:

$$T_{umo} = 43.9\ \rho_{wb}^{720}, \tag{5.1a}$$

$$T_{umo} = 39.2\ \rho_{wb}^{660} - 105.2\ \rho_{eb}^{1080}, \tag{5.1b}$$

$$T_{umo} = 42.4\ \rho_{wb}^{660} - 102.0\ \rho_{eb}^{1080} + 105.8\ \rho_{wb}^{1960}, \tag{5.1c}$$

$$T_{umo} = 43.6\ \rho_{wb}^{660} - 101.4\ \rho_{eb}^{1080} + 102.7\ \rho_{wb}^{1900} + 37.2\ \rho_{wb}^{3500}, \tag{5.1d}$$

where ρ is the in situ density anomaly (kg/m^3), the superscript shows the depth of the anomaly, and the subscript shows whether it is from the western (wb) or eastern (eb) boundary profile. The regressions were again generalised least squares models with auto-regressive errors of lag 1. Comparing the selected density anomaly depths with the regression equations in Eq. 3.10, the thermocline layer density anomaly depth is consistently shallower than that selected for the RAPID training data, but that the depths of the other three NEMO anomalies are all deeper.

Linear regression model	**Adj. R^2**	**SE [Sv]**	**Mean error [Sv]**	**σ error [Sv]**	**RMS error [Sv]**	**AIC**	**BIC**
Thermocline	0.366	2.03	1.78	1.25	2.17	498.4	501.2
Therm. + AAIW	0.563	1.69	1.47	1.19	1.89	456.3	461.8
Therm. + AAIW + UNADW	0.616	1.59	1.35	1.18	1.80	442.9	451.2
Therm. + AAIW + UNADW + LNADW	0.625	1.57	1.35	1.19	1.79	441.6	452.7

Table 5.2: Comparison of model statistics from the generalised least squares linear regressions in Eq. 5.1. Statistics shown are the adjusted R^2 ; standard error of the regression (SE); the mean, standard deviation (σ) and root-mean-square (RMS) of the error, where the error is the difference between the actual and model-predicted values; Akaike's Information Criterion (AIC); and Bayes Information Criterion (BIC).

The time series in Fig. 5.6 show the regression model UMO transport anomalies compared to those calculated from NEMO data by the RapidMoc package, using the RAPID dynamic height method of calculating geostrophic transport. The increase in adjusted R^2 values show that although adding the second and third density anomalies

consistently improve the model, adding the fourth increases the adjusted R^2 by only 0.009. Table 5.2 shows that although the Akaike's Information Criterion (AIC) of the fourth regression model increases slightly, the Bayes Information Criterion (BIC), which penalises the addition of spurious variables more harshly, decreases. The p-value of the fourth regressor, ρ_{wb}^{3500} in Eq. 5.1d, is 0.058, indicating it is not statistically significant at the 95% confidence level. The standard error and RMS error, very similar statistics in assessing model fit, also show diminishing benefits in adding each variable in turn.

The LNADW regression trained on NEMO density anomalies explains 68% of the variance in the RapidMoc-estimated LNADW transport anomaly, and Eq. 5.2 shows both coefficients and the depth selected for the deep density anomaly are similar to the RAPID-trained regression:

$$T_{\text{lnadw}} = -167.6\rho_{wb}^{3020} - 0.52\, T_{ekman}. \tag{5.2}$$

The adjusted R^2 value of 68% and standard error of 1.08 Sv show again that the NEMO-trained regression is not quite as good as the RAPID-trained one, but it is maybe surprisingly good considering the non-significance of the deepest density anomaly in the UMO regression. The mean prediction interval for the NEMO-trained regression is 2.2 Sv, again slightly more than the RAPID-trained equivalent of 1.9 Sv.

5.3.2 Empirical orthogonal function analysis

Fig. 5.8 shows how the standard deviation of the NEMO regression density anomaly training data varies with depth, and that the upper 1200-1300 dbar of the western boundary profile is much noisier than the RAPID training data (Fig. 3.20). Below the EOF analysis cutoff depths, both RAPID and NEMO density anomalies show two depth bands with high standard deviations, at around 500 and 800 dbar, but they are much more distinctive in the NEMO data.

Fig. 5.9 shows the variance on the western boundary is more evenly split between the first and second EOF modes for NEMO density anomalies (49.8% and 34.9%) than for RAPID (72.5% and 13.6%, Fig. 3.21). All three leading modes show more variance in the upper 1000 dbar, and the third mode shows much less variance at depth.

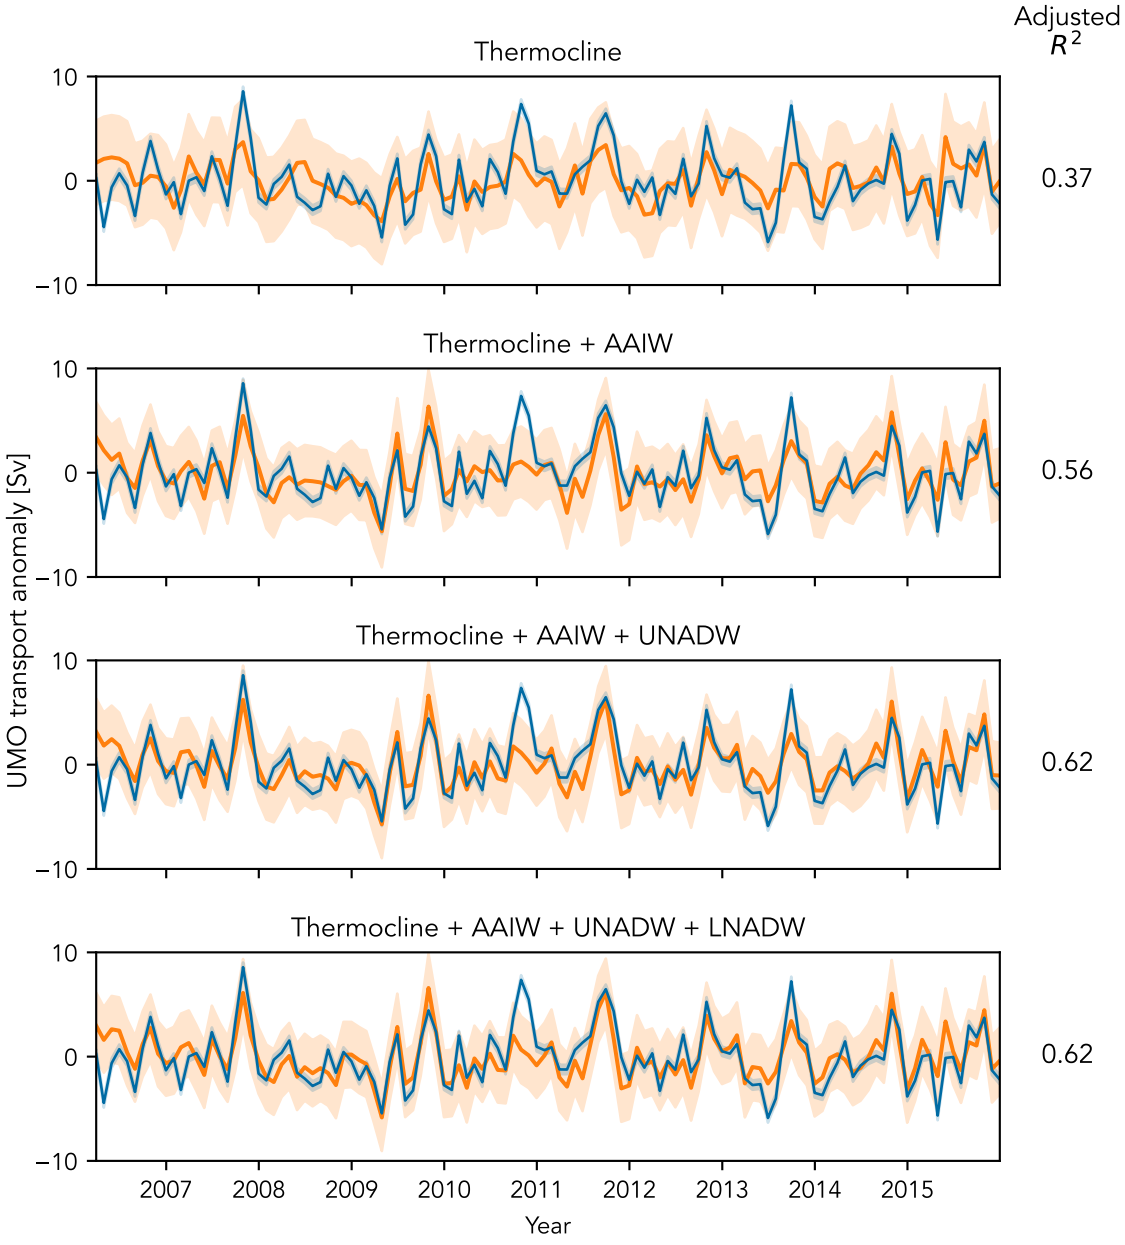

Figure 5.6: Time series of NEMO UMO transport anomalies predicted by GLSAR(1) linear regression models (orange line), and estimated by RapidMoc (blue line). The models have one, two, three and four explanatory variables – western and eastern boundary NEMO density anomalies – representing different water masses at 26°N, which are named above each plot. The model R^2 value, adjusted for the degrees of freedom of each model, is shown to the right of each plot. The orange shaded area around each model prediction line shows the 95% prediction interval of the model.

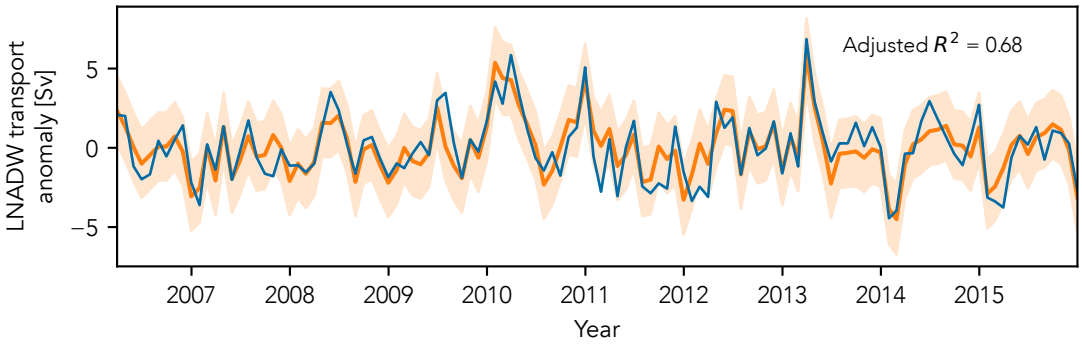

Figure 5.7: Time series of NEMO LNADW transport anomaly predicted by GLSAR(1) linear regression models (orange line), and estimated by RapidMoc (blue line). The models have two explanatory variables; a deep (3020 dbar) NEMO western boundary density anomaly, and the RapidMoc Ekman transport. The model R^2 value, adjusted for the degrees of freedom, is shown. The orange shaded area around the model prediction line shows the 95% prediction interval of the model.

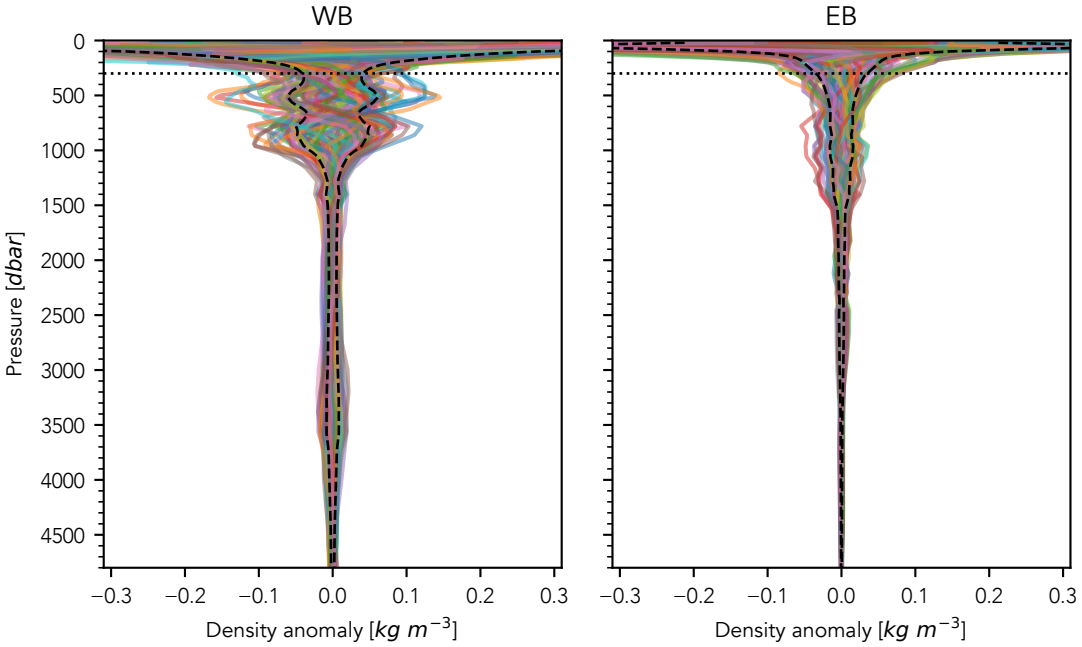

Figure 5.8: Variability of NEMO density anomalies over the full depth of the western (left) and eastern (right) profiles used to train the regressions. The standard deviation is shown by black dashed lines. the horizontal dotted line shows the upper depth limit of EOF analysis.

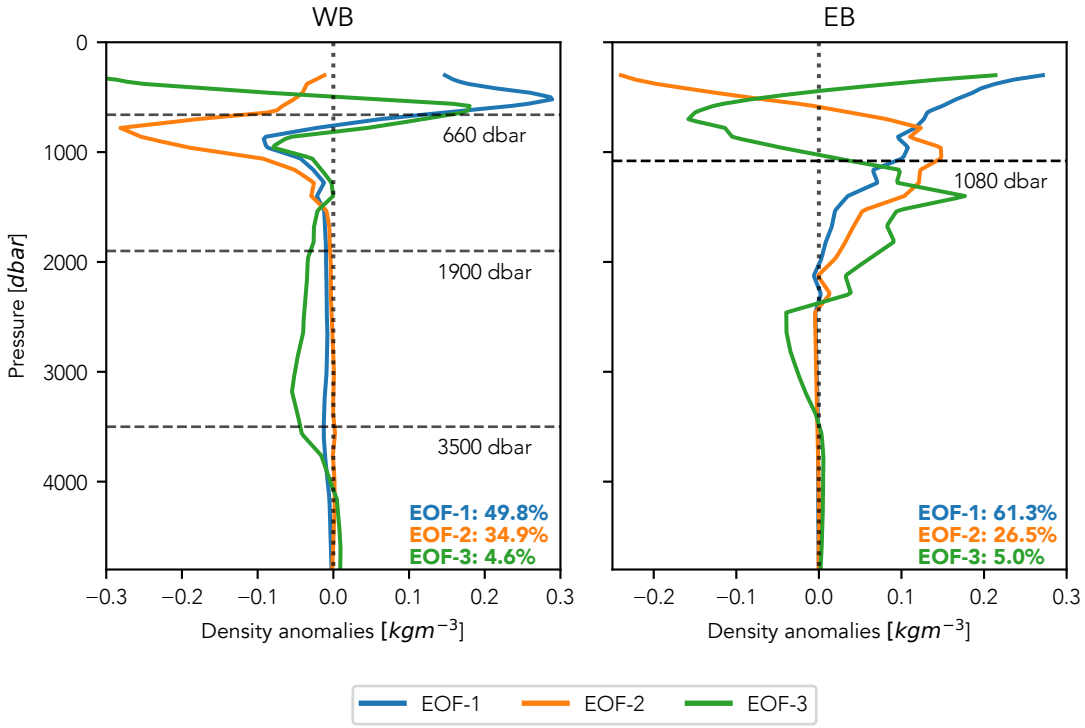

Figure 5.9: Depth profiles of NEMO density anomalies associated with the first three leading EOF modes of density variability for the western (left) and eastern (right) boundary profiles, and the variance explained by each mode. The top 300 dbar have been excluded. Horizontal dashed lines show the depths of density anomalies used in the UMO regression.

The correlation between the principal components (PC) of the first three EOF modes and the density anomalies (Fig. 5.10), does not show the same striking agreement between the depth of maximum correlation and the depth of each algorithm-selected density anomaly seen for the RAPID data (Fig. 3.22). The PC of EOF1 (PC-1) shows maximum correlation at a depth just above the 660 dbar of the shallowest western boundary density anomaly. The other two density anomalies at 1900 and 3500 dbar coincide at depths where the correlation of PC-1 and PC-3 with the density anomalies are the same, rather than at a maximum for any single mode. The correlation on the eastern boundary is more comparable to RAPID, with the depth of the selected density anomaly at the depth of maximum correlation with PC-1.

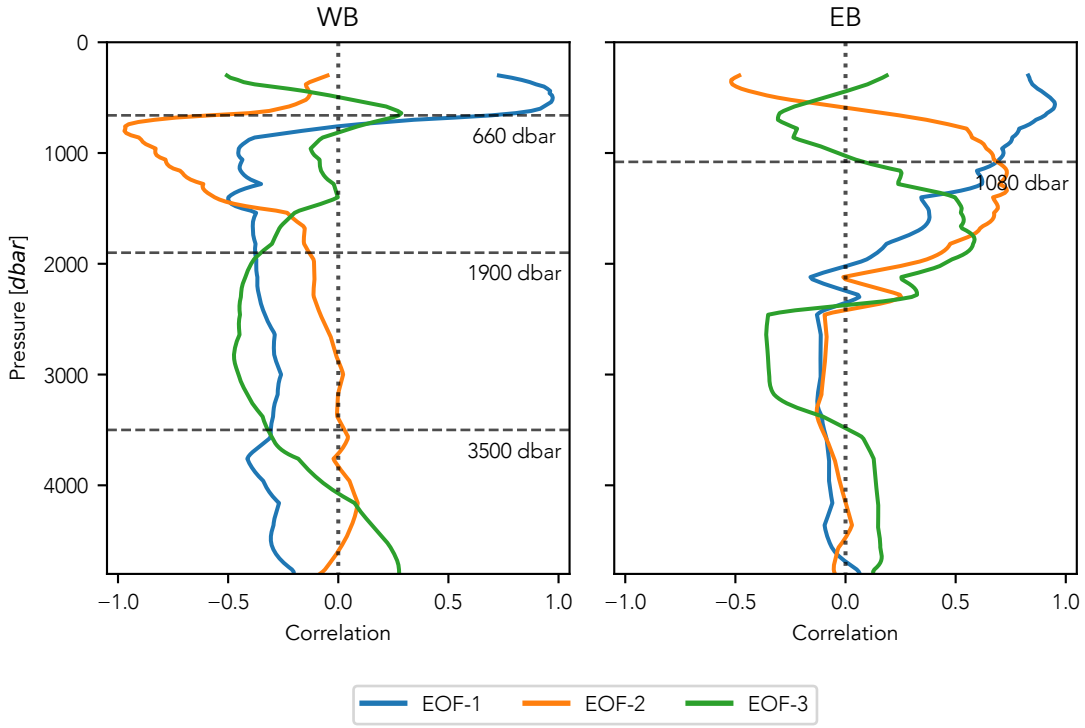

Figure 5.10: Correlation between the principal components associated with each of the three leading EOF modes, and the NEMO density anomalies at each depth for the western (left) and eastern (right) boundary profiles.

5.3.3 Comparison of regression models trained on RAPID and NEMO data

Fig. 5.11 compares the 30-day smoothed UMO transports from the RapidMoc estimations with the predictions of two regressions: one trained on RAPID data and described in Section 3.4.3.1; and one trained on NEMO data and described in Section 5.2.2.3. The time series is limited to 1980 onwards as a comparison with the AMOC reconstruction period in Chapter 4. Even though the NEMO regression explains less UMO transport anomaly variance, it is slightly more effective at reconstructing the RapidMoc UMO transport, with a Pearson's correlation coefficient of r = 0.69 between the smoothed time series. The RAPID regression, despite having been trained on observational data rather than a portion of the NEMO data, still predicts a UMO transport that correlates almost as well (r = 0.64), showing that the regression is robust and not overly sensitive to different depths being used for the density anomaly variables. The unsmoothed UMO transports predicted by each regression show reduced correlation (r = 0.57 for NEMO, r = 0.56 for

RAPID) with the RapidMoc UMO transport.

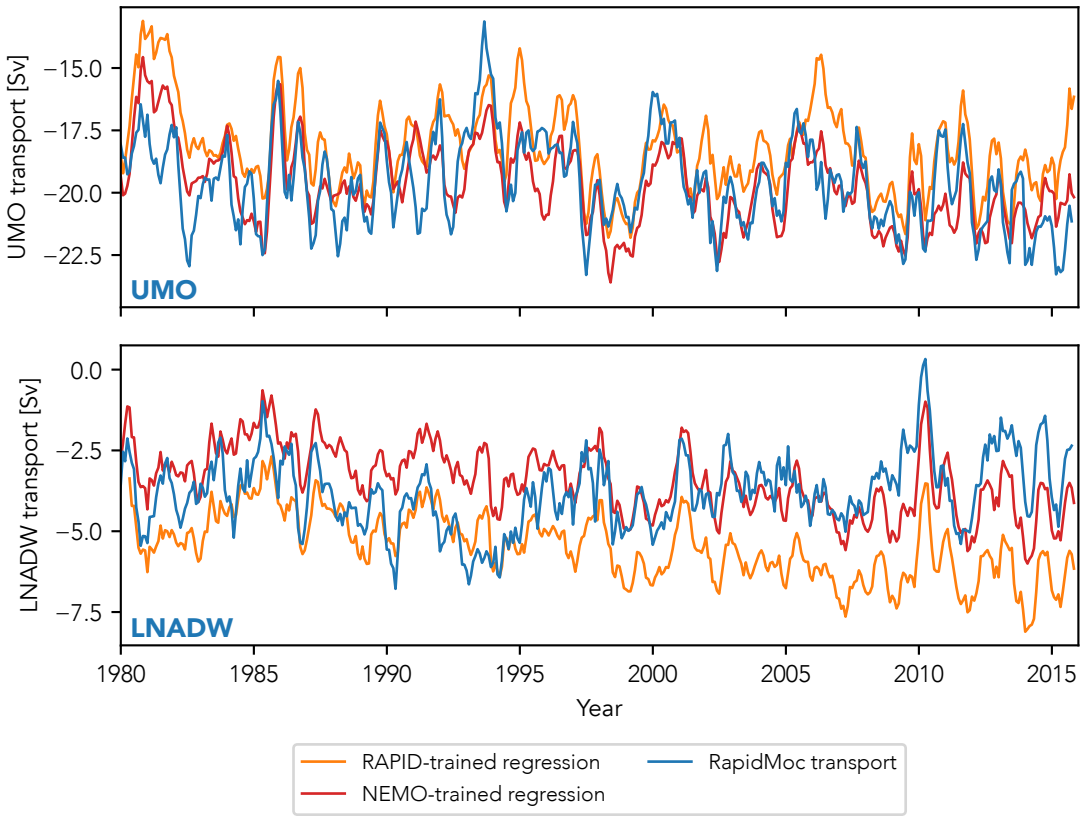

Figure 5.11: Predicted UMO (top) and LNADW (bottom) transports from two empirical models, one trained on RAPID data and described in Chapter 3 (orange) and one on NEMO data (red). The blue line is the UMO estimated from NEMO data by RapidMoc using a RAPID-style methodology. All time series are smoothed with a 30-day rolling mean.

The LNADW transport is reconstructed much less effectively by both the NEMO-trained (r = 0.35) and RAPID-trained (r = 0.27) regression models. In contrast to the UMO transports, the unsmoothed LNADW transports show higher correlation (r = 0.46 for NEMO, r = 0.40 for RAPID). This is surprising given the relatively high adjusted R^2 value of the model.

The difference in effectiveness of the two regression models is further highlighted by the full 1958-2015 transports in Fig. 5.12, which are smoothed with a 90-day rolling mean. Only the NEMO-trained regressions are shown for clarity, and show that the UMO transports are in good agreement (r = 0.69), with most of the interannual and longer period variability generally captured well, although there are periods when the model-predicted magnitude is poorly reconstructed, for example between 1972 and 1975. The agreement

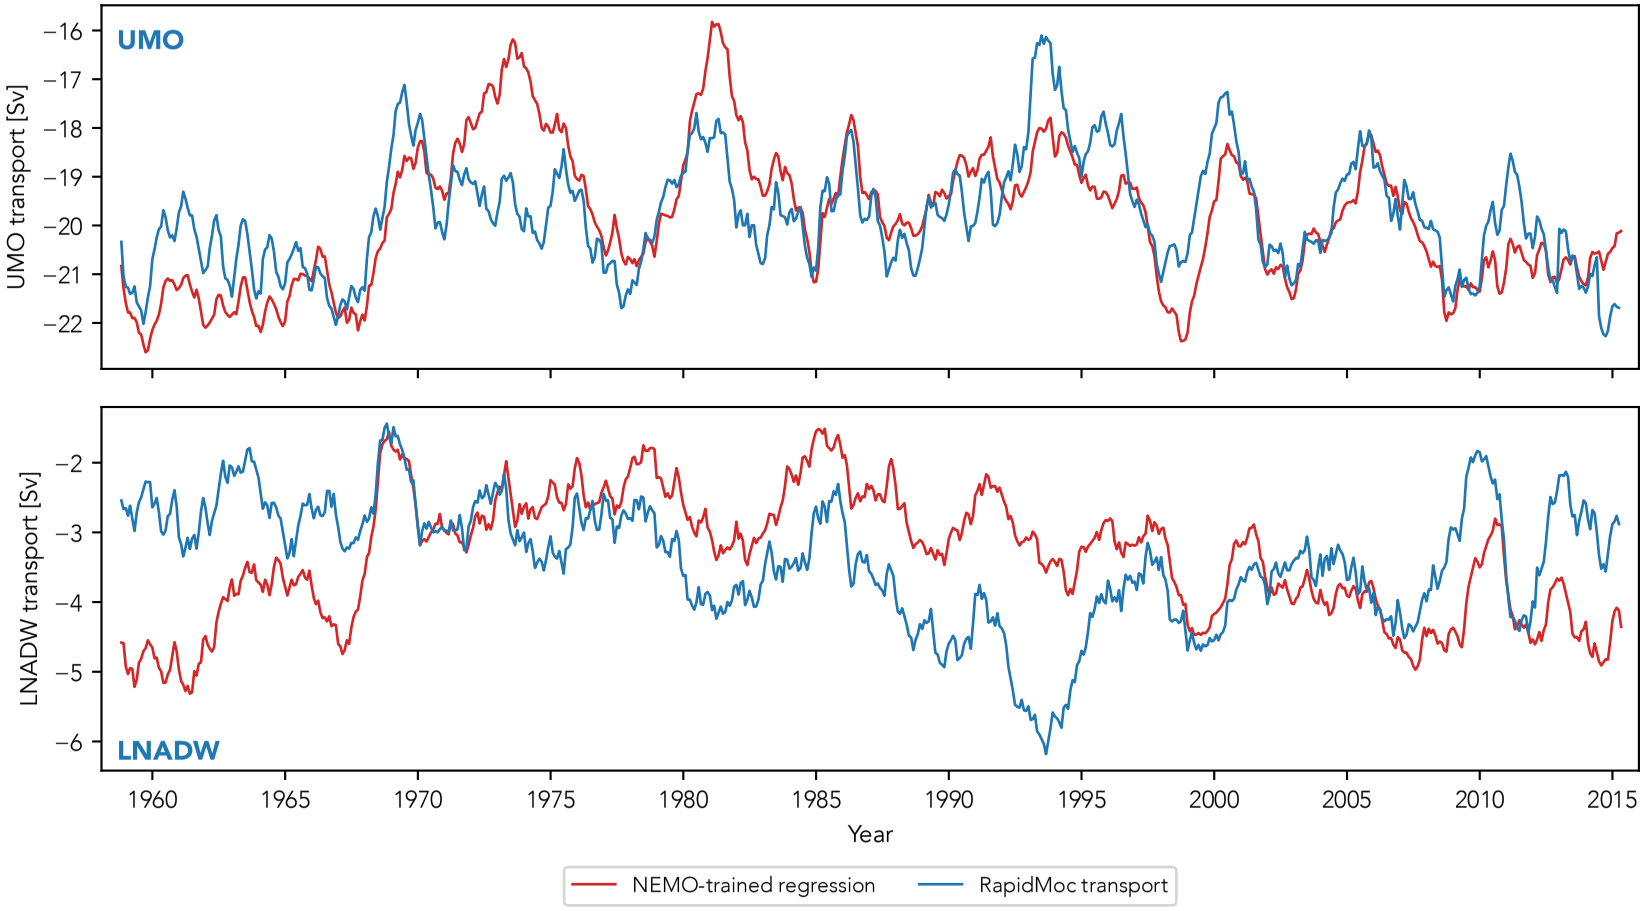

Figure 5.12: NEMO UMO (top) and LNADW (bottom) transports predicted by the regression model trained on NEMO data (red line), and estimated by RapidMoc (blue line). The full NEMO N06 run from 1958 to 2015 was used, and all time series were smoothed with a 90-day rolling mean.

for the 30-day smoothed time series is very similar for both the 1958-2015 and 1980-2015 time series (r = 0.66 and r = 0.69 respectively), and Fig. 5.13 shows that although the rolling four-year Pearson's correlation coefficient varies over time, there is no overall trend. This linear regression is robust over a period much longer than the training data. In contrast, the agreement between the 90-day smoothed LNADW transports drops to 0.11 for the 1958-2015 period. This poor performance meant the LNADW regression was not analysed any further within the NEMO ocean model. We had enough confidence in the UMO regression, however, to use it to ask questions that could not be answered using observational data.

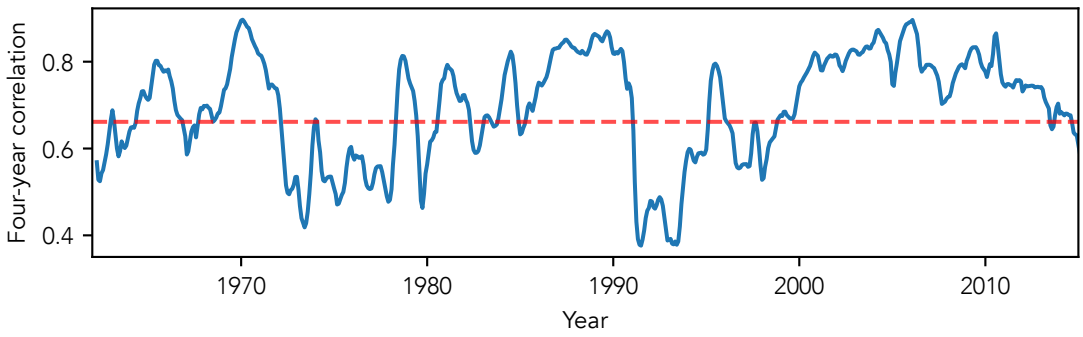

Figure 5.13: Rolling four-year Pearson's correlation coefficient between the 30-day rolling mean UMO transports predicted by the NEMO-trained regression and the RapidMoc estimate, for 1958-2015. The dashed red line shows the mean correlation for the whole time series (r = 0.69)

5.3.4 Uncertainty of using CTD profiles to predict monthly AMOC

One such question was asked in Section 3.4.4. If a regression model was trained on monthly mean data, then how well do 'snapshots' such as individual CTD profiles represent the monthly mean state? The uncertainty in annual mean estimates of the AMOC that are based upon single hydrographic sections has been estimated as ±6 Sv, and is in part due to the effects of eddy variability (Ganachaud, 2003; Longworth, 2007).

The uncertainty in monthly mean predicted UMO transports was estimated from the bootstrapped time-random replica NEMO CTD profiles created in Section 5.2.2.4. The error was the difference between the UMO transport anomaly predicted by the regression model from the density anomalies of each replica CTD profile, and the monthly mean RapidMoc UMO transport anomaly at 26°N. Fig. 5.14 shows that although the

distribution of these errors is close to normal, the peak and the mean do not coincide due a positive bias of 1.5 Sv. The root mean square error (RMSE) is 3.2 Sv, which is much greater than the RMSE of 0.9 (±0.45) Sv for 10-day geostrophic transports found by McCarthy et al. (2015), but less than the ±6 Sv error in upper and lower transports quoted by Bryden et al. (2005) and due to sampling in or out of eddies. The mean monthly range for 10-day filtered RAPID observations between April 2004 and September 2018 is 9.5 Sv, so the RMSE and even 2 standard deviations (5.6 Sv) are much less than this mean variability. The cause of the positive error bias remains unclear.

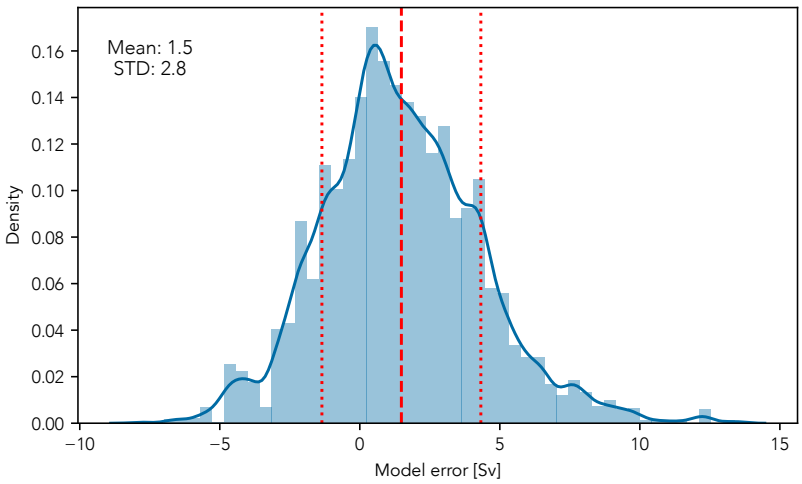

Figure 5.14: Distribution of regression model error (difference between predicted and actual UMO transport anomaly) for 62 replica hydrographic CTD profiles at 500 randomly selected times between 1980 and 2015. The vertical red lines are the mean (dashed) and one standard deviation (dotted).

Given that the regression model performs better when trained on observations than on NEMO data, and that some of the 3.2 Sv RMSE will come from regression model error rather than sampling issues, we would expect the confidence intervals for the AMOC reconstruction from observed CTD profiles to be within ±3.2 Sv.

5.3.5 Meridional coherence of the UMO regression model

Fig. 5.15 shows how the UMO regression model performs as the latitude and longitude of the density anomaly data applied to it changes. The correlation at each grid point is between the density profile created at that grid point (by merging if necessary)

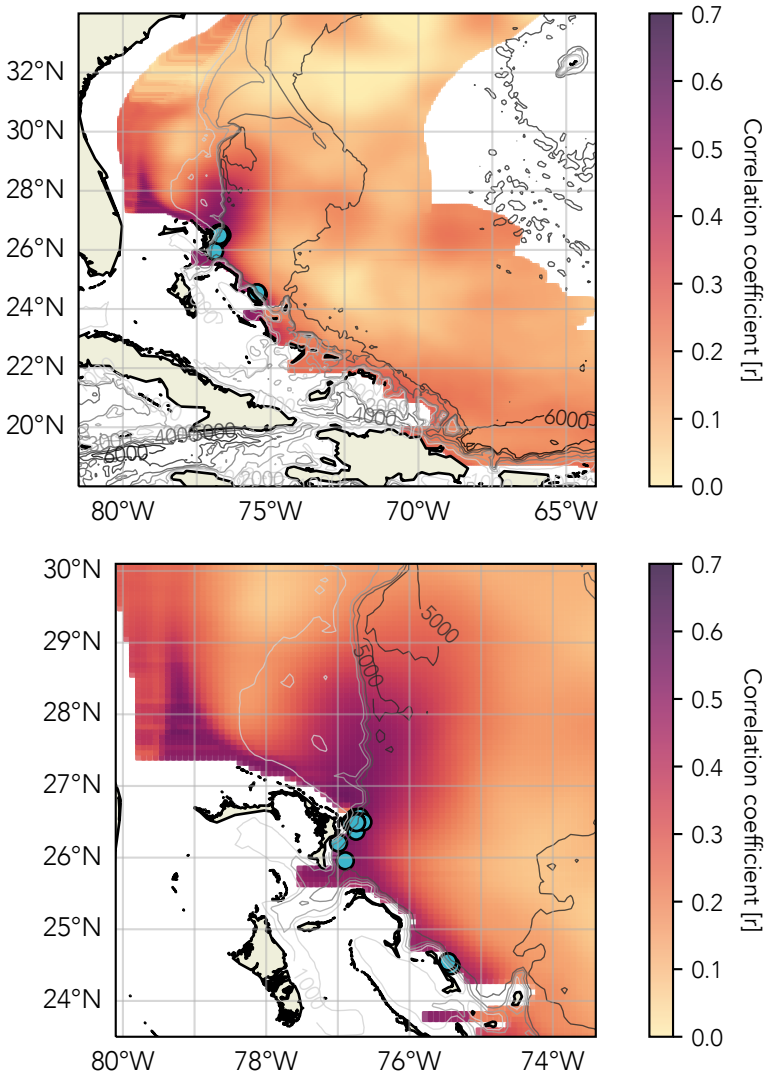

Figure 5.15: (Top) Correlation between the RapidMoc estimated UMO transport anomaly at 26°N, and the UMO transport anomaly predicted by empirical model for the density anomaly profiles at each ORCA grid point between 18°N and 34°N, and approximately within 10° of the western boundary. Correlation (r) is the Pearson's correlation coefficient between the unsmoothed time series. The light blue circles show the positions of the hydrographic data profiles used to reconstruct the AMOC in Chapter 4. ORCA-12 V3.3 bathymetry is shown by isobaths every 1000 m. (Bottom) The same but focused on the western boundary region near 26.5°N.

and the NEMO UMO transport anomaly at 26°N, calculated using the RapidMoc RAPID-analogous method. The correlation r is the Pearson's correlation coefficient between the two monthly mean 1958-2015 time series. From 26.5°N, the correlation remains higher up to over 200 km north and south along the steeper bathymetry. To the north, the correlation decreases rapidly after around 28.5°N, where the bathymetry becomes more complex eastwards of the steeper slope. To the south, the correlation remains higher along the slope until the latitude approaches 24°N, where there are islands and their associated bathymetry to the east.

Moving longitudinally from 26.5°N, the higher level of correlation is lost around 50 km to the east. The band of higher correlation remains narrow along the steep continental slope to the east of The Bahamas, but as the slope stretches northwards of 27°N, it broadens, getting closer to 100 km in width. Part of the broadening is due to the existence of profiles to the west of the steep bathymetry, rather than land, but the correlation also drops more slowly eastwards of the slope.

The narrowing triangular plume of high correlation north of 27.5°N at around 79°W may correspond to the exit of the Florida Current from the Florida Straits. As this region is shallower than 1000 m, the lower sections of the profiles in this region were taken from the east of the continental slope, a distance of up to 300 km from the position shown by the correlation.

5.4 Conclusions

Our results showed that the UMO linear regression performed well in predicting the UMO transport when applied to NEMO data, and justified using NEMO to ask questions that we could not with the observational data. The regression model fit was better when trained on NEMO boundary density anomalies, but was surprisingly good when trained on RAPID observational data. The LNADW regression model however performed poorly, and may suggest that the NEMO 1/12° model does not represent changes in the deep return limb of the overturning circulation particularly well. The poor performance of the LNADW model meant further investigation of it within NEMO was abandoned. The UMO model was however effective enough to allow us to quantify the uncertainty associated with using individual and isolated CTD profiles to predict a monthly mean UMO transport. A root

mean square error of 3.2 Sv was slightly more than the 2.7 Sv obtained from more recent RAPID observations (Section 3.4.4), but included the simulation of different CTD locations. Since the RAPID regression model is a better fit than the NEMO equivalent, we would expect its uncertainty to be less than 3.2 Sv.

We also used NEMO to assess whether our assumptions in selecting the locations of CTD profiles were correct, and if we could extend the selection of hydrographic data for use in the regression models. It showed that the regression model performed well for profiles close to the steep bathymetry of the continental slope, so vindicated the use of distance from the boundary in selecting CTD locations. The width of the higher correlation band along the slope varies between around 50 and 100 km, with steeper bathymetry corresponding to a narrower band. This is in agreement with the decrease in variability within 100 km of the Bahamas shelf found in an eddy-permitting model by Kanzow et al. (2009).

Profiles selected along the bathymetry between 24°N and 28°N, when applied to the UMO regression model, predicted UMO transports that correlated well with the 26°N UMO transport, suggesting that the density anomalies at the model depths show meridional coherence over at least 200 km. Polo et al. (2020) showed in a lower resolution ocean general circulation model (NEMO 1°) that density anomalies at 2000-3000 m propagate southwards along the western boundary at an average speed of 0.3 m/s. For a distance of 200 km this would equate to around a week, however anomaly propagation speeds are dependent on model resolution (Getzlaff, 2005). The small spatial and temporal scale of high correlation seen in our results suggests that our model is mostly representing the higher frequency AMOC variability associated with wind forcing. This small region of coherence may nevertheless allow us to select observational data to apply to the model from an expanded region compared to Chapter 4, and increase the resolution of our reconstruction further.

Chapter 6

Discussion and Future Work

6.1 Discussion

The importance of the Atlantic meridional overturning circulation to the transport of ocean heat, freshwater and carbon, and thus to global climate and carbon storage has been well documented (Gruber et al., 2019; Kostov et al., 2014; McDonagh et al., 2015; Sabine et al., 2004; Winton et al., 2013). This importance motivated the deployment in 2004 of the RAPID permanent mooring array at 26°N in the North Atlantic. The resulting observations have provided a wealth of information on AMOC variability over a range of timescales (Cunningham et al., 2007; Kanzow et al., 2010; McCarthy et al., 2012), however the length of the RAPID time series remains too short to distinguish a possible climate change-related decline from natural decadal variability (Roberts et al., 2014; Robson et al., 2014; Smeed et al., 2014). Estimates of the length of time series required to detect low frequency changes are as long as 40 to 60 years (Baehr et al., 2008; Roberts and Palmer, 2012), but fingerprinting methods suggest that including deep signals, such as changes in deep density gradients, reduces this time significantly (Baehr et al., 2007; Vellinga and Wood, 2004).

Although direct AMOC observations are high resolution and have low uncertainty, their brevity means most estimates of the AMOC prior to 2004 are provided by proxy reconstruction methods. These proxies can be approximately divided into implicitly 'single-layer' reconstructions of 20-30 years that may not capture all the observed variability (Caesar et al., 2018; Frajka-Williams, 2015; e.g., Longworth et al., 2011; Rahmstorf et al., 2015; Willis, 2010; Zhang, 2008); and reconstructions from tens to thousands of years long with high uncertainty and, for the most part, not based on hydrographic observations (e.g., Caesar et al., 2021; Thornalley et al., 2018).

This thesis addressed the following questions. Firstly, could the AMOC at 26°N, with its different water masses and changing direction of flow with depth, be represented by a linear regression trained on RAPID boundary data? Secondly, if this were successful, could the regression be used to reconstruct the AMOC from pre-2004 hydrographic data and capture lower frequency changes more effectively than 'single-layer' models, and with higher resolution and lower uncertainty than palaeoclimatic reconstructions? Lastly, could this method of reconstruction be applied to hydrographic data away from the western boundary at 26°N, and how would this affect the associated uncertainty?

To answer the first question, we assessed a number of different linear regressions trained on RAPID observations. Following the principle of parsimony, we started with the simplest ordinary least squares linear model, informed by Longworth et al. (2011), and regressed a western boundary temperature anomaly on the upper mid-ocean transport anomaly. Iteratively increasing the complexity of the models, our algorithmic process found the best fit was supplied by a multiple linear regression combining four boundary density anomalies, each representing a RAPID depth-defined layer: thermocline; AAIW; UNADW; and LNADW. As is common with linear regressions based on time series data, our regression models failed the assumption of having no autocorrelation of residuals, showing significant correlation between errors with a lag of one month. The ordinary least squares regression was thus replaced by a generalised least squares regression that modelled the first-order auto-correlative error process. We also created an additional model that regressed a combination of a deep (3040 dbar) density anomaly and Ekman transport on the LNADW transport anomaly. These two models explained 79% of UMO and 74% of LNADW transport anomaly variances respectively. This was a considerable improvement on the 53% variance of thermocline transport anomaly that the simple linear regression in Longworth et al. (2011) found, and showed that multiple linear regression models could be used to effectively reconstruct the UMO and LNADW transports from historical hydrographic data.

Although the results of Chapter 3 showed that adding the two deepest western boundary density anomalies – representing changes in the upper and lower NADW layers – improved the adjusted R^2 measure of regression model fit, the improvements were only slight. This naturally begs the question, why include them at all? To address this, we initially showed

that adding these two explanatory variables had an effect significantly different from adding statistically similar noisy signals. The regression coefficient p-values for each of these variables also indicated that the relationships between them and the UMO transport anomaly response variable were significant. The coefficients of each of the three western boundary density anomaly variables are of similar magnitude, so changes to the shallow and deep densities have a similar effect on the upper mid-ocean transport within the model. In addition, the regression's variance inflation factors proved that the deep variables were independent of each other and those representing the shallower layers, while other statistics of model fit also showed significant improvement when they were included, as did the mean prediction interval of the regression, giving increased confidence in the model's predicted values. Smeed et al. (2014) showed that the LNADW had a larger relative change in transport associated with the AMOC 2004-2012 decline than the UMO, while our coherence analysis showed that the 3000 dbar (LNADW) density anomaly was significantly coherent with the UMO transport anomaly for periods of 1 year and longer, highlighting its importance in lower frequency changes. EOF analysis of boundary density anomalies with depth also showed the strongest correlation between the principal components of the second and third leading modes, and the density anomalies at the depths of the UNADW and LNADW variables. Finally, Baehr et al. (2007) showed that by including deep density differences in a fingerprinting method, the detection time of a climate change-related AMOC decline was halved. The fingerprinting method uses properties that are not directly affected by global warming, such as sub-surface properties. Our method also has the advantage over reconstruction methods using sea-surface temperatures, as our sub-surface densities are relatively isolated from direct atmospheric effects.

What are the limitations of this reconstruction method? We have argued for the inclusion of variables representing the upper and lower NADW by citing Baehr et al. (2007), but whereas they used cross-basin density differences, we have included only deep density anomalies on the western boundary. We have assumed little variability in deep eastern boundary densities, since AMOC variability is dominated by the western boundary (Bryden et al., 2009; Elipot et al., 2014), and eastern boundary variability is mainly observed at intermediate depths (Chidichimo et al., 2010; Pérez-Hernández et al., 2015). Lastly, there are the uncertainties related to the regression model itself, which will be described following the discussion of Chapter 4 results.

Chapter 4 addressed our second question – could a successful regression model reconstruct the pre-2004 AMOC, and would this reconstruction improve on other proxy methods? Our empirical models require only measurements of temperature, salinity and pressure or depth, so potentially could reconstruct the AMOC back decades, and certainly longer than methods that rely on satellite data and/or glider technology. In practice, we found data within our defined western boundary region only between 1981 and 2016. The sampling of the 62 merged CTD profiles created from this hydrographic data was sufficient to reconstruct the AMOC with approximately annual resolution, although some periods (e.g., the 1980s) were much better sampled than others (e.g., the 1990s). This is sufficient to show pentadal and decadal AMOC variability, and detect the AMOC recovery during the late 1990s that is observed in other reconstructions (Caesar et al., 2018; Thornalley et al., 2018) and may be linked to the weakened sub-polar gyre circulation in the 1990s (Häkkinen and Rhines, 2004). It also contradicts ensemble mean of coupled-climate CMIP5 and CMIP6 model historical runs that show an AMOC increase during the 1960s and a decrease from the 1990s onwards (Menary et al., 2020). Our model results have not revealed a long-term AMOC decline during this period indicative of anthropogenic climate change (Stocker et al., 2013), but this is in agreement with many reconstructions that also show a steep decline between the 1950s and 1980s (Caesar et al., 2021).

Our AMOC reconstruction has accurately reproduced the multi-year variability observed in the RAPID data, and showed that the 2008-2012 downturn (McCarthy et al., 2012) was not only the weakest AMOC since RAPID observations began, but the weakest since the mid-1980s. Our model also captures the recent recovery in AMOC strength observed by RAPID (Moat et al., 2020). Comparison with Frajka-Williams (2015) single-layer sea-surface height reconstruction shows that only by including all four variables in our regression do we see the expected decadal variability in upper mid-ocean transport.

Although our model shows the importance of including deep density measurements in AMOC reconstructions, shown to be key in detecting anthropogenic AMOC change over shorter timescales (Baehr et al., 2007), our 35-year reconstruction is only just greater than the 30 years required by their fingerprinting method. The temporal resolution of our AMOC reconstruction, while an improvement on paleoclimatic methods, remains much coarser than RAPID observations. There are periods where the sampling resolution is

sub-annual (e.g., late 1980s to early 1990s), however there are others where it is much larger, for example the data gap between February 1998 and April 2001, and the rolling four-year means for these sections will be more influenced by outlying data points.

Chapters 3 and 4 of this thesis generated a number of different estimates of error and uncertainty. The standard error of a multiple linear regression model describes how well the model fits the data, analogous to how well a simple regression line fits scattered data points. Our UMO model had a standard error of 1.3 Sv, and our LNADW model had a standard error of 0.9 Sv. In contrast, the prediction interval of a regression model quantifies the uncertainty of its predicted values. For example, the UMO regression model will predict the same UMO transport anomaly for a particular combination of density anomalies, but it also provides an associated prediction interval, where we have 95% confidence that the actual UMO transport anomaly falls with this range. The prediction interval varies for each combination of explanatory variables, but the mean prediction intervals for the RAPID training data were 2.4 Sv for the UMO model, with a range of 0.2 Sv, and 1.9 Sv for the LNADW model, with a range of 0.1 Sv. Testing the UMO model against more recent RAPID 10-day filtered data gave another opportunity to quantify the uncertainty for data with a shorter timescale. The RMS error between the RAPID-observed and model-predicted UMO transports was 2.16 Sv, which increased slightly to 2.20 Sv when the eastern boundary density anomaly was replaced with a climatology. The mean prediction interval for these data was again 2.4 Sv.

To quantify the uncertainty in applying the models to single CTD profiles, we bootstrapped random samples from the newer, independent 10-day RAPID data to simulate individual profiles. The normally-distributed errors had a RMS of 2.7 Sv. Finally, applying the model to hydrographic data allowed direct comparison with overlapping RAPID observations. The 28 data points from the overlapping period showed an RMS error of 2.0 Sv, and a mean prediction interval of 2.45 Sv. The mean model prediction intervals remained between 2.4 and 2.5 Sv whether applied to monthly mean, 10-day or single profile data. The RMS errors are between 2.0 and 2.2 Sv except for the bootstrapping results of 2.7 Sv. Data with shorter timescales and fewer data points have lower RMS errors. Given these results, the maximum uncertainty estimate of ± 2.7 Sv is conservative.

The success of this regression model in reconstructing the UMO and AMOC transports raises the question of whether it could be used as an alternative to the current RAPID mooring array. The RAPID estimates have an RMS error of ±0.9 Sv for annual transports and ±1.5 Sv for 10-day filtered transports (McCarthy et al., 2015). A ±6 Sv uncertainty has been associated with using single data points to determine the mean annual AMOC (Atkinson et al., 2012; Ganachaud, 2003; Longworth, 2007). Our conservative estimate of uncertainty is ±2.7 Sv, less than half of ±6 Sv, but still nearly double the ±1.5 Sv for the 10-day RAPID transports. We could assume that model predictions going forward would have similar uncertainty, but this depends on the linear regression not changing over time. The results of Chapter 5 show that the UMO model is stable over 58 years in an ocean simulation, so this is not an unreasonable assumption. Since the main expense of maintaining the RAPID array is in servicing via ship, having only three instruments on the main WB2 mooring would not give much of a cost advantage over the current method. As of early 2022, the array has already been reduced in scope by the removal of the Mid-Atlantic Ridge and deep eastern boundary moorings. Where the regression method may prove useful is if the eastern boundary moorings fail, or it is decided to no longer maintain them, although this also would depend on the eastern boundary density climatology not changing significantly over time. Our results show that whether transports are predicted from monthly mean data, 10-day filtered data or single hydrographic profiles, we are confident the estimated transport is within ±2.7 Sv of the predicted value. However, the higher resolution data currently provided by RAPID would be reduced to monthly means, and the increased uncertainty over the current ±1.5 Sv for 10-day transports would need a longer time series to distinguish low frequency changes from natural variability. However if RAPID is ever replaced, for example if ocean reanalyses become sufficiently reliable and accurate to calculate the geostrophic components, then hydrography-based model reconstructions could provide an additional constraint.

The sparseness of hydrographic data limited further investigations into our reconstruction method. To answer our final questions about the uncertainty inherent in the method, and about the influence of hydrographic data sampling location, we moved to a high resolution, eddy-resolving ocean model, NEMO-ORCA 1/12°, using a hindcast run between 1958 and 2015. The UMO regression model did not perform as well when trained and tested on NEMO data as the models in Chapter 3, but was good enough to justify

further investigation. The LNADW regression model on the other hand performed poorly, despite reasonably good model fit statistics, which may suggest that this NEMO model does not represent the lower limb of the overturning circulation as well as the upper. This was reflected in the EOF leading modes, which showed over 89% of the density anomaly variance was concentrated in the upper 1300 m of the water column, in contrast to the RAPID data, where the third leading mode had a maximum at 3000 dbar. We used the UMO model within NEMO to further constrain the uncertainty associated with predictions based on 'snapshot' CTD profiles, and showed the root mean square error was 3.2 Sv, although the errors showed a slight positive bias. This was a slight increase on the 2.7 Sv estimated in Chapter 3, but given that the fit of the NEMO-trained model was not as good as for the RAPID-trained model, 2.7 Sv remains a reasonable estimate of uncertainty for our AMOC reconstruction.

We also showed that our criteria for selecting hydrographic data were justified, both the limits of the defined region and the distance from the steep western boundary. The correlation between UMO transport at 26°N and that predicted from density anomalies was highest when those anomalies were selected from model grid points within approximately 50 km of the boundary. This is in good agreement with the sharp decline of sea-surface height variability towards the boundary observed by Kanzow et al. (2009) in both an eddy-resolving model and satellite altimetry. The higher correlation between UMO transports was also shown to persist along the steep bathymetry at least 200 km north and south of 26.5°N, and showed that any additional available hydrographic data from within this higher correlation region could contribute to our AMOC reconstruction. The relatively small scale meridional coherence of density anomalies that we found agrees with other general circulation model (GCM) studies that show wind-forced, higher frequency variability processes overwhelm coherent signals of low frequency buoyancy forcing (e.g., Bingham et al., 2007; Polo et al., 2020, 2014). Overall, the results of assessing this regression method within NEMO suggests that it would be a useful first step if the method were to be applied to other regions.

In summary, this thesis has successfully created the longest continuous AMOC reconstruction based on both shallow and deep sub-surface hydrographic data, and has created an independent AMOC index to go alongside those such as Caesar et al.

(2018). The length of this AMOC reconstruction approaches the length of time series for which a climate change-related decline is detectable. It is longer than current RAPID observations and satellite altimetry-based reconstructions; has much lower uncertainty than palaeoclimatic proxy reconstructions; and by uniquely including deep density anomalies, increases the likelihood of detecting low frequency changes in the AMOC. This reconstruction bridges the gap between short, low uncertainty observations and long, high uncertainty proxy reconstructions, and shows good agreement with sea-surface temperature indices. The resolution of our reconstruction is sufficient to show multi-year and decadal changes, and shows an AMOC recovery in the late 1990s and downturn around 2010 seen in other reconstructions and observations, and an apparent recovery in recent years in line with observations (Moat et al., 2020). Finally, we showed that there was no overall AMOC decline for the period between 1981 and 2016.

6.2 Future work

Our NEMO ocean GCM investigations showed we could select hydrographic data from a slightly enlarged spatial area, which would be an easy and obvious next step. We could also extend our reconstruction temporally by using data from the International Geophysical Year of 1957/58, and possibly even earlier, although we would need to account for uncertainties related to different salinity calculation methods.

In terms of improving the model, applying machine learning would be one option, especially using techniques that are not limited to modelling linear processes. We co-supervised the thesis of Tasneem Ahmed, a Masters student who applied a non-linear random forest technique to our regression data, and although this improved the model marginally, her results were qualitatively similar. This approach could be extended to use machine learning techniques beyond simple algorithms to select the depths of density anomalies, for example. Applying machine learning methods to ocean or coupled-climate models before more limited hydrographic data may highlight an alternate method that would be more effective. Using GCMs would also allow easier investigation of similar methods over different timescales, for example regression models trained on annual mean anomalies to reduce the influence of wind-forced variability.

Future studies could apply this method – using temperature and salinity mooring data

to train a transport regression model – to other locations and mooring datasets, such as OSNAP. Care would need to be taken with how the water mass layers are represented by independent variables within the regression equation, as isopycnals are much more steeply sloped at the latitude of the OSNAP array compared to 26°N. Transport calculations are typically made in density rather than depth space as a result. The OSNAP array is also divided into western and eastern arrays, either side of Greenland. The majority of variability is associated with the eastern array, suggesting that its data would be more suitable for a regression model. The utility of an OSNAP regression model would depend on how many historical CTD data are available from the region to reconstruct transport variability. The results of Chapter 5 suggest a feasibility study using a high resolution ocean or climate GCM may be worthwhile in defining the locations of suitable hydrographic data.

Another study could evaluate the model's performance within a more sophisticated coupled-climate ensemble, such as CMIP6. This was beyond the scope of this study, but would make an interesting CMIP6 AMOC proxy to add to those more commonly used such as sea-surface temperature fingerprints and sea level on the North American seaboard. It would also provide more insight into how the CMIP6 simulations model low-frequency deep density changes, compared with the surface/single-layer proxies described above.

In summary, this thesis has proven the worth of this method in extending observational time series, and provided a foundation for its adaptation by future studies in other locations, and a basis for improved machine learning methods.

Bibliography

Aitken, A.C., 1936. IV.—On Least Squares And Linear Combination of Observations. Proc. R. Soc. Edinb. 55, 42–48. https://doi.org/10.1017/S0370164600014346

Atkinson, C.P., Bryden, H.L., Cunningham, S.A., King, B.A., 2012. Atlantic transport variability at 25° N in six hydrographic sections. Ocean Sci. 8, 497–523. https://doi.org/10.5194/os-8-497-2012

Baehr, J., 2011. Detecting changes in the meridional overturning circulation at 26°N in the Atlantic from observations. Deep Sea Res. Part II Top. Stud. Oceanogr. 58, 1833–1836. https://doi.org/10.1016/j.dsr2.2010.10.062

Baehr, J., Haak, H., Alderson, S., Cunningham, S.A., Jungclaus, J.H., Marotzke, J., 2007. Timely Detection of Changes in the Meridional Overturning Circulation at 26°N in the Atlantic. J. Clim. 20, 5827–5841. https://doi.org/10.1175/2007JCLI1686.1

Baehr, J., Keller, K., Marotzke, J., 2008. Detecting potential changes in the meridional overturning circulation at 26°N in the Atlantic. Clim. Change 91, 11–27. https://doi.org/10.1007/s10584-006-9153-z

Baringer, M.O., Larsen, J.C., 2001. Sixteen years of Florida current transport at 27°N. Geophys. Res. Lett. 28, 3179–3182. https://doi.org/10.1029/2001GL013246

Baringer, M.O., Molinari, R., 1999. Atlantic Ocean baroclinic heat flux at 24 to 26° N. Geophys. Res. Lett. 26, 353–356. https://doi.org/10.1029/1998GL900323

Bendat, J.S., Piersol, A.G., 2010. Random Data: Analysis and Measurement Procedures, Wiley Series in Probability and Statistics. John Wiley & Sons, Inc., Hoboken, NJ, USA. https://doi.org/10.1002/9781118032428

Biastoch, A., Böning, C.W., Lutjeharms, J.R.E., 2008. Agulhas leakage dynamics affects decadal variability in Atlantic overturning circulation. Nature 456, 489–492. https://doi.org/10.1038/nature07426

Bingham, R.J., Hughes, C.W., Roussenov, V., Williams, R.G., 2007. Meridional coherence of the North Atlantic meridional overturning circulation. Geophys. Res. Lett. 34. https://doi.org/10.1029/2007GL031731

Boyer, T.P., Baranova, O.K., Coleman, C., Garcia, H.E., Grodsky, A., Locarnini, R.A.,

Mishonov, A.V., O'Brien, T.D., Paver, C.R., Reagan, J.R., Seidov, D., Smolyar, I.V., Weathers, K., Zweng, M.M., Antonov, J.I., Johnson, D.R., 2018. NCEI Standard Product: World Ocean Database (WOD).

Breusch, T.S., Pagan, A.R., 1979. A Simple Test for Heteroscedasticity and Random Coefficient Variation. Econometrica 47, 1287. https://doi.org/10.2307/1911963

Broecker, W.S., Peteet, D.M., Rind, D., 1985. Does the ocean–atmosphere system have more than one stable mode of operation? Nature 315, 21–26. https://doi.org/10.1038/315021a0

Broeker, W., 1991. The Great Ocean Conveyor. oceanog 4, 79–89. https://doi.org/10.5670/oceanog.1991.07

Bryden, H.L., King, B.A., McCarthy, G.D., McDonagh, E.L., 2014. Impact of a 30% reduction in Atlantic meridional overturning during 2009–2010. Ocean Sci. 10, 683–691. https://doi.org/10.5194/os-10-683-2014

Bryden, H.L., Longworth, H.R., Cunningham, S.A., 2005. Slowing of the Atlantic meridional overturning circulation at 25° N. Nature 438, 655–657. https://doi.org/10.1038/nature04385

Bryden, H.L., Mujahid, A., Cunningham, S.A., Kanzow, T., 2009. Adjustment of the basin-scale circulation at 26 N to variations in Gulf Stream, deep western boundary current and Ekman transports as observed by the Rapid array. Ocean Sci. 5, 421–433. https://doi.org/10.5194/osd-6-871-2009

Buckley, M.W., Marshall, J., 2016. Observations, inferences, and mechanisms of the Atlantic Meridional Overturning Circulation: A review. Rev. Geophys. 54, 5–63. https://doi.org/10.1002/2015RG000493

Caesar, L., McCarthy, G.D., Thornalley, D.J.R., Cahill, N., Rahmstorf, S., 2021. Current Atlantic Meridional Overturning Circulation weakest in last millennium. Nat. Geosci. https://doi.org/10.1038/s41561-021-00699-z

Caesar, L., Rahmstorf, S., Robinson, A., Feulner, G., Saba, V., 2018. Observed fingerprint of a weakening Atlantic Ocean overturning circulation. Nature 556, 191–196. https://doi.org/10.1038/s41586-018-0006-5

Casanova-Masjoan, M., Pérez-Hernández, M.D., Vélez-Belchí, P., Cana, L., Hernández-Guerra, A., 2020. Variability of the Canary Current Diagnosed by Inverse Box Models. J. Geophys. Res. Oceans 125. https://doi.org/10.1029/2020JC016199

Center, N.G.D., 2006. ETOPO2v2 2-minute Global Relief Model. https://doi.org/10.7289/V5J1012Q

Chatterjee, S., Simonoff, J.S., 2013. Handbook of Regression Analysis. John Wiley & Sons, Inc., Hoboken, NJ, USA.

Chidichimo, M.P., Kanzow, T., Cunningham, S.A., Johns, W.E., Marotzke, J., 2010. The contribution of eastern-boundary density variations to the Atlantic meridional overturning circulation at 26.5°N. Ocean Sci. 6, 475–490. https://doi.org/10.5194/os-6-475-2010

Cochrane, D., Orcutt, G.H., 1949. Application of Least Squares Regression to Relationships Containing Auto-Correlated Error Terms. Journal of the American Statistical Association 44, 32–61. https://doi.org/10.1080/01621459.1949.10483290

Cunningham, S.A., 2005. A transatlantic hydrography section at 24.5°N: RRS Discovery Cruise D279, 04 April - 10 May 2004. Southampton Oceanography Centre Cruise Report 54, Southampton, UK 199.

Cunningham, S.A., Kanzow, T., Rayner, D., Baringer, M.O., Johns, W.E., Marotzke, J., Longworth, H.R., Grant, E.M., Hirschi, J.J.-M., Beal, L.M., Meinen, C.S., Bryden, H.L., 2007. Temporal Variability of the Atlantic Meridional Overturning Circulation at 26.5°N. Science (80-.). 317, 935–938.

Cunningham, S.A., Roberts, C.D., Frajka-Williams, E., Johns, W.E., Hobbs, W., Palmer, M.D., Rayner, D., Smeed, D.A., McCarthy, G., 2013. Atlantic Meridional Overturning Circulation slowdown cooled the subtropical ocean. Geophys. Res. Lett. 40, 6202–6207. https://doi.org/10.1002/2013GL058464

Dawson, A., 2016. Eofs: A Library for EOF Analysis of Meteorological, Oceanographic, and Climate Data. Journal of Open Research Software 4, e14. https://doi.org/10.5334/jors.122

Durbin, J., Watson, G.S., 1950. Testing for Serial Correlation in Least Squares Regression: I. Biometrika 37, 409. https://doi.org/10.2307/2332391

Elipot, S., Frajka-Williams, E., Hughes, C., Willis, J., 2014. The Observed North Atlantic Meridional Overturning Circulation: Its Meridional Coherence and Ocean Bottom Pressure. J. Phys. Oceanogr. 44, 517–537. https://doi.org/10.1175/JPO-D-13-026.1

Evans, D.G., Frajka-Williams, E., Naveira Garabato, A.C., 2022. Dissipation of mesoscale eddies at a western boundary via a direct energy cascade. Sci Rep 12, 887. https://doi.

org/10.1038/s41598-022-05002-7

Frajka-Williams, E., 2015. Estimating the Atlantic overturning at 26°N using satellite altimetry and cable measurements. Geophys. Res. Lett. 42, 3458–3464. https://doi.org/10.1002/2015GL063220

Frajka-Williams, E., Lankhorst, M., Koelling, J., Send, U., 2018. Coherent Circulation Changes in the Deep North Atlantic From 16°N and 26°N Transport Arrays. J. Geophys. Res. Oceans 123, 3427–3443. https://doi.org/10.1029/2018JC013949

Frajka-Williams, E., Meinen, C.S., Johns, W.E., Smeed, D.A., Duchez, A., Lawrence, A.J., Cuthbertson, D.A., McCarthy, G.D., Bryden, H.L., Baringer, M.O., Moat, B.I., Rayner, D., 2016. Compensation between meridional flow components of the Atlantic MOC at 26°N. Ocean Sci. 12, 481–493. https://doi.org/10.5194/os-12-481-2016

Fraser, N.J., Cunningham, S.A., 2021. 120 Years of AMOC Variability Reconstructed From Observations Using the Bernoulli Inverse. Geophysical Research Letters 48. https://doi.org/10.1029/2021GL093893

Fuglister, F.C., 1960. Atlantic Ocean atlas of temperature and salinity profiles and data from the International Geophysical Year of 1957-1958.

Ganachaud, A., 2003. Error Budget of Inverse Box Models: The North Atlantic. J. Atmos. Oceanic Technol. 20, 1641–1655. https://doi.org/10.1175/1520-0426(2003)020%3C1641:EBOIBM%3E2.0.CO;2

Ganopolski, A., Rahmstorf, S., 2001. Rapid changes of glacial climate simulated in a coupled climate model. Nature 409, 153–158. https://doi.org/10.1038/35051500

Getzlaff, J., 2005. Signal propagation related to the North Atlantic overturning. Geophys. Res. Lett. 32, L09602. https://doi.org/10.1029/2004GL021002

Gordon, A.L., 1986. Interocean exchange of thermocline water. J. Geophys. Res. 91, 5037. https://doi.org/10.1029/JC091iC04p05037

Graybill, F.A., 1976. Theory and application of the linear model. Duxbury Press, North Scituate, Mass.

Gruber, N., Clement, D., Carter, B.R., Feely, R.A., van Heuven, S., Hoppema, M., Ishii, M., Key, R.M., Kozyr, A., Lauvset, S.K., Lo Monaco, C., Mathis, J.T., Murata, A., Olsen, A., Perez, F.F., Sabine, C.L., Tanhua, T., Wanninkhof, R., 2019. The oceanic sink for anthropogenic CO_2 from 1994 to 2007. Science 363, 1193–1199. https://doi.org/10.1126/science.aau5153

Guerra, L.A.A., Paiva, A.M., Chassignet, E.P., 2018. On the translation of Agulhas rings to the western South Atlantic Ocean. Deep Sea Research Part I: Oceanographic Research Papers 139, 104–113. https://doi.org/10.1016/j.dsr.2018.08.005

Häkkinen, S., Rhines, P.B., 2004. Decline of Subpolar North Atlantic Circulation During the 1990s. Science 304, 555–559. https://doi.org/10.1126/science.1094917

Hallam, S., Marsh, R., Josey, S.A., Hyder, P., Moat, B., Hirschi, J.J.-M., 2019. Ocean precursors to the extreme Atlantic 2017 hurricane season. Nat Commun 10, 896. https://doi.org/10.1038/s41467-019-08496-4

Hernández-Guerra, A., Espino-Falcón, E., Vélez-Belchí, P., Dolores Pérez-Hernández, M., Martínez-Marrero, A., Cana, L., 2017. Recirculation of the Canary Current in fall 2014. Journal of Marine Systems 174, 25–39. https://doi.org/10.1016/j.jmarsys.2017.04.002

Hirschi, J.J.-M., Lynch-Stieglitz, J., 2006. Ocean margin densities and paleoestimates of the Atlantic meridional overturning circulation: A model study. Geochemistry, Geophys. Geosystems 7. https://doi.org/10.1029/2006GC001301

Hughes, C.W., Williams, J., Blaker, A., Coward, A., Stepanov, V., 2018. A window on the deep ocean: The special value of ocean bottom pressure for monitoring the large-scale, deep-ocean circulation. Prog. Oceanogr. 161, 19–46. https://doi.org/10.1016/j.pocean.2018.01.011

Jackett, D.R., Mcdougall, T.J., 1995. Minimal Adjustment of Hydrographic Profiles to Achieve Static Stability. J. Atmos. Oceanic Technol. 12, 381–389. https://doi.org/10.1175/1520-0426(1995)012%3C0381:MAOHPT%3E2.0.CO;2

Johns, W.E., Baringer, M.O., Beal, L.M., Cunningham, S.A., Kanzow, T., Bryden, H.L., Hirschi, J.J.M., Marotzke, J., Meinen, C.S., Shaw, B., Curry, R., 2011. Continuous, array-based estimates of Atlantic ocean heat transport at 26.5°N. J. Clim. 24, 2429–2449. https://doi.org/10.1175/2010JCLI3997.1

Johns, W.E., Beal, L.M., Baringer, M.O., Molina, J.R., Cunningham, S.A., Kanzow, T., Rayner, D., 2008. Variability of shallow and deep western boundary currents off the Bahamas during 2004–05: Results from the 26°N RAPID–MOC Array. J. Phys. Oceanogr. 38, 605–623. https://doi.org/10.1175/2007JPO3791.1

Josey, S.A., Kent, E.C., Taylor, P.K., 2002. Wind Stress Forcing of the Ocean in the SOC Climatology: Comparisons with the NCEP–NCAR, ECMWF, UWM/COADS, and

Hellerman and Rosenstein Datasets. J. Phys. Oceanogr. 32, 1993–2019. https://doi.org/10.1175/1520-0485(2002)032%3C1993:WSFOTO%3E2.0.CO;2

Kanzow, T., Cunningham, S.A., Johns, W.E., Hirschi, J.J.-M., Marotzke, J., Baringer, M.O., Meinen, C.S., Chidichimo, M.P., Atkinson, C., Beal, L.M., Bryden, H.L., Collins, J., 2010. Seasonal Variability of the Atlantic Meridional Overturning Circulation at 26.5°N. J. Clim. 23, 5678–5698. https://doi.org/10.1175/2010JCLI3389.1

Kanzow, T., Johnson, H.L., Marshall, D.P., Cunningham, S.A., Hirschi, J.J.-M., Mujahid, A., Bryden, H.L., Johns, W.E., 2009. Basinwide Integrated Volume Transports in an Eddy-Filled Ocean. J. Phys. Oceanogr. 39, 3091–3110. https://doi.org/10.1175/2009JPO4185.1

Kendall, M.G., 1948. Rank Correlation Methods. Charles Griffin and Co. Ltd, 42 Drury Lane, London.

Koenker, R., 1981. A note on studentizing a test for heteroscedasticity. Journal of Econometrics 17, 107–112. https://doi.org/10.1016/0304-4076(81)90062-2

Koltermann, K., Jancke, V.V., Gouretski, K.P., 2011. Volume 3: Atlantic Ocean. International WOCE Project Office, Southampton, UK. https://doi.org/10.21976/C6RP4Z

Kostov, Y., Armour, K.C., Marshall, J., 2014. Impact of the Atlantic meridional overturning circulation on ocean heat storage and transient climate change. Geophys. Res. Lett. 41, 2108–2116. https://doi.org/10.1002/2013GL058998

Lee, T.N., Johns, W.E., Zantopp, R.J., Fillenbaum, E.R., 1996. Moored Observations of Western Boundary Current Variability and Thermohaline Circulation at 26.5° in the Subtropical North Atlantic. Journal of Physical Oceanography 26, 962–983. https://doi.org/10.1175/1520-0485(1996)026%3C0962:MOOWBC%3E2.0.CO;2

Longworth, H.R., 2007. Constraining Variability of the Atlantic Meridional Overturning Circulation at 25°N from Historical Observations, 1980 to 2005. (PhD thesis). University of Southampton.

Longworth, H.R., Bryden, H.L., Baringer, M.O., 2011. Historical variability in Atlantic meridional baroclinic transport at 26.5°N from boundary dynamic height observations. Deep. Res. Part II 58, 1754–1767. https://doi.org/10.1016/j.dsr2.2010.10.057

Lorenz, E.N., 1956. Empirical Orthogonal Functions and Statistical Weather Prediction.

Lozier, M.S., De Jong, M.F., Li, F., Bacon, S., Holliday, N.P., Bower, A.S., Pickart, R.S.,

Straneo, F., Weller, R.A., Yang, J., Zhao, J., Cunningham, S.A., Gary, S.F., Houpert, L., Inall, M.E., De Steur, L., De Young, B., Fischer, J., Karstensen, J., Greenan, B.J.W., Heimbmbach, P., Johns, W.E., Johnson, H.L., Marshall, D.P., Lin, X., Mackay, N., Williams, R.G., Mercier, H., Thierry, V., Myers, P.G., Pillar, H.R., Wilson, C., Zika, J.D., 2017. Overturning in the Subpolar north Atlantic program: A new international ocean observing system. Bull. Am. Meteorol. Soc. 98, 737–752. https://doi.org/10.1175/BAMS-D-16-0057.1

Lozier, M.S., Li, F., Bacon, S., Bahr, F., Bower, A.S., Cunningham, S.A., de Jong, M.F., de Steur, L., deYoung, B., Fischer, J., Gary, S.F., Greenan, B.J.W., Holliday, N.P., Houk, A., Houpert, L., Inall, M.E., Johns, W.E., Johnson, H.L., Johnson, C., Karstensen, J., Koman, G., Le Bras, I.A., Lin, X., Mackay, N., Marshall, D.P., Mercier, H., Oltmanns, M., Pickart, R.S., Ramsey, A.L., Rayner, D., Straneo, F., Thierry, V., Torres, D.J., Williams, R.G., Wilson, C., Yang, J., Yashayaev, I., Zhao, J., 2019. A sea change in our view of overturning in the subpolar North Atlantic. Science 363, 516–521. https://doi.org/10.1126/science.aau6592

Lozier, S., 2011. Atlantic Meridional Overturning. Ann. Rev. Mar. Sci. 4, 110301100421037. https://doi.org/10.1146/annurev-marine-120710-100740

Madec, G., and the NEMO team, 2008. NEMO ocean engine. Note du Pôle de modélisation, Institut Pierre-Simon Laplace (IPSL), France 27.

Mann, H.B., 1945. Nonparametric Tests Against Trend. Econometrica 13, 245. https://doi.org/10.2307/1907187

Mann, M.E., Steinman, B.A., Brouillette, D.J., Miller, S.K., 2021. Multidecadal climate oscillations during the past millennium driven by volcanic forcing. Science 371, 1014–1019. https://doi.org/10.1126/science.abc5810

McCarthy, G.D., Menary, M.B., Mecking, J.V., Moat, B.I., Johns, W.E., Andrews, M.B., Rayner, D., Smeed, D.A., 2017. The importance of deep, basinwide measurements in optimized Atlantic Meridional Overturning Circulation observing arrays. J. Geophys. Res. Ocean. 122, 1808–1826. https://doi.org/10.1002/2016JC012200

McCarthy, G.D., Smeed, D.A., Johns, W.E., Frajka-Williams, E., Moat, B.I., Rayner, D., Baringer, M.O., Meinen, C.S., Collins, J., Bryden, H.L., 2015. Measuring the Atlantic Meridional Overturning Circulation at 26°N. Prog. Oceanogr. 130, 91–111. https://doi.org/10.1016/j.pocean.2014.10.006

McCarthy, G., Frajka-Williams, E., Johns, W.E., Baringer, M.O., Meinen, C.S., Bryden, H.L., Rayner, D., Duchez, A., Roberts, C., Cunningham, S.A., 2012. Observed interannual variability of the Atlantic meridional overturning circulation at 26.5°N. Geophys. Res. Lett. 39. https://doi.org/10.1029/2012GL052933

McDonagh, E.L., King, B.A., Bryden, H.L., Courtois, P., Szuts, Z., Baringer, M., Cunningham, S.A., Atkinson, C., McCarthy, G., 2015. Continuous Estimate of Atlantic Oceanic Freshwater Flux at 26.5°N. J. Clim. 28, 8888–8906. https://doi.org/10.1175/JCLI-D-14-00519.1

McKinney, W., Perktold, J., Seabold, S., 2019. Time Series Analysis in Python with statsmodels. Proceedings of the 10th Python in Science Conference.

Meinen, C.S., Baringer, M.O., Garcia, R.F., 2010. Florida Current transport variability: An analysis of annual and longer-period signals. Deep. Res. Part I 57, 835–846.

Meissner, K.J., 2007. Younger Dryas: A data to model comparison to constrain the strength of the overturning circulation. Geophys. Res. Lett. 34, L21705. https://doi.org/10.1029/2007GL031304

Menary, M.B., Robson, J., Allan, R.P., Booth, B.B.B., Cassou, C., Gastineau, G., Gregory, J., Hodson, D., Jones, C., Mignot, J., Ringer, M., Sutton, R., Wilcox, L., Zhang, R., 2020. Aerosol-Forced AMOC Changes in CMIP6 Historical Simulations. Geophys. Res. Lett. 47. https://doi.org/10.1029/2020GL088166

Moat, B.I., Josey, S.A., Sinha, B., Blaker, A.T., Smeed, D.A., McCarthy, G.D., Johns, W.E., Hirschi, J.J.-M., Frajka-Williams, E., Rayner, D., Duchez, A., Coward, A.C., 2016. Major variations in subtropical North Atlantic heat transport at short (5 day) timescales and their causes. J. Geophys. Res. Oceans 121, 3237–3249. https://doi.org/10.1002/2016JC011660

Moat, B.I., Smeed, D.A., Frajka-Williams, E., Desbruyères, D.G., Beaulieu, C., Johns, W.E., Rayner, D., Sanchez-Franks, A., Baringer, M.O., Volkov, D., Bryden, H.L., 2020. Pending recovery in the strength of the meridional overturning circulation at 26°N. Ocean Science 16. https://doi.org/10.5194/os-2019-134

Parrilla, G., Lavín, A., Bryden, H., García, M., Millard, R., 1994. Rising temperatures in the subtropical North Atlantic Ocean over the past 35 years. Nature 369, 48–51. https://doi.org/10.1038/369048a0

Patricola, C.M., Chang, P., Saravanan, R., 2015. Impact of Atlantic SST and high

frequency atmospheric variability on the 1993 and 2008 Midwest floods: Regional climate model simulations of extreme climate events. Climatic Change 129, 397–411. https://doi.org/10.1007/s10584-013-0886-1

Percival, D.B., Walden, A.T., 1998. Spectral Analysis for Physical Applications : Multitaper and Conventional Univariate Techniques. Cambridge ; New York : Cambridge University Press, 1998.

Pérez-Hernández, M.D., McCarthy, G.D., Vélez-Belchí, P., Smeed, D.A., Fraile-Nuez, E., Hernández-Guerra, A., 2015. The Canary Basin contribution to the seasonal cycle of the Atlantic Meridional Overturning Circulation at 26°N. J. Geophys. Res. Ocean. 120, 7237–7252. https://doi.org/10.1002/2015JC010969

Petit, T., Lozier, M.S., Josey, S.A., Cunningham, S.A., 2020. Atlantic Deep Water Formation Occurs Primarily in the Iceland Basin and Irminger Sea by Local Buoyancy Forcing. Geophys. Res. Lett. 47. https://doi.org/10.1029/2020GL091028

Pickart, R.S., Spall, M.A., Ribergaard, M.H., Moore, G.W.K., Milliff, R.F., 2003. Deep convection in the Irminger sea forced by the Greenland tip jet. Nature 424, 152–156. https://doi.org/10.1038/nature01729

Polo, I., Haines, K., Robson, J., Thomas, C., 2020. Can the boundary profiles at 26° N be used to extract buoyancy-forced Atlantic Meridional Overturning Circulation signals? Ocean Sci. 16, 1067–1088. https://doi.org/10.5194/os-16-1067-2020

Polo, I., Robson, J., Sutton, R., Balmaseda, M.A., 2014. The Importance of Wind and Buoyancy Forcing for the Boundary Density Variations and the Geostrophic Component of the AMOC at 26°N. Journal of Physical Oceanography 44, 2387–2408. https://doi.org/10.1175/JPO-D-13-0264.1

Prerau, M.J., Brown, R.E., Bianchi, M.T., Ellenbogen, J.M., Purdon, P.L., 2017. Sleep Neurophysiological Dynamics Through the Lens of Multitaper Spectral Analysis. Physiology 32, 60–92. https://doi.org/10.1152/physiol.00062.2015

Rahmstorf, S., 2002. Ocean circulation and climate during the past 120,000 years. Nature 419, 207–214. https://doi.org/10.1038/nature01090

Rahmstorf, S., Box, J.E., Feulner, G., Mann, M.E., Robinson, A., Rutherford, S., Schaffernicht, E.J., 2015. Exceptional twentieth-century slowdown in Atlantic Ocean overturning circulation. Nat. Clim. Chang. 5, 475.

Rayner, D., Hirschi, J.J.-M., Kanzow, T., Johns, W.E., Wright, P.G., Frajka-Williams, E.,

Bryden, H.L., Meinen, C.S., Baringer, M.O., Marotzke, J., Beal, L.M., Cunningham, S.A., 2011. Monitoring the Atlantic meridional overturning circulation. Deep Sea Res. Part II Top. Stud. Oceanogr. 58, 1744–1753. https://doi.org/http://dx.doi.org/10.1016/j.dsr2.2010.10.056

Rhines, P.B., Hakkinen, S., Josey, S.A., 2008. Is oceanic heat transport significant in the climate system?, in: Arctic-Subarctic Ocean Fluxes. Springer Verlag, pp. 87–109.

Roberts, C.D., 2017. Cdr30/Rapidmoc: Rapidmoc V1.0.1. Zenodo. https://doi.org/10.5281/ZENODO.1036387

Roberts, C.D., Garry, F.K., Jackson, L.C., 2013. A Multimodel Study of Sea Surface Temperature and Subsurface Density Fingerprints of the Atlantic Meridional Overturning Circulation. Journal of Climate 26, 9155–9174. https://doi.org/10.1175/JCLI-D-12-00762.1

Roberts, C.D., Jackson, L., McNeall, D., 2014. Is the 2004 - 2012 reduction of the Atlantic meridional overturning circulation significant? Geophys. Res. Lett. 41, 3204–3210. https://doi.org/10.1002/2014GL059473

Roberts, C.D., Palmer, M.D., 2012. Detectability of changes to the Atlantic meridional overturning circulation in the Hadley Centre Climate Models. Clim. Dyn. 39, 2533–2546. https://doi.org/10.1007/s00382-012-1306-3

Robson, J., Hodson, D., Hawkins, E., Sutton, R., 2014. Atlantic overturning in decline? Nat. Geosci. 7, 2. https://doi.org/10.1038/ngeo2050

Roemmich, D., Wunsch, C., 1985. Two transatlantic sections: Meridional circulation and heat flux in the subtropical North Atlantic Ocean. Deep Sea Research Part A. Oceanographic Research Papers 32, 619–664. https://doi.org/10.1016/0198-0149(85)90070-6

Sabine, C.L., Feely, R.A., Gruber, N., Key, R.M., Lee, K., Bullister, J.L., Wanninkhof, R., Wong, C.S., Wallace, D.W.R., Tilbrook, B., Millero, F.J., Peng, T.-H., Kozyr, A., Ono, T., Rios, A.F., 2004. The Oceanic Sink for Anthropogenic CO2. Science 305, 367–371. https://doi.org/10.1126/science.1097403

Savin, N.E., White, K.J., 1977. The Durbin-Watson Test for Serial Correlation with Extreme Sample Sizes or Many Regressors. Econometrica 45, 1989. https://doi.org/10.2307/1914122

Seager, R., Battisti, D.S., Yin, J., Gordon, N., Naik, N., Clement, A.C., Cane, M.A., 2002.

Is the Gulf Stream responsible for Europe's mild winters? Q. J. R. Meteorol. Soc. 128, 2563–2586. https://doi.org/10.1256/qj.01.128

Shapiro, S.S., Wilk, M.B., 1965. An Analysis of Variance Test for Normality (Complete Samples). Biometrika 52, 591. https://doi.org/10.2307/2333709

Sinha, B., Smeed, D.A., McCarthy, G., Moat, B.I., Josey, S.A., Hirschi, J.J.M., Frajka-Williams, E., Blaker, A.T., Rayner, D., Madec, G., 2018. The accuracy of estimates of the overturning circulation from basin-wide mooring arrays. Prog. Oceanogr. 160, 101–123. https://doi.org/10.1016/j.pocean.2017.12.001

Smeed, D.A., Josey, S.A., Beaulieu, C., Johns, W.E., Moat, B.I., Frajka-Williams, E., Rayner, D., Meinen, C.S., Baringer, M.O., Bryden, H.L., McCarthy, G.D., 2018. The North Atlantic Ocean Is in a State of Reduced Overturning. Geophys. Res. Lett. 45, 1527–1533. https://doi.org/10.1002/2017GL076350

Smeed, D.A., McCarthy, G.D., Cunningham, S.A., Frajka-Williams, E., Rayner, D., Johns, W.E., Meinen, C.S., Baringer, M.O., Moat, B.I., Duchez, A., Bryden, H.L., 2014. Observed decline of the Atlantic meridional overturning circulation 2004 - 2012. Ocean Sci. 10, 29–38. https://doi.org/10.5194/os-10-29-2014

Smeed, D.A., McCarthy, G., Rayner, D., Moat, B.I., Johns, W.E., Baringer, M.O., Meinen, C.S., 2017. Atlantic meridional overturning circulation observed by the RAPID-MOCHA-WBTS (RAPID-Meridional Overturning Circulation and Heatflux Array-Western Boundary Time Series) array at 26N from 2004 to 2017. https://doi.org/10.5285/5ACFD143-1104-7B58-E053-6C86ABC0D94B

Smeed, D.A., Moat, B.I., Rayner, D., Johns, W.E., Baringer, M.O., Volkov, D.L., Frajka-Williams, E., 2019. Atlantic meridional overturning circulation observed by the RAPID-MOCHA-WBTS (RAPID-Meridional Overturning Circulation and Heatflux Array-Western Boundary Time Series) array at 26N from 2004 to 2018. https://doi.org/10.5285/8CD7E7BB-9A20-05D8-E053-6C86ABC012C2

Stocker, T.F., Qin, D., Plattner, G.-K., Tignor, M., Allen, S.K., Boschung, J., Nauels, A., Xia, Y., Bex, V., Midgley, P.M. (Eds.), 2013. Climate change 2013 : The physical science basis : Working Group I contribution to the Fifth assessment report of the Intergovernmental Panel on Climate Change. Cambridge ; New York : Cambridge University Press, 2014.

Thomson, D.J., 1982. Spectrum estimation and harmonic analysis. Proc. IEEE 70,

1055–1096. https://doi.org/10.1109/PROC.1982.12433

Thomson, R.E., Emery, W.J., 2014. Data Analysis Methods in Physical Oceanography. Elsevier. https://doi.org/10.1016/C2010-0-66362-0

Thornalley, D.J.R., Oppo, D.W., Ortega, P., Robson, J.I., Brierley, C.M., Davis, R., Hall, I.R., Moffa-Sanchez, P., Rose, N.L., Spooner, P.T., Yashayaev, I., Keigwin, L.D., 2018. Anomalously weak Labrador Sea convection and Atlantic overturning during the past 150 years. Nature 556, 227–230. https://doi.org/10.1038/s41586-018-0007-4

Vélez-Belchí, P., Pérez-Hernández, M.D., Casanova-Masjoan, M., Cana, L., Hernández-Guerra, A., 2017. On the seasonal variability of the Canary Current and the Atlantic Meridional Overturning Circulation. J. Geophys. Res. Oceans 122, 4518–4538. https://doi.org/10.1002/2017JC012774

Vellinga, M., Wood, R.A., 2004. Timely detection of anthropogenic change in the Atlantic meridional overturning circulation. Geophys. Res. Lett. 31. https://doi.org/10.1029/2004GL020306

Vellinga, M., Wood, R.A., 2002. Global climatic impacts of a collapse of the atlantic thermohaline circulation. Clim. Change 54, 251–267. https://doi.org/10.1023/A:1016168827653

Walker, 1931. On periodicity in series of related terms. Proc. R. Soc. Lond. A 131, 518–532. https://doi.org/10.1098/rspa.1931.0069

Willis, J.K., 2010. Can in situ floats and satellite altimeters detect long-term changes in Atlantic Ocean overturning? Geophys. Res. Lett. 37, n/a–n/a. https://doi.org/10.1029/2010GL042372

Winton, M., Griffies, S.M., Samuels, B.L., Sarmiento, J.L., Licher, T.L.F., 2013. Connecting changing ocean circulation with changing climate. J. Clim. 26, 2268–2278. https://doi.org/10.1175/JCLI-D-12-00296.1

Worthington, E.L., Frajka-Williams, E., McCarthy, G.D., 2019. Estimating the Deep Overturning Transport Variability at 26°N Using Bottom Pressure Recorders. J. Geophys. Res. Oceans 124, 335–348. https://doi.org/10.1029/2018JC014221

Worthington, E.L., Moat, B.I., Smeed, D.A., Mecking, J.V., Marsh, R., McCarthy, G.D., 2021. A 30-year reconstruction of the Atlantic meridional overturning circulation shows no decline. Ocean Sci. 17, 285–299. https://doi.org/10.5194/os-17-285-2021

Yule, G.U., 1927. On a Method of Investigating Periodicities in Disturbed Series, with

Special Reference to Wolfer's Sunspot Numbers. Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character 226, 267–298.

Zhang, R., 2008. Coherent surface-subsurface fingerprint of the Atlantic meridional overturning circulation. Geophys. Res. Lett. 35, L20705. https://doi.org/10.1029/2008GL035463

Zhang, R., Delworth, T.L., 2006. Impact of Atlantic multidecadal oscillations on India/Sahel rainfall and Atlantic hurricanes. Geophys. Res. Lett. 33, L17712. https://doi.org/10.1029/2006GL026267

Zhang, R., Sutton, R., Danabasoglu, G., Kwon, Y.-O., Marsh, R., Yeager, S.G., Amrhein, D.E., Little, C.M., 2019. A Review of the Role of the Atlantic Meridional Overturning Circulation in Atlantic Multidecadal Variability and Associated Climate Impacts. Rev. Geophys. 57, 316–375. https://doi.org/10.1029/2019RG000644