UNIVERSITY OF SOUTHAMPTON

Faculty of Environmental and Life Sciences
School of Ocean and Earth Sciences

Characterising water mass and circulation change using heat and carbon covariability during the anthropogenic era

DOI: 10.1002/0470841559.ch1

by

Charles Turner
MSci
ORCiD: 0000-0002-3262-4972

A thesis for the degree of
Doctor of Philosophy

February 2023

University of Southampton

Abstract

Faculty of Environmental and Life Sciences
School of Ocean and Earth Sciences

Doctor of Philosophy

Characterising water mass and circulation change using heat and carbon covariability during the anthropogenic era

by Charles Turner

The ocean is a major sink of both CO_2 and heat, having absorbed approximately a third of cumulative carbon emissions to data and 93% of the additional heat contained in the climate system. Whilst not straightforward, it is possible to identify this additional carbon in the ocean, allowing the quantification of the global and regional ocean carbon sinks. However, for a number of reasons, it is far more difficult to decompose heat changes in an analogous fashion in order to identify the 'excess' heat.

In this work, two new and related techniques are developed for the identification of excess heat. The first removes the 'anthropogenic' carbon signal from total carbon changes, leaving changes in carbon which are not driven by increased atmospheric CO_2. By relating the remaining changes in temperature and carbon, the redistributed temperature is identified, and the excess isolated by residual. This technique is applicable to additional tracers, for example salinity. This technique is demonstrated in the NEMO GCM, finding significant excess salinity changes generally precede excess temperature changes, but that excess temperature changes dominate later in the model run.

Previous work has also shown that changes in anthropogenic carbon and excess heat are linked by a transient response coupling, and are therefore linearly related. By combining the previous technique and this transient response coupling, a second technique, which does not require an explicit carbon decomposition is developed. This technique is then applied in the Subtropical North Atlantic, and to the full GLODAP dataset, to produce global fields of excess heat and salinity accumulation. As expected, excess heat content increases smoothly with time, with the majority of excess heat accumulation in the upper thousand metres. Additionally, the rate of excess heat storage is higher in the Atlantic than in other ocean basins, in agreement with previous studies. Patterns of excess salinity storage are less spatially uniform and exert a strong influence on excess density changes, suggesting that changes to the water cycle may impact ocean circulation to a similar or greater degree than additional heat content.

Contents

List of Figures

Listings

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has 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 or 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. Where 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: C. E. Turner, P. J. Brown, K. I. C. Oliver, and E. L. McDonagh. Decomposing oceanic temperature and salinity change using ocean carbon change. *Ocean Science*, 18(2):523–548, 2022. . URL https://os.copernicus.org/articles/18/523/2022/

Signed:..... Date:.....

Acknowledgements

I'd like to thank everyone who has managed to put up with me over the past four years. It's been an unusually difficult time and I'm sure I've been a nightmare at times but we've somehow made it through. In particular, I'd like to thank my parents for raising me to be optimistic (or maybe arrogant) enough to believe I'm capable of doing anything I stick at, and my brother for always being up for a drink when I needed one. I'd like to thank Mike, Spencer, Phil, Pablo, Cath and Vanesa for getting me through my first year in Southampton, Gabby and Phil for getting me through lockdown without tearing my hair out, Alice and Louis for being there for me when I was really in a pickle, and Katie for letting me stay at hers to write up my thesis.

But finally and most importantly, I'd like to thank Pete, Elaine and Kevin. Without their insight, support, guidance and patience (and it was really quite a lot of patience), the work in this thesis never would have materialised, and I wouldn't be the person I am today.

"The first gulp from the glass of natural sciences will turn you into an atheist, but at the bottom of the glass God is waiting for you" - Werner Heisenberg

Chapter 1

Introduction

1.1 Motivation

As a result of continuing anthropogenic CO_2 emissions, atmospheric CO_2 levels continue to rise, driving an increase in global temperature. However, the global ocean acts as a brake on both of these processes, ameliorating both increases in atmospheric CO_2 concentrations and increases in global mean temperatures.

The global ocean is a major sink of anthropogenic CO_2 emissions, currently absorbing approximately a quarter (Le Quéré et al., 2018), and having absorbed approximately a third of all anthropogenic CO_2 emissions to date (Khatiwala et al., 2013). Furthermore, it is expected to absorb up to 80% of all anthropogenic carbon emissions on multi-centennial timescales (Archer et al., 1997). It is therefore hard to overestimate its importance to the global carbon cycle and atmospheric CO_2 concentrations.

The ocean is similarly important for the global heat budget, having stored over 93% of the heat accumulated in the earth system over 1961-2008 due to top of atmosphere flux imbalances (Church et al., 2011), significantly ameliorating the effects of global warming. However, as a result of the added heat content of the ocean, sea levels are expected to rise significantly over the coming centuries (Pardaens et al., 2011). This is anticipated to cause significant disruption to many societies: over 10% of the worlds population live in the Low Elevation Coastal Zone (LECZ), the contiguous area along the coast with less than 10m of elevation (McGranahan et al., 2007). Compounding this, sea levels are not projected to rise uniformly: previous modelling studies indicate low latitude sea level rise (SLR) will exceed the global mean by up to 20% (Perrette et al., 2013), in large part due to the thermal expansion of seawater (thermosteric sea level rise) (Pardaens et al. (2011), Church et al. (2013)), with enormous implications for future economic development (Hinkel et al., 2014). As at least 14% of the population in developing economies live in the LECZ, where SLR is projected to be greatest, ocean

heat uptake and SLR therefore have enormous implications for global development (Hinkel et al., 2014), coastal nations, and their populations. It also has important implications for the future of marine ecosystem health: ocean warming has a direct effect on marine life as a driver of deoxygenation (Oschlies et al., 2018), as well as through increased stratification (Gruber, 2011). The uptake of carbon similarly affects marine life through its role in ocean acidification (Gruber, 2011).

In its preindustrial state, ocean heat and carbon are closely coupled (Sarmiento and Gruber, 2006), mostly due to the strong temperature dependence of CO_2 solubility in water. CO_2 is more soluble in cooler waters, and therefore the highest concentrations of DIC (Dissolved Inorganic Carbon, the total carbon concentration of seawater excluding biological carbon), is highest in cool waters. This leads to a close linear relationship between the temperature of water and its carbon concentration: the carbon content of water decreases by approximately 9μmol/kg for every degree. However, CO_2 does not simply dissolve in water: dissolved CO_2 also reacts with water, forming carbonate and bicarbonate ions. As a result of this unique chemistry, the concentration of DIC in seawater is far higher than might be naively be expected from equilibration with the atmosphere: in fact, the preindustrial ocean carbon pool (the total amount of carbon contained in the ocean preindustrially) was approximately sixty times larger than that of atmospheric CO_2 (38 thousand petagrams of carbon, PgC for the ocean versus 610 PgC for the atmosphere) (Graven, 2016). In fact, the oceanic carbon pool is so much larger than the atmospheric carbon pool that it was thought until the 1950's that virtually all anthropogenic CO_2 emissions would be absorbed by the ocean almost instantaneously: it was largely due to the interest of the US military in detecting surreptitious Soviet nuclear tests it was discovered that atmospheric CO_2 was in fact increasing (Weart, 1997). It was subsequently realised this was due to the oceans limited ability to sequester the additional CO_2 released to the atmosphere by the burning of fossil fuels (Weart, 1997). However, the large oceanic pool of DIC, coupled with the sensitivity of CO_2 solubility to temperature, means that carbon and global climate are closely linked: previous studies have found changes in CO_2 concentrations as a response to ocean circulation change to be a primary driver of glaciation and deglaciation (Shakun et al. (2012), Shackleton (2000)), although others suggest a more limited influence (Imbrie et al. (1993), Alley and Clark (1999)), or indeed the opposite causality (Toggweiler and Lea, 2010). Regardless of the mechanisms linking preindustrial ocean carbon content and global climate, it is well known that ocean circulation presently exerts a strong influence on the ocean's uptake of anthropogenic carbon (DeVries et al., 2017).

In the same way that the ocean and its circulation is important for global CO_2 concentrations due to its large pool of carbon, the enormous heat capacity of the ocean in comparison to the atmosphere (over 1000 times larger), means that it also plays a key role in regulating global climate through its storage and transport of heat (Schmitt, 2018). For example, though both the ocean and atmosphere contribute similarly to the transport of heat from low to high latitudes (Ferrari and Ferreira, 2011), a large fraction of atmospheric heat transport is in fact achieved through the transport of moisture by the atmosphere, which is supplied by the ocean (Yang et al., 2015).

There are numerous indications that both the patterns and scale of heat and carbon sequestration by the ocean are tightly coupled. The uptake of both are strongly influenced by ocean circulation (Sarmiento et al. (1992), DeVries et al. (2017)), with greatest sequestration in regions of deep and intermediate water formation, notably the North Atlantic and Atlantic sectors of the Southern Ocean. This is the simple result of the physics of ocean heat and carbon uptake: once a parcel of water is entrained below the mixed layer, the layer of water in the upper ocean that can be considered to be in contact with the atmosphere, its properties are essentially fixed, save for mixing, radioactive decay and biological activity.

Changes to surface air temperature (SAT) as a result of anthropogenic activities are known to be linearly related to cumulative carbon emissions, for both local and global temperature changes (Leduc et al. (2016), Gillett et al. (2013)), with this linear relationship controlled predominantly by the uptake of heat and carbon by the ocean (Goodwin et al. (2015), Ehlert et al. (2017), Katavouta et al. (2018), Katavouta et al. (2019)). This relationship is known as the transient climate response to cumulative carbon emissions (TCRE), and its uses include, for example, calculating the carbon emissions remaining to limit warming to 1.5 degrees relative to preindustrial (Damon Matthews et al., 2021). Understanding the relationship between ocean heat and carbon uptake therefore represents a valuable opportunity to better quantify the effects anthropogenic CO_2 emissions on the global climate and to understand how to limit climate change to safe levels.

Although the coverage, both spatial and temporal, of ocean observations continues to increase rapidly (Davis et al., 2019), observational data remain relatively sparse spatially and temporally, particularly at depth. In CMIP3 models, deep ocean heat uptake was shown to be a major contributor to the uncertainty in transient global warming (Boé et al., 2009), with this uncertainty compounded by the lack of observational constraints (Karspeck et al., 2017). An improved understanding of deep ocean heat uptake contributing to a better understanding of transient climate change again underscores the importance of understanding the global ocean uptake of heat and carbon for global climate change.

Work extending the TCRE to a regional transient response (RTCRE) (Leduc et al., 2016) indicates that although a large range of SAT increase rates are observed depending

on location, there is a great deal of geographic coherence. This would indicate, assuming tight coupling between SST and SAT changes due to the large heat capacity of water, that we would expect similar geographic coherence in any relationship between changes in ocean temperature and carbon content. Bronselaer and Zanna (2020) investigated such a relationship, finding a near-linear global relationship between increases in ocean heat and carbon content (Bronselaer and Zanna, 2020). They labelled the parameter linking the two the carbon-heat coupling, α. This can be observed at a range of scales: both increases in global ocean heat and carbon inventories, and in local ocean excess temperature (temperature changes propagated from surface temperature changes) and anthropogenic carbon (DIC changes resulting from equilibration with increased atmospheric CO_2) are linearly related. Bronselaer and Zanna (2020) argues that this relationship ought be globally uniform thanks to the responses of seawater carbonate chemistry and radiative forcing to rising atmospheric CO_2. However, there are are indications that there may be some regional variability in this ratio (Turner et al. (2022), Bopp et al. (2015), Davila et al. (2022)), and that it may be necessary to better understand the drivers of this regional variability in order to properly understand how ocean heat and carbon uptake are linked. It appears to be the case that at high latitudes, ocean heat uptake is increased relative to carbon uptake. It is unclear from previous research why this is the case. It may be due to the polar amplification of warming (the Arctic is warming significantly faster than the global mean): for example, Leduc et al. (2016) finds greater SAT changes above the ocean at high latitudes than low. Alternatively, it may be driven by differing mechanisms of ocean heat and carbon uptake at high and low latitudes. Bopp et al. (2015) found that total subduction of anthropogenic carbon through the base of the mixed layer was significantly more variable than air-sea fluxes, and was generally reduced at high latitudes, relative to the global mean. Such a mechanism would indicate that the different processes contributing to the transport of anthropogenic carbon into the ocean interior are a key factor in driving regional variability in this analogous relationship between changes in ocean heat and carbon content.

In addition to transient warming, SLR and global energy budgets also rely on estimates of deep ocean heat uptake for closure, and so improving our understanding of deep ocean heat uptake is key to improving our estimates of these quantities (Purkey and Johnson (2012), Domingues et al. (2008), Murphy et al. (2009)). Linking carbon uptake to heat uptake as an additional constraint therefore represents an opportunity to improve our estimates of SLR and global energy balance: the field of decomposing the increases in carbon in the ocean is far more developed than analogous decompositions to heat changes, and so linking the two represents an avenue for improving our understanding of ocean heat uptake.

Though previous research into the TCRE has emphasised the importance of the ocean both in setting its magnitude and controlling the linear relation between warming and

emissions, how this transient response manifests in the ocean remains an understudied problem. Bronselaer and Zanna (2020) use this transient response framework in order to explain whether ocean heat content change is due to changes in surface forcing (excess heat), or due to changes in circulation (redistributed heat), but do not consider the implications for transient climate change, nor the potential for spatial variability in the transient link between changes in ocean heat and carbon content. Similarly, Williams et al. (2021) uses the correlation imparted by this transient response between changes in ocean temperature and carbon, along with the anticorrelation of the background temperature and carbon fields, to identify regions where changes in ocean temperature and carbon result predominately from either changes in ocean circulation or from the uptake of excess heat and anthropogenic carbon by the ocean. An improved framework in which to understand this oceanic response therefore represents an enormous opportunity, with implications for understanding sea level rise, transient global warming and ocean circulation change. It may also represent an opportunity for improved model evaluation by providing a new metric with which to understand model performance and biases.

1.2 Ocean Tracers

Ocean tracers are, at the most basic level, properties of a water parcel which may be used to ascertain the location at which the parcel was formed: that is, subducted from the ocean surface. They also enable the study the changes such a parcel has undergone as it is transported throughout the global ocean: advection, mixing, biological activity, etc. Tracers include physical properties such as temperature and salinity, and biogeochemical properties such as silicate and nitrate, amongst many others. Tracers have been used to infer and quantify ocean circulation for centuries, with Capt. Henry Ellis first discovering the deep waters of the Subtropical Atlantic to be much cooler than the surface in 1751 (Ellis and Hales, 1751). It was later hypothesised by Count Rumford that this water must have come from much higher latitude, pioneering the use of temperature as a tracer to determine ocean circulation. Over the years, the use of tracers has allowed us, to a large extent, to describe the general circulation of the ocean. One commonly used description is that of the "Great Ocean Conveyor" of Broecker (1991): an adaptation of this is shown in Figure 1.1. Though extremely oversimplified, it captures several key features of global ocean circulation: the formation of deep and bottom water in the North Atlantic and Southern Oceans, and the return of water from depth to the surface in the Indian and Pacific oceans.

The use of tracers in the determination of ocean circulation is extensive and well developed, and as such there are a number of classifications of tracers, though categories are not necessarily mutually exclusive. Examples include:

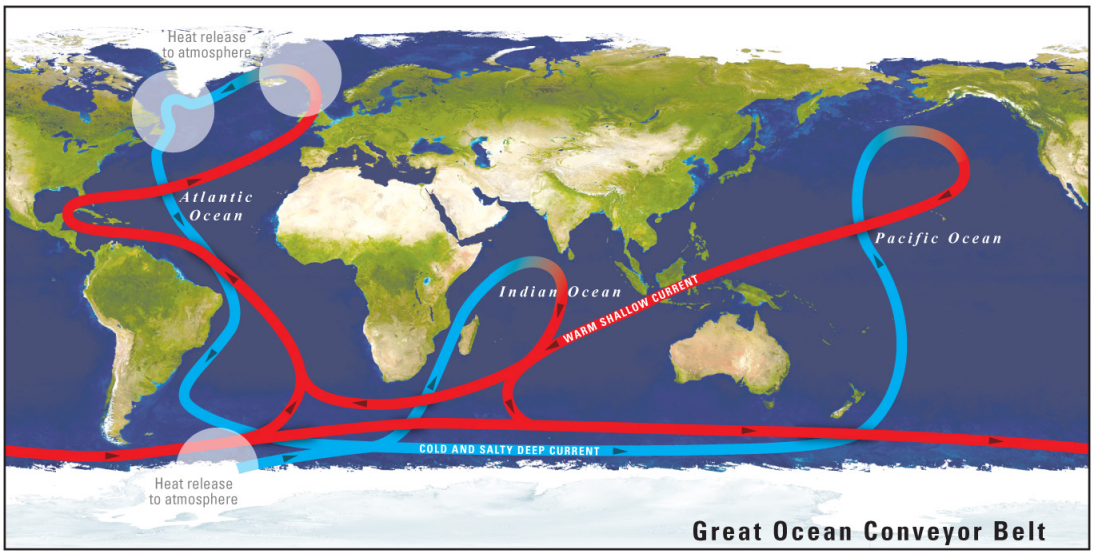

FIGURE 1.1: The Great Ocean Conveyor belt, a simplified schematic of ocean circulation. Deep water forms in regions where heat is lost to the atmosphere: these are shown by shading. Deep water formation is indicated by a change of colour from blue to red, and deep water returning to the surface by a change of colour from red to blue. Reproduced from https://pubs.usgs.gov/pp/p1386a/gallery2-fig31.html.

- Passive: Tracers which do not affect ocean circulation. Chemical tracers such as phosphate, silicate, nitrate etc. are all passive, whereas physical tracers such as temperature, salinity and potential vorticity are generally not passive, though it is common to treat perturbations to physical tracers as passive. Non passive tracers are referred to as active.
- Conservative: Tracers which do not exhibit time or state (pressure, temperature) dependence. Tracers such as chlorofluorocarbons (CFCs) are conservative, as they exhibit no in situ interactions or decay. However, CFCs can also be treated as a transient tracer, as they have a time dependent surface input.
- Transient: Tracers which exhibit time dependence, often in the form of radioactive decay, but also due to time dependent surface conditions as noted above. Radiocarbon (Carbon-14) is a commonly used transient tracer.

Other classifications of tracers exist, though are not discussed further.

1.2.0.1 Carbon in the ocean: Dissolved Inorganic Carbon, Anthropogenic and Natural Carbon

Unlike in the atmosphere, Dissolved Inorganic Carbon (DIC) in the ocean is not well mixed at global scales, nor are changes in its distribution entirely attributable to increased atmospheric CO_2. In its preindustrial state, atmospheric CO_2 is thought to have

been well mixed to within a few parts per million (ppm), compared to a global mean of approximately 286ppm (Meinshausen et al., 2017). Although anthropogenic activities have caused an increase in both the amplitude of the seasonal CO_2 cycle (Graven et al., 2013) and the difference between Northern and Southern hemisphere concentrations (Meinshausen et al., 2017), atmospheric CO_2 is still well mixed to within a few ppm, compared to a global mean of approximately 410ppm presently. As such, atmospheric CO_2 is well described, both presently and historically, by a single number, the annual global mean CO_2.

Such a simple description of carbon in the ocean is not possible, for a number of reasons. Atmospheric CO_2 records extend back millions of years thanks to ice core records, such as those used in Meinshausen et al. (2017). No such record exists for DIC, and so its preindustrial state must be estimated indirectly. In addition to DIC, organic carbon is also found in the ocean: Particulate Organic Carbon (POC), for example, which comprises the remains and excretions of small marine organisms, which sink upon their death. As these remains sink, they are consumed by heterotrophs, creating a cascade to smaller excretions, eventually returning to inorganic nutrients. Thus, POC is transformed into DIC, an interaction which must be accounted for to understand the global distribution of DIC. This particular process is known as the soft tissue pump, although other interactions occur in the ocean carbon system: for example the Carbonate Pump, in which the hard carbonate shells of marine organisms (Particulate Inorganic Carbon, PIC) are similarly remineralised at depth due to the increased pH of deeper waters. The soft tissue and carbonate pumps are collectively referred to as the biological carbon pump, without which preindustrial CO_2 concentrations would have been approximately 50% higher (Sarmiento and Gruber, 2006) than they were.

Another important control on the distribution of DIC in the ocean interior is the solubility pump (Williams and Follows, 2011). The solubility of CO_2 is a strong inverse function of temperature. However, as previously noted, the deep waters of the global ocean are formed at high latitude, where water masses are cool and therefore dense enough to subduct to depth. Thus, DIC generally increases with depth, even discounting the effects of the biological pumps: this effect is known as as the solubility pump.

As a result of these complex interactions between physical, chemical and biological processes, the distribution of DIC in the global ocean is far more heterogeneous than that of atmospheric CO_2, with concentrations ranging from less than 2000μmol/kg in surface waters to more than 2400μmol/kg in the deep North Pacific. Figure 1.2 gives an indication of this complex distribution. In general, concentrations of DIC increase from the surface to depth. Concentrations are also visibly higher in the deep North Pacific than the deep North Atlantic. This is due to the patterns of deep water formation: deep water in the Pacific can be hundreds to thousands of years old, and so has accumulated more DIC through the biological carbon pumps. This pattern of older waters having higher sDIC (DIC normalised to a salinity of 35) concentrations can also be seen in the

North Atlantic: here, Antarctic Bottom Water (AABW) sits below North Atlantic Deep Water (NADW) at depths greater than approximately 4000m. Above NADW, at depths of approximately 1000m, Antarctic Intermediate Water (AAIW) is found. Both AAIW and AABW take longer to reach this region than NADW: thus, due to its greater age, higher DIC concentrations are seen in AAIW and AABW than NADW (Touratier et al., 2005).

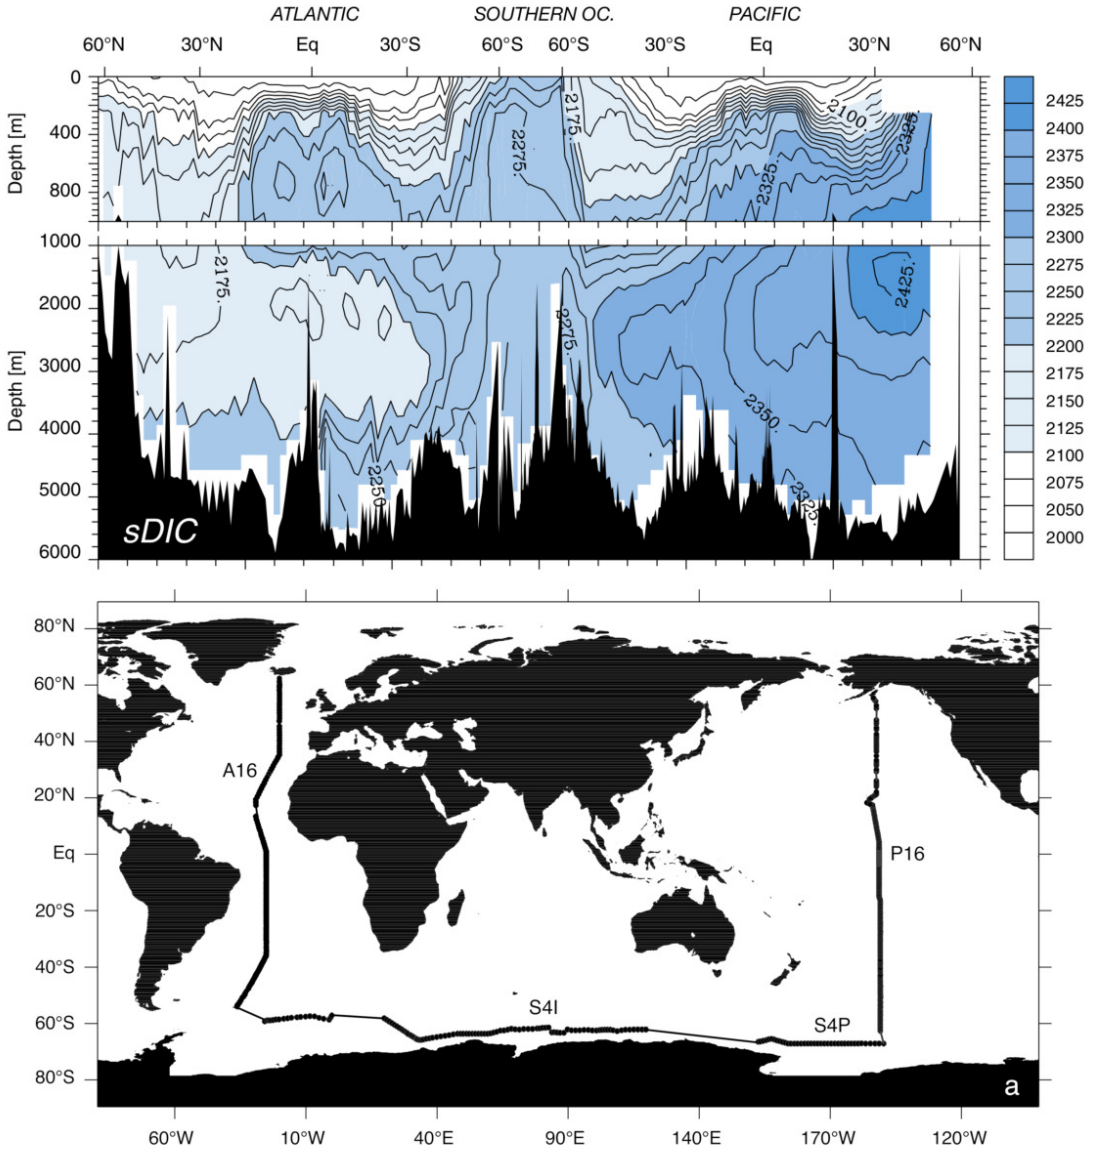

FIGURE 1.2: A map of sDIC, DIC normalised to a salinity of 35, from shipboard observations during WOCE (the World Ocean Circulation Experiment (Thompson et al., 2001)) (top), and the location of the shipboard observations (bottom). Normalising to constant salinity helps to remove the effect of freshwater fluxes on DIC distributions. Reproduced from Sarmiento and Gruber (2006).

If we wish to understand what controls the distribution of DIC in the ocean, we must first separate it into its different constituent components. This is a crucial step in the

identification of anthropogenic carbon, C_{anth}. Due to heterogeneity of preindustrial DIC distributions, the complexity of processes controlling the exchange of CO_2 between atmosphere and ocean (Wanninkhof, 2014), the chemical indistinguishably of C_{anth} from other carbon pools, and the complex in situ biogeochemical interactions of the ocean carbon cycle, a number of techniques exist which aim to separate present day DIC into the component due to increased atmospheric CO_2 (C_{anth}) and changes in the distribution of the background DIC distribution due to changing ocean dynamics (C_{nat}). These techniques, though differing in their details, fall broadly into two categories: Back calculation techniques and Transit Tracer based techniques (Khatiwala et al., 2013).

Back calculation techniques rely on an understanding of the ocean carbon cycle, and the different reservoirs of carbon which DIC may be decomposed into. Several different combinations and definitions of the carbon reservoirs may be used (see Williams and Follows (2011)), though all are essentially different combinations of a commonly used 'base set', which is given here. We may decompose DIC into 5 'pools' as follows:

$$\mathrm{DIC} = \mathrm{DIC}_{sat}^{PI} + \mathrm{DIC}_{carb} + \mathrm{DIC}_{soft} + \mathrm{DIC}_{diseq.} + \mathrm{DIC}_{anth} \tag{1.1}$$

Each term is now described in turn:

- DIC_{sat}^{PI}: The DIC content which a parcel of water would have, if allowed to equilibrate with the preindustrial atmosphere, at its present potential temperature and salinity. This comprises the bulk of all DIC (approximately 2000μmol/kg), and is referred to as Saturation Carbon.

- DIC_{soft}: The DIC content due to the remineralisation of soft tissue, referred to as Soft Tissue Carbon. It ranges approximately 0-200μmol/kg, depending on depth and the age of a water parcel: surface waters and the deep North Atlantic have concentrations near to zero, whilst concentrations of over 200μmol/kg are seen in the intermediate North Pacific.

- DIC_{carb}: DIC content due to the remineralisation of calcium carbonate. The distribution is broadly similar to that of DIC_{soft}, but with lower concentrations, ranging from zero up to approximately 60μmol/kg.

- DIC_{diseq}: DIC content due to the disequilibrium of a parcel of water with the overlying atmosphere, when subducted away from the surface. It is referred to as Disequilibrium Carbon, and unlike other carbon reservoirs, may be positive or negative. Calculating its local magnitude can be extremely difficult, and so it is often combined with the DIC_{anth} term as the term DIC_{res}.

- DIC_{anth}: Anthropogenic Carbon, usually referred to simply as C_{anth} or C_{ant}. This is broadly similar to DIC_{sat}^{PI}, but accounts for the increased CO_2 content of the

atmosphere: we may write $DIC_{sat} = DIC_{sat}^{PI} + C_{anth}$. Concentrations range from zero in deep, unventilated waters, up to around 80 μmol/kg in surface, well ventilated waters.

In back calculation techniques, observations of the current ocean state are used to remove the soft tissue, carbonate and saturation carbon, leaving DIC_{res}. This is then decomposed for C_{anth} and the disequilibrium carbon. Back calculation techniques include ΔC^* (Gruber et al., 1996), TrOCA (Touratier and Goyet, 2004), φC_T^0 (Vázquez-Rodríguez et al., 2009) and eMLR(C^*) (Clement and Gruber, 2018) (though the eMLR(C^*) technique actually calculates changes in C_{anth} and so requires an initial distribution estimate).

Unlike back calculation techniques, transient tracer based techniques instead begin with estimates of ocean transport calculated using transient tracers, and aim to propagate the history of atmospheric conditions into the ocean to obtain estimates of C_{anth} inventories. Such techniques generally assume C_{anth} to be a conservative tracer, since C_{anth} concentrations are only a fraction of that of DIC. As such, its interaction with the biological carbon pump is a minimal, 2nd order term. Examples include the TTD (Transit Time Distribution) (Hall et al., 2002), and Green's Function based implementations of surface histories, such as the Transport Matrix Method (TMM) (Khatiwala et al., 2005).

Like DIC, the distribution of C_{anth} is also very heterogeneous, predominately due to its time dependent history and strong dependence on ocean circulation. An estimated distribution of the additional C_{anth} uptaken by the ocean from 1994-2007 (using the eMLR(C^*) technique) is shown in Figure 1.3, which is broadly representative of the total C_{anth} distribution in the global ocean.

As a result of the myriad of C_{anth} estimation techniques, there are numerous discrepancies between techniques, though globally integrated inventories tend to agree much more closely (Khatiwala et al., 2013), (Wang et al., 2012). This is shown in Figure 1.4. Near global inventories from data based methods agree to within 20% (experimental uncertainty), but regional estimates can differ by up to 50%, particularly in the North Atlantic and Southern Oceans. Importantly, C_{anth} accumulation tends to be greatest in regions of deep water formation, as illustrated in Figure 1.1. In fact, ocean circulation is well known to be a leading order control on regional patterns of CO_2 fluxes and C_{anth} accumulation (Sarmiento et al., 1992).

1.2.1 Changing ocean dynamics: Excess and Redistributed Tracers

1.2.1.1 Excess and Redistributed Temperature

Though the art of decomposing changes in DIC into C_{anth} accumulation and changes in other carbon pools remains imperfect, it is relatively well developed, largely thanks

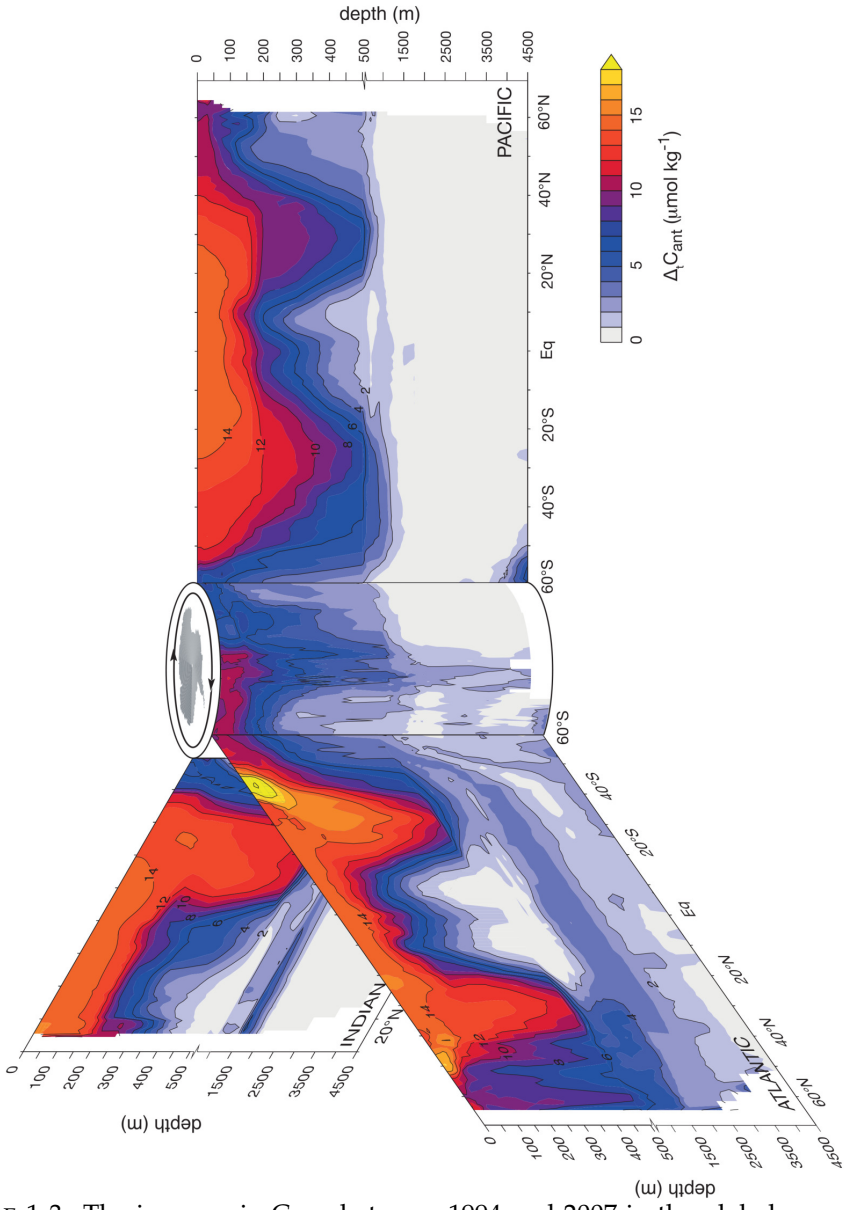

FIGURE 1.3: The increase in C_{anth} between 1994 and 2007 in the global ocean. Concentrations are highest in the subtropical gyres, and decay rapidly at depth in both the Indian and Pacific Oceans. Concentrations of C_{anth} at abyssal depths in the Atlantic are much higher due to it's importance in deep water formation: deep waters in the Atlantic tend to be ventilated more recently than those in the Pacific and Indian, and thus show faster increase rates in their inventories of C_{anth}. Adapted from Gruber et al. (2019).

to the well understood biogeochemical processes controlling it. Attempting to analogously decompose temperature changes is more difficult: reliably identifying whether a change in temperature is due to dynamical variability or warming of a water mass can be extremely difficult. A framework which allows us to identify the drivers of these changes is known as the excess-redistribution decomposition, where 'excess' changes are changes in an ocean tracer resulting from changes in surface forcing, whereas 'redistributed' changes are the changes in a tracer resulting from perturbations to ocean circulation. The notion of an 'excess' tracer is closely related to the principles used

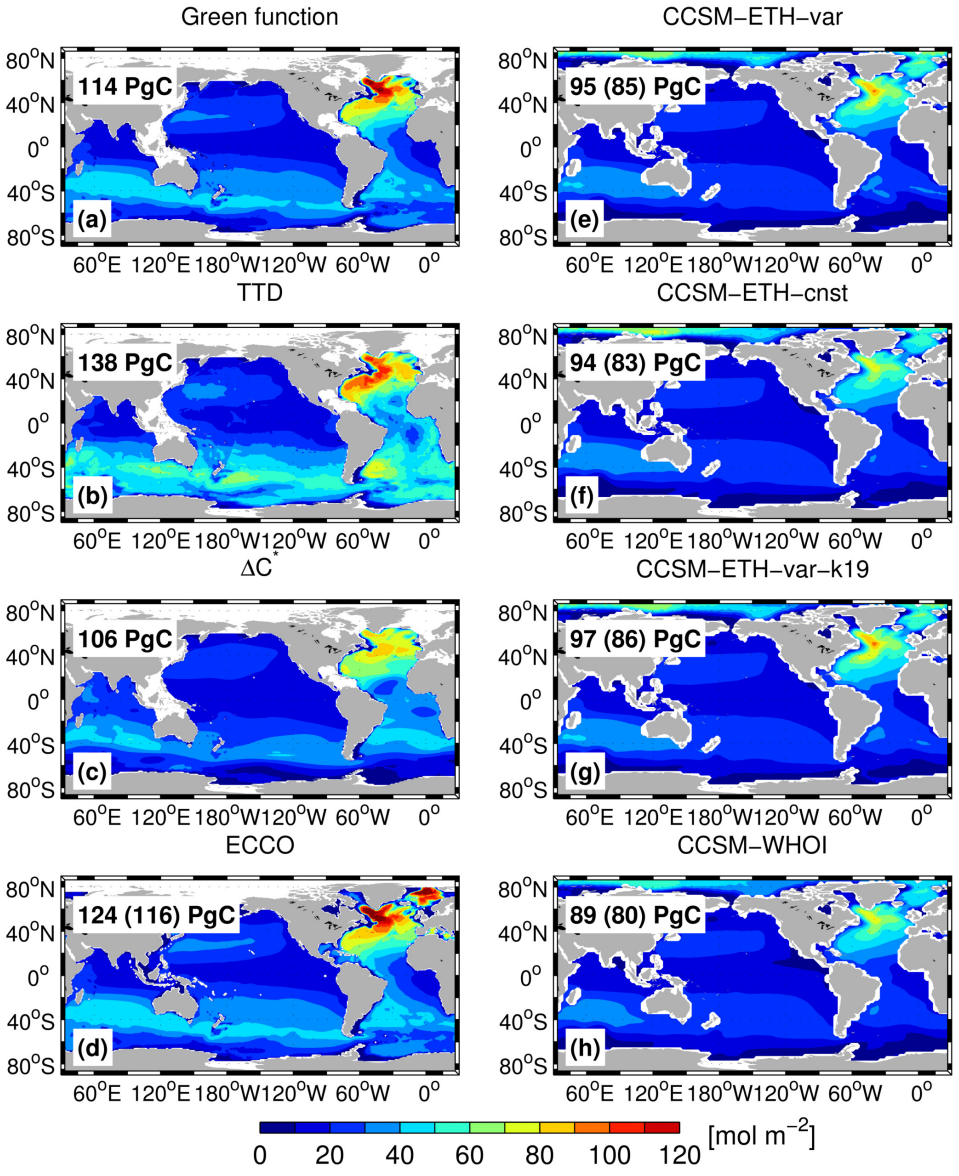

FIGURE 1.4: Estimates of C_{anth} column inventories in 1990, reproduced from Khatiwala et al. (2013). Panels (a)-(c) show estimates from data based techniques, (d) estimates from the TMM method using an MITgcm ECCO configuration, and (e)-(h) ESM CMIP historical run estimates.

in the transient tracer based estimates of C_{anth}, in particular the GF technique: by beginning with a description of ocean circulation, we may propagate changes in surface conditions into the ocean interior: changes in a given tracer due to the propagation of the surface signal by the steady state ocean circulation are described as 'excess'. Thus, the anthropogenic carbon estimated by transient tracer based estimation techniques is closely related to the excess carbon content of the ocean. The utility of this framework is that it can be extended to any ocean tracer, not just carbon: the tracer best studied under this framework thus far (after carbon) is temperature.

As an illustration of this, consider a velocity-temperature field, $\vec{v}T$, which may be written as two initial fields, plus small perturbations:

$$\vec{v} = \vec{v}_0 + \vec{v}', \tag{1.2}$$

where $\vec{v}_0$ is the initial field and $\vec{v}'$ the perturbation, and likewise for T. Then, the total perturbation of this velocity-temperature field may be written as

$$(\vec{v}T)' = \vec{v}T - \vec{v}_0 T_0 = \vec{v}_0 T' + \vec{v}' T_0 + \vec{v}' T' \tag{1.3}$$

In Equation 1.3, the term $\vec{v}_0 T'$ is referred to as excess temperature, and $\vec{v}' T_0$ referred to as redistributed temperature. The $\vec{v}' T'$ term is small and therefore may be neglected. Note that $\vec{v}'$ is nonzero as it is a function of T', and therefore T' is not a passive tracer. However, it is often useful to treat T' as a passive tracer, as this approximation remains valid in a number of circumstances (Church et al., 1991), (Jackett et al., 2000), (Russell et al., 2006). In others, for example estimates of regional ocean heat uptake, this simplification cannot be made (Banks and Gregory, 2006), (Xie and Vallis, 2012).

As heat is not a passive tracer, heat uptake will act to modify ocean circulation, and so $\vec{v}'$ is in general nonzero. This is problematic in understanding heat uptake as the terms $\vec{v}' T_0$ and $\vec{v}_0 T'$ are of similar magnitude, with the redistributed temperature term $\vec{v}_0 T'$ often dominating changes in ocean temperature, particularly on decadal and shorter timescales. Observationally identifying excess and redistributed heat changes can therefore be very difficult, and is an active area of research.

One solution is to use models and impose artificial conditions in order to deconvolve excess and redistributed heat change. Two main conceptual approaches are used to this end. The first is to force model dynamics to remain the same throughout a climate change run, effectively setting $\vec{v}' = 0$ in Equation 1.3, despite the uptake of heat. The second is to use a "Passive Anomalous Tracer" (PAT), which observes the same physics as the temperature field, but has no preindustrial field, and does not itself affect the evolving ocean circulation (instead, model dynamics respond to the co-evolving temperature field). This is analogous to forcing $T_0 = 0$ in Equation 1.3. Importantly, the two techniques are subtly different: the PAT technique will incorporate a second order, $v'T'$ term, whilst the fixed circulation technique will not.

Figure 1.5 shows both heat and PAT uptake, from CMIP5 *1pctCO2* runs. Of note are the significantly higher PAT uptake relative to heat in the North Atlantic and Arctic oceans, and lower uptakes across much of the Pacific and South Atlantic than heat. The patterns also strongly resemble those of C_{anth} uptake from 1.4 as ocean circulation (and deep water formation) is a leading control on the uptake of both heat and carbon. The differences in the Atlantic are largely representative of the effect of heat uptake on ocean dynamics: the Atlantic Meridional Overturning Circulation (AMOC), transports heat northward throughout the entire Atlantic, and weakens as the North Atlantic warms

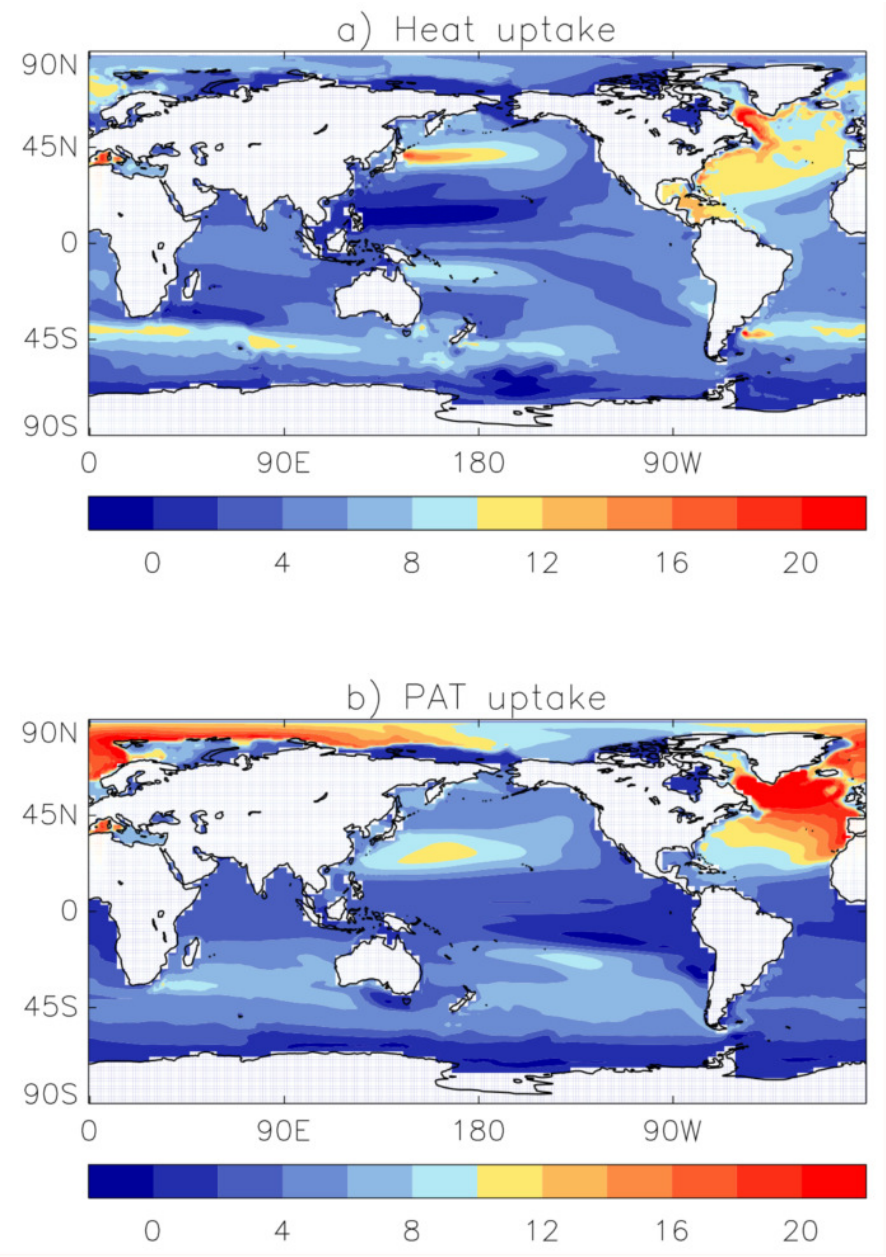

FIGURE 1.5: Column inventories of total heat and PAT (Passive Anomalous Tracer) uptake: climate change - control runs, reproduced from Banks and Gregory (2006).

(Buckley and Marshall). As such, heat uptake in the North Atlantic is associated with a reduction in northwards heat transport. The temperature change resulting from this heat transport change can be loosely thought of as a warming in the South Atlantic, and a cooling in the North Atlantic. Thus, ocean circulation changes act to reduce horizontal gradients in heat uptake. Whilst the response of dynamical change can not be described in such a simple fashion in other ocean basins, the processes are heuristically similar: see for example Newsom et al. (2022).

1.2.1.2 A general excess tracer

Extending Equation 1.3 to a general tracer, Q, and velocity field, $\vec{v}$, we obtain

$$(\vec{v}Q)' = \vec{v}Q - \vec{v}_0 Q_0 = \vec{v}_0 Q' + \vec{v}' Q_0 + \vec{v}' Q', \tag{1.4}$$

where again the subscript 0 indicates the stable equilibrium component, and the primed component the perturbation, with the sum of the two denoting the perturbed field:

$$Q = Q_0 + Q', \tag{1.5}$$

and likewise for $\vec{v}$. Following Winton et al. (2013), we may define a Redistribution Number for the tracer Q, denoted R_Q, as follows:

$$R_Q = \frac{v' Q_0}{v_0 Q'}, \tag{1.6}$$

where v' and v_0 represents magnitudes of our initial and perturbation velocity vectors, respectively. Thus, R_Q is an estimate of what fraction of the change in Q is due to the change of the velocity field acting to redistribute the background field of the initial equilibrium, and what fraction is due to the addition of the perturbed component of Q, advected by the equilibrium velocity field. An illustration of this is shown in Figure 1.6, in which the model dynamics are either fixed to preindustrial conditions ($\vec{v}' = 0$), or allowed to vary freely. It can easily be seen that the pattern of heat uptake is changed substantially, particularly in the Atlantic. Fixed current heat uptake strongly resembles that of PAT in Figure 1.5. Though carbon shows some differences between runs, patterns are largely similar. Note this is the total carbon field, DIC, and not anthropogenic carbon, C_{anth}.

When estimating R_Q, scales for each quantity must be used. For perturbations, this is approximately the size of the largest perturbation at any point in the ocean. For the background quantity, a natural scale to use is the range of values seen preindustrially. Estimating the ratio v'/v_0 is somewhat harder but can be done: details are given in Winton et al. (2013).

Computing redistribution numbers for Temperature, DIC and C_{anth} give redistribution numbers of $R_T \sim 3.6$, $R_{DIC} \sim 0.6$ and $R_{C_{anth}} = 0$. This implies that on a global scale, the majority of observed temperature changes are due to redistribution, which is not the case for DIC. The redistribution for C_{anth} is 0 as C_{anth} is constructed in order to be zero preindustrially, and so Q_0 is by definition zero for C_{anth}. This differing impact on heat and carbon as ocean circulation changes has important implications for understanding the link between heat and carbon inventory changes: whilst the distribution of excess heat and carbon are expected to be similar, this is not the case for redistribution.

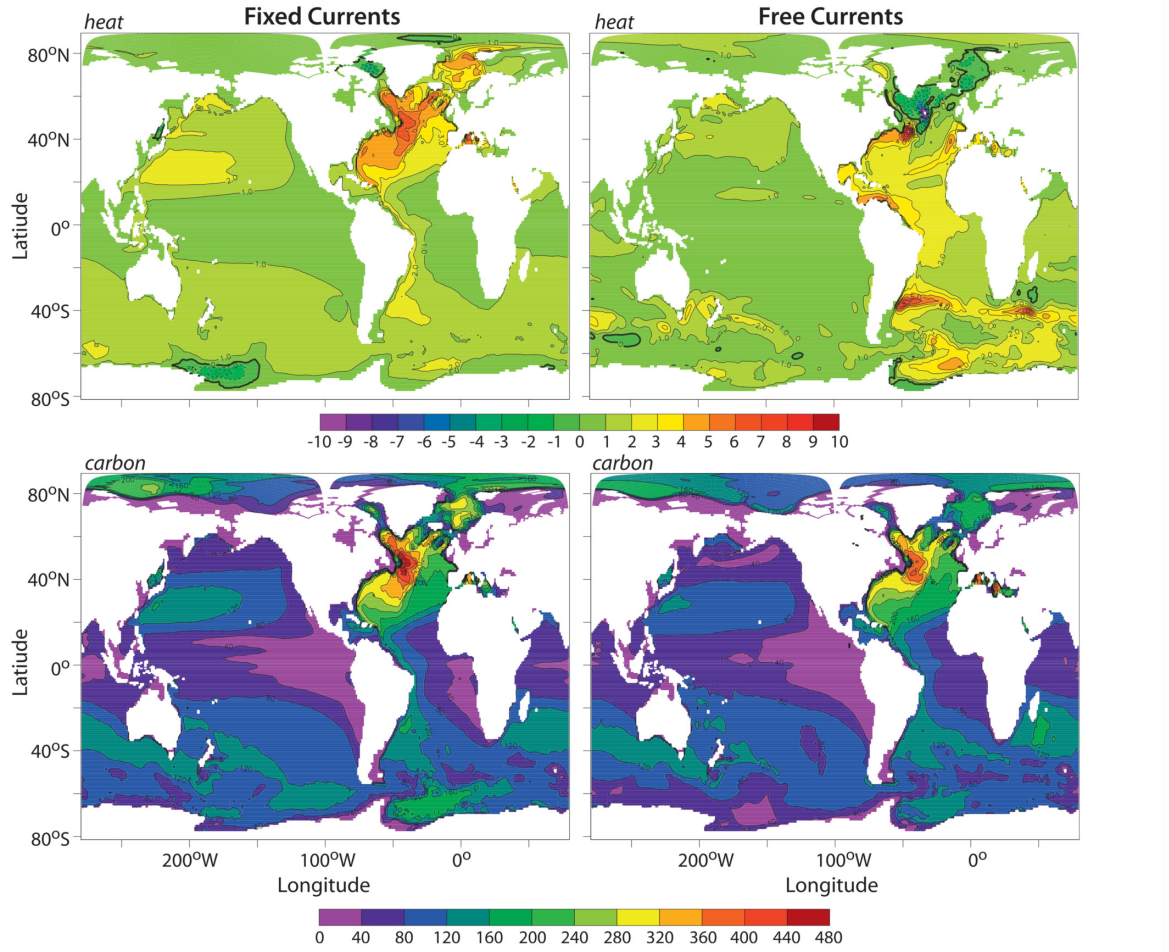

FIGURE 1.6: The effect of fixing currents or allowing them to vary with changing conditions under climate change model simulations, for DIC and heat content. Reproduced from Winton et al.. Thick black lines denote the zero contour.

1.2.1.3 Excess and redistributed salinity

Though Section 1.2.1 has focussed predominantly on how the preindustrial fields of ocean temperature and carbon respond to changes in circulation and the implications for understanding ocean heat and carbon uptake, the formulation is applicable to any ocean tracer. Temperature is not the only variable that impacts ocean circulation, with salinity also playing a fundamental role in controlling ocean circulation. So far, understanding 'excess' and 'redistributed' salinity in the same way remains a gap in our understanding of the response of ocean dynamics to anthropogenic global warming: Zika et al. (2021) notes that their reconstruction of material warming from observations (closely related to excess heat) also produces a material salinity change, but do not analyse this material change, and Turner et al. (2022) investigates whether modelled salinity change is due to either excess or redistribution of the preindustrial salinity field, but beyond this, the problem remains fundamentally unstudied.

As with temperature, salinity may change due to changes in surface fluxes, or due to changes in circulation. Whereas for temperature these surface fluxes are predominately the result of increased radiative forcing due to strengthening of the greenhouse effect, excess salinity is controlled by the response of atmospheric freshwater transport to the warming atmosphere and surface ocean. Previous model experiments, as well as observational studies of historical sea surface salinity (SSS) changes, indicate that the general pattern of ocean salinity change is that the fresh gets fresher and the salty gets saltier, as patterns of SSS become amplified (Stott et al. (2008), Durack and Wijffels (2010), Terray et al. (2012), Pierce et al. (2012), Skliris et al. (2014), Zika et al. (2018)). This is demonstrated in Figure 1.7: the patterns of SSS trends (Figure 1.7c) are reasonably well captured by the response of SSS to water cycle amplification (Figure 1.7d). Net freshwater fluxes are typically out of the ocean in the subtropical gyres (Panel a), and the highest sea surface salinities are generally found in these regions (Panel b), with the North Atlantic significantly saltier than the North Pacific. Generally, observations of trends in SSS show increasing salinity in these already salty regions, and freshening in fresher regions. Panels (d)-(h) show the response of sea surface salinity to idealised modelled forcings, for more details see Zika et al. (2018).

Unlike temperature, for which there is a large net flux of excess heat into the global ocean, globally integrated excess salinity changes are expected to be small: the only significant potential sources of a global net flux of excess salinity being the melting of the Greenland and Antarctic ice sheets, which act to freshen the global ocean by the addition of freshwater. This is a key distinction between excess heat and salinity: whilst excess heat accumulates in or is removed from the ocean due to to the addition or removal of heat, excess salinity changes are due to the addition or removal of freshwater, not salt (excess salt and excess salinity are theoretically distinct, unlike excess heat and excess temperature). As such, changes in excess salinity can be considered to be driven by the atmospheric response to climate change, rather than the ocean response as is the case for temperature. As the ocean has a much larger heat capacity than the atmosphere and longer timescales for response, we therefore expect excess salinity to be more easily detected than excess temperature, particularly for historical observations for which the global mean warming is small. This is already well known in the context of changes in sea surface salinity: Stott et al. (2008), Terray et al. (2012), Pierce et al. (2012) and Skliris et al. (2014), who suggest the salinification of the subtropical North Atlantic and freshening of the Western Pacific Warm Pool may constitute an early fingerprint of anthropogenic forcings.

In addition, it might be expected that the contribution of redistributed salinity to salinity change is smaller than that of redistributed temperature to temperature change: returning to the redistribution number 1.6, we may note that the range of preindustrial salinity values spans a smaller range than that of temperature. However, the size of the perturbations to surface salinity values is also different, and so it is not necessarily

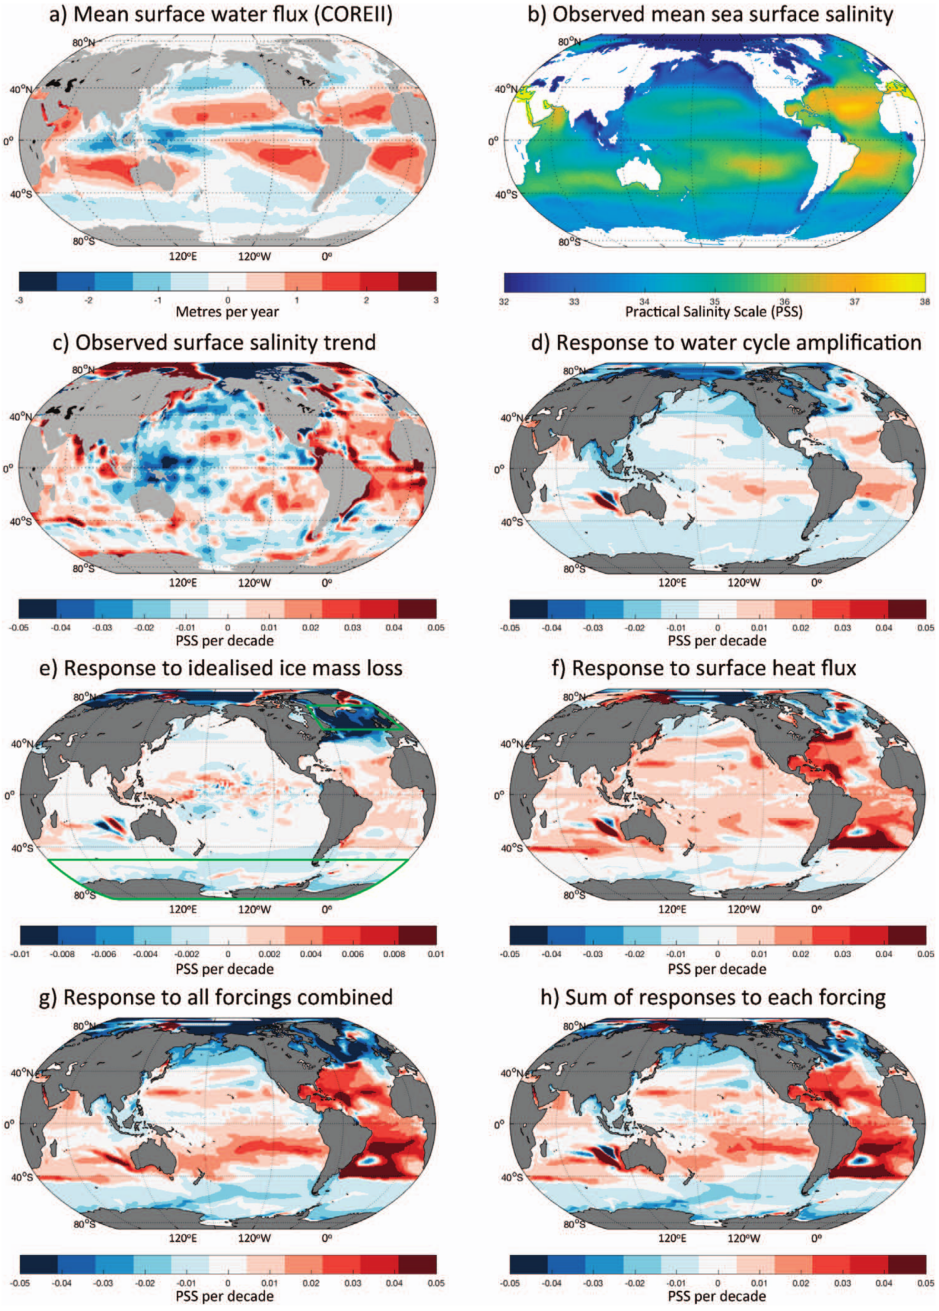

FIGURE 1.7: Surface water fluxes, sea surface salinity, and the response of sea surface salinity to various forcings, reproduced from Zika et al. (2018).

the case that the redistribution number is smaller for salinity than for temperature. We do expect negligible trends in the globally integrated excess salinity, except for a slight freshening due to the melting of the Greenland and Antarctic ice sheets (Antonov et al. (2002), Wadhams and Munk (2004)).

1.2.1.4 Excess DIC and Anthropogenic Carbon

In Section 1.2.1.2, it was noted that the redistribution number for C_{anth} is zero as anthropogenic carbon is defined to be zero preindustrially. From this, it may appear than excess DIC and anthropogenic carbon are the same quantity. Whilst closely analagous, this is not precisely the case: these differences are now explored. Returning to Equation 1.4, and including labels for the excess, redistributed, and preindustrial components:

$$\vec{v}Q = (\vec{v}_0 + \vec{v}')(Q_0 + Q') = \underbrace{\vec{v}_0 Q_0}_{Preindustrial} + \underbrace{\vec{v}' Q_0}_{Redistributed} + \underbrace{\vec{v}_o Q' + \vec{v}' Q'}_{Excess}, \tag{1.7}$$

we see that, for our arbitrary tracer Q, our definition of excess and redistributed Q is purely transport based: redistributed Q is simply the change in Q due to anomalous transport of the preindustrial field. Following Williams et al. (2021), we may therefore write the change in excess and redistributed Q explicitly as

$$Q_e(t) = \int_{t_0}^{t} \left(F'_Q - \left(\vec{v}_0 + \vec{v}'\right) \cdot \nabla Q' \right) dt, \tag{1.8}$$

and

$$Q_r(t) = -\int_{t_0}^{t} \left(\vec{v}' \cdot \nabla Q_0 \right) dt, \tag{1.9}$$

where F'_Q is the anomalous surface flux in Q, t_0 is a preindustrial time, and t is a generic time.

Now consider the global mean value of Q_r: the net velocity of the global ocean is zero, and so this quantity must be zero at all times. This result can also be intuited from considering the redistribution of a parcel of water from one location to another: warming at one location must be compensated for by equivalent cooling at another when the locations of two parcels of water are exchanged. As a result, all changes in the global inventory of Q must be the result of excess Q, rather than redistributed Q.

Now instead consider the response of the ocean carbonate system to both increasing atmospheric CO_2 and the consequent warming of the ocean. As noted in Section 1.2.0.1, anthropogenic carbon is the carbon content of a parcel of water resulting from its equilibrium with the increased atmospheric carbon content. However, at the same time, the warming of a parcel of water will reduce the amount of CO_2 which may dissolve in it due to the inverse temperature dependence of CO_2 solubility. Thus, saturation carbon will outgas in response to ocean warming. This effect constitutes a net loss of DIC in

response to ocean warming, separate from the changes in the ocean DIC inventory due to anthropogenic carbon: thus, anthropogenic carbon and excess DIC must differ. This effect is demonstrated in Figure 1.8, where the effect of increasing C_{anth} and outgassing C_{sat} on the global mean DIC inventory is shown. Note that the size of the preindustrial pool of DIC is not to scale: it is in fact significantly larger than the anthropogenic contribution.

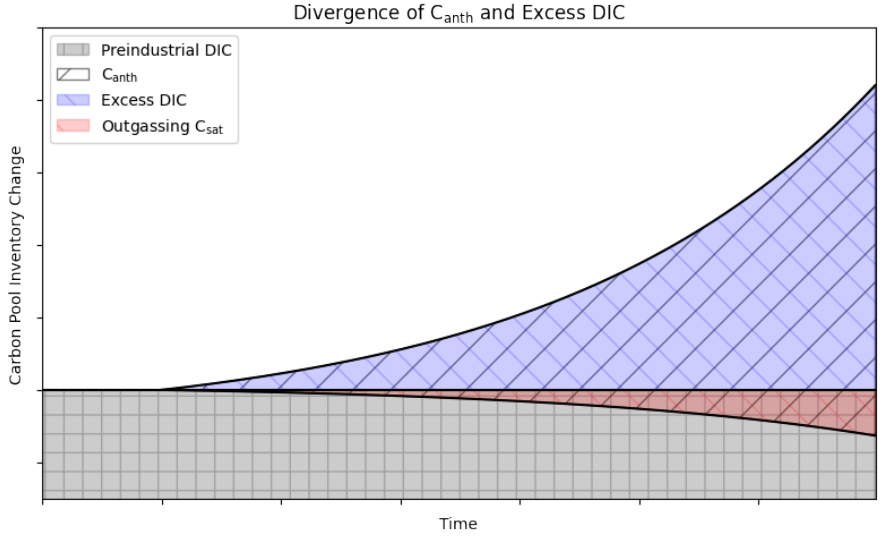

FIGURE 1.8: The difference between excess and anthropogenic carbon: the global C_{anth} inventory slightly exceeds that of excess DIC, as some C_{sat} outgasses in response to the ocean warming. This diagram is not to scale: preindustrial carbon concentrations far exceed those of anthropogenic carbon.

As a result of this outgassing of C_{sat}, excess DIC and anthropogenic carbon do differ systematically. However, on short timescales (sub decadal), the change in the two can be considered to be approximately equivalent. This is also the case for historic observations, where the reduction in C_{sat} is small, but this approximate equivalence will become less accurate over time as the ocean continues to warm.

1.2.1.5 Redistributed DIC and Natural Carbon

As a result of the close analogue between excess DIC and C_{anth}, there also exists a natural analogue between redistributed DIC and natural carbon (C_{nat}), the pool of carbon which remains when anthropogenic carbon is subtracted from total carbon.

Returning to Equation 1.1,

$$DIC = DIC^{PI}_{sat} + DIC_{carb} + DIC_{soft} + DIC_{diseq} + DIC_{anth}, \tag{1.1}$$

we may group the first four pools into 'Natural Carbon':

$$DIC = C_{nat} + C_{anth}, \tag{1.10}$$

where C_{nat} is therefore defined as

$$C_{nat} = DIC^{PI}_{sat} + DIC_{carb} + DIC_{soft} + DIC_{diseq}. \tag{1.11}$$

However, unlike C_{anth}, there are 2 issues with C_{nat} as an analogue to redistributed DIC:

1. Redistributed DIC is defined to globally integrate to zero, however this is not the case for C_{nat} (see Figure 1.8).
2. C_{anth} is defined to not contain any biologically driven carbon, whereas C_{nat} is.

For the issue of redistributed DIC integrating globally to zero but changes in C_{nat} not doing so, it might be noted that this is also an issue with the use of C_{anth} as an analogue to excess DIC. However, whilst C_{anth} will always be an overestimate (though not necessarily a large overestimate) of excess DIC, as in both will generally increase with atmospheric CO_2, this is not an issue: the patterns of both will be largely similar, with C_{anth} generally slightly exceeding excess carbon. However, this is not the case for redistributed DIC and C_{nat}: as redistributed carbon and temperature are defined to globally integrate to zero, patterns of reduced or increased natural carbon due to redistribution will be substantially and meaningfully altered if we find a trend in the global mean value of redistributed carbon content: for example, regions which should exhibit a slight decrease will instead show a slight increase. In general, in a warming climate, we expect the global pool of C_{nat} to decrease as saturation carbon outgasses due to the reduced solubility of CO_2 in warmer waters. One solution to this problem is to 'repartition' some C_{anth} back into our natural carbon pool to obtain an 'adjusted natural carbon' pool, correcting for this outgassing of saturation carbon. This process is described in detail in Chapter 3.

The second issue is more subtle, as it requires an understanding of the degree to which changes in biology alter the distribution of C_{nat} on a variety of timescales. Whilst changes in the strength of the biological carbon pump control the strength of the ocean carbon sink on millennial timescales and are therefore key to regulating atmospheric CO_2 concentrations, for our purposes, we are interested in their response to anthropogenically forced climate change, as well as their natural variability, on much shorter timescales: yearly to centennial. For the analogy between redistributed carbon and natural carbon to remain valid, we do not necessarily require that the strength of the biological carbon pumps remains constant: only that the trend in the carbon content of a parcel of water due to changes in the strengths of the biological carbon pumps remains small compared to the trend in the carbon content of a parcel of water due to the change in its saturation carbon content. We may write this as

$$\left\langle \left| \frac{d}{dt} C_{sat}(\mathbf{r}) \right| \right\rangle \gg \left\langle \left| \frac{d}{dt} C_{carb}(\mathbf{r}) + \frac{d}{dt} C_{soft}(\mathbf{r}) \right| \right\rangle, \tag{1.12}$$

where angle brackets represent an expected scale of change over a time period of interest, and modulo signs represent the absolute value of a quantity at a fixed spatial location, **r**. Therefore, so long as the scale of changes due to saturation carbon change are much larger than those due to biologically driven carbon change, over our period of interest, we can consider C_{nat} to be an appropriate approximation for redistributed carbon.

Finally, we may also note that changes in the soft tissue and carbonate DIC content at fixed spatial location may change not only due to changes in the soft tissue and carbonate pumps, but due to changes in ocean circulation: that is, soft tissue and carbonate carbon may be redistributed in the same way as saturation carbon. Thus, the condition stated in Equation 1.12 are can be relaxed to the component of the change in the soft tissue and carbonate carbon inventories due to time integrated changes in the strength of the biological carbon pumps, rather than the total change in the soft tissue and carbonate carbon. As the size of these carbonate pools is generally much smaller than that of saturation carbon, and that the magnitude of the change in the biological carbon pumps is generally thought to be small, we can consider the approximation of redistributed DIC with natural carbon change to be generally appropriate (Couldrey (2018), Chapter 4).

1.2.2 Thesis Structure and Scope

Due to the importance of the global ocean for both heat and carbon uptake under transient climate change (Khatiwala et al. (2012), Zanna et al. (2019)), and its importance in setting both the magnitude and form of the relationship between carbon emissions and global warming (Goodwin et al. (2015), Ehlert et al. (2017)), an improved understanding of the link between changes in ocean heat and carbon content represents an opportunity for improved understanding of the global climate response to continued anthropogenic carbon emissions. In addition, recent work has shown that the covariability of the background and added/excess heat and carbon in the ocean can be leveraged in order to understand whether regional ocean changes in heat content are due to circulation change or due to the accumulation of excess/added heat (Bronselaer and Zanna (2020), Williams et al. (2021)).

In this work, both the mechanisms coupling changes in ocean heat and carbon content, as well as the potential uses of these couplings are explored in detail. This is performed both in a Ocean General Circulation Model (OGCM), and also applied to observations. Additionally, the techniques which will be developed, though initially designed for the decomposition of temperature, will also be applied to salinity, as it is also a key driver of ocean circulation, and responding strongly to anthropogenic climate change.

This first chapter has provided a brief introduction to ocean tracers, and their decomposition in an excess/redistribution framework. It has also given a brief introduction into the decomposition of temperature, salinity, and carbon, in both the excess/redistribution framework, and also the more traditional decomposition of DIC into its constituent pools.

In the second chapter, theories and mathematical structures which relate temperature and carbon in the ocean, as well as the atmosphere, will be introduced. This will show why excess temperature and anthropogenic carbon are so closely linked, as well as why this link is expected to be temperature dependent. The covariability of the background temperature and carbon fields will then be explored, introducing the notion of a redistribution coefficient: a number which informs how we expect two tracers to covary as a result of ocean circulation change. The two relationships between the excess/anthropogenic components, and the background components, are then combined, showing how the excess and redistributed temperature, salinity, and DIC can be obtained directly in temperature-carbon space.

In the third chapter, the relationship between redistributed temperature and salinity and natural carbon are explored in depth, using the NEMO OGCM. Relevant theory introduced in the second chapter is expanded in greater detail, and it is shown how this approach can be used to determine the mechanisms of ocean temperature and salinity change. This chapter was previously published in Ocean Sciences (Turner et al., 2022), and the material is broadly the same; however, some changes to the order of the content have been made, due to the extensive appendices in the published work.

In the fourth chapter, the assumptions made in Chapter 3 are explored in greater depth, in order to explore the limitations of the validity of the decomposition. Techniques which would enable this the extension of this method are proposed, and the validity of some assumptions are tested explicitly. However, implementation of these proposed techniques is left for future work.

In the fifth chapter, the temperature-carbon space approach described in Chapter 2 is explored in detail, for the A05 hydrographic section: a well observed ocean location, which is known to be important for both heat and carbon uptake and storage. The theory introduced in Chapter 2 is expanded, and the sensitivity of results to the choice of parametrisations utilised is explored. Uncertainties are explored in detail, and it is shown how this technique can reliably identify excess carbon through comparison with previously existing anthropogenic carbon estimates.

In the sixth chapter, the approach explored in detail in Chapter 5 is applied to the global ocean, to produce global estimates of excess heat and salinity accumulation, using the GLODAP dataset (Lauvset et al., 2021). The implications of these global excess heat and salinity accumulation fields are then explored in detail. Chapter 7 then presents some final conclusions, as well as investigating how the work presented here might be

extended in future to improve understanding of ocean heat and carbon uptake. Finally, Appendix A describes the mapping software used in chapters five and six.

Chapter 2

Theory: Transient Climate Responses, Carbon-Heat Coupling, Redistribution Coefficients and Temperature-Carbon Space

In this Chapter, the notion of a Transient Climate Response is introduced and explored. Transient Climate Responses refer to the reliable ways in which certain climatic parameters change in response to a change in forcing. The Transient Climate Response to Emissions (TCRE) is first introduced, followed by a less commonly used emergent relationship between changes in ocean heat and carbon content under anthropogenic climate change, known as the Carbon-Heat coupling. This Carbon-Heat coupling allows for the identification of excess heat from anthropogenic carbon. Following this, a new derivation for this Carbon-Heat coupling, which allows for this coupling to vary in space is presented. Biogeochemical reasons for this spatial varying coupling are explored, before an approach which allows for the explicit identification of this spatial variability is then introduced and explored. This method uses orthogonal assumptions, aiming to understand how the spatial distribution of two tracers covary in the ocean: in this case, to understand the redistribution of heat from natural carbon. Finally, these two methods for understanding how heat and carbon covary in the ocean are combined into a single method in Temperature-Carbon space.

2.1 Transient Climate Responses

2.1.1 The Transient Climate Response to cumulative carbon Emissions (TCRE)

Equilibrium climate sensitivity is a metric which has been used in a number of climate studies in order to estimate the response of the climate system to an instantaneous doubling of atmospheric CO_2 from preindustrial levels (Cubasch et al.). It is defined as the global mean temperature change after the climate equilibrates to this doubling of atmospheric CO_2. This measure was of historical utility as it enabled estimates of the response of the climate system to anthropogenic forcing, at a time when computational power was significantly more limited than it is presently.

A more realistic measure of climate change is the Transient Climate Response (TCR). This is defined as the global mean surface air temperature change for an AOGCM (Atmosphere-Ocean General Circulation Model) at the time of CO_2 doubling in a 1%/yr CO_2 increase experiment.

Beginning with CMIP5 (the Coupled Model Intercomparison Project, phase 5), a global coordination effort to produce robust estimates of climate change using a range of models, climate change experiments have been performed using emission driven scenarios to estimate atmospheric CO_2 and climate change, as well as just specifying atmospheric CO_2 as was the case in CMIP3 (Taylor et al., 2012) (N.B. There was no CMIP4). This allowed for the extension of the Transient Climate Response to that of the Transient Climate Response to Cumulative Carbon Emissions (TCRE), incorporating a more realistic response of the terrestrial biosphere and ocean to emissions. Note that some early literature regarding the TCRE is referred to as the 'Carbon Climate Response' (CCR), such as Matthews et al. (2009).

The TCRE is known to be both approximately constant over time, and approximately independent of emissions scenario, (Matthews et al., 2009), (Zickfeld et al., 2009), (Allen et al., 2009), with a value of 0.7-2.0K/1000PgC (Gillett et al., 2013): this is shown in Figure 2.1. This linearity is known to be predominantly a result of ocean processes controlling the planetary response to carbon emissions (Goodwin et al., 2015), (Ehlert et al., 2017), (Katavouta et al., 2018), (Katavouta et al., 2019).

Following Goodwin et al. (2015), we may write a perturbation to the ocean, atmosphere and terrestrial carbon system as

$$M\delta\mathrm{CO}_2 + V\delta\overline{\mathrm{DIC}}(t) + \delta I_{ter} = I_{em}(t), \tag{2.1}$$

where M is the molar volume of the atmosphere, V the volume of the ocean, I_{ter} the amount of carbon in the terrestrial biosphere and I_{em} cumulative carbon emissions.

Assuming biological activity remains unchanged, we may write

$$\delta\overline{\mathrm{DIC}}(t) = \delta\overline{\mathrm{DIC}_{sat}}(t) + \delta\overline{\mathrm{DIC}_{res}}(t), \tag{2.2}$$

where $\mathrm{DIC}_{sat}(t)$ is the saturated DIC concentration of a water parcel brought to equilibrium with the atmosphere at time t, and $\mathrm{DIC}_{res}(t)$ the residual DIC concentration: the total DIC concentration, minus $\mathrm{DIC}_{\mathrm{sat}}$ and DIC due to the biological pump, which may be positive or negative. Now, defining the buffered carbon inventory of the ocean, I_B, as

$$I_b = I_A + \frac{V\overline{\mathrm{DIC}}}{B}, \tag{2.3}$$

following Goodwin et al. (2009), where is I_A total preindustrial atmospheric CO_2, $V\overline{\mathrm{DIC}}$ the total ocean carbon inventory and B the global mean Revelle/Buffer factor (Sarmiento and Gruber, 2006). B is defined as

$$B = \frac{\mathrm{DIC}}{\mathrm{CO}_2}\frac{\partial \mathrm{CO}_2}{\partial \mathrm{DIC}}, \tag{2.4}$$

and relates the change in DIC for a parcel of water at the oceans surface for a given change in CO2. Here, CO_2 refers to atmospheric CO_2 in ppm (parts per million), and DIC global mean DIC in μmol/kg. Combining Equations 2.1-2.4, we obtain:

$$I_b\delta\ln \mathrm{CO}_2(t) = I_{em}(t) - V\delta\overline{\mathrm{DIC}}(t) - \delta I_{ter}(t). \tag{2.5}$$

Since $I_b \approx$ const, $\overline{\mathrm{DIC}_{res}}(t_{PI}) \approx 0$ and $\delta I_{ter} \approx 0$, where t_{PI} denotes a preindustrial time, we may write

$$\Delta\ln \mathrm{CO}_2 = \frac{I_{em}(t) + I_{usat}(t)}{I_B}, \tag{2.6}$$

where $I_{usat}(t) = -V\overline{\mathrm{DIC}_{res}}(t)$. We therefore see that the change in atmospheric CO_2 levels is set predominantly by the response of the ocean. Goodwin et. al. then note that the radiative forcing due to atmospheric CO_2 is a logarithmic function of atmospheric CO_2 concentrations (Myhre et al., 1998), and so obtain a theoretical linear relationship between emissions and global mean surface warming.

Additional work has been done on extending this warming to emissions approach to local surface area temperature (SAT) change (Leduc et al., 2016), and to a range of other variables, such as surface ocean pH (Steinacher and Joos, 2016). A result of particular note from Leduc et al. (2016) is that in general, high latitude, ocean RTCRE's deviate more from the relatively uniform value seen above much of the ocean surface, typically showing greater warming: this is shown in Figure 2.2. Ocean values are only approximately 2/3 of land values, though the fractional spread between models is similar. However, both the expected warming and their uncertainties are highest, in general, over the Arctic ocean. An implication of this is that we may expect to see greater warming, relative to carbon uptake, in these high latitude regions. As these regions are

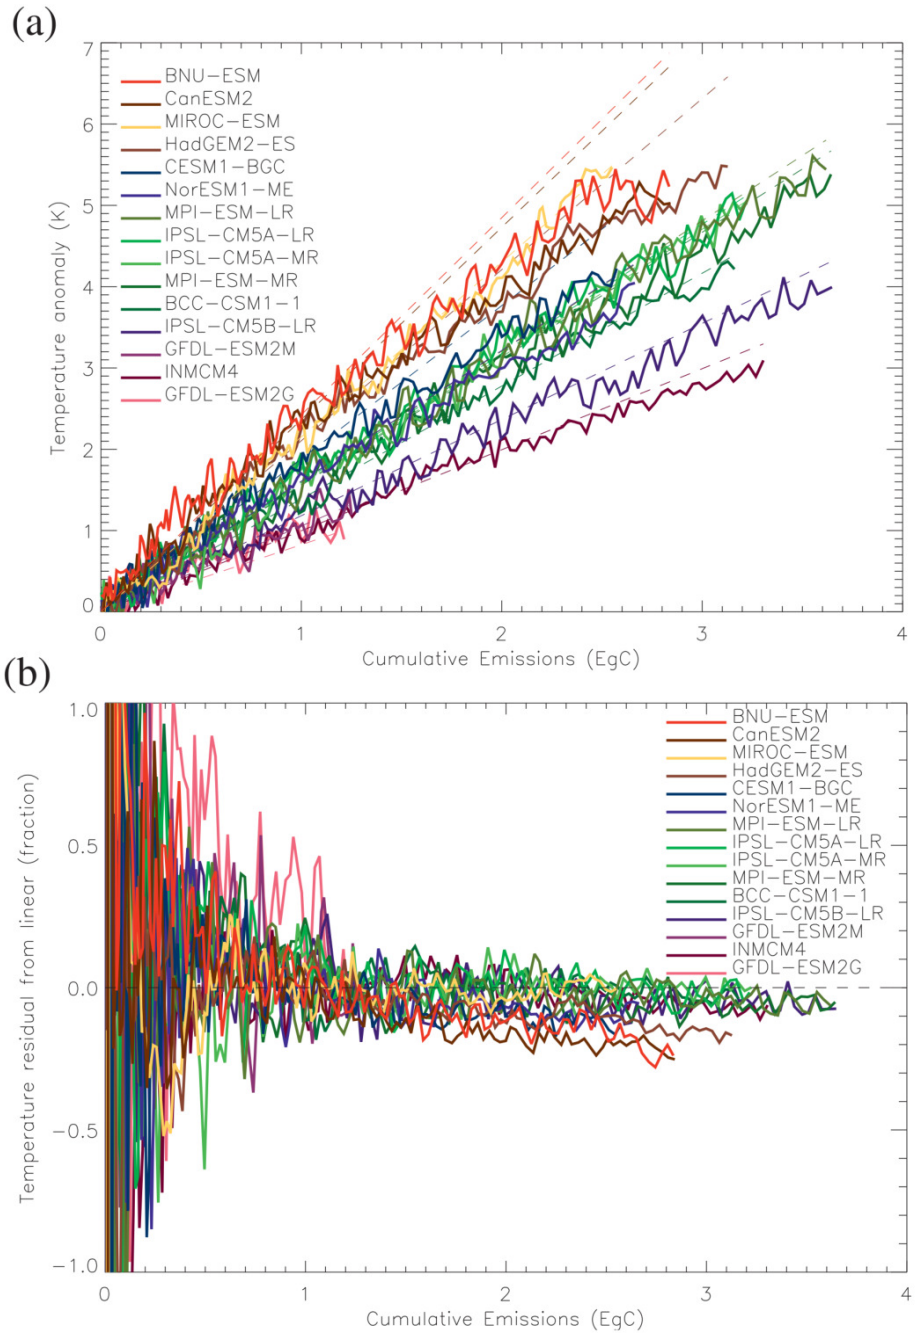

FIGURE 2.1: Global mean warming (a) in response to a 1%/year CO_2 increase, and deviations from a linear relationship (b), from 15 CMIP5 models. Here, CO_2 emissions are calculated rather than prescribed. All models show approximately linear behaviour over a large range of emissions (3EgC $\equiv$ 3000PgC): current emissions total approximately 600PgC (Stocker et al., 2013). Reproduced from Gillett et al. (2013).

crucial for deep ocean heat uptake, we may therefore expect to see a greater ratio of warming to C_{anth} uptake in the high latitude and deep ocean than we do in the upper ocean and at lower latitudes.

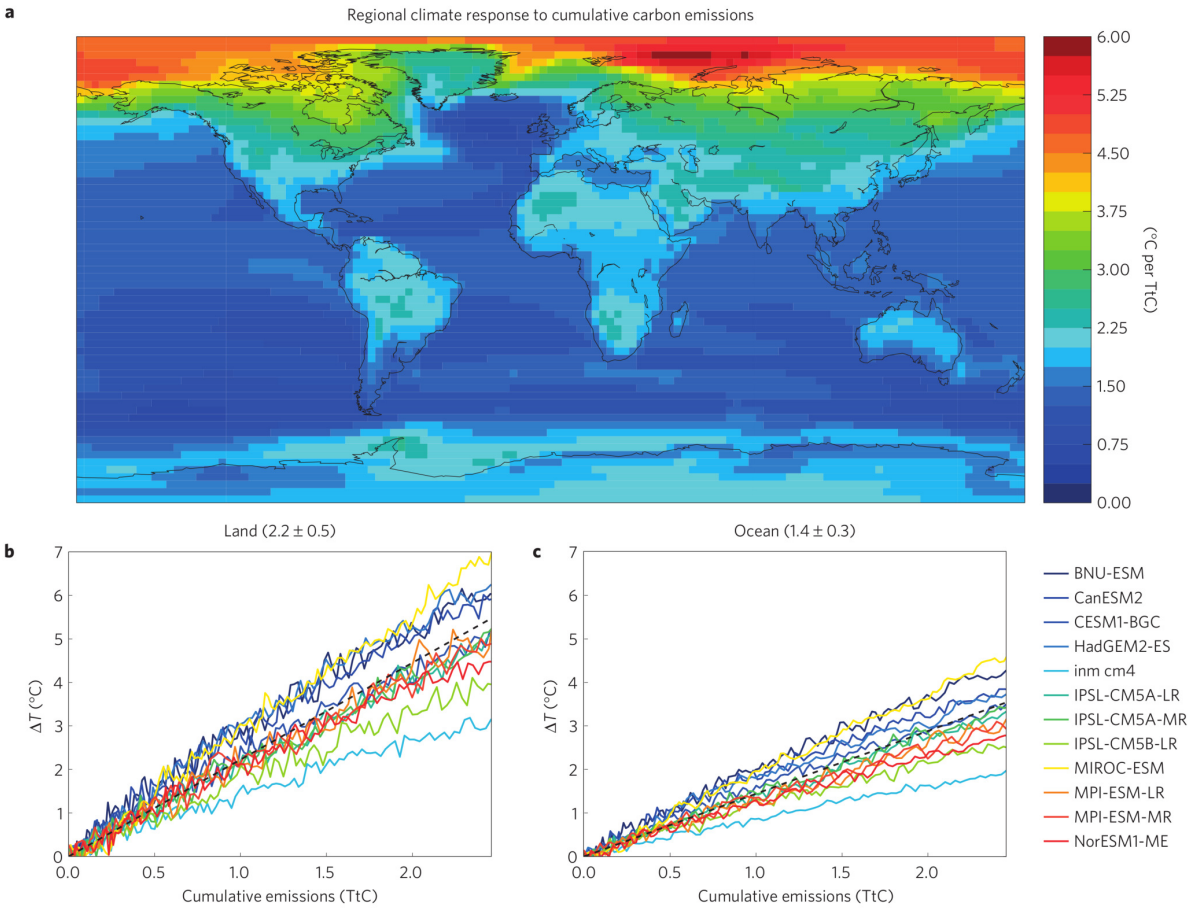

FIGURE 2.2: Maps of the ensemble mean RTCRE (a), for the ensemble shown to the left of panel (c). Panels (b) and (c) show land and ocean mean RTCRE values for each model. Reproduced from Leduc et al. (2016).

2.1.2 An oceanic analogue to the TCRE

So far, little research has been performed into understanding an ocean analogue to the TCRE. Bronselaer and Zanna (2020), the only paper to investigate such an analogue, labelled it the carbon-heat coupling, $\alpha(t)$. Similarly to Goodwin et al. (2015), they derive a theoretical value from the interaction of ocean carbonate chemistry with increased atmospheric CO_2, and the radiative forcing of CO_2. By approximating this ratio as globally uniform, they are able to use patterns of C_{anth} storage to approximate those of excess heat storage. However, Turner et al. (2022) was able to show that this approximation may not be entirely appropriate due to variability in the ratio of heat and carbon uptake at high latitude, as discussed in the previous section. In this section, the derivation of $\alpha(t)$ by Bronselaer and Zanna (2020) is presented. In the following section, an alternative derivation which produces a slightly different result but allows for spatial variability in the carbon-heat coupling. The two approaches are then contrasted. Note that a slightly different notation to that presented in Bronselaer and Zanna (2020) is used here, however, the formulations are equivalent.

From Myhre et al. (2013) we may express the change in global mean surface air temperature (SAT) as

$$\Delta T(t) = \frac{R(t) - N(t)}{\lambda}, \tag{2.7}$$

where $R(t)$ is TOA radiative forcing at time t, $N(t)$ the net heat flux into the earth system and λ the climate feedback parameter. λ is defined as $T_{eq} = 1/\lambda$, where T_{eq} is equilibrium climate sensitivity. We may then rearrange for

$$N(t) = R(t) - \lambda \Delta T(t), \tag{2.8}$$

and note virtually all heat flux into Earth system goes into the ocean, such that

$$N(t) \approx \frac{\hat{H}(t)}{\Delta t} = R(t) - \lambda \Delta T(t), \tag{2.9}$$

where $\hat{H}(t)$ is cumulative global heat uptake and Δt the time elapsed since preindustrial time. Note that there is a factor of A, the Earth's surface area, in the derivation given by Bronselaer and Zanna (2020), which has been absorbed into $N(t)$ for convenience. We then use the form of the dependence of $R(t)$ on atmospheric CO_2:

$$R(t) = a\Delta \ln CO_2(t) = a \ln \left(\frac{CO_2(t)}{CO_2(t_0)} \right) \tag{2.10}$$

where t_0 refers to preindustrial time, usually taken to be either 1765 or 1850. a is a constant with value $5.35 W/m^2$.

Now, using 2.10, we can rewrite the above as

$$N(t) \approx \frac{\hat{H}(t)}{\Delta t} = a\Delta \ln CO_2(t) - \lambda \Delta T(t), \tag{2.11}$$

and rearrange for

$$\hat{H}(t) = \Big(a\Delta \ln CO_2(t) - \lambda \Delta T(t) \Big) \Delta t \tag{2.12}$$

Now, we take the Buffer/ Revelle Factor, which was previously defined as:

$$\mathrm{B} = \frac{\partial CO_2}{CO_2} \frac{DIC}{\partial DIC}. \tag{1.16}$$

We may rewrite this in the form

$$\mathrm{B} = \frac{\partial \ln CO_2}{\partial \ln DIC}, \tag{2.13}$$

using the identity

$$\partial \ln x = \frac{\partial x}{x} \tag{2.14}$$

Assuming that the ratio DIC/B remains approximately constant, we can take $\delta \rightarrow \Delta$ and rearrange to obtain

$$\Delta \ln CO_2(t) = \frac{B}{DIC}\Delta DIC. \tag{2.15}$$

Inserting this into Equation 2.12, we get

$$\hat{H}(t) = \left(a\frac{B}{DIC}\Delta DIC - \lambda \Delta T(t)\right)\Delta t \tag{2.16}$$

Now, we need to make another assumption: that total DIC change is the same as the global C_{anth} inventory: $\hat{C}_{anth} = V_o \Delta DIC$, where V_0 is the global ocean volume. The limitations of this assumption have been discussed in Section 1.2.1.4. This allows us to write

$$\hat{H}(t) = \left(\frac{aB}{I_o}\hat{C}_{anth} - \lambda \Delta T\right)\Delta t, \tag{2.17}$$

where I_o is the global ocean preindustrial carbon inventory. Now, we can simply divide by the C_{anth} uptake to get the ratio of cumulative heat uptake to carbon uptake, obtaining $\alpha\hat{(t)}$:

$$\hat{\alpha}(t) = \left(a\frac{B}{I_o} - \lambda(t)\frac{\Delta T(t)}{\hat{C}_{anth}(t)}\right)\Delta t. \tag{2.18}$$

$\hat{\alpha}(t)$ thus relates the global inventory of excess heat and anthropogenic carbon:

$$\hat{\alpha}(t) = \frac{\hat{H}(t)}{\hat{C}_{anth}(t)} \tag{2.19}$$

In this approximation $\hat{\alpha}$ increases linearly with time (ignoring the C_{anth} term), and so $\hat{\alpha}$ may be estimated as

$$\hat{\alpha}(t) = \frac{1}{2}\frac{d\hat{H}}{dt} \Big/ \frac{d\hat{C}_{anth}}{dt}, \tag{2.20}$$

where estimates of the rates of change of the global ocean heat and carbon inventories are at the present time. To obtain the factor of $1/2$, note that under a linear approximation, we may write

$$\hat{\alpha}(t) = \frac{1}{2}\left[\frac{d\hat{H}}{dt} \Big/ \frac{d\hat{C}_{anth}}{dt}\bigg|_{t_0} + \frac{d\hat{H}}{dt} \Big/ \frac{d\hat{C}_{anth}}{dt}\bigg|_{t}\right], \tag{2.21}$$

where t_0 is a preindustrial time and t is the present day. Noting that the preindustrial rates of change are zero, we obtain the expression given by Equation 2.20.

Bronselaer and Zanna (2020) then use a value for $d\hat{C}_{anth}/dt$ from DeVries (2014) of 2.6±2% PgC/yr, and 17.44±3.64 ZJ/yr from Zanna et al. (2019) to obtain a contemporary estimate of $\hat{\alpha}$ of 41.7± 9.6 MJ/mol (their Figure 2a, observations, white dot, reproduced in Figure 2.3). This corresponds to approximately 0.01±0.0025 K/μmol/kg.

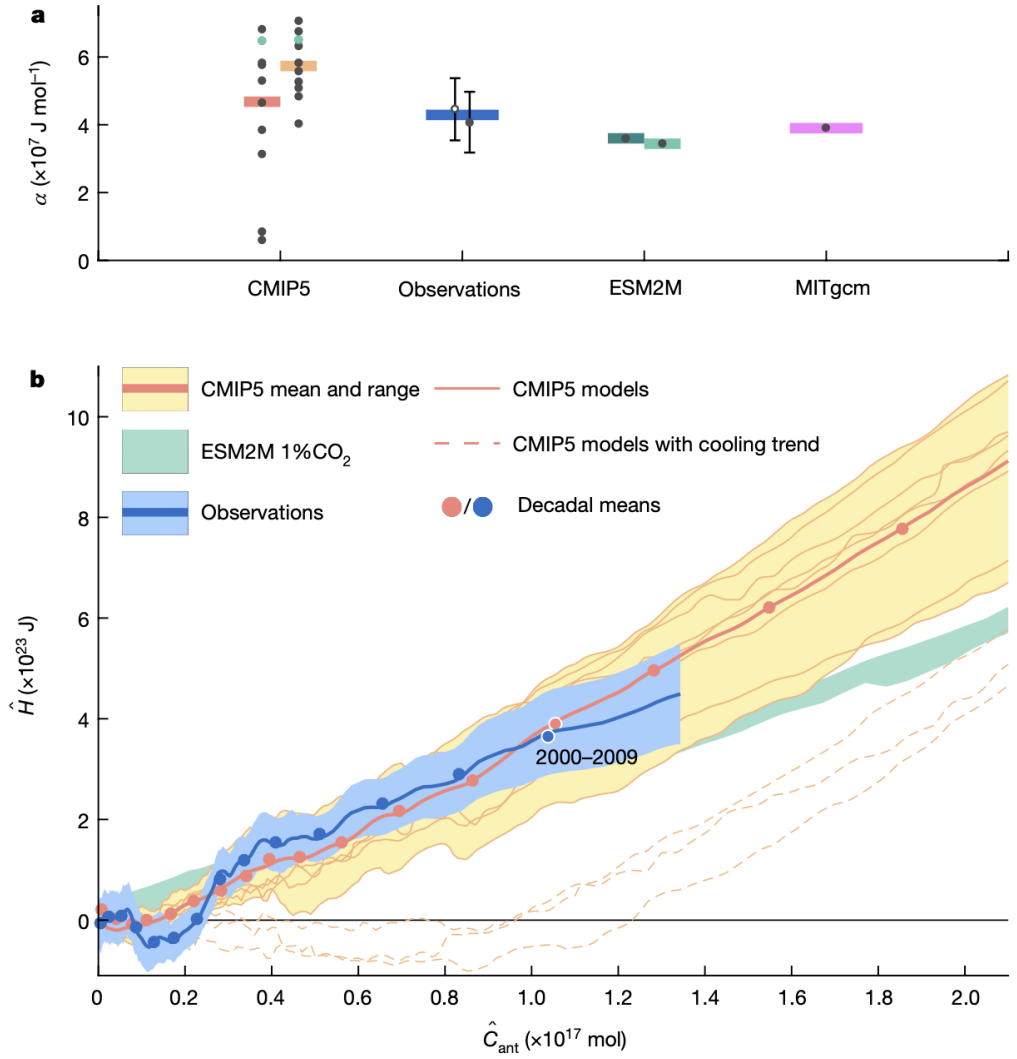

FIGURE 2.3: Obeservational and model estimates of $\hat{\alpha}$, reproduced from Bronselaer and Zanna (2020).

2.2 Deriving the carbon-heat coupling

In Section 2.1.2, the derivation of the carbon-heat coupling by Bronselaer and Zanna (2020) was presented. Here, a derivation which allows for spatial variability by making slightly different assumptions is also given. The two approaches are then contrasted.

Equations 2.7 and 2.10 may be combined to obtain

$$\Delta T(t) = \frac{1}{\lambda} a \Delta \ln CO_2(t) - \varepsilon N(t). \tag{2.22}$$

Then, rearranging Equation 2.15 into the form

$$\Delta \mathrm{DIC} = \frac{\mathrm{DIC}}{\mathrm{B}} \Delta \ln CO_2, \tag{2.23}$$

we may divide Equation 2.22 by Equation 2.23 to obtain

$$\frac{\Delta T}{\Delta \mathrm{DIC}} = \frac{\mathrm{B}}{\mathrm{DIC}} \frac{1}{\lambda} \left\{ a - \frac{\varepsilon N(t)}{\Delta \ln \mathrm{CO_2}} \right\}. \tag{2.24}$$

Note that in Equation 2.24, $\Delta T(t)$ refers to the change in global mean surface air temperature (SAT), not sea surface temperature (SST). We therefore need to apply a coupling factor, labelled k:

$$\frac{\Delta T_{SS}}{\Delta \mathrm{DIC}} = \frac{\mathrm{B}}{\mathrm{DIC}} \frac{k}{\lambda} \left\{ a - \frac{\varepsilon N(t)}{\Delta \ln \mathrm{CO_2}} \right\}. \tag{2.25}$$

Here, the subscript SS has been added to ΔT in order to make it clear this is a change in SST. k is defined as

$$k = \frac{\Delta \mathrm{SST}}{\Delta \mathrm{SAT}}, \tag{2.26}$$

where ΔSST and ΔSAT refer to global mean changes.

In order to estimate a value for $\Delta T_{SS}/\Delta \mathrm{DIC}$, we need estimates of all the quantities used in Equation 2.25. Steinacher and Joos (2016) investigated the transient response of a number of earth system parameters, including global mean SST and SAT, to carbon emissions in a large ensemble modelling experiment. They found that over this ensemble, global mean SST changes by 1.30K/1000PgC, and global mean SAT by 1.75K/1000PgC, with both quantities having uncertainties of approximately 35% (in general, greater SAT changes are seen over land than sea, and thus this does not indicate a change in the air-sea temperature difference - it merely represents the different regions considered in these metrics). A naive calculation therefore gives $k = 0.74 \pm 49\%$. However, due to the large heat capacity of seawater and extent of coverage of the global ocean, global mean SST and global mean SAT show strong covariability (see for example Li et al. (2020)). A better estimate of the uncertainty on k is therefore obtained by comparison of the change in the ratio k over the different periods for which global mean SST and SAT changes are computed in Steinacher and Joos (2016): this gives $k = 0.72 \pm 0.03$.

Estimates for all quantities other than the ratio B/DIC may be found Winton et al. (2010), Table 2, and Myhre et al. (1998). They are as follows:

- $1/\lambda$ = 3.4 ± 0.17K
- a = 5.35 ± 0.54 $\mathrm{Wm^{-2}}$
- ε = 1.34 ± 0.07
- $N(t) \approx$ const = 1.29 ± 0.06 $\mathrm{Wm^{-2}}$
- $\Delta \ln \mathrm{CO_2}$ = 0.41

Note that although $N(t)$ is time dependent, it can be treated as a constant. Values of 280 and 420 ppm have been used for preindustrial and present day CO_2, and uncertainty

in this ratio is assumed to be negligible. Winton et al. (2010) gives these values at the time of CO_2 doubling, but as they are all time independent, except $N(t)$, which we may approximate as constant, their values should hold.

The ratio B/DIC may be estimated using GLODAPv2 data: this gives 2034μmol/kg for DIC and 11.2 for the Buffer Factor in the upper 20m of the ocean, with a combined uncertainty of approximately 30%. This comprises the largest source of uncertainty in our estimate of $\Delta T_{SS}/\Delta$DIC.

Using these values, we obtain an estimate for α of 0.0195$\pm$ 0.006 K/(μmol/kg), at current atmospheric CO_2 levels. This is, as expected, double the value obtained for $\hat{\alpha}$ by Bronselaer and Zanna (2020) (after transforming units): the surface ocean is in direct contact with the overlying atmosphere, and so we expect this time evolving ratio to be exactly that obtained by comparing the rate of heat and carbon uptake by the ocean, whereas the global ocean mean value will be half this.

In this form, provided that the ratio B/DIC remains constant, we see that we expect the value of $\Delta T_{SS}/\Delta$DIC to asymptote. This approximate asymptotic behaviour is also seen in Bronselaer and Zanna (2020) (Extended Data, Figure 2), but is less clear from their formulation, which has a linear time dependence that is balanced by a increasingly large negative contribution. Assuming this constant B/DIC ratio, Equation 2.25 asymptotes to

$$\frac{\Delta T_{SS}}{\Delta \mathrm{DIC}} = \frac{\mathrm{B}}{\mathrm{DIC}} \frac{k}{\lambda} a. \tag{2.27}$$

We may therefore assume that the ratio of changes in DIC and temperature will, as changes in atmospheric CO_2 become large, be linked by a constant ratio, provided that the ratio B/DIC is fixed. From Egleston et al. (2010) we see that we may alternatively define the Buffer Factor as

$$\mathrm{B} = \frac{\mathrm{DIC}}{\gamma_{\mathrm{DIC}}}, \tag{2.28}$$

where γ_{DIC} is defined as

$$\gamma_{\mathrm{DIC}} = \left(\frac{\partial \ln \mathrm{CO_2}}{\partial \mathrm{DIC}} \right)^{-1}. \tag{2.29}$$

Therefore, provided γ_{DIC} is a linear function of DIC, the ratio B/DIC remains constant. From Egleston et al. (2010) (Figure 2, reproduced in Figure 2.4) we see that we may approximate changes in γ_{DIC} from present values of around 2000 μmol/kg as being linearly dependent on DIC until global mean surface concentrations reach approximately 2200 μmol/kg. This is not expected to happen until approximately the end of the 21$^{\mathrm{st}}$ century following a RCP8.5 scenario (Riahi et al., 2011) and so whilst the asymptotic behaviour of Equation 2.25 does break down, it remains a valid approximation both historically and over the coming decades (Schwalm et al., 2020).

It is clear that these two equations describing the ratio of temperature to carbon change are intimately linked, with Equation 2.25 relating the instantaneous ratio of warming to

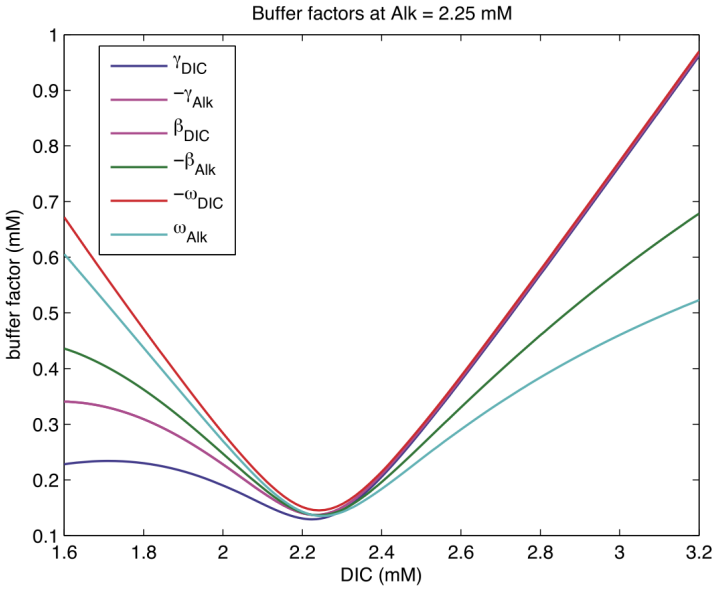

FIGURE 2.4: The various buffer factors explored in Egleston et al. (2010), and reproduced here. The DIC buffer factor (γ_{DIC}, purple line) can be approximated as linearly dependent on DIC until global mean surface DIC values of approximately 2000μmol/kg.

carbon increase for a 'global mean parcel of water' at the oceans surface, and Equation 2.19 relating the ratio of global heat and carbon content change.

An advantage of the form given by Equation 2.25 is that it straightforwardly allows for variability in the ratio $\Delta T_{SS}/\Delta\mathrm{DIC}$: consider an arbitrary parcel of water, rather than the global mean parcel of water. For such a parcel, k_p (the parcel value of k) may vary significantly: for example, as a result of the polar amplification of warming or the latitudinal variability of B/DIC, we might expect parcels of water subducted at high latitudes to have intensified warming or decreases $\mathrm{C_{anth}}$ uptake and therefore higher values of k_p than the global mean. Thus, we may expect a greater ratio of heat uptake to carbon uptake in water masses formed at high latitude. However, Bronselaer and Zanna (2020) provide an argument for a globally uniform ratio $\alpha(t)$: this argument is now explored.

Consider the change in the surface ocean partial pressure of $\mathrm{pCO_2}$, $\Delta\mathrm{pCO_2}$, which is a function of temperature (T), salinity (S), alkalinity (Alk), and DIC:

$$\Delta\mathrm{pCO_2}(\mathrm{DIC},\mathrm{Alk},\mathrm{S},\mathrm{T}) = \frac{\partial\mathrm{pCO_2}}{\partial\mathrm{DIC}}\Delta\mathrm{DIC} + \frac{\partial\mathrm{pCO_2}}{\partial\mathrm{Alk}}\Delta\mathrm{Alk} + \frac{\partial\mathrm{pCO_2}}{\partial\mathrm{T}}\Delta\mathrm{T} + \frac{\partial\mathrm{pCO_2}}{\partial\mathrm{S}}\Delta\mathrm{S} \tag{2.30}$$

$$\approx \frac{\partial\mathrm{pCO_2}}{\partial\mathrm{DIC}}\Delta\mathrm{DIC} + \frac{\partial\mathrm{pCO_2}}{\partial\mathrm{Alk}}\Delta\mathrm{Alk} + \frac{\partial\mathrm{pCO_2}}{\partial\mathrm{T}}\Delta\mathrm{T}, \tag{2.31}$$

where the contribution from changes in salinity (addition or removal of freshwater) is negligible compared to changes due to changes in temperature, alkalinity or DIC.

Over the timescales for which anthropogenically driven climate change is significant (Bronselaer and Zanna (2020) show changes at the time of CO_2 doubling), these changes are dominated by changes in DIC. Using the definition of the Buffer Factor,

$$B = \frac{\partial pCO_2}{pCO_2}\frac{DIC}{\partial DIC}, \tag{1.16}$$

we may rewrite the DIC term as

$$\frac{\partial pCO_2}{\partial DIC}\Delta DIC = pCO_2\frac{B_{DIC}}{DIC}, \tag{2.32}$$

where I have relabelled $B \to B_{DIC}$. Now, similarly defining an Alkalinity Buffer Factor (following for example Egleston et al. (2010)) as

$$B_{Alk} = \frac{\partial pCO_2}{pCO_2}\frac{Alk}{\partial Alk}, \tag{2.33}$$

and using the empirical relationship derived by Takahashi et al. (1993):

$$\frac{\partial pCO_2}{\partial T} = 0.0423 \times pCO_2, \tag{2.34}$$

we may rewrite Equation 2.31 in the following form:

$$\Delta pCO_2 = pCO_2\big(0.0423\Delta T + \frac{B_{DIC}}{DIC}\Delta DIC + \frac{B_{Alk}}{Alk}\Delta Alk\big). \tag{2.35}$$

However, Bronselaer and Zanna (2020) instead provide the form

$$pCO_2 = pCO_2\big(0.0423\Delta T + \frac{B_{DIC}}{DIC}\Delta DIC + \frac{B_{Alk}}{DIC}DIC + \frac{B_{Alk}}{Alk}Alk\big), \tag{2.36}$$

or equivalently,

$$pCO_2 = pCO_2\big(0.0423\Delta T + \frac{B_{DIC}}{DIC}\Delta DIC + 2B_{Alk}\big). \tag{2.37}$$

Note also the change on the left hand side of Equation 2.35 to Equation 2.36: this is reproduced from Equation 7 of their paper, and appears to be simply a typographical error. However, it is unclear whether the change in variables from ΔDIC and ΔAlk to DIC and Alk, respectively, is also a typographical error or intentional: that the presented equation contains two uncancelled fractions (DIC/DIC and Alk/Alk) and an uncollected term ($2B_{Alk}$) suggests they are typographical errors. It therefore appears to be the case that Equation 2.36 is in error. Bronselaer and Zanna (2020) then goes on to discuss how patterns of $\partial pCO_2/\partial DIC$ are set by background ocean chemistry, and so changes in DIC are controlled predominately by background ocean chemistry.

The study then states that patterns of ΔDIC are strongly correlated with added temperature, with a correlation coefficient of 0.64, but that CFC changes are not strongly

correlated with added temperature, and that this discrepancy arises from the lack of chemical buffering of CFC's. Whilst a correlation of 0.64 between changes in added temperature and ΔDIC at the time of CO_2 doubling is a relatively strong correlation, it is plausible that imperfect correlations result from spatial variability in the local $\hat{\alpha}$ value. Without an explicit plot of changes in surface temperature and DIC in their fixed simulation circulations, it is impossible to know whether this is the result of systematic variability, or simply due to noise. Personal communications with Laure Zanna indicate investigation into the patterns of α did find higher values at high latitudes. However, the study did not attempt to quantify this variation, finding the assumption of a globally uniform carbon-heat coupling to be sufficient for their purposes.

As such, whilst the assumption of a globally uniform $\hat{\alpha}$ is a useful approximation, it is insufficient for the purposes of this study, which will aim to quantify how the carbon-heat coupling varies explicitly. In Chapter 5, it will be shown at A05, in the subtropical North Atlantic, that a temperature dependent α_T and a uniform α_T value do not produce meaningfully different results. In addition, Chapter 6 will also show that results at I08S, in the Southern Ocean, are not strongly affected by the choice of a uniform or temperature dependent α value either. However, a better understanding of variability in this coupling ought to produce better estimates of excess heat accumulation from changes in anthropogenic carbon.

The arguments provided above would provide a natural and simple explanation for the near-linear relationship between changes in ocean heat and carbon content, were the ocean simply a parcel of water in continuous contact and equilibrium with the atmospheric forcing. However, this is not the case for the real ocean: both heat and carbon are transported into the ocean by a number of processes. In the subtropical gyres, a major storage region for both heat and carbon (Sabine et al. (2004), Zika et al. (2021), the approximation of the ocean as being in continuous contact with the atmosphere is reasonable. It is less clear that this is the case in the deep ocean and at high latitude. Deep waters are typically formed from convective events, for example in the Labrador Sea, constituting a major pathway for the transport of heat and carbon to the deep ocean Lazier et al. (2002). Here, deep convective events typically require severe winters, thus, we might expect the near-linear relationship between excess heat and carbon uptake to be broken in the deep ocean.

There are, however, a number of pieces of evidence which suggest that we might expect this relationship between ocean heat and carbon uptake to persist into the deep ocean. It is well known that the patterns of ocean heat and carbon storage are strongly linked by ocean circulation: model simulations indicate storage of both excess heat and anthropogenic carbon by North Atlantic Deep Water (NADW) (Sarmiento et al. (1992), Banks and Gregory (2006), Winton et al. (2013)). Using observations, Mauritzen et al. (2012) found that density compensated anomalies ocean temperature are important for ocean heat uptake in the Subpolar North Atlantic. This requires anomalously

high salinities, in addition to anomalously high temperatures, in order to propagate a warming signal into the deep ocean. They also note that such warm and salty conditions have been found since approximately the mid 1990s. Thus we may well expect warming and salinification of NADW formed over the past 25 years, in addition to increased C_{anth} concentrations.

Finally, though these deep convective events act on short timescales and are typically associated with rapid cooling of surface waters, these events typically form well mixed water: Labrador Sea Water is characterised by a minimum in the vertical density gradient (Talley and McCartney, 1982). Thus, this strong mixing is likely to smooth out variability in local conditions, leading to a relatively homogeneous water mass with coherent signals in the temperature and carbon content change. Though we may not expect these changes to be related in a way which perfectly aligns with the global mean ratio α (as discussed in the following section), in combination, these pieces of evidence suggest we do indeed expect to see a broad and consistent linkage between temperature and carbon increase in the high latitude and deep ocean. Indeed, the results of Bronselaer and Zanna (2020) depend on this linkage between temperature and carbon existing in the high, as well as low latitude ocean.

2.2.1 Evidence for a temperature dependent carbon-heat coupling

As noted in the previous section, though Bronselaer and Zanna (2020) makes an argument for a globally uniform α value, this appears to be a useful approximation rather than a robust result. Several lines of evidence which suggest that is may be more accurate to approximate α as temperature dependent are now considered. Importantly, these arguments do not require α to depend on temperature explicitly: only for the intrinsic variability in α to be able to be parametrised usefully as temperature dependent. Potential issues with the global uniformity of Bronselaer and Zanna (2020) have already been discussed, and so will not be further explored. Some results from Turner et al. (2022), which basis for Chapter 3, are also alluded to: as such, whilst an overview of the arguments presented there are given here, they will not be explored in detail at this point. However, some details which underlie the arguments in Chapter 3 will be expanded upon here, as the arguments presented there are relatively terse.

The work discussed here was performed by Dr. Pete Brown for the TICTOC project proposal (`https://projects.noc.ac.uk/tictoc/`), and implies that we may expect a near linear relationship between α_T and background water temperature. The previously discussed approaches which aim to understand the relationship between warming and carbon content changes all begin with a description of climate system parameters, and aim to establish a relationship between changes in global ocean heat and carbon content from these general climate parameters. Alternatively, a semi-empirical approach might be taken: this is the approach discussed here.

CO2sys is a software package which allows for the computational determination of carbonate system parameters (Van Heuven et al., 2011). It has a number of uses, however, here only one is considered: calculating DIC change from projected temperature and CO_2 changes. By providing CO2sys with the relevant parameters of a number of water masses and a projected warming and CO_2 change, the associated DIC change may be evaluated. These projected changes are shown in Figure 2.5. Though there is a notable difference in the expected coefficients for the higher and lower end of the expected global mean warming, two patterns are clearly apparent. Firstly, there is an extremely good near linear dependence on water mass temperature of the coefficient relating the change in temperature and DIC. Secondly, all these coefficients increase with time.

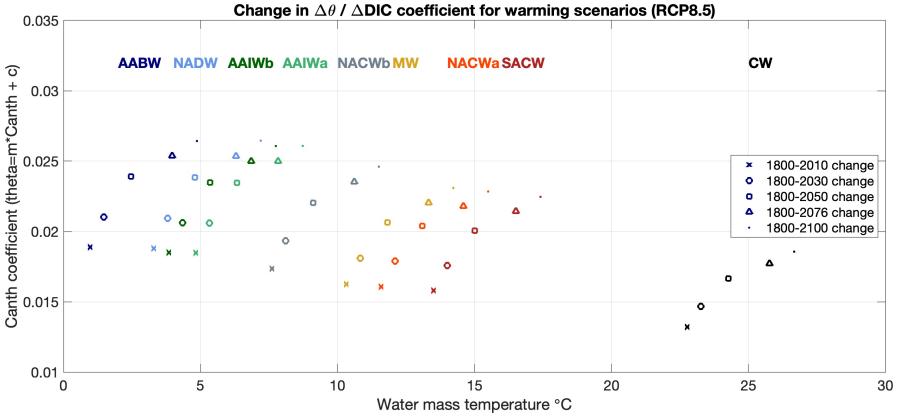

FIGURE 2.5: The projected ratio of DIC change to warming, for a number of water masses. For each bar, the middle point represents the expected coefficient for the expected global mean temperature change associated with a given atmospheric CO_2 concentration, the top point the expected coefficient for the upper limit of uncertainty in the global mean temperature change associated with the same atmospheric CO_2 concentration, and vice versa for the lowest points.

The linear dependence is striking: for the ratios of change over 1980:2050, a linear fit gives a R^2 value of 0.98. This strongly suggests a linear relationship between background temperature and the constant α. In addition, the arguments of Bronselaer and Zanna (2020) rely on making a number of assumptions about carbonate chemistry in order to derive a theoretical global uniformity in α. CO2sys does not make such assumptions, instead directly calculating changes associated with this warming. It therefore ought to better capture the response of the carbonate system to increasing atmospheric CO_2.

Subsequent idealised CO2sys investigations reveal that this apparent temperature dependence appears to be due to the relationship between the background carbonate chemistry and temperature. CO_2 is more soluble in cooler waters, and thus they exhibit higher DIC values: this is the solubility pump. As a result, surface DIC concentrations tend to vary latitudinally, with higher surface DIC values found in the cooler waters near the poles. However, surface alkalinity concentrations are less latitudinally variable, and so the surface DIC to alkalinity ratio varies with latitude. This leads to

a strong correlation between the surface buffer factor and temperature due to the influence of temperature on DIC concentrations (Sarmiento and Gruber, 2006). However, this is just a correlation: temperature does not directly drive buffer factor variations, instead influencing the buffer factor through its influence on surface DIC concentrations.

Due to their (generally) higher buffer factors, cool waters tend to have a reduced capacity to sequester additional carbon. This leads to a reduction in the total carbon change, which is anticorrelated with the initial carbon content. Thus, ΔDIC is proportional to the background temperature of water. As a result, the carbon change content of warmer waters is larger than that of cooler waters, and so the coefficient $\Delta\theta/\Delta$DIC is smaller in warmer waters than in cooler ones. This dependence is therefore well approximated as being temperature dependent (though this temperature dependence is not direct), and is shown in Figure 2.6, where the initial distribution of DIC values as a function of temperature (bottom edge of the shaded region) exhibits a stronger dependence on temperature than the distribution at a concentration of 1000ppm (top edge of shaded region). The distributions shown in Figure 2.6 are calculated using CO2sys in order to estimate DIC from a temperature, pCO_2 and using a fixed alkalinity of 2300μmol/kg. Global mean warming due to climate change is included in the calculation: it is parametrised as a linear function of pCO_2.

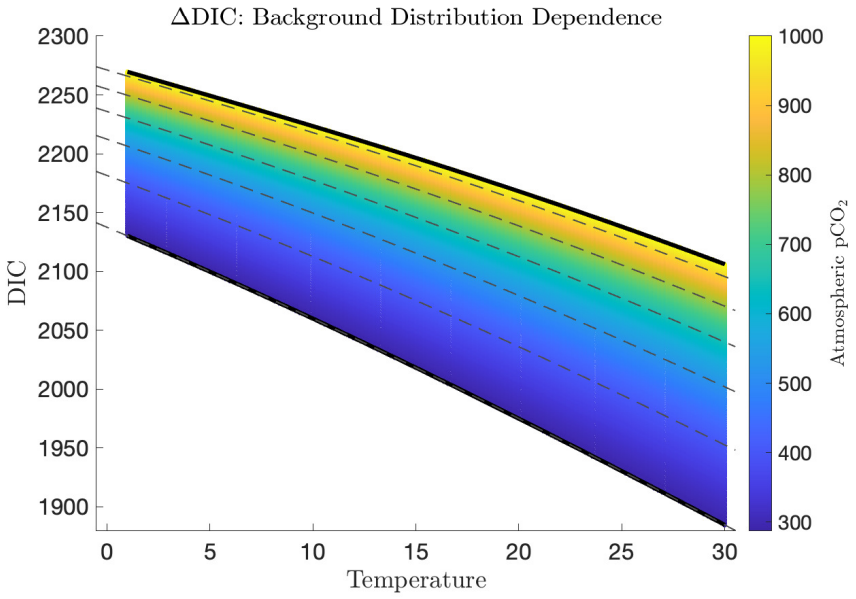

FIGURE 2.6: The dependence on initial temperature and DIC of DIC changes: cooler, higher carbon waters are less able to take up additional carbon, and so their carbon content changes less for a given change in pCO_2. Contours of constant pCO_2 are indicated with both thick black and dashed grey lines: the thick black lines show the initial and final contours of constant pCO_2 (280ppm and 1000ppm respectively).

I now return to the pattern scaling argument, which we may observe from the results of Leduc et al. (2016). Recall Equation 2.25:

$$\frac{\Delta T_{SS}}{\Delta \mathrm{DIC}} = \frac{\mathrm{B}}{\mathrm{DIC}} \frac{k}{\lambda} \left\{ a - \frac{\varepsilon N(t)}{\Delta \ln \mathrm{CO}_2} \right\}, \tag{1.36}$$

which importantly refers to temperature change for a global mean parcel of water. For simplicity, we may define

$$\hat{\alpha}_T = \frac{\Delta T_{SS}}{\Delta \mathrm{DIC}}, \tag{2.38}$$

where the label $\hat{\alpha}_T$ reflects the analogy between this quantity and the carbon-heat coupling of Bronselaer and Zanna (2020). Now, we may note that as Equation 2.25 refers to a global mean parcel of water, we may straightforwardly extend this to a local relationship by carbon and temperature changes by allowing the SST:SAT coupling constant k to vary spatially:

$$\frac{\Delta T_{SS}}{\Delta \mathrm{DIC}}(\vec{r}) = \frac{\mathrm{B}}{\mathrm{DIC}} \frac{k(\vec{r})}{\lambda} \left\{ a - \frac{\varepsilon N(t)}{\Delta \ln \mathrm{CO}_2} \right\}, \tag{2.39}$$

where $\vec{r}$ reflects an arbitrary location in the surface ocean. By spatially resolving k, we can then consider the pattern scaling approach utilised by Leduc et al. (2016) to instruct our understanding of how we expect changes in sea surface temperature to vary spatially, and therefore where we might expect to patterns to emerge. Additionally, we may consider k to include all spatial variability in the ratio B/DIC, allowing the use of global mean values in Equation 2.39.

Returning to Figure 1a from Leduc et al. (2016) (reproduced in Figure 2.7), we see again that in general, RTCRE values exceeding approximately 1.5 degrees warming per teratonne of carbon emissions are typically found only at high latitudes. This implies, all other surface changes being constant, a greater warming at high latitudes than low latitudes. We would therefore expect larger values of $\hat{\alpha}_T$ in cooler waters in the ocean, which form at these high latitudes. One notable exception to the greater warming seen at higher latitude in Figure 2.7 is in the Gulf Stream and Grand Banks region, where high RTCRE values are seen despite being at mid latitude. This likely reflects a shift in the position of the Gulf Stream: SAT's are strongly controlled by SST's, and so a northward shift in the position of the Gulf Stream will lead to significant warming at the locations it shifts to.

This plot from Leduc et al. (2016) provides one plausible explanation for why we might expect to see larger values of $\hat{\alpha}_T$ in cooler waters: intensified heat uptake at high latitudes. An alternative explanation is now explored: reduced carbon uptake at high latitudes. It is important to note at this point that the previous explanation of intensified heat uptake does not explain the near linear dependence of α_T values seen in Figure 2.5, as only a global mean warming was applied to each water mass here. Thus, whilst

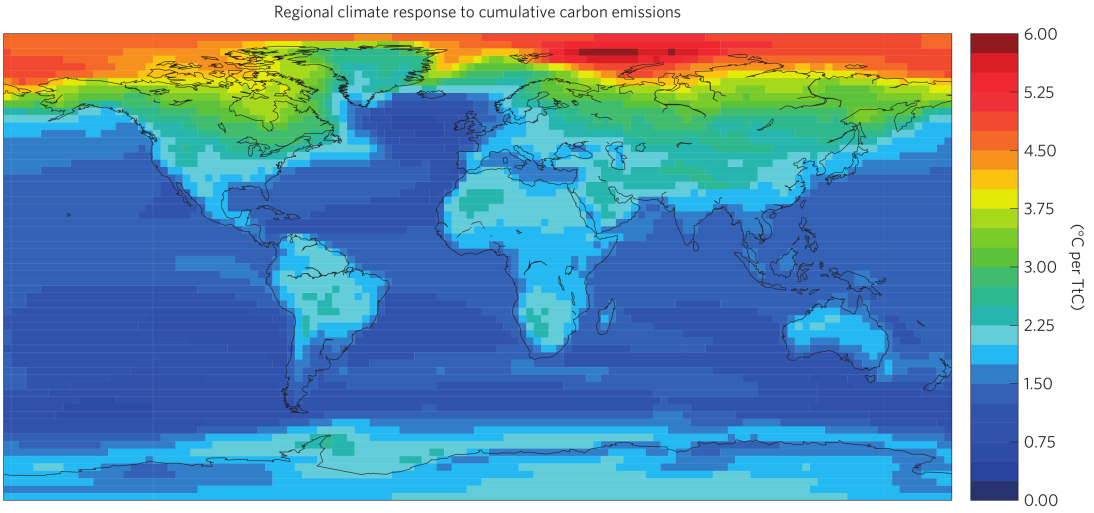

FIGURE 2.7: The map of RTCRE values from Leduc et al. (2016), also reproduced in Figure 2.2. RTCRE values refer to SAT change at a given spatial location per teratonne of carbon emissions.

variations in the background carbonate chemistry response are considered there, if intensified heat uptake at high latitude does in fact drive larger values of α_T, we would expect to see a stronger temperature dependence than suggested by Figure 2.5.

Consider the patterns shown in Figure 2.8. Whilst patterns of air-sea C_{anth} fluxes are relatively uniform (panel a), the patterns of C_{anth} subduction through the base of the mixed layer (panel b) are far less so. In general, C_{anth} subduction through the base of the mixed layer is lower than air-sea C_{anth} fluxes at high latitude, as well as at the Equator. These regions of reduced C_{anth} subduction are, in general, where the effects of vertical mixing and advection are in opposition. As a result of this, we would expect a reduced C_{anth} flux into the ocean interior at high latitudes, relative to surface fluxes. Such a mechanism would act to increase the variability of heat and carbon uptake: a globally uniform surface distribution could become far more variable in the ocean interior due to such a mechanism. As the arguments of Bronselaer and Zanna (2020) concern only the ocean surface, it is plausible that their arguments are correct here, but become less valid after propagation into the ocean interior.

Davila et al. (2022) directly investigated the propagation of C_{anth} though the mixed layer into the ocean interior, and its spatial variability. They found that ventilation of C_{anth} into the ocean interior was significantly less efficient at high latitude to low, but that purely Buffer Factor driven effects are insufficient to account for this: thus, whilst the simple CO2sys experiments presented above may partially explain the indicated temperature dependence, there are likely other contributing factors.

Finally, the findings of Turner et al. (2022), presented in Chapter 3, indicate similarly that higher values of α are found in high latitude waters. The method of diagnosing

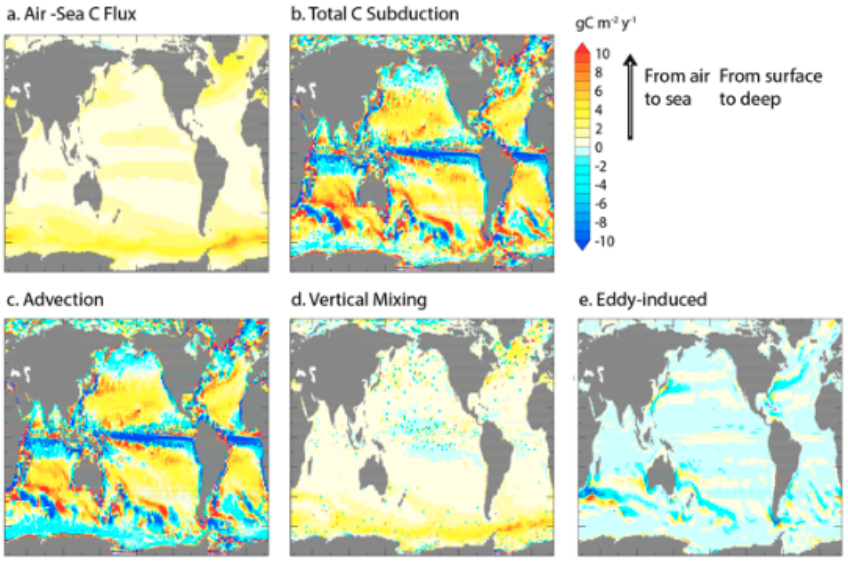

FIGURE 2.8: Maps of anthropogenic carbon fluxes (a), total C_{anth} subduction through the base of the mixed layer (b), and the contributions to subduction from advection (c), vertical mixing (d) and eddy induced transport (e). Reproduced from Bopp et al. (2015).

excess temperature presented there does not rely on the assumption of a globally uniform α value, instead diagnosing changes in redistributed temperature from changes in natural carbon, and calculating excess temperature by residual. It therefore allows for the direct calculation of local α values. The results there strongly imply that higher values of α are found at high latitudes, as well as potentially in regions of deep water formation. This result is in general agreement with the CO2sys calculation presented in Figure 2.5, and the results of Bopp et al. (2015) and Davila et al. (2022) suggest a plausible mechanism for this. In addition, the results of Leduc et al. (2016), Davila et al. (2022) and the CO2sys calculations imply together that this temperature dependence may in fact be even stronger than estimated here. As such, it appears to be very likely that α is not in fact a globally uniform constant, but exhibits spatial variability which may be accurately parametrised as a temperature dependence. However, whether it is necessary to account for this temperature dependence is less clear, and is investigated in Chapters 5 and 6 in more detail.

2.3 Redistribution and Natural Carbon

If the local ratio of excess heat to anthropogenic or excess carbon uptake is unknown, it is not possible to readily identify excess heat from anthropogenic carbon. However, provided that we retain the analogy between anthropogenic carbon and excess heat, we must also retain an analogy between natural carbon changes and redistributed heat, and it therefore ought to be possible to relate these two quantities in order to estimate

the redistribution of heat. By differencing total changes in heat content and redistribution, we should therefore be able to estimate the excess heat.

A benefit of this redistribution first approach is that provided the assumptions used are orthogonal to those relating excess temperature and anthropogenic carbon, it allows the explicit validation of the assumptions necessary to use anthropogenic carbon to estimate excess heat. Additionally, the redistribution of a parcel of water fundamentally differs from the addition of excess tracer to this parcel of water, as the redistribution of a parcel of water ought not to alter its properties, save for the effects of mixing. In fact, it is this property which motivated the development and use of the water mass framework: by considering the properties of water in tracer space, the effects of ocean circulation can be removed. This is the approach of Zika et al. (2021), who applied these water mass considerations to directly estimate material heat content change, a quantity closely analogous to excess heat. This means that by understanding the properties of redistribution, we may also understand the effects of ocean circulation change on multiple tracers.

In other words, there is no (known) clear relationship between changes in excess salinity and anthropogenic carbon. However, because redistribution does not alter the properties of a parcel of water, there should be a relationship between the redistributed salinity and the redistributed carbon. If we can understand the nature of this relationship, we ought to be able to estimate the redistribution of salinity by changes in ocean transport - and therefore also estimate the excess salinity. Even more usefully, this is not limited to salinity - it is in principle extensible to any ocean tracer.

I now explore how we may decompose the changes between two general tracers at a point in the ocean, before showing how we may apply this to temperature and carbon. This principle will be explored in greater detail and in a more intuitive fashion in Chapter 3. However, it is useful to understand the mathematical underpinning of relationship between the two, and how the relationship between excess temperature and anthropogenic carbon and the relationship between redistributed temperature and natural carbon change can be considered to be two limits of the same equation, and in some sense analogous.

Recall the definition of the excess and redistributed components of an arbitrary tracer, Q, following Winton et al. (2013), who break the tracer Q and the transport v up as follows (note now I use a slightly diferent notation to the previous definition and collect terms explicitly):

$$\vec{v}Q = (\vec{\bar{v}} + \vec{v}')(\bar{Q} + Q') = \underbrace{\vec{\bar{v}}\bar{Q}}_{Preindustrial} + \underbrace{\vec{v}'\bar{Q}}_{Redistributed} + \underbrace{\vec{\bar{v}}Q' + \vec{v}'Q'}_{Excess}, \tag{2.40}$$

where $\vec{\bar{v}}$ and $\bar{Q}$ refer to the preindustrial, steady state components of $\vec{v}$ and Q, and $\vec{v}'$ and Q' the perturbations. Note that these components may be either on short timescales (for

example seasonal variability), or longer timescales (for example centennial timescale AMOC decline). We may now wish to show how changes in two arbitrary tracers, Q_1 and Q_2, are related. Following Williams et al. (2021), we may write changes in Q in the following fashion:

$$F = \frac{\partial Q}{\partial t} + \vec{v} \cdot \nabla Q. \tag{2.41}$$

We now note that we have defined $\vec{v}$ and $\bar{Q}$ as the steady state, preindustrial components: thus

$$\frac{\partial \bar{Q}}{\partial t} = 0, \tag{2.42}$$

and therefore

$$\bar{F} = \vec{v} \cdot \nabla Q, \tag{2.43}$$

where the surface forcing F may also be decomposed as $F = \bar{F} + F'$. Now expanding Equation 2.41 using the decomposition of Equation 2.40, and collecting terms, we find

$$\frac{\partial Q'}{\partial t} = \underbrace{F' - \vec{v} \cdot \nabla Q'}_{Excess} \underbrace{-\vec{v'} \cdot \nabla \bar{Q}}_{Redist} \tag{2.44}$$

Now, consider the case where the excess term can be considered negligible: we have

$$\frac{\partial Q}{\partial t} = \frac{\partial Q'}{\partial t} = -\vec{v'} \cdot \nabla \bar{Q} \tag{2.45}$$

In this case, the relationship between changes in Q_1 and Q_2 due to only to redistribution, as measured at fixed spatial location, is given by

$$\frac{\partial Q_1}{\partial t} \Big/ \frac{\partial Q_2}{\partial t} = \frac{\vec{v'} \cdot \nabla \bar{Q}_1}{\vec{v'} \cdot \nabla \bar{Q}_2} \tag{2.46}$$

I will label this quantity a redistribution coefficient, κ_r: it relates the changes in Q_1 and Q_2 due to redistribution:

$$\kappa_r = \frac{\vec{v'} \cdot \nabla \bar{Q}_1}{\vec{v'} \cdot \nabla \bar{Q}_2}. \tag{2.47}$$

As we have defined our terms, $\vec{v'}$ is the non steady state component of ocean transport, comprising all transport except mean flow. As previously stated, it therefore includes short timescale velocities, such as seasonal perturbations to mean velocities, and longer timescales velocities. A natural question might therefore be: how sensitive is κ_r to our choice of velocity perturbation, $\vec{v'}$. In other words, does knowing κ_r for seasonal perturbations to velocity inform us of the relevant value of κ_r for perturbations to the oceans velocity field on centennial timescales: is the coupling between the redistribution of Q_1 and Q_2 timescale invariant?

We know that, from the way we have defined our terms, $\bar{Q}_1$ and $\bar{Q}_2$ are constants: they have no time dependence. Therefore, all potential time dependence in κ_r must

come from the time dependence in $\vec{v'}$. To make this explicit, we may split our velocity perturbation into short timescale perturbations and longer timescale perturbations:

$$\vec{v'} = \vec{v}_s + \vec{v}_l, \tag{2.48}$$

where $\vec{v}^s$ represents short timescale (hourly to subdecadal) velocity perturbations, and $\vec{v}^l$ longer timescale velocity perturbations (decadal and longer) (N.B: This cutoff timescale is entirely arbitrary, the choice of cutoff will be justified later). We can therefore write

$$\kappa_r = \frac{(\vec{v}_s + \vec{v}_l) \cdot \nabla \bar{Q}_1}{(\vec{v}_s + \vec{v}_l) \cdot \nabla \bar{Q}_2} \tag{2.49}$$

From this, we may see that κ_r is time invariant in two situations: if the velocities $\vec{v}_s$ and $\vec{v}_l$ are parallel, or if the normal vectors $\nabla \bar{Q}_1$ and $\nabla \bar{Q}_2$ are parallel. The first case is trivial: simply let $\vec{v}_l = a\vec{v}_s$, where a is an arbitrary scalar constant. The factor of $1 + a$ in the numerator and the denominator cancels, and the expression for κ_r now must be time invariant.

To demonstrate the second case is more involved. For any velocity vector $\vec{v}$, we may decompose it as the component perpendicular to the gradient of a scalar tracer field Q, and the change parallel to the gradient of this field:

$$\vec{v} = \vec{v}_{\parallel} + \vec{v}_{\perp}, \tag{2.50}$$

where $\vec{v}_{\parallel}$ is defined as

$$\vec{v}_{\parallel} = \vec{v} \cdot \nabla Q \frac{\nabla Q}{|\nabla Q||\vec{v}|} \tag{2.51}$$

and $\vec{v}_{\perp}$ is defined as

$$\vec{v}_{\perp} = \frac{\vec{v} \times \nabla Q}{|\vec{v} \times \nabla Q|} \tag{2.52}$$

As, by definition,

$$\vec{v}_{\perp} \cdot \nabla Q = 0, \tag{2.53}$$

we have that

$$\kappa_r = \frac{\vec{v'}_{\parallel} \cdot \nabla \bar{Q}_1}{\vec{v'}_{\parallel} \cdot \nabla \bar{Q}_2}. \tag{2.54}$$

If the short and long timescale velocities contain different contributions from the parallel and perpendicular components of velocities, it is still possible that κ_r varies on longer timescales. However, let us write

$$\nabla \bar{Q}_1 = b \nabla \bar{Q}_2 + \vec{\delta}, \tag{2.55}$$

where b is an arbitrary scalar constant, and $\vec{\delta}$ is an arbitrary vector, which is by definition perpendicular to $\nabla \bar{Q}_1$:

$$\vec{\delta} \perp \nabla \bar{Q}_1. \tag{2.56}$$

We can therefore rewrite Equation 2.54 in the form

$$\kappa_r = b\frac{\vec{v'}_{\parallel} \cdot \nabla \bar{Q}_2}{\vec{v'}_{\parallel} \cdot \nabla \bar{Q}_2} + \frac{\vec{\delta}}{\vec{v'}_{\parallel} \cdot \nabla \bar{Q}_2}. \tag{2.57}$$

Noting that if $\nabla Q_1 \parallel \nabla Q_2$, $\vec{\delta} = \vec{0}$, we have

$$\kappa_r = b\frac{\vec{v'}_{\parallel} \cdot \nabla \bar{Q}_2}{\vec{v'}_{\parallel} \cdot \nabla \bar{Q}_2} = \frac{\vec{v'}_{\parallel} \cdot \nabla \bar{Q}_1}{\vec{v'}_{\parallel} \cdot \nabla \bar{Q}_2}, \tag{2.58}$$

an equation which now contains only constant terms. It is therefore the case that κ_r is time independent either if our short or long timescale velocities are parallel, or if the normal directions to our two tracer fields, ∇Q_1 and ∇Q_2 are a parallel. The first case, whilst simpler, is less likely to be strictly obeyed throughout the ocean on a large range of timescales. However, if we can demonstrate the second condition is true, it remains true on all timescales: this makes it the more useful condition.

The question is now whether we can apply this to relate temperature and carbon changes. As noted previously, we may approximate redistributed DIC with changes in natural carbon. We may therefore link the redistribution of temperature with changes in natural carbon, provided the gradient of the background fields of temperature and natural carbon are either parallel or antiparallel. We know this to be generally the case as a result of the strong temperature dependence of CO_2 solubility: C_{nat} and temperature are therefore strongly anticorrelated throughout the ocean. This is shown explicitly in Figure 2.9, using data from the NEMO simulations used in Chapter 3. Cross products of the gradients of the two fields are reliably 0, and dot products equal to ± 1, indicating that the diathermal and diacarbon directions are almost perfectly parallel throughout the global ocean.

As a result of this strong spatial covariability, we may write

$$\Delta\theta_r = \kappa_r \Delta C_{nat} \tag{2.59}$$

by rearranging Equation 2.47 and substituting $Q_1 = \Delta\theta$ and $Q_2 = \Delta C_{nat}$. This is analogous to the approach of Bronselaer and Zanna (2020): in the same fashion they associate C_{anth} with excess temperature, we may associate C_{nat} with redistributed temperature.

It was previously stated that these two relationships: the α relationship between excess temperature and anthropogenic carbon, and the κ_r relationship between redistributed temperature and natural carbon, can be considered to be two limits of the same relationship. This is now demonstrated. Returning to our arbitrary tracer Q, we may write the total derivative of Q with respect to time as

$$\frac{dQ}{dt} = \frac{\partial Q}{\partial t} + \vec{v} \cdot \nabla Q, \tag{2.60}$$

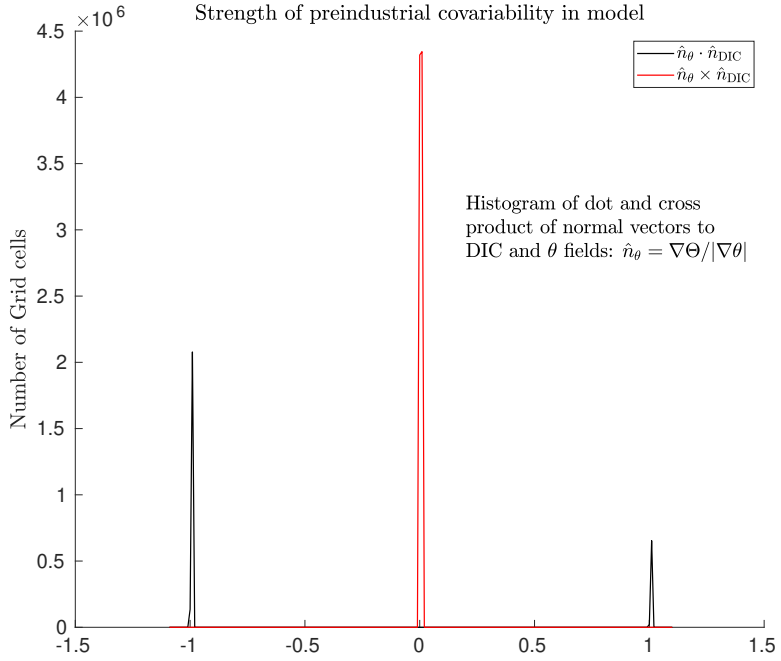

FIGURE 2.9: The covariability of the preindustrial temperature and carbon fields. Dot products of the two fields are shown in black, and cross products in red.

which is simply the standard relationship between Eulerian and Lagrangian derivatives. Noting that we may simply cancel the time dependence in the ratio of total derivatives, we may therefore write the relationship between changes in two arbitrary tracers as

$$\frac{dQ_1}{dQ_2} = \frac{\partial_t Q_1 + \vec{v} \cdot \nabla Q_1}{\partial_t Q_2 + \vec{v} \cdot \nabla Q_2}, \tag{2.61}$$

where $\partial_t \equiv \partial/\partial t$. Now, consider the change in the global mean value of Q_1 and Q_2: in the global mean $\vec{v} = \vec{0}$, and so

$$\frac{dQ_1}{dQ_2} = \frac{\partial_t Q_1}{\partial_t Q_2}. \tag{2.62}$$

Replacing Q_1 with temperature and Q_2 with C_{anth}, we obtain the formulation of Bronselaer and Zanna (2020). If we instead assume that we are inspecting the relationship between steady state components, we may take $\partial_t Q = 0$ (from Equation 2.42), and so

$$\frac{dQ_1}{dQ_2} = \frac{\vec{v} \cdot \nabla Q_1}{\vec{v} \cdot \nabla Q_2}, \tag{2.63}$$

which is the relationship we used to define our redistribution coefficients. Thus, the two relationships may both be derived from considering derivatives in temperature carbon space, and are both limits of the equation relating changes in temperature and carbon.

2.3.1 Redistribution Coefficients: Analogy to Spice and Heave

Equation 2.59, in isolation, may seem somewhat novel and relatively unsubstantiated. However, it is in fact closely related the spice-heave decomposition, a conventional, well defined and widely used tool to interpret changes in the properties of water. In this section, it is show how the two are related, and how the use of C_{nat} produces a decomposition that is more readily interpretable than a spice-heave decomposition: an excess-redistribution decomposition.

The spice-heave decomposition was introduced by Bindoff and Mcdougall (1994), which looked at changes in water properties along isopycnals, and changes in water properties due to the vertical migration of isopycnals. They note that the properties of water at a fixed point in the global ocean may change in one of three orthogonal ways:

- Pure warming: The introduction of warmed mixed layer water into the ocean interior.
- Pure freshening: The introduction of water which has been freshened through changes in the ocean evaporation and precipitation balance at the ocean surface into the ocean interior.
- Pure heave: The vertical displacement of the water column either due to changes in water mass formation rates or dynamical changes.

From these properties, it is tempting to assume that pure heave is directly analagous to the redistribution of water masses, pure warming corresponds to excess temperature, and pure freshening corresponds to excess salinity. However, the picture is more complicated. On short timescales, we expect changes in water properties at a point to be dominated by the heaving and shoaling of the water column. To remove the effects of this heaving and shoaling of the water column, we might follow the trajectory of an isopycnal, looking at changes in water mass properties along an isopycnal. These changes are referred to as spice: as warm and salty water may have the same density as cool and fresh water, we will still be able to observe changes in water mass properties at the same density. For water of a given density, warm and salty water is referred to as 'spicy', and cool and fresh water as 'minty'.

Consider the case of pure warming. When water is warmed at constant salinity, it will move to a new isopycnal, before being subducted along that isopycnal into the ocean interior. However, this isopycnal does not necessarily follow the same trajectory in θ-S space as the profile of water properties here: if this is not the case, then pure warming will act to cause a change in the depth of the isopycnal here. In the case that the slope of the profile of water properties is steeper than the slope of the isopycnal, then when viewed on isopycnals, water will appear to have cooled. This behaviour is counter intuitive, and so a diagram is shown in Figure 2.10.

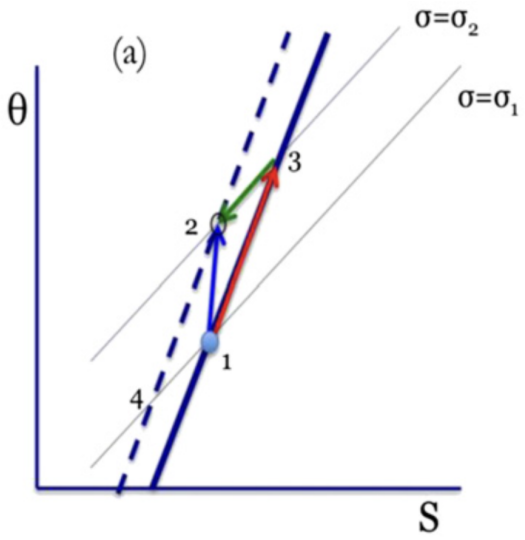

FIGURE 2.10: The warming along isopycnals due to pure spice. A parcel at the surface is warmed from point 1 to point 2, which lie on isopycnals σ_1 and σ_2, respectively. This motion can be viewed as the movement along the original θ-S profile without changing it ($1 \rightarrow 3$, heave), followed by movement along an isopycnal ($3 \rightarrow 2 \equiv 1 \rightarrow 4$). The effect of pure warming is therefore to cool water when viewed on isopycnals. This Figure is reproduced from Häkkinen et al. (2015).

Mathematically, a spice-heave decomposition may be expressed as a Taylor expansion of a water mass property about an isopycnal surface. For a general tracer Q, we may write

$$dQ|_\sigma = dQ|_z + \frac{\partial Q}{\partial z} dz|_\sigma + \mathcal{O}(dz|_\sigma)^2. \tag{2.64}$$

Rearranging and ignoring higher order terms, we have

$$\underbrace{dQ|_z}_{\textit{Eulerian Change}} = \underbrace{dQ|_\sigma}_{\textit{Spice}} - \underbrace{\frac{\partial Q}{\partial z} dz|_\sigma}_{\textit{Heave}}. \tag{2.65}$$

Note at this point that Equation 2.65 bears a resemblance to 2.44. In fact, if we consider the relationship between two tracers, Q_1 and Q_2, resulting solely from heave driven changes, we see

$$\frac{dQ_1}{dQ_2} = \frac{\partial Q_1}{\partial z} \Big/ \frac{\partial Q_2}{\partial z}, \tag{2.66}$$

which is simply the one dimensional (vertical only) approximation to κ_r. Assuming $\vec{v'} = v_z \hat{\vec{z}}$ and $\nabla Q = \partial_z Q \hat{\vec{z}}$, we see that we may reduce Equation 2.47 in the following way:

$$\kappa_r = \frac{\vec{v'} \cdot \nabla \bar{Q}_1}{\vec{v'} \cdot \nabla \bar{Q}_2} = \frac{v_z \partial_z Q_1}{v_z \partial_z Q_2} = \frac{\partial Q_1}{\partial z} \Big/ \frac{\partial Q_2}{\partial z}. \tag{2.67}$$

Thus, the relationship between variables as the result of heave is the one dimensional approximation to the redistribution coefficient.

There are two important features to this result. First, we may approximate the redistribution coefficient linking two tracers as the ratio of vertical derivatives of two tracers if the spatial distribution of the two can be considered essentially one dimensional. This explains intuitively why we may consider the redistribution coefficient to be independent of timescale in cases where the normal vectors to the fields Q_1 and Q_2 are parallel. Second, the heave driven relationship between two variables is independent of the distance over which an isopycnal heaves, as a result of our linear approximation. As a corollary to this, the heave driven relationship between the two variables is also independent of our choice density surfaces as a surface to follow: we might instead choose natural carbon surfaces. This leads very naturally to the concept of a redistribution depth. In the same way that the quantity $dz|_\sigma$ tells us the distance an isopycnal heaves, we can define a redistribution depth, Δz_r as

$$\Delta z_r = \frac{\Delta Q_r}{\partial Q / \partial z} \tag{2.68}$$

The difference between a spice-heave decomposition and an excess-redistribution decomposition, under this one dimensional approximation, is therefore fundamentally different only in the choice of depth surface to follow. For the case of temperature, in an excess-redistribution decomposition, there is no velocity imposed upon the variable depth surface we choose to follow as a result of imposed surface warming. However, this is not the case for a spice-heave decomposition: here, an imposed surface warming leads to an imposed velocity on our isopycnal surface, leading to the difficulty in interpreting the results of a spice-heave decomposition, as this imposed velocity is a function of the background T-S profile.

2.4 Temperature-Carbon Space: Estimating circulation change without decomposing DIC

From Sections 2.1.2, 2.2 and 2.3, we have two clear relationships linking changes in ocean temperature and carbon: excess temperature is closely linked to anthropogenic carbon, and redistributed temperature is closely linked to natural carbon. In this section, it is shown how these two relationships may be leveraged to directly calculate excess and redistributed temperature, without the need for an explicit decomposition of carbon. In fact, it is the case that this decomposition also yields a decomposition into excess and redistributed carbon, fields which can be compared to the independently calculated anthropogenic and natural carbon fields.

Decomposing temperature and carbon change as follows:

$$\Delta\theta = \Delta\theta_e + \Delta\theta_r \tag{2.69}$$

$$\Delta\text{DIC} = \Delta\text{DIC}_e + \Delta\text{DIC}_r, \tag{2.70}$$

where the subscript e refers to changes due to excess temperature or carbon, and the subscript r refers to changes due to the redistribution of background temperature or carbon.

Here, I choose to partition DIC changes into an excess and redistributed component, instead of anthropogenic and preindustrial (or natural) components, as I do not believe that the relationship I derive maps precisely onto previous definitions. However, excess DIC and C_{anth}, as well as redistributed DIC and C_{nat} changes map well onto one another. As such, this decomposition can be considered both a temperature and a carbon decomposition. This point is discussed in Williams et al. (2021): I note at this point that the technique I present here may be considered a localised formulation of the principles they use to identify regions where temperature and carbon changes are dominated by either the excess or redistributed components.

Though Bronselaer & Zanna relate changes in anthropogenic carbon and excess heat, and here I relate changes in excess carbon and temperature, in this formulation, I follow their naming convention, rather than that of Turner et al. (2022) or the nomenclature used in Section 2.3. Thus, in the following equations, the redistribution coefficient κ_r will instead be referred to as β_T. I relate changes in the local excess components by a constant, α_T:

$$\Delta\theta_e = \alpha_T \Delta\text{DIC}_e \tag{2.71}$$

and changes in the redistributed components by β_T:

$$\Delta\theta_r = \beta_T \Delta\text{DIC}_r \tag{2.72}$$

Both are given the subscript T to identify that they relate changes in temperature and DIC: later these definitions will be extended to include salinity. Rearranging these Equations, we obtain expressions for the excess and redistributed DIC in terms of the change in the temperature and DIC:

$$\Delta\text{DIC}_e = \frac{\Delta\theta - \beta_T \Delta\text{DIC}}{\alpha_T - \beta_T} \tag{2.73}$$

and

$$\Delta\text{DIC}_r = \frac{\alpha_T \Delta\text{DIC} - \Delta\theta}{\alpha_T - \beta_T} \tag{2.74}$$

Equations 2.73 and 2.74 may be combined in a vector equation as follows:

$$\begin{pmatrix} \Delta\text{DIC}_e \\ \Delta\text{DIC}_r \end{pmatrix} = \frac{1}{\alpha_T - \beta_T} \begin{pmatrix} 1 & -\beta_T \\ -1 & \alpha_T \end{pmatrix} \cdot \begin{pmatrix} \Delta\theta \\ \Delta\text{DIC} \end{pmatrix} \tag{2.75}$$

Here, approximating C_{anth} as excess DIC allows us to use this decomposition to identify anthropogenic carbon, producing estimates consistent with established techniques for identifying anthropogenic carbon: this will be shown in Chapters 5 and 6. However, as we are concerned primarily with identifying excess and redistributed temperature and salinity, we instead use the relationship

$$\begin{pmatrix} \Delta\theta_e \\ \Delta\theta_r \end{pmatrix} = \begin{pmatrix} \alpha_T & 0 \\ 0 & \beta_T \end{pmatrix} \cdot \begin{pmatrix} \Delta\text{DIC}_e \\ \Delta\text{DIC}_r \end{pmatrix} \tag{2.76}$$

to obtain an expression for excess and redistributed temperature:

$$\begin{pmatrix} \Delta\theta_e \\ \Delta\theta_r \end{pmatrix} = \frac{1}{\alpha_T - \beta_T} \begin{pmatrix} \alpha_T & -\alpha_T\beta_T \\ -\beta_T & \alpha_T\beta_T \end{pmatrix} \cdot \begin{pmatrix} \Delta\theta \\ \Delta\text{DIC} \end{pmatrix} \tag{2.77}$$

Conceptually, this decomposition may be thought of as assigning two axes in $\theta - \text{DIC}$ space: α_T and β_T. A change in the $\theta - \text{DIC}$ space position of water at a given geographical location can thus be thought of as the sum of its motion along the transient axis (α_T, associated with excess) and it's motion along the redistribution axis (β_T, associated with redistribution). We may then estimate excess and redistributed temperature from Equation 2.77, provided the values of α_T and β_T are known. This decomposition is illustrated in Figure 2.11.

The displacement of water at fixed geographic location being decomposed in terms of the sum of vector displacements along two predefined axes is the same principle as that applied by the spice-heave decomposition: the diagram shown in Figure 2.10 can be straightforwardly relabelled to describe this excess-redistribution diagram.

This matrix inversion is closely related to the technique used by Clément et al. (2022), which similarly projects changes in temperature and salinity into changes along and across isopycnals: indeed, this decomposition can be thought of straightforwardly as the analogue of their analysis in Temperature-Carbon space. I also note that the arctan of the quantity β_T is analogous to the Turner angle (Ruddick, 1983) in Temperature-Carbon space.

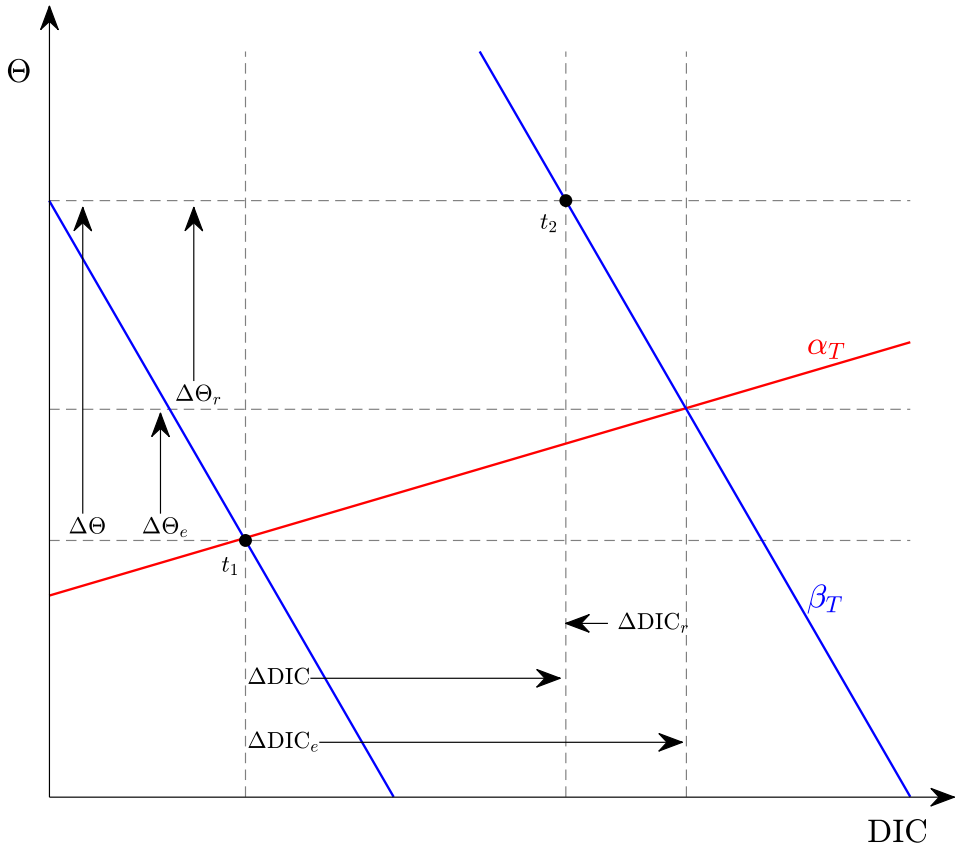

FIGURE 2.11: A diagram illustrating the matrix inversion decomposition of temperature and DIC presented here. Changes in temperature and DIC components are indicated with arrows, and the preferred axes in red (transient axis, α_T) and blue (redistribution axis, β_T). For a change in temperature and DIC, $(\Delta\theta, \Delta\text{DIC})$, between times t_1 and t_2, the excess components are linked by the transient axis α_T and the redistributed components by the redistribution axis β_T.

2.5 Discussion

In this Chapter, the transient climate response to cumulative carbon emissions (TCRE) has been introduced, and it's implications for ocean heat and carbon uptake discussed. In particular, it has been shown how an analagous transient response relationship, known as the carbon-heat coupling (α), between excess temperature and anthropogenic carbon can be derived using similar considerations. Two derivations for this carbon-heat coupling have been presented: the derivation of Bronselaer and Zanna (2020), and a new one, which allows for spatial variability in the constant α. Carbonate chemistry considerations have shown that we expect α to depend linearly on background temperature, and this temperature dependence will be explored explicitly in Chapters 5 and 6.

In addition, it has been shown how the background temperature and carbon fields of the ocean spatially covary, and how this spatial covariability can be leveraged using

a redistribution coefficient in order to predict the changes in one tracer from another. This approach, which shares some key characteristics with the more commonly used spice-heave decomposition, allows us to identify excess temperature changes without relying on a prescribed relationship between excess temperature and anthropogenic carbon - thus, it will be used in Chapter 3 to directly study variability in the carbon-heat coupling.

Finally, it has been shown how these two approaches may be combined to estimate excess and redistributed temperature directly from changes in DIC and temperature, without any decomposition of carbon. This approach will be applied to data from the GLODAP dataset in Chapters 5 and 6 to produce both local and global estimates of excess heat uptake by the global ocean.

Chapter 3

Decomposing ocean temperature and salinity change using ocean carbon change in the NEMO OGCM

3.1 Introduction

This section is based upon the paper Turner et al. (2022), in which I explore how we can use the principles illustrated in Section 2.3 in order to understand changes in global ocean temperature and carbon.

Local ocean heat content changes are contributed to by both the addition and removal of heat from the surface due to perturbed radiative forcing (*excess heat*), and from the rearrangement of the preindustrial temperature field from circulation variability (*redistributed heat*). Ocean salinity changes can also result from perturbations to air-sea freshwater fluxes (*excess salinity*), as well as due to the rearrangement of the preindustrial salinity field (*redistributed salinity*).

The redistribution of temperature and salinity as a result of ocean circulation variability acts on much shorter timescales than the accumulation of excess heat and salinity. Circulation-related variability comprises the majority of temporal variability in contemporary ocean temperature and salinity, (Bindoff and Mcdougall (1994), Desbruyères et al. (2017)) and regional sea level (Church et al., 2013). However, the excess component is anticipated to dominate in the future (Bronselaer and Zanna (2020), Zika et al. (2021)). Thus the evolution of excess temperature and patterns of excess salinity as well as changes in ocean circulation comprise a key source of uncertainty in estimates of regional sea level rise (Church et al., 2013).

One approach to determine excess temperature is to use a Passive Anomalous Tracer (PAT), which obeys the same physics as temperature, but is defined to have a preindustrial field which is zero everywhere: the preindustrial field therefore cannot contribute to redistribution (Banks and Gregory (2006), Gregory et al. (2016)), and so PAT reveals the distribution and evolution of the excess temperature field. Alternatively, it is possible in simulations to force ocean circulation to obey preindustrial dynamics despite increasing radiative forcing: this gives a similar result, though differing by a second order term to the PAT implementation (Winton et al., 2013).

While these methods have been very informative they are only applicable to models: no real world PAT tracer exists and while transient tracers such as chlorofluorocarbons are a close analogue their interpretation in terms of excess temperature necessitates the determination of an excess temperature boundary condition. This motivates the development of proxy methods which aim to diagnose the excess and redistributed temperature from other tracers and might be more generally applied, in particular to observations. The approach of Bronselaer and Zanna (2020) is an example of this: by approximating the distribution of excess temperature with that of anthropogenic carbon, they are able to leverage the mechanistic coupling relating excess heat accumulation to anthropogenic carbon accumulation to produce estimates of the scale and patterns of excess heat uptake in models and observations.

Using an alternative carbon based methodology, Williams et al. (2021) explains differences in storage of heat and carbon in terms of two components: 1) the correlation of excess heat and carbon (both increase over time), and 2) anticorrelation of redistributed heat and carbon (the preindustrial distributions of temperature and carbon are inverted due to the inverse temperature dependence of carbon dioxide solubility). They use this to diagnose excess and redistributed heat (note Williams et al. (2021) refer to this as added heat, though the definitions used are identical). Bronselaer and Zanna (2020) can therefore be thought of as specifying the character of this positive correlation between excess heat and anthropogenic carbon, in order to estimate excess heat directly from anthropogenic carbon. In this Chapter, the approach described in Chapter 1, Section 2.3 is expanded upon, and its uses explored. This approach uses similar principles to those presented in Williams et al. (2021) and Bronselaer and Zanna (2020), specifying the character of the anticorrelation between redistributed heat and natural carbon locally via the preindustrial ocean state.

Instead specifying the relationship between the excess components of temperature and DIC change, as done by Bronselaer and Zanna (2020), relies on a global biogeochemical relationship derived from the radiative forcing of CO_2 and the ocean carbon buffer factor, making their approach applicable only to temperature (Chapter 1, Section 2.1.2). In contrast, in the absence of perturbations to mixing, redistribution leaves the properties of a parcel of water unchanged. As a result, the redistribution first approach applied here is more generally applicable: for example, if we identify a clear spatial relationship

between natural carbon and salinity, we may use the redistribution of natural carbon to estimate the redistribution of salinity. This allows us to not only produce estimates of temperature redistribution, but also estimates of salinity, and by extension density, redistribution. Using these, the patterns of storage of excess and redistributed temperature and salinity by the global ocean are investigated in this Chapter.

3.2 Data and Methods

3.2.1 Model set up

In this Chapter, the NEMO v3.2 OGCM (Ocean General Circulation Model) is used, (Madec et al., 2017) coupled to the MEDUSA-2 biogeochemical model (Yool et al., 2013) and the Louvain-la-Neuve (LIM2) dynamic sea ice model (Timmermann et al., 2005). The model was configured with the ORCA1 grid with a nominal 1 degree resolution and 64 vertical levels (Madec and Imbard, 1996). The model was spun up for 900 years, before three 240 year simulations spanning 1860-2099 were spawned: a control run (CTR), coupled climate change run (COU), and a 'warming only' run (RAD), following the convention of Schwinger et al. (2014), Rodgers et al. (2020). The ocean model was forced with output from the HadGEM2-ES (Collins et al., 2011), an earth system model driven using prescribed greenhouse gas, land use and atmospheric chemistry forcing following the RCP8.5 scenario over the 1860-2099 time period. In this scenario, atmospheric CO_2 increases to over 900ppm by the end of the simulations (Riahi et al. (2011), atmospheric CO_2 in these simulations is shown in Couldrey et al. (2016), Figure 1a). Surface heat, momentum, freshwater fluxes, and atmospheric chemistry from HadGEM2-ES were used to force NEMO at 6 hourly intervals, and no restoring was used.

The CTR run is forced with 8 repetitions of the first 30 years of these fluxes from the HadGEM2-ES forcing, with a fixed atmospheric CO_2 of 286ppm: no significant climate change occurs in these 30 years. The 900 year spinup for all 3 model runs was also forced using this 30 year repeat forcing.

The COU run is forced with the full 240 year output from HadGEM2-ES. The RAD run has the same physical variability as in COU including that driven by atmospheric carbon increases but the atmospheric carbon is artificially relaxed to preindustrial conditions. As the RAD run only includes changes in DIC due to physical change (circulation change and warming), rather than the ocean biogeochemical response to increased atmospheric CO_2, we can calculate this response, namely anthropogenic carbon or C_{anth}, directly from the difference of the COU and RAD runs:

$$C_{anth}(x,y,z,t) = DIC^{COU}(x,y,z,t) - DIC^{RAD}(x,y,z,t). \tag{3.1}$$

Natural carbon, or C_{nat}, is then defined to be the total DIC content with the anthropogenic carbon contribution removed: it is therefore calculated as

$$C_{nat} = DIC^{COU} - C_{anth} - \Delta DIC^{CTR} = DIC^{RAD} - \Delta DIC^{CTR}, \tag{3.2}$$

where ΔDIC^{CTR} is control run drift, equivalent to $DIC^{CTR}(x, y, z, t) - DIC^{CTR}(x, y, z, t_0)$, where t_0 is the beginning of the three simulations, 1860, and Δ refers to change since 1860 and t is an arbitrary time. Therefore by definition all DIC is natural carbon at the beginning of our simulations, as the DIC fields are identical at the beginning of all 3 runs. DIC changes are then the sum of natural and anthropogenic carbon change. As such, the local DIC content may be decomposed at any time in the following way (note as C_{anth} is defined to be zero at time $t = t_0$, $C_{anth} = \Delta C_{anth}$ here):

$$DIC(x, y, z, t) = DIC(x, y, z, t_0) + \Delta C_{nat}(x, y, z, t) + C_{anth}(x, y, z, t). \tag{3.3}$$

Changes in natural carbon, ΔC_{nat}, are thus given by the difference in DIC between the RAD and CTR runs:

$$\Delta C_{nat} = DIC^{RAD} - DIC^{CTR}. \tag{3.4}$$

For further detail on model setup, see Couldrey et al. (2016) and Couldrey et al. (2019): the same simulations are used in these papers. Additionally, Couldrey et al. (2019) compared the representation of DIC and alkalinity in these models runs to GLODAPv2 observations (Lauvset et al., 2016), finding the modelled carbon cycle to be representative of observations, and so expect the carbon derived identification of excess temperature and salinity in this chapter ought to also be representative.

3.2.2 Relating the redistribution of temperature and carbon

In Chapter 1, Section 2.3, the mathematics relating changes in two arbitrary tracers was explored in detail. Here, this framework is presented in a more intuitive fashion, and its implementation in the NEMO OGCM detailed.

Following Williams et al. (2021), the preindustrial temperature and carbon fields of the ocean are broadly anticorrelated as a result of the strong inverse temperature dependence of carbon solubility. In contrast, the excess temperature and anthropogenic carbon fields are correlated due to the radiative forcing of atmospheric CO_2. Bronselaer and Zanna (2020) specify this correlation between excess heat and anthropogenic carbon using a time varying, globally uniform constant, which they refer to as the carbon-heat coupling or α. Here, the aim is to similarly relate the redistribution of temperature and natural carbon using an analogous redistribution coefficient, referred to as κ_r. As salinity will also be decomposed, superscripts are used to denote the variable which

will be related to natural carbon: the temperature redistribution coefficient, κ_r^T, refers to the preindustrial spatial covariability of natural carbon and temperature, whereas the salinity redistribution coefficient, κ_r^S refers to the preindustrial spatial covariability of natural carbon and salinity. Decomposing the total temperature change,

$$\Delta\theta(x,y,z,t) = \Delta\theta_e(x,y,z,t) + \Delta\theta_r(x,y,z,t), \tag{3.5}$$

where θ is in situ potential temperature, and the subscripts e and r refer to the excess and redistributed components, respectively. Again following the definitions of Winton et al. (2013) for the excess and redistributed temperature:

$$\vec{v}\theta = (\vec{v}_0 + \vec{v}')(\theta_0 + \theta') = \underbrace{\vec{v}_0\theta_0}_{Preindustrial} + \underbrace{\vec{v}'\theta_0}_{Redistributed} + \underbrace{\vec{v}_o\theta' + \vec{v}'\theta'}_{Excess}, \tag{3.6}$$

where $\vec{v}_o$ and θ_0 refer to the preindustrial components of the velocity field, $\vec{v}$, and the temperature field, θ, and $\vec{v}'$ and θ' the perturbations. Salinity, DIC and C_{anth} (or indeed any other tracer) changes may be decomposed in the same fashion. The excess component of a tracer can therefore be interpreted as changes in a tracer due to changes in surface forcing, and the redistributed component as changes in a tracer resulting from circulation change: redistribution driven changes in a tracer should therefore globally integrate to zero. At this point, note that the preindustrial distribution of C_{anth} is defined to be zero everywhere: thus C_{anth} well approximates excess carbon. However, as discussed in Chapter 1, Section 1.2.1.4, C_{nat} changes are not constrained to globally integrate to zero, and thus C_{anth} and excess carbon, though similar, are not the same entity.

The approach of Bronselaer and Zanna (2020) parameterises $\Delta\theta_e$ as

$$\Delta\theta_e(x,y,z,t) = \alpha_T(\Delta t) \times C_{anth}(x,y,z,t), \tag{3.7}$$

where α_T is their coefficient α, expressed in units of temperature rather than heat. α is estimated as the ratio of global heat to DIC accumulation, over the time period $\Delta t = t - t_0$, where t_0 is a preindustrial time (1860 here). Alternatively, we might parameterise the redistribution of temperature, $\Delta\theta_r$ in terms of the natural carbon change:

$$\Delta\theta_r(x,y,z,t) \approx \kappa_r^T(x,y,z) \times \Delta C_{nat}(x,y,z,t). \tag{3.8}$$

Unlike α_T, κ_r^T is not a function of time: it is instead a function of position, as it relates the spatial covariability of the preindustrial temperature and carbon fields at a given point. This method is equally applicable to any property for which we aim to estimate redistribution although each property pair will have a distinct distribution of κ_r: we could instead choose to relate the spatial covariability of the preindustrial salinity and

carbon fields. Redistributed salinity, ΔS_r, would therefore be estimated as

$$\Delta S_r(x,y,z,t) \approx \kappa_r^S(x,y,z) \times \Delta \mathrm{C_{nat}}(x,y,z,t). \tag{3.9}$$

In Equations 3.8 and 3.9, no constraint is made such that the global integral of redistributed heat is zero (or equivalently the global mean redistributed temperature is zero). If $\mathrm{C_{nat}}$ were a perfect tracer for redistribution, then its global integral would be zero. However, we would expect the global integral of $\Delta \mathrm{C_{nat}}$ to be nonzero, predominantly as a result of the outgassing of saturation carbon, $\mathrm{C_{sat}}$ (the DIC content of the ocean resulting from equilibrium with the preindustrial atmosphere), in response to ocean warming. Thus the quantities $\Delta \mathrm{DIC}_r$ (redistributed DIC) and $\Delta \mathrm{C_{nat}}$ will differ, particularly over timescales of multiple decades to centuries (Williams et al., 2021): this is reflected through the approximate rather than exact equality in Equations 3.8 and 3.9. In general, when integrating over the global ocean,

$$\frac{d}{dt} \iiint \mathrm{C_{nat}} dV < 0, \tag{3.10}$$

so it is necessary to correct for the divergence of $\mathrm{C_{nat}}$ and the ideal behaviour of a redistributed preindustrial carbon field using a repartitioning factor, referred to as γ. The corrected quantity is referred to as adjusted natural carbon, $\mathrm{C_{nat}^{adj}}$. This repartitions a fraction of anthropogenic carbon into the adjusted natural carbon in order to correct for $\mathrm{C_{sat}}$ outgassing.

This repartitioning allows us to force the global integral of adjusted natural carbon changes to zero. However, because globally integrated biology driven changes in $\mathrm{C_{nat}}$ may be nonzero, it is instead enforced that globally integrated redistributed heat, not adjusted natural carbon, is zero. The redistributed temperature field is therefore estimated as

$$\begin{aligned} \Delta\theta_r(x,y,z,t) &= \kappa_r^T(x,y,z) \times \Delta \mathrm{C_{nat}^{adj}}(x,y,z,t) \\ &= \kappa_r^T(x,y,z) \times \Big(\Delta \mathrm{C_{nat}}(x,y,z,t) + \gamma(t) \mathrm{C_{anth}}(x,y,z,t) \Big), \end{aligned} \tag{3.11}$$

where $\gamma(t)$ is a factor between 0 and 1 such that over the global ocean

$$\iiint \Delta\theta_r dV = 0 \tag{3.12}$$

at all times. γ must be less than 1 or $\mathrm{C_{nat}^{adj}}$ would exceed DIC, and so would not be physically meaningful. It is constrained to be positive as historically, atmospheric CO_2 has increased from preindustrial levels and so the global $\mathrm{C_{anth}}$ inventory is positive. However, if the global $\mathrm{C_{anth}}$ inventory were negative, γ could also be negative (though the magnitude is always less than 1).

As with redistributed temperature, redistributed salinity will be estimated as

$$\begin{aligned}\Delta S_r(x,y,z,t) &= \kappa_r^S(x,y,z) \times \Delta C_{nat}^{adj}(x,y,z,t) \\ &= \kappa_r^S(x,y,z) \times \Big(\Delta C_{nat}(x,y,z,t) + \gamma(t) C_{anth}(x,y,z,t)\Big).\end{aligned} \tag{3.13}$$

Also note that Equations 3.11 and 3.13 may be combined in order to directly estimate salinity redistribution from temperature redistribution and vice versa. This follows from the property that (assuming no perturbation to mixing) redistribution does not alter the properties of a parcel of water, and so the redistribution of natural carbon, temperature and salinity are related by the spatial covariability of their preindustrial fields. Alternatively, it might be observed that the choice of C_{nat} is not unique as a tracer for which to estimate redistribution: as previously noted, we only a require a tracer which can be considered to change only through redistribution. The sum of the preindustrial temperature or salinity fields and their redistributed components both satisfy this, and so can be used to estimate redistribution of other tracers themselves.

Relating the change in adjusted natural carbon to changes in temperature (κ_r^T) and salinity (κ_r^S) now requires an estimate of the redistribution coefficents. To estimate these, a statistical method was used, examining how the model temperature or salinity and C_{nat} fields covary on subdecadal timescales in our control run, in order to estimate the covariability of their preindustrial state. It is well known that when making repeated observations at a fixed spatial location, the majority of observed changes in temperature and salinity are due to circulation variability, rather than material changes in water mass properties on subdecadal timescales (for example Bindoff and Mcdougall (1994), Firing et al. (2017)). The dominance of circulation variability on these timescales can therefore be exploited, assuming that the correlation between deviations in temperature, salinity and DIC from their mean state on subdecadal timescales are due entirely to circulation variability. The correlations obtained allow us to estimate how circulation acts to couple changes in temperature and salinity to changes in natural carbon, at every point in the ocean. Thus, by looking at the relationship between temperature or salinity and DIC on subdecadal timescales, the spatial covariability of the background fields may be identified, with this approach not requiring an explicitly decomposed temperature or salinity field. Though these spatial correlations will change due to the addition of excess temperature (salinity), excess and redistributed temperature (salinity) are defined such that these preindustrial correlations correctly capture the relationship between redistributed temperature (salinity) and natural carbon change throughout the COU simulation (Equation 3.6)).

The calculation is performed as follows: in each grid cell, the full 240 years (1860-2099) of yearly mean temperature, salinity and DIC from our CTR run is used, binned into 24 decades. In each decadal bin, the mean tracer (θ or S and C_{nat}) values are subtracted, giving yearly θ, S and C_{nat} anomalies from the decadal mean in that grid cell. This

decadal binning is performed in order to preclude the possibility of any excess temperature or excess DIC contaminating the relationship as the result of model drifts or surface forcing driven variability due to the 30 year repeated forcing: though these effects should be small, they are both partitioned by the excess/redistribution decomposition into excess.

The correlations between the yearly anomalies from decadal means, for the entire 240 years of data, are then used to establish an intermediate value, which is labelled κ_i, at each grid cell. This is done using a total least squares linear fit, implemented as two dimensional PCA (Principal Component Analysis): κ_i is estimated as the gradient of the slope obtained. A total least squares fit, rather than an ordinary least squares fit is performed, as whilst the two variables are correlated, it is not a causal relationship: total least squares is therefore more appropriate, as the relationship obtained should not be affected by the choice of dependent variable.

A suppression factor, ϕ, is then calculated, based on the quality of the correlations to estimate κ_r for each variable: this process is detailed in Section 3.2.3, along with a visualisation of the estimation process. As with κ_r, this will be unique to each variable. ϕ is designed such that where the correlations we obtain between the θ or S and C_{nat} anomalies from decadal means are poor or nonexistent, no estimate of redistribution is made. As a result of this, if local C_{nat} changes due to biological processes but temperature or salinity due to circulation variability, this method will misclassify these changes as excess temperature or salinity: this also will occur at maxima/minima of temperature or salinity. However, due to the implementation of the γ correction, these misclassifications will globally integrate to zero. Over the full simulation, adjusted natural carbon increases by approximately 2μmol/kg, 0.1% of the mean preindustrial DIC concentration. This implies the net global divergence of C_{nat}^{adj} and ΔDIC_r is approximately 0.1%.

The full calculation is therefore performed as

$$\begin{aligned}\Delta\theta_r(x,y,z,t) &= \kappa_r^T(x,y,z) \times \Delta C_{nat}^{adj}(x,y,z,t) \\ &= \phi_\theta(x,y,z) \times \kappa_i^T(x,y,z) \times \Delta C_{nat}^{adj}(x,y,z,t)\end{aligned} \tag{3.14}$$

for temperature, and

$$\begin{aligned}\Delta S_r(x,y,z,t) &= \kappa_r^S(x,y,z) \times \Delta C_{nat}^{adj}(x,y,z,t) \\ &= \phi_S(x,y,z) \times \kappa_i^S(x,y,z) \times \Delta C_{nat}^{adj}(x,y,z,t)\end{aligned} \tag{3.15}$$

for salinity. Excess temperature was then calculated as

$$\begin{aligned}\Delta\theta_e(x,y,z,t) &= \Delta\theta(x,y,z,t) - \Delta\theta_r(x,y,z,t) \\ &= \mathrm{TMP}^{\mathrm{COU}}(x,y,z,t) - \mathrm{TMP}^{\mathrm{CTR}}(x,y,z,t) \\ &- \kappa_r^T(x,y,z) \times \Delta\mathrm{C}_{\mathrm{nat}}^{\mathrm{adj}}(x,y,z,t),\end{aligned} \tag{3.16}$$

and likewise for salinity.

Excess and redistributed density fields were then built from the decomposed temperature and salinity fields. To do this, the redistributed fields were added to the initial fields, and redistributed density calculated using TEOS-10 (McDougall and Barker, 2011). Initial density was then subtracted for density redistribution. Excess density fields were then calculated as the difference between the redistributed density field and the total density change.

3.2.3 Uncertainty in estimates of local redistribution

A local gradient, $\Delta\theta/\Delta\mathrm{C}_{\mathrm{nat}}$ or $\Delta S/\Delta\mathrm{C}_{\mathrm{nat}}$, was estimated by applying two dimensional PCA to the timeseries of yearly deviations of the two variables from their decadal mean values at each grid cell. This is equivalent to performing a total least squares fit to obtain a linear relationship between the two variables.

The data are then scaled to normalise the ranges of θ/S and $\mathrm{C}_{\mathrm{nat}}$ before again performing 2D PCA on these timeseries at each grid cell to estimate the fraction of the covariance contained within each principle component. This yields the fraction of the total variance explained by each principal component, referred to as ε_1 and ε_2: these can be thought of as the axes of an ellipse describing a scatter cloud relating the two variables. A fit which is a perfect line can be thought of as the limit of this ellipse where $\varepsilon_1 \to 1$ and $\varepsilon_2 \to 0$. Conversely, an essentially random fit through a spherical cloud of points can be thought of as the case where $\varepsilon_1 = \varepsilon_2$.

The eccentricity of this ellipse is then used as a suppression factor, ϕ_u:

$$\phi_u = \sqrt{1 - \left(\frac{\varepsilon_2}{\varepsilon_1}\right)^2} \tag{3.17}$$

The need for conservative estimates of confidence in the fit is particularly important for fits in which no discernable correlation can be drawn: for these, gradients associating minor changes in $\mathrm{C}_{\mathrm{nat}}$ with large changes in θ or S can be obtained, effectively at random, and so the suppression factor must remove these effectively. As this Chapter is concerned primarily with inventories, this approach was found to be preferable to including large uncertainties due to a small number of spurious points, or simply setting a threshold below which no attempt to diagnose the redistribution of heat and salinity

is made. Only 6% of ε_1 values are scaled by a factor of $1/2$ or less: this was found to be a suitable compromise, with only the most unreliable estimates strongly suppressed.

The final redistribution coefficient κ_r is then calculated as

$$\kappa_r = \phi_u \times \kappa_i \tag{3.18}$$

The implementation of this is demonstrated in Figure 3.1, for two points in the North Atlantic at approximately 24°N, 30°W and 850m and 1950m. The poorly correlated point, Figure 3.1a and 3.1c, is an extreme outlier, chosen for demonstrative purposes. Here the fit is essentially random, and so estimates of temperature redistribution are scaled to reflect this uncertainty: the eccentricity of the ellipse described by the cloud of points in $\theta - \mathrm{C_{nat}}$ space is used as a scale factor. For the strongly correlated point, shown in panels (b) and (d), temperature and $\mathrm{C_{nat}}$ variability are almost perfectly anticorrelated, representing the dominance of vertical structure in determining the redistribution coefficient κ_r^T. Here, $\kappa_r^T = -0.0210$, $\partial_z\theta/\partial_z\mathrm{DIC} = -0.0208$.

3.2.4 Merging one and two step estimates

The estimation of redistribution coefficients using the technique described above assumes that the relationship between short timescale changes is dominated by circulation variability. However, in the top 200m of the ocean, changes in salinity and $\mathrm{C_{nat}}$ are instead dominated by freshwater fluxes: an excess of evaporation over precipitation will increase concentrations of salt and $\mathrm{C_{nat}}$, coupling changes in the two. This leads to changes which are properly described as excess salinity being partitioned into redistributed salinity. To resolve this issue, salinity redistribution is recalculated using the same statistical approach to locally estimate the salinity redistribution from the redistributed temperature field: this will be referred to as a two step estimation. This calculation is performed as

$$\begin{aligned}\Delta S_r^2(x,y,z,t) &= \kappa_r^{T-S}(x,y,z) \times \Delta\theta_r(x,y,z,t) \\ &= \kappa_r^{T-S}(x,y,z) \times \phi_T(x,y,z) \times \kappa_i^T(x,y,z) \times \Delta\mathrm{C}_{\mathrm{nat}}^{\mathrm{adj}}(x,y,z,t),\end{aligned} \tag{3.19}$$

where the superscript 2 in $\Delta S_r^2(x,y,z,t)$ refers to the two step estimation. κ_r^{T-S} is an estimate of the T-S curve angle, and is estimated in the same way as κ_i^T and κ_i^S: no new suppression factor is calculated.

The two estimates are then merged using a sigmoidal weighting scheme based on depth. The grid in the simulations used here has 64 vertical levels, with the 20th level corresponding to approximately 200m. Denoting the $\mathrm{i^{th}}$ vertical level z_i, the one step estimate as S^1 and the two step estimate as S^2, the final estimate of salinity redistribution,

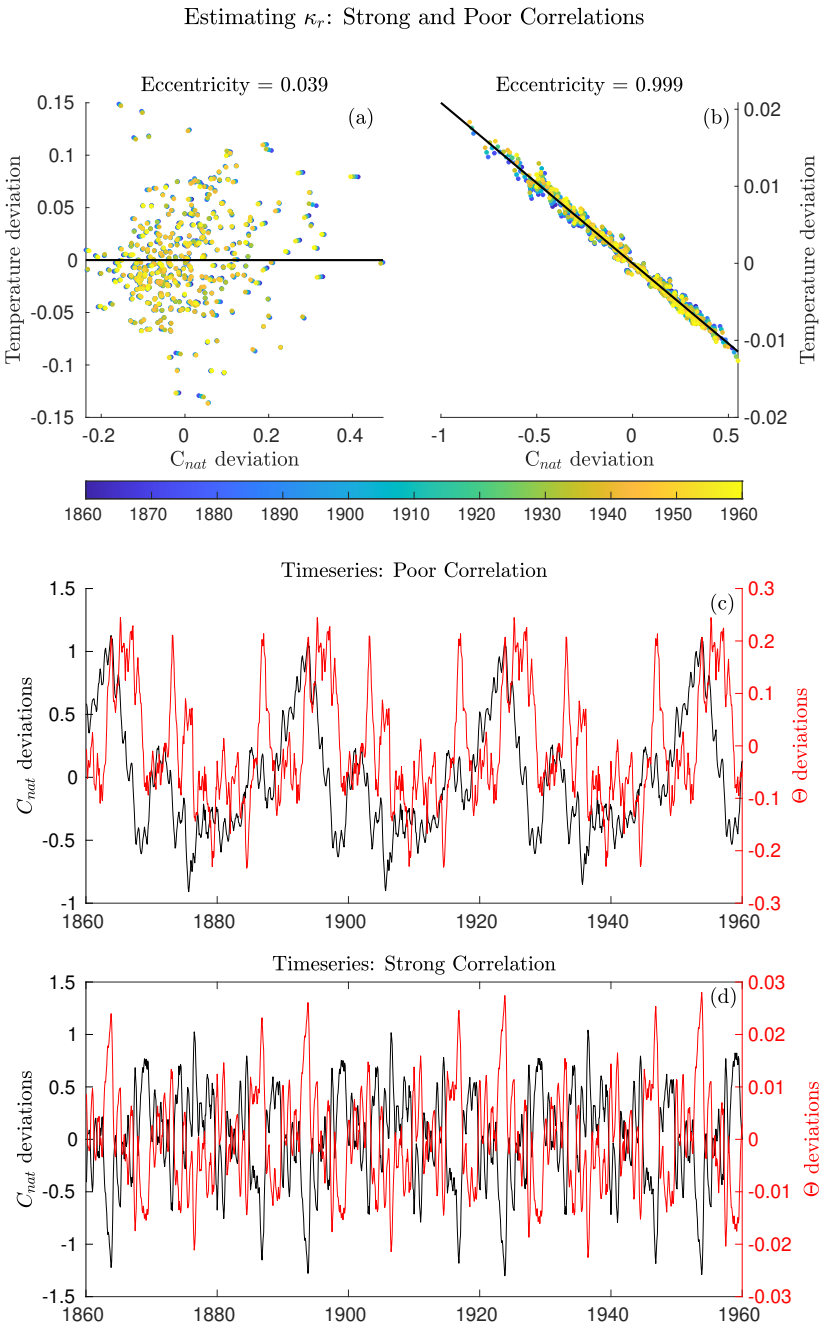

FIGURE 3.1: The correlations between C_{nat} and θ used to establish a κ_r value, for a poorly correlated point ((a),(c)), and a well correlated point ((b),(d)), in $\theta - C_{nat}$ space ((a), (b)), and timeseries of both ((c),(d)). These two points are located at 24N,30W in the Atlantic, at depths of 850 and 1950m. The major axis of the covariance ellipse in panels (a) and (b) is shown in black. Note these two points are chosen to demo

Sr, is then calculated as

$$S = S^1 \times \sigma\Big(\frac{z_i + 20}{2}\Big) + S^2 \times \left(1 - \sigma\Big(\frac{z_i + 20}{2}\Big)\right), \tag{3.20}$$

where $\sigma(z)$ is the sigmoid function:

$$\sigma(z) = \frac{1}{1 + e^{-z}} \tag{3.21}$$

For temperature, approximately 80% of grid cells globally have a scale factor of 0.8-1, and by the end of the model run, the suppression factor ϕ alters the redistributed temperature of 93% grid cells globally by less than 0.04 degrees, and 60% by less than 0.02 degrees, though the RMS mean redistributed temperature is reduced by 5%. However, the small number of grid cells producing extremely large estimates (10's of degrees of change) are effectively suppressed. It is therefore estimated that the statistical nature of this method introduces a minimum uncertainty of approximately 5% into the inventories calculated here.

3.2.5 Gamma Factor

γ was calculated for each year using Equation 3.11 to satisfy Equation 3.12: a fraction of C_{anth} was added to C_{nat} to ensure the global integral of redistributed heat is zero in each year, with the fraction representing the γ value that year. The value of γ is then smoothed over a 10 year period, before the fraction of C_{anth} each year given by the smoothed γ value is added back to C_{nat} to obtain the C_{nat}^{adj} field. Figure 3.2 shows the γ factor over our full run. It increases from 0 at the beginning of the run to 0.117 by 2099. Smoothing is performed as in the late 19th and early 20th century, C_{anth} inventories are small and so large corrections are necessary to perfectly correct a small amount of C_{sat} outgassing: smoothing is an effective way to remove this. By the 21st century, C_{anth} inventories are large enough that smoothing has little effect.

The γ factor does not begin to increase significantly until the late 20th century, approximately the same time that globally integrated ocean heat content begins to increase. Thus, to first order, γ corrects for C_{sat} outgassing due to ocean warming.

Once the C_{nat}^{adj} field had been built, it was used to generate both the redistributed temperature and salinity fields: we did not recalculate a new γ value to force a zero integral of salinity redistribution in our salinity decomposition. This approach was chosen for 3 reasons. Calculating a new γ for salinity would mean a new C_{nat}^{adj} field, and so the evolution of the redistributed temperature and salinity fields would not be linked by the same adjusted C_{nat}^{adj} field. In addition, the salinity of sea ice in the model (practical salinity of 6) and reduced carbon content of sea ice, relative to ocean DIC content,

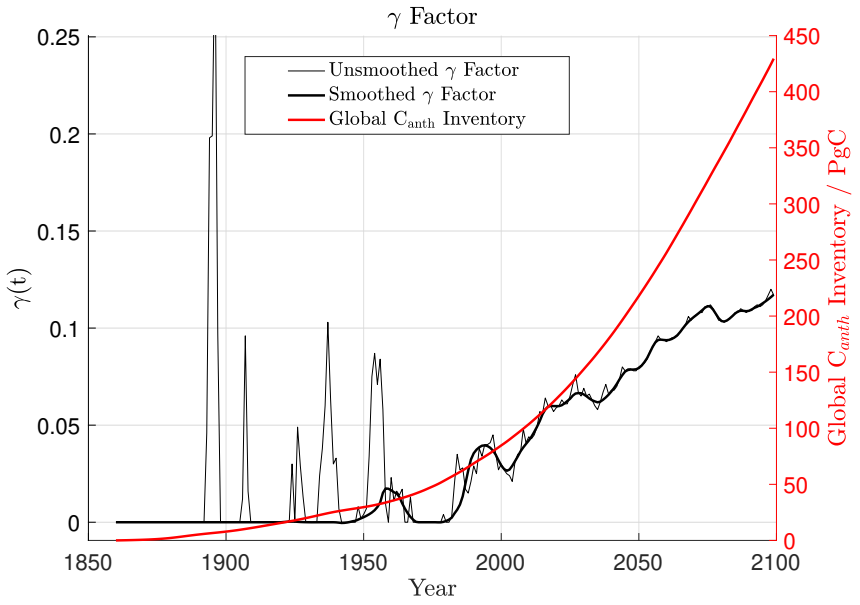

FIGURE 3.2: The calculated γ factor (thin black line), smoothed γ factor (thick black line), and global anthropogenic carbon inventory (red line).

causes some ice melt to be captured as redistributed salinity, rather than excess. This means that globally integrated salinity redistribution is not expected to sum to zero as is the case for temperature. However, as globally integrated redistributed salinity is not independently constrained to be zero, this allows this global integral to be used as a check on the validity of the method.

3.3 Results

3.3.1 Methodology Validation

In order to validate these results, a comparison with previous carbon proxy based methods is performed. The method of Bronselaer and Zanna (2020) relies on a globally uniform α value, linking carbon and heat changes at all scales, which they refer to as the carbon-heat coupling. In comparison, the technique presented here does not enforce global uniformity of this carbon-heat coupling, with a local carbon-heat coupling, $\Delta\theta_e/\Delta C_{anth}$ instead being an output of our method. For the rest of this Chapter, the global mean carbon-heat coupling will be referred to as α_T, and the local carbon-heat coupling as $\Delta\theta_e/\Delta C_{anth}$: specifically, the local carbon-heat coupling links the anthropogenic carbon and excess heat.

As we expect the correlations between the excess components of temperature and DIC changes to be positive, and between the redistributed components of temperature and

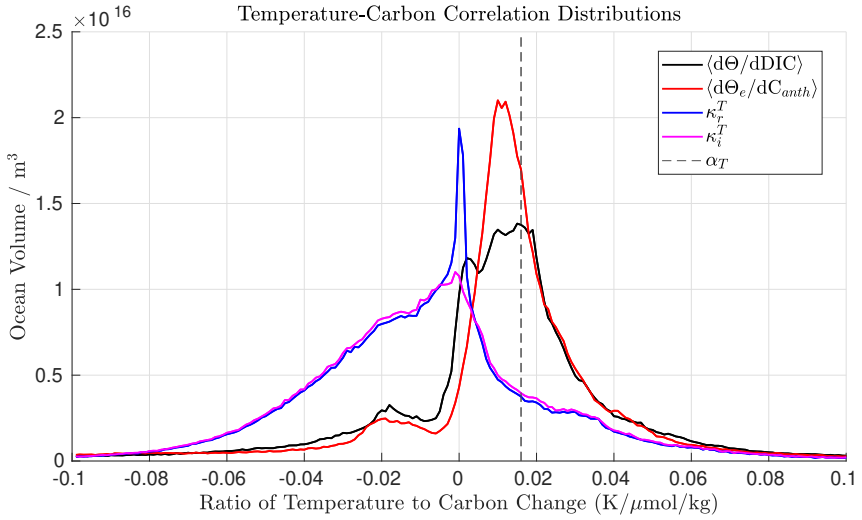

FIGURE 3.3: Histograms of the distribution of correlations relating different components of the temperature and carbon fields, over the full simulation (1860-2099). The global mean value α_T is shown by the dashed line. We include both the final redistribution coefficient, κ_r^T (blue), and its intermediate estimate, κ_i^T (magenta), as well as the ratio of total temperature change to DIC change ($\langle d\theta/d\text{DIC}\rangle$, black), and local excess temperature to anthropogenic carbon change ($\langle d\theta_e/d\text{C}_{\text{anth}}\rangle$, red).

DIC changes to be negative, we can infer whether the redistribution coefficient technique reliably estimates excess heat by comparing histograms of correlations between the different components of temperature and carbon change. To do this, the total temperature change to DIC change, the excess temperature change to C_{anth} change, and the redistributed temperature change to $\text{C}_{\text{nat}}^{\text{adj}}$ change (equivalent to κ_r^T), are compared for each grid cell at depths of less than 2000m in our simulations (volume weighted histograms of each of these quantities over the upper 2000m of the ocean is shown in Figure 3.3). Depths greater than 2000m are excluded due to the negligible ventilation and C_{anth} beyond this depth horizon. The total change and excess component correlations are calculated as the ratio of decadal mean temperature and carbon at each grid cell for the period 2090-2099 minus the initial values in 1860. Taking the assumption of a globally uniform α_T to be accurate, we expect to find a broad distribution of ratios of total temperature change to DIC change with both positive and negative correlations, and a narrower distribution of ratios of excess temperature change to C_{anth} change, centred about the global mean value α_T. We would also expect the correlations between redistributed temperature changes and $\text{C}_{\text{nat}}^{\text{adj}}$ to generally be negative.

In Figure 3.3, distribution of the ratio of total temperature change to DIC change (black line) is generally positive, indicating the dominance of excess temperature and DIC over redistribution over this time period and region, but is broad and encompasses both positive and negative values. Its mode occurs at the global mean value α_T: 0.016 K/μmol/kg. The mode of the ratio of excess temperature change to C_{anth} accumulation (red line) is slightly lower (0.012-0.014 K/μmol/kg), but the magnitude of the peak at the mode is approximately 50% greater than that of total change (2.1×10^{16} m^3

vs $1.4 \times 10^{16}\,m^3$). This implies that the assumption of a globally uniform carbon-heat coupling, α, is broadly appropriate, although a large spread in values exists, and that the redistribution coefficient method reliably identifies excess heat.

The distribution of the ratio of redistributed temperature change to C_{nat}^{adj} change (κ_r^T, blue line), is also generally negative, as expected, with a much broader distribution than the distribution of the ratio of excess temperature and C_{anth}. Generally, the intermediate value histogram (κ_i^T, magenta line) resembles the final ratio (κ_r^T, blue line), with the exception of the large peak at zero, resulting from the suppression factor, ϕ_T. Figure 3.4 shows patterns of κ_r^T for the Atlantic (a), Indian (b), and Pacific (c) oceans, and Figure 3.5 the patterns of effective κ_r^S for these basins. The positive tail of κ_r^T values is predominantly due to the inversion of the DIC field with depth in the North Pacific.

Figure 3.4 shows patterns of κ_r^T, representing the spatial covariability of temperature and salinity background distributions throughout the global ocean. As redistributed salinity estimates are produced from the combination of a one step and two step estimate (as described in Section 3.2.4), the patterns of κ_r^S shown in Figure 3.5 are instead a map of effective values. These are calculated by diagnosing the mean redistributed salinity for the decade 2090-2099 (this was chosen to maximise changes to adjusted C_{nat} and thus avoid numerical issues), before dividing by mean changes in adjusted C_{nat} to obtain a map of effective κ_r^S values. Thus, these values are identical (to numerical precision) to those calculated directly in the mid-depth and deep ocean, but represent a best estimate of the spatial coupling between salinity and natural carbon in the upper 200m of the ocean, and avoid the complicating effects of freshwater dilution.

That the correlation between the redistribution of temperature and carbon is positive in the North Pacific implies that the method of Williams et al. (2021) may not be appropriate in this location. However, the shape of the distributions presented are in clear agreement with their method: the redistribution coefficient decomposition generally infers negative correlations between redistributed temperature and natural carbon, and positive correlations between excess temperature and anthropogenic carbon. As the redistribution coefficient method identifies correlations between excess temperature and anthropogenic carbon, and between redistributed temperature and natural carbon changes that are consistent with both the assumptions of Williams et al. (2021) and Bronselaer and Zanna (2020), despite not enforcing this to be the case, we may be confident that it is reliably separating excess and redistributed temperature.

Estimates of excess temperature from the redistribution coefficient method and the method of Bronselaer and Zanna (2020) are now compared, with the assumption that both methods producing consistent estimates indicates an accurate identification of the excess temperature field. Figure 3.6 shows the zonally averaged excess and total temperature fields obtained for the Atlantic and Indo-Pacific, for the final decade of the simulations, 2090-2099. In the Atlantic and Indo-Pacific, the estimate using the method

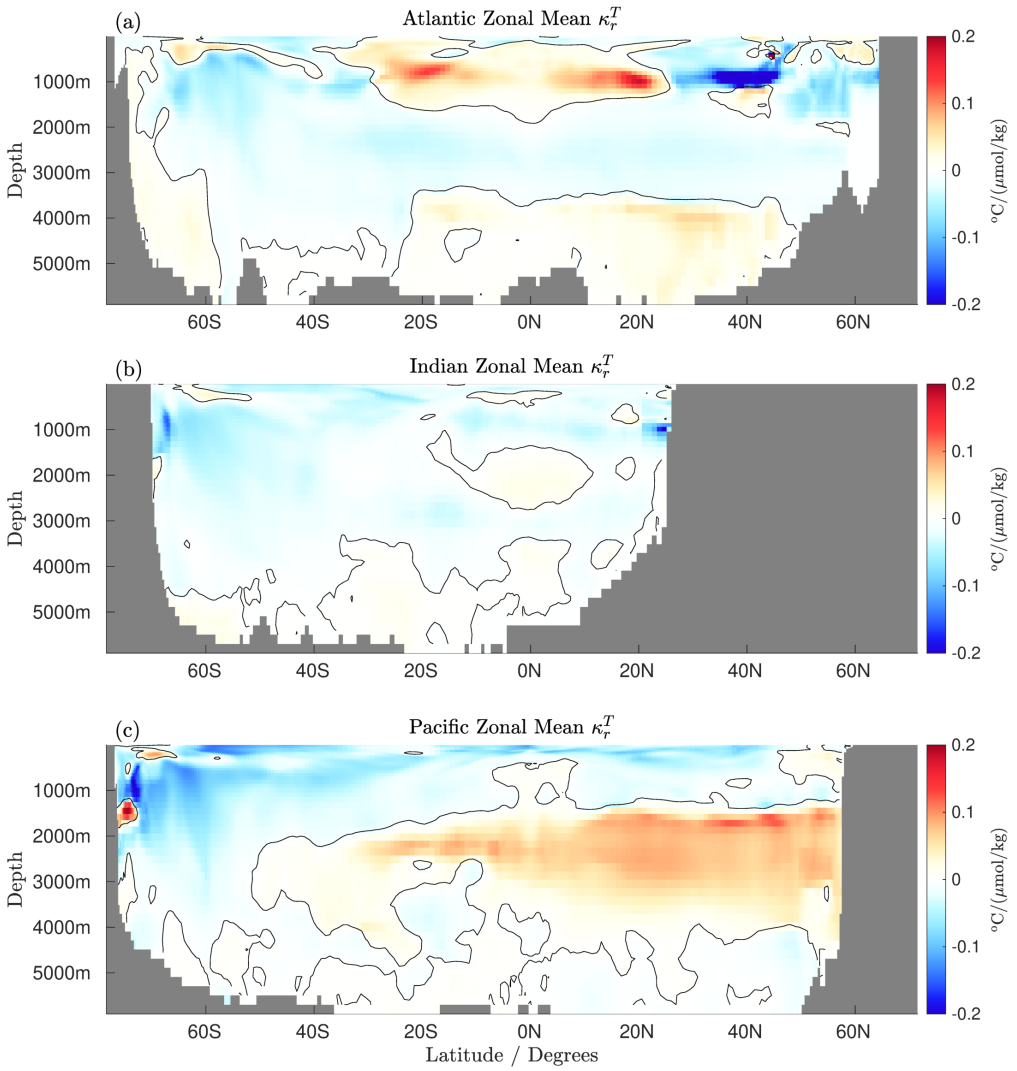

FIGURE 3.4: Zonal mean κ_r^T values for the Atlantic (a), the Indian ocean (b) and the Pacific (c).

of Bronselaer and Zanna (2020) (Figure 3.6a,b) produces smoother estimates than the redistribution coefficient technique (Figure 3.6c,d), but there are a number of common features which both techniques identify that are not due to the accumulation of excess heat. In the Atlantic, the tongue of warming at 2000-2500m depth, extending from approximately 40°N to 30°S is identified by both techniques as redistribution of the preindustrial temperature field, rather than excess heat. In addition, both techniques identify the region of warming extending from approximately 2000-4000m depth between 60°S and 40°S as redistributed, rather than excess heat. In the Indo-Pacific, both methods identify the cooling at approximately 1000m at 20°S as redistribution, rather than excess temperature. However, the redistribution coefficient method identifies the penetration of excess temperature to depth in the Southern Ocean, unlike the method of Bronselaer and Zanna (2020).

In the upper 1000m, there are significant divergences between the two techniques. To

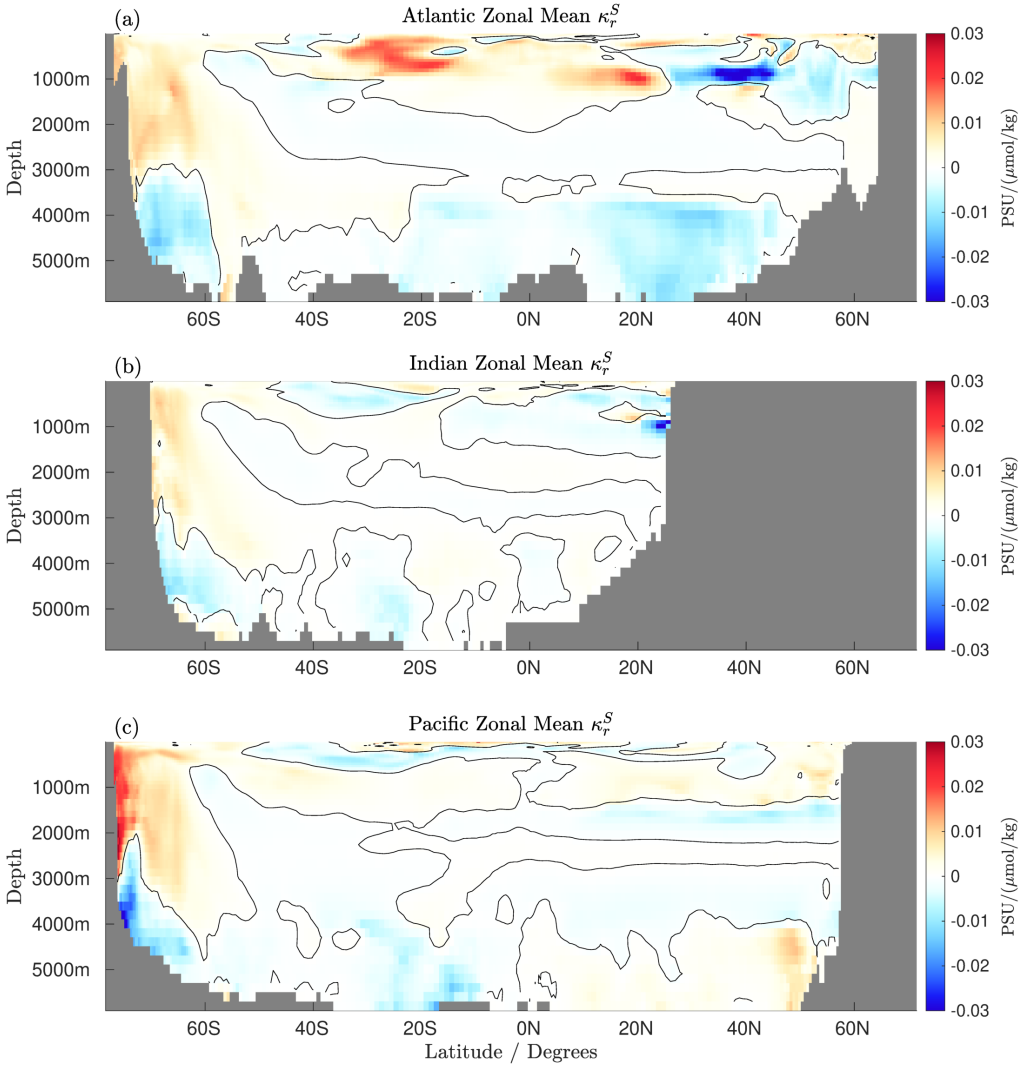

FIGURE 3.5: Zonal mean effective κ_r^S values for the Atlantic (a), the Indian ocean (b) and the Pacific (c).

explore the sources of these differences, local estimates of the quantity $\Delta\theta_e/\Delta C_{anth}$ are computed from estimates of $\Delta\theta_e$ obtained and the model C_{anth} output. By comparing the locally obtained estimates with the patterns of excess heat and anthropogenic carbon uptake estimated by assuming a globally uniform α_T, we are able to show how the relaxation of the assumption of a globally uniform α_T causes our estimates to differ. This is demonstrated in Figure 3.7.

Figure 3.7a and 3.73b show local values of $\Delta\theta_e/\Delta C_{anth}$, presented as the zonal mean of the ratio of total excess temperature accumulated to total anthropogenic carbon accumulated, averaged over the decade 2090-2099. Figure 3.7c and 3.7d show the differences between the excess temperature estimated using our technique, and estimated using the technique of Bronselaer and Zanna (2020), and Figure 3.7e, 3.7f shows the total C_{anth} accumulated over the same period and domain.

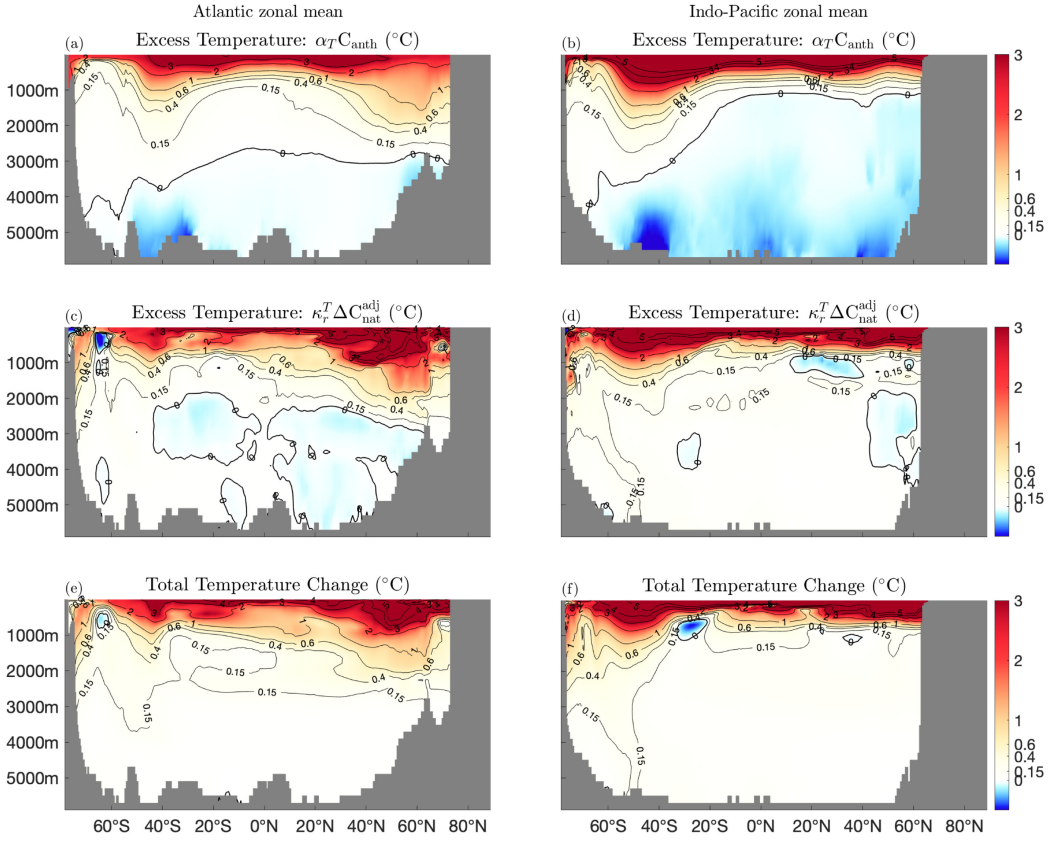

FIGURE 3.6: Atlantic and Indo-Pacific zonal, decadal mean excess temperature estimates, for the decade 2090-99, and total temperature change. The method of Bronselaer and Zanna (2020) is shown in panels (a,b), our method in panels (c,d), and the total temperature change in panels (e,f). The thick black contour indicates the zero contour, and temperature changes are indicated by thin contours, which are also indicated on the colour axes.

At depths of below 2000m in the Atlantic and 1000m in the North Pacific, ventilation is negligible and so despite large $\Delta\theta_e / \Delta C_{anth}$ estimates, the two methods produce similar estimates of excess temperature. In the Southern Ocean, North Atlantic and North Pacific, we see large $\Delta\theta_e / \Delta C_{anth}$ values, as well as nontrivial accumulation of excess temperature. As a result, in these regions, the two estimates diverge.

In general, estimates of $\Delta\theta_e / \Delta C_{anth}$ produced using the redistribution coefficient method show a large degree of spatial coherence, despite no constraints being imposed to enforce this. This suggests that these variations are likely real, rather than an artifact of the estimation technique. An implication of this is that heat uptake is intensified, relative to C_{anth} uptake, in the high latitude Northern Hemisphere, and reduced in the low latitude Northern Hemisphere and Southern Hemisphere. This may be explained in terms of a reduction of carbon export through the mixed layer at high latitudes. Bronselaer and Zanna (2020) make an argument for a globally uniform α value based on surface carbonate chemistry. However, Bopp et al. (2015) found total C_{anth} subduction through the base of the mixed layer to be significantly more variable than air-sea C_{anth} fluxes,

and generally reduced at high latitudes (their Figure 3c): this mechanism could potentially act to reduce the spatial uniformity of α below the base of the mixed layer. In particular, water masses where the effects of advection and vertical mixing on carbon subduction are in opposition (namely high latitudes) tend to produce higher values of $\Delta\theta_e/\Delta C_{anth}$. This is in agreement with the semi-empirical estimate of $\Delta\theta_e/\Delta C_{anth}$ produced using CO2sys in Chapter 2.

To test whether these variations in local values of $\Delta\theta_e/\Delta C_{anth}$ may constitute a source of error in the method of Bronselaer and Zanna (2020), the column inventories of excess heat uptake over the top 2000m of the ocean obtained using both methods in these simulations are compared: this is shown in Figure 3.8. Bronselaer and Zanna (2020) were able to directly compare their estimates of excess heat and the simulated excess heat (their Figure 3f, reproduced in Figure 3.9). Although the two estimates do differ, these differences (Figure 3.8c) closely resemble those between their method and the simulated excess (Figure 3.9): their method underestimates excess heat content in the North Atlantic, North Pacific and Southern Oceans, whilst overestimating it in the South Atlantic, South Pacific and Indian Oceans: patterns which are broadly reproduced by our method. The zonally integrated difference in upper 2000m excess heat content (Figure 3.8d) is again consistent with a reduction of carbon export through the mixed layer base at high latitudes.

Finally, due to the absence of simulations explicitly including PAT, the accuracy of the decomposition and the appropriateness of Equations 3.11 is investigated by examining the total and redistributed heat content change of a region of ocean for which ventilation is negligible; here, excess heat content is assumed to be zero and heat content change and redistributed heat content change are expected be identical. Thus, the accuracy of the decomposition can be examined, without the need for simulations including PAT.

The unventilated region of the ocean is defined as the fixed volume of all grid cells in our model runs for which C_{anth} concentrations do not exceed a given threshold ($|1|\mu$mol/kg) at any time during the model run. In addition, only grid cells in the depth range 2000-3000m in this region are considered: this excludes grid cells for which large spurious, negative C_{anth} values are seen (See Figure 3.6). Changes in the heat content of this region of ocean are therefore the result of changes to ocean circulation. For example, a slowdown in the AMOC, even absent any excess heat content in entering the ocean, would result in a reduction in the mean temperature of the North Atlantic, as the northward heat transport is reduced. Were we to a pick a volume of ocean in the North Atlantic as our 'unventilated ocean' region in such a case, we would find a redistribution driven decrease in the mean temperature of this volume.

Figure 3.10 shows the total heat inventory change, and our reconstruction redistributed heat content for this region of unventilated ocean: for a perfect reconstruction, these

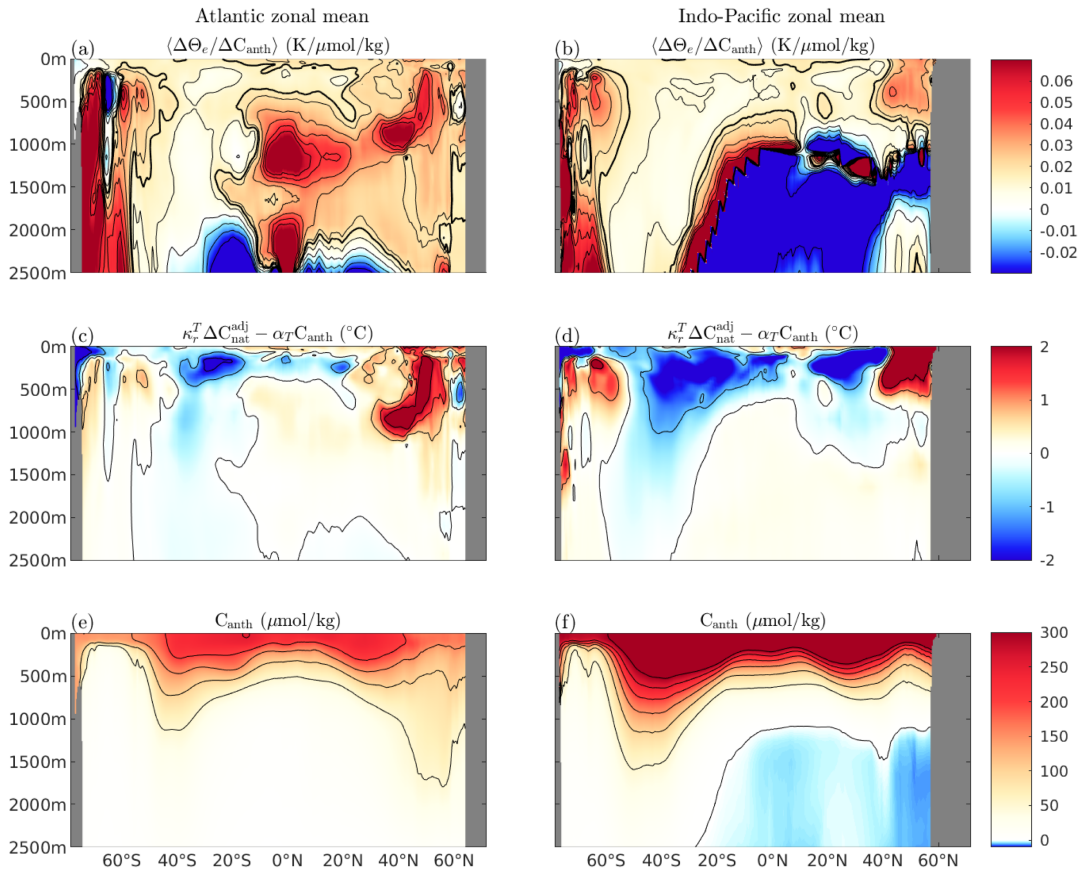

FIGURE 3.7: Atlantic and Indo-Pacific zonal mean ratio of excess temperature to C_{anth} accumulation, calculated as the 2090's decadal, zonal mean temperature divided by 2090's decadal, zonal mean C_{anth} (Panels (a), (b)). Panels (c) and (d) show the difference between our excess temperature estimate and the excess temperature estimate produced using the method of Bronselaer and Zanna (2020), and Panels (e) and (f) the zonal mean C_{anth} accumulation, calculated as the 2090's decadal mean. The thick black contour in Panels (a), (b) indicate the global mean value of α_T of 0.016K/μmol/kg, and the thin contours are indicated on the colour axes.

quantities would be nearly identical, and this is indeed the case, in particular for higher frequency variability (3.10a), and for longer term variability before approximately 2050 (3.10b). However, over the full model run, the reconstruction begins to systematically diverge, with the estimate of mean temperature change due to redistribution slightly lower than the mean temperature change.

To determine the quality of this reconstruction, Taylor Skill Scores are computed (Taylor, 2001) for the periods 1890-1950, and 1890-2099. Following Hirota et al. (2011), these are calculated as

$$S = \frac{(1+R)^4}{4(\hat{\sigma_f} + 1/\hat{\sigma_f})}, \tag{3.22}$$

where S is the Taylor Skill Score, R the Pearson correlation coefficient of the two timeseries, and $\hat{\sigma_f}$ the ratio of the model output timeseries standard deviation (σ_f) and the

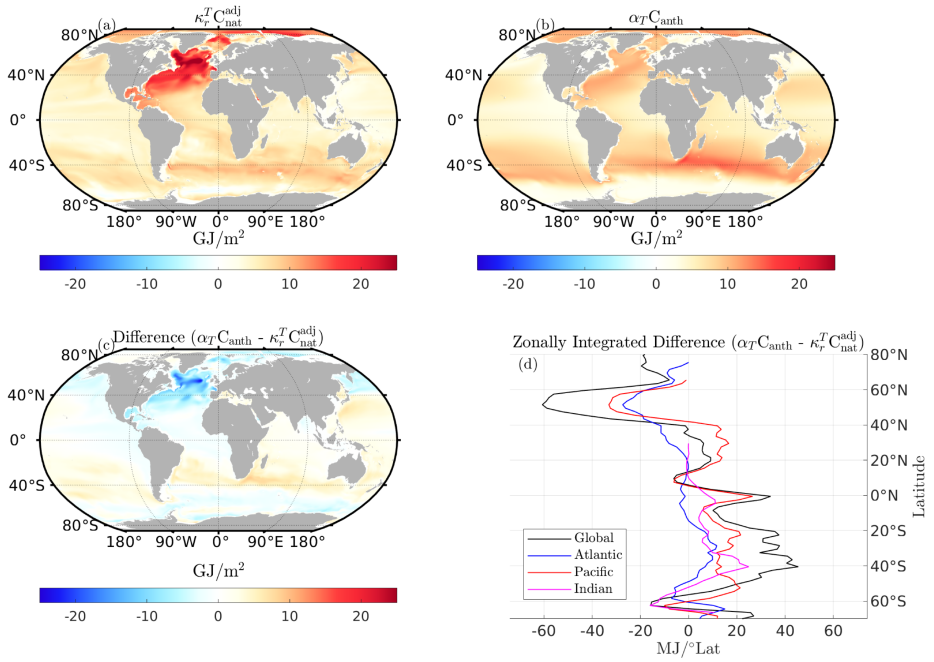

FIGURE 3.8: Column inventories of excess heat (0-2000m), calculated as the 2090-2099 decadal mean excess temperature relative to preindustrial. Panel (a) shows our method, panel (b) the method of Bronselaer and Zanna (2020), and panel (c) the difference between the two estimates. Panel (d) shows the zonally integrated upper 2000m excess heat content difference.

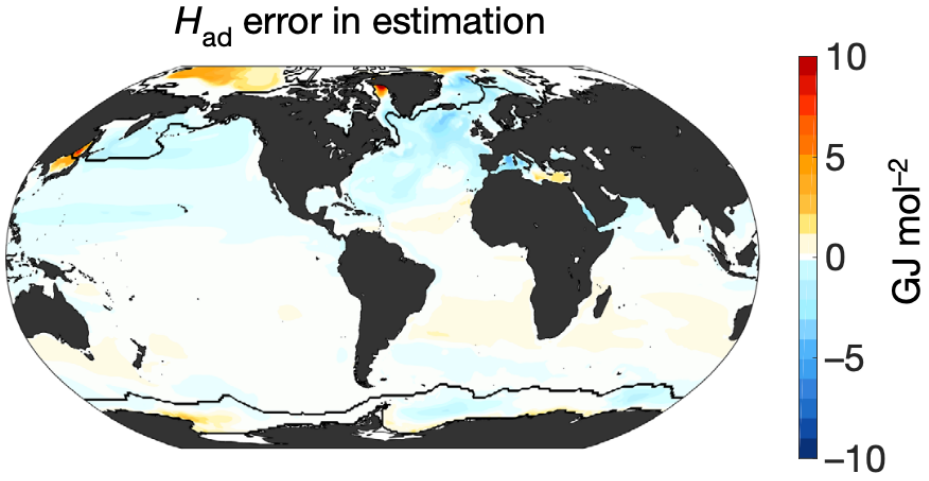

FIGURE 3.9: Column inventories of the difference in estimated and simulated excess heat (0-2000m), reproduced from Bronselaer and Zanna (2020).

reconstruction timeseries standard deviation (σ_r):

$$\hat{\sigma_f} = \frac{\sigma_f}{\sigma_r}. \tag{3.23}$$

For the period 1890-1950 (Figure 3.10a), before meaningful climate change, we obtain a Taylor Skill Score of 0.838: thus representing skillful reconstruction of the redistributed heat content.

Over the full time period (Figure 3.10b), the reconstruction again captures the higher frequency variability in heat content of the unventilated region well, though the rate of accumulation is somewhat lower: the ratio of mean temperature change to mean redistributed temperature asymptotes to approximately 80% over the full model run. Over this period, we obtain a slightly higher Taylor Skill Score of 0.964, indicating the reconstruction is capturing the redistribution of heat into this region accurately on longer timescales, although again biased low as with the shorter timescale changes. The decomposition is thus thought to capture approximately 80% of the forced circulation change.

However, the definition of unventilated waters used here (Canth $\ll 1$ μmol/kg) may be too expansive, and thus the 20% uncertainty detailed here represents an upper estimate of the error introduced by our method: as our cutoff is nonzero, some excess heat will enter this region, which will act to systematically increase the total heat content in this region beyond the redistributed heat content.

In addition, truncating the period considered to 1890-2050, a Taylor Skill Score of 0.987 is obtained, and over the period 2000-2050, the mean ratio of total temperature changes to redistribution driven temperature changes is 0.994, suggesting a highly accurate reconstruction.

As the redistribution coefficient method produces results broadly consistent with both the method of Bronselaer and Zanna (2020) and Williams et al. (2021), it appears to reliably identify excess temperature. In conjunction with the high Taylor Skill Scores and plausible explanation for differences between the results of the two methods, these results imply that the redistribution coefficient method accurately identifies excess heat.

3.3.2 Inventory Changes

Global mean excess and redistributed salinity change, as well as globally integrated excess and redistributed heat content change are shown in Figure 3.11. The global mean excess and redistributed salinity (thick lines, Figure 3.11a) begin to show a small net decrease in 1891, when the RAD and CTR forcing ceases to be identical, though this sea ice melt driven decrease is much smaller than the scale of either the positive and negative only excess or redistributed salinity components (thin dashed lines): global mean

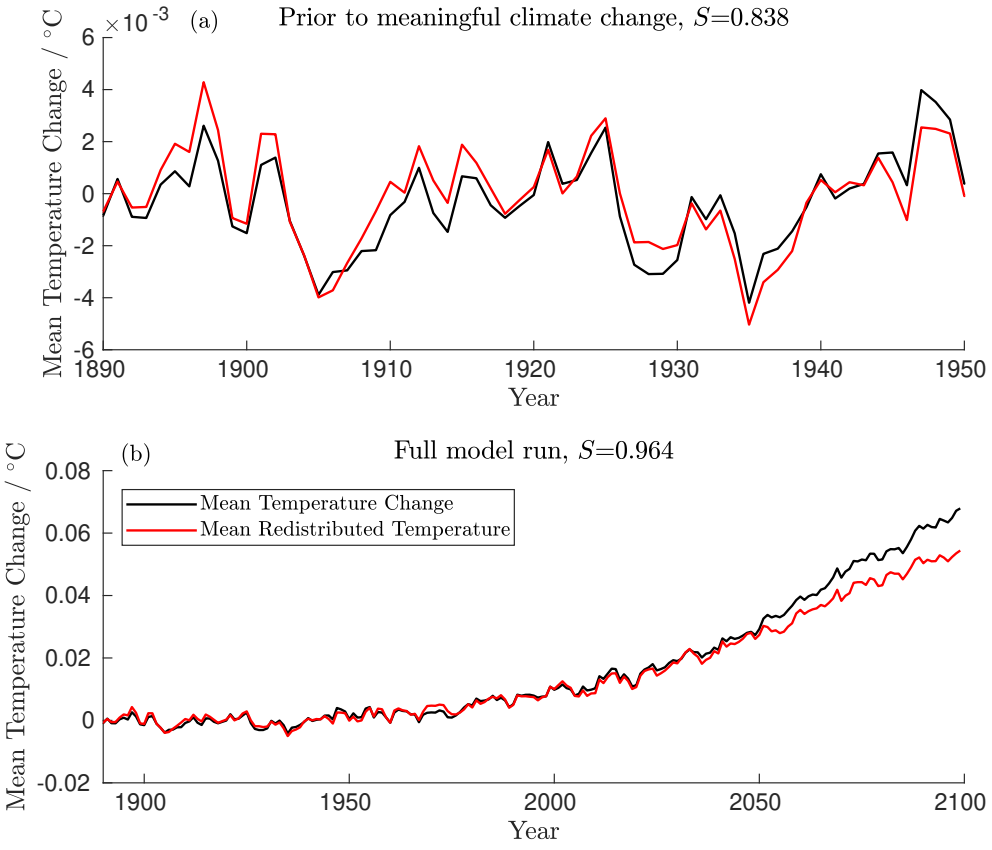

FIGURE 3.10: Total and redistribution driven changes in the mean temperature of the unventilated ocean for the period before meaningful climate change (a), and the full model run (b). Taylor Skill Scores for the periods presented are shown in subplot titles.

excess and redistributed salinity both decrease by approximately 0.001 over the full run. Globally integrated excess heat does not begin to accumulate significantly until approximately 2000: until this point, both positive only (global integral of excess heat content only in regions where excess temperature is positive) and negative only excess and redistributed heat are of similar scales. Positive only excess heat and globally integrated excess heat are approximately the same by 2050, and negative only excess heat increases from approximately -200 ZJ in 2000 to approximately -50 ZJ by 2100: some negative excess heat due to cooling in the first half of the run remains throughout the full simulation.

The global integral of positive and negative only regions are useful for assessing the extent of redistribution: whilst the global integral of redistributed temperature is constrained to be zero, this is the result of the cancellation of the positive and negative regions. Whilst excess heat begins to dominate during the mid 21st century, the extent of temperature (and salinity) redistribution continually increases: there is no indication of 'settling' into a new circulation state, where redistribution ceases to increase, on the timescale of the full simulation. This can be seen from the continued and accelerating

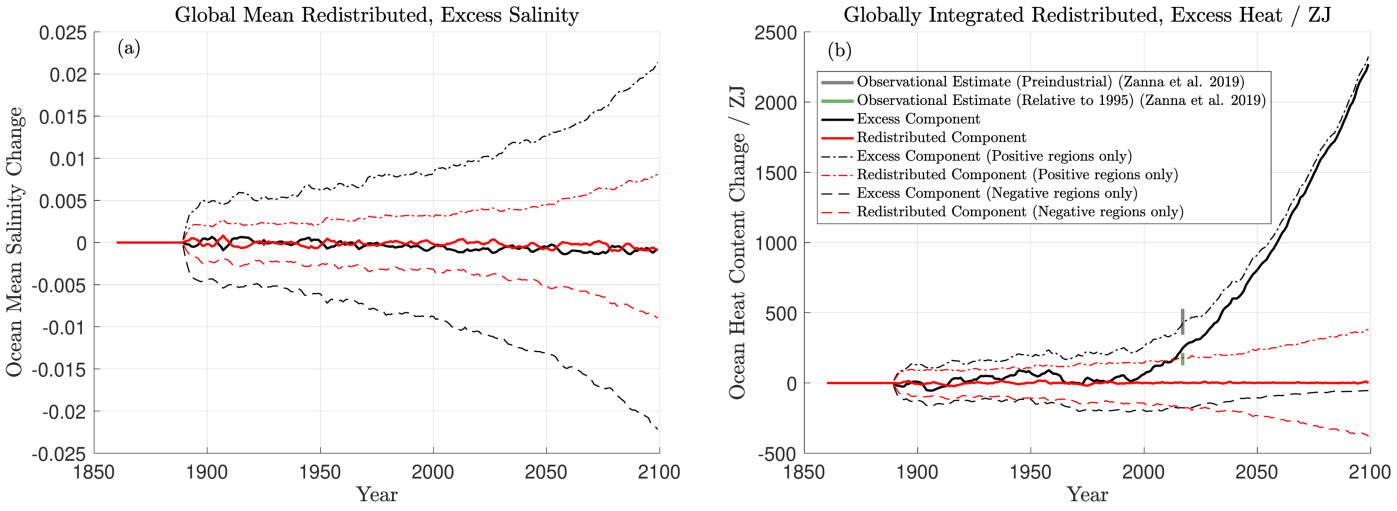

FIGURE 3.11: Global mean excess and redistributed salinity (a), and globally integrated excess and redistributed heat (b). Excess components are shown in black, redistributed components in red. The integrals of only the positive and negative regions are also shown (thin dashed lines). Climate change and control runs use the same first 30 years of forcing, so values are by definition zero here: the jump in 1890 represents the initial divergence of states. Observational estimates of global ocean heat uptake from Zanna et al. (2019) are also shown in panel (b).

increases in positive and negative only redistributed heat and salinity. It is also the case that whilst the magnitude of positive and negative only redistributed heat are similar until approximately 2000, excess salinity is significantly larger than redistributed salinity at all times. This indicates that during the full course of our simulations, salinity changes are dominated by changes in the freshwater cycle, rather than changes in circulation.

For comparison, observational estimates of ocean heat uptake from Zanna et al. (2019) (Figure 3.11b) are also included: cumulative heat uptake over 1871-2015 in grey (436$\pm$91 ZJ) and over 1995-2015 in green (153$\pm$44 ZJ). Over the period 1871-2015, our simulations produce less cumulative heat uptake (249 ZJ), but more heat uptake over 1992-2015 (232ZJ) than observations.

Figure 3.12 shows the integrated redistributed and excess temperature, salinity, and densities for each ocean basin. As with the global mean, excess salinity begins to accumulate almost immediately in most ocean basins (Figure 3.12c), particularly the North Atlantic and South Pacific: trends here are distinct from noise at 2σ in 1893 and 1911, respectively. Excess temperature does not begin to accumulate until the 21st century, at which point it begins to rapidly accumulate in all ocean basins; the exception to this is the South Atlantic (Figure 3.12a, dashed black line) which cools in the 20th century, its excess heat signal emerging from noise at 2σ in 1918. In contrast, the excess heat signal in the North Atlantic and South Pacific do not emerge from noise at 2σ until 2023 and 2021, respectively. Over the period 2023-2099, for which the excess heat signal of

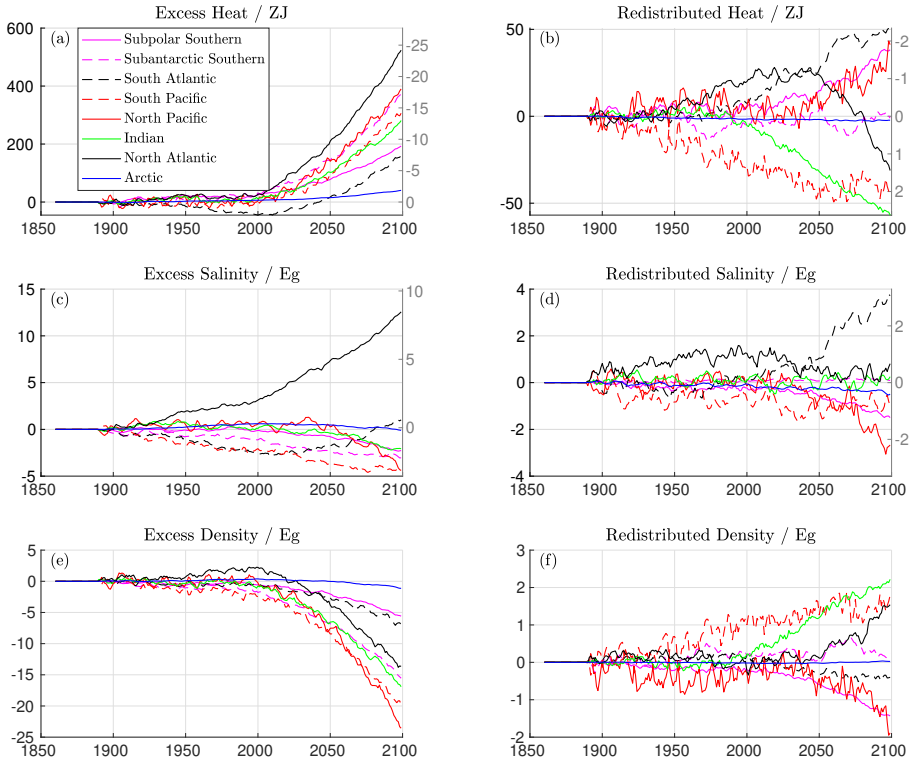

FIGURE 3.12: Excess (left column) and redistributed (right) heat, salinity and density integrals for each ocean basin over the full model run. For the changes in heat and salinity (Panels (a)-(d)), the equivalent integrated density change in units of Pg are given in grey on the right. Scales differ for excess and redistributed components, and changes in salinity and density are given as mass changes rather than volumes.

the North Atlantic is distinct from noise, 25±2% of global excess heat accumulated is located in the North Atlantic.

The accumulation of negative excess density is dominated by the accumulation of excess temperature, rather than salinity: the grey scales on the right hand side of panels (a)-(d) in Figure 3.12 show the density change associated with heat and salinity change. In the North Atlantic, changes in the excess heat and salinity compensate to reduce density anomalies: a reduction of almost 25Pg associated with excess heat is compensated for by an increase of approximately 8Pg associated with increased salinity. Similar compensation, though much weaker, is seen in the South Atlantic, which cools and freshens during the 20th century before warming and salinifying in the 21st. This is not the case in other basins, where the changes in excess heat and salinity both act to decrease density and therefore increase stratification (Figure 3.12a, 3.12c).

The redistribution of density is less dominated by heat, with heat and salinity contributing similarly to redistributed density. In the North Atlantic, the redistribution of heat and salinity are approximately density compensated until around 2050, at which

point the redistributed density inventory begins to increase rapidly (Figure 3.12f, black line). Good density compensation in the redistributed component is also seen in the subantarctic Southern Ocean, with minimal accumulation of redistributed density.

In the COU run, AMOC strength (calculated as peak depth integrated meridional volume transport at 26°N) increases until 1990 before declining continually thereafter. The cumulative transport anomaly (time integrated difference between COU and CTR AMOC volume transport) peaks in 2035 before also declining continually for the rest of the simulation. The signal of AMOC decline is visible in the redistributed heat content of the North Atlantic, which peaks in 2037 before declining rapidly, as well as the redistributed salinity content of the South Atlantic, which begins to increase at approximately the same time: consistent with previous studies (Zhu and Liu, 2020) which find a 'pile up' of salinity in the South Atlantic as a result of AMOC slowdown. The AMOC in these simulations is too weak, with a preindustrial mean of approximately 7.5 Sv at 26°N, and a maximum value of 13Sv in our COU run, declining to approximately 4.5 Sv by 2099, as compared to approximately 15 Sv in HadGEM2-ES (Martin et al., 2011) and 18±4.9 Sv observationally (Johns et al., 2011). This AMOC strength at 26°N in HadGEM2-ES itself is towards the weaker end of estimates from CMIP5 models (Weaver et al., 2012). However, the heat transport is realistic, with a control run heat transport of 0.075 PW/Sv at 26°N, as compared to observations of 0.079 PW/Sv (Johns et al., 2011). The decline in AMOC strength in our ocean only simulations and HadGEM2-ES simulations are also proportional: over an RCP8.5 scenario, Sgubin et al. (2015) found a decline of AMOC strength at 26°N from approximately 15.5 to 8Sv at 26°N in HadGEM2-ES.

To explicitly test whether the redistribution of heat from the North Atlantic, and salinity to the South Atlantic, can be explained in terms of a changing AMOC, we calculate the redistribution of heat and salinity through the Equator in the Atlantic. This is calculated as the difference in meridional velocities between the COU and CTR runs, multiplied by the control run temperature and salinity fields (this analysis is conceptually similar to that performed by Williams et al. (2021) in order to calculate the redistribution into/out of a volume, though here we consider only the equatorial boundary between the North and South Atlantic). For the period 1950-2099, for which there are non-negligible changes in the redistributed heat content of the North Atlantic, we find the correlation between the redistributed heat content of the North Atlantic and the redistribution of heat through the Equator due to AMOC change has an R^2 value of 0.58, suggesting that the change in overturning circulation plays a key role redistributing heat out of the North Atlantic, and into the South Atlantic. A slightly weaker correlation is found between the non AMOC driven redistribution of heat past the equator and the North Atlantic heat inventory, with a R^2 value of 0.45. These R^2 values are reduced to 0.50 and 0.38, respectively, when considering the period 1890-2099.

The picture is similar for salinity: for the period 2000-2099, for which there are non-negligible changes in the redistributed salinity content of the South Atlantic, the correlation between the redistributed salinity content of the South Atlantic and redistribution of salinity through the equator due to AMOC change has an R^2 value of 0.61, which is reduced to 0.04 when considering the period 1890-2099. Changes due to gyre circulation driven redistribution have R^2 values of 0.09 and 0.33, respectively, suggesting that the large scale mechanisms of salinity redistribution differ from those of heat.

Given the reliability of the outputs of the decomposition technique presented here, we may confidently assert this decomposition is reliably identifying excess heat and salinity. It therefore appears that the redistribution of heat out of the North Atlantic and salinity into the South Atlantic are driven predominantly by AMOC variability, with non AMOC circulation changes influencing the redistribution of temperature and salinity differently. Identifying whether the lack of correlation between our estimates and the explicitly calculated redistribution when there is no appreciable accumulation of either is due to inaccuracies in the redistribution coefficient approach or the dominance of other factors in the redistribution of heat and salinity would likely improve understanding of the strengths and weaknesses of this method, but is beyond the scope of this study.

As with the global inventories, there is little evidence of 'settling' into a new circulation state: in most basins, redistributed heat and salinity inventories continue to grow during our simulations, and AMOC strength declines continually throughout the 21st century. A notable exception is the South Pacific, for which the redistributed heat inventory increases to approximately -50 ZJ by 2050, before remaining at a similar value for the rest of the simulation.

One way of assessing the interaction of excess and redistributed heat is to plot changes in their accumulation against each other, with emergent relationships consistent with coupling between the two (Figure 3.13).

In the North Atlantic, there is an acceleration of the accumulation of redistributed heat with respect to the excess heat inventory (Figure 3.13g). However, in all other basins for which relationships emerge clearly, the accumulation of excess and redistributed heat are either linearly related (Subpolar Southern (Figure 3.13a), North Pacific (Figure 3.13e)), or sublinear. This is as expected: the acceleration of the accumulation of redistributed heat is unique to the North Atlantic. In all basins other than the North Atlantic, the rate of accumulation of redistributed heat with respect to excess heat slows over the timeseries.

Despite this slowing, the redistributed heat inventories continue to grow, except in the Subantarctic Southern, South Atlantic, South Pacific and Arctic. In all other basins, there is continued accumulation of redistributed temperature, indicating the continual dynamic readjustment of the ocean, at an inter-basin scale: the lack of growth at a

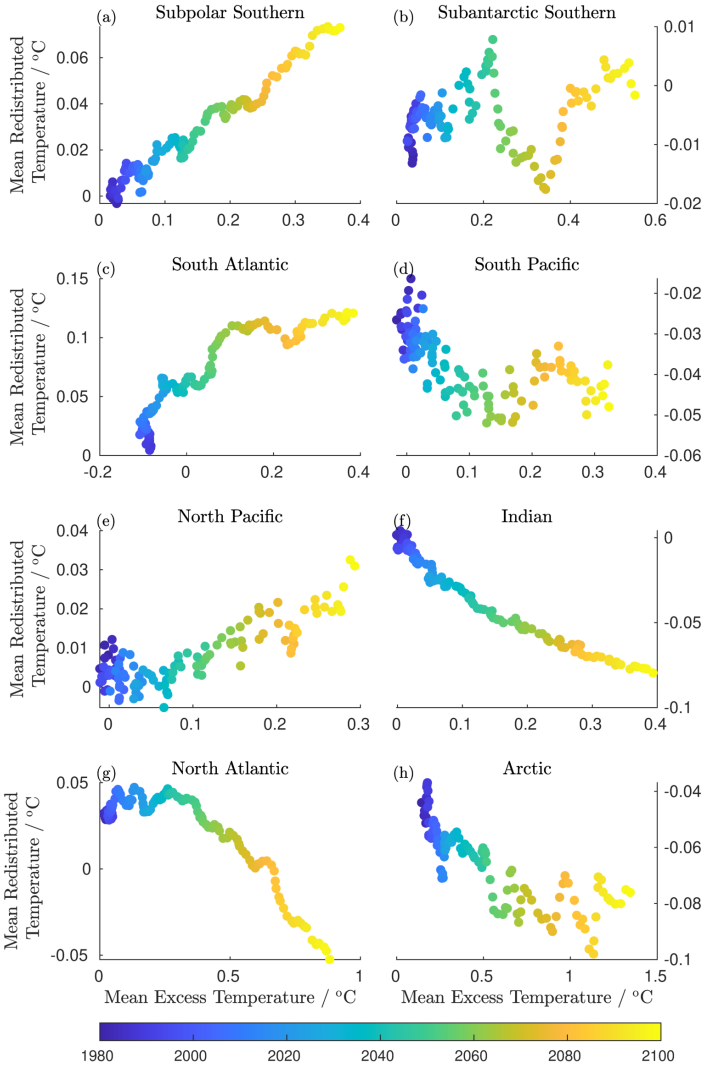

FIGURE 3.13: The emergent relationships (if any) observed between excess and redistributed heat in each of the 8 ocean basins shown in Figure 3.12, presented in terms of the mean redistributed and excess temperature changes for the basin. Timeseries begin in 1980 as there is no appreciable accumulation of excess or redistributed heat in the first half of the run. Scales differ for each basin.

basin scale imposes no constraints on intra-basin redistribution. Of these, the Subpolar Southern and North Atlantic are the most striking, with heat redistribution increasing linearly and with the square of excess heat accumulation, respectively.

Previous studies have found AMOC strength to be proportional to SST anomalies in the North Atlantic (Caesar et al., 2018), and SST anomalies are thought to be proportional to excess heat (MacDougall and Friedlingstein, 2015). Though it would initially appear that this would act to linearly couple the excess heat content of the North Atlantic to the redistribution of heat out of the North Atlantic, the redistributed heat inventory will be proportional to the time integrated changes in AMOC strength. The excess heat inventory of the North Atlantic increases monotonically with time, and so the rate of change of the redistributed heat inventory will be proportional to the excess

heat inventory. The proportionality of the redistributed heat inventory of the North Atlantic to the excess heat inventory can therefore be explained in terms of the unique circulation of the North Atlantic.

3.3.3 Mapping storage of excess and redistributed temperature and salinity

The regional patterns of decadal mean excess and redistributed temperature for the 2090s at the surface and at 2000m is shown in Figure 3.14 and the regional patterns of the 2090s decadal mean excess and redistributed salinity in Figure 3.15. For both temperature and salinity, surface changes are dominated in most locations by the excess component. Excess temperatures are positive nearly everywhere, whilst excess salinity is generally positive in the South Atlantic, Subtropical North Atlantic and Indian Oceans, with the Pacific generally negative. This is consistent with increased evaporation over the Atlantic and increased atmospheric freshwater transport from the Atlantic to the Pacific.

It is generally expected that in a warming climate, the hydrological cycle will become amplified, with increased evaporation (precipitation) in regions of net evaporation (precipitation) (Durack and Wijffels (2010), Zika et al. (2018), Gould and Cunningham (2021)). Thus, salty regions of the ocean surface become saltier, and fresh regions fresher. As these changes result from changing surface fluxes, hydrological amplification should be captured by the excess salinity at the surface, rather than redistributed salinity: this is consistent with the results presented here.

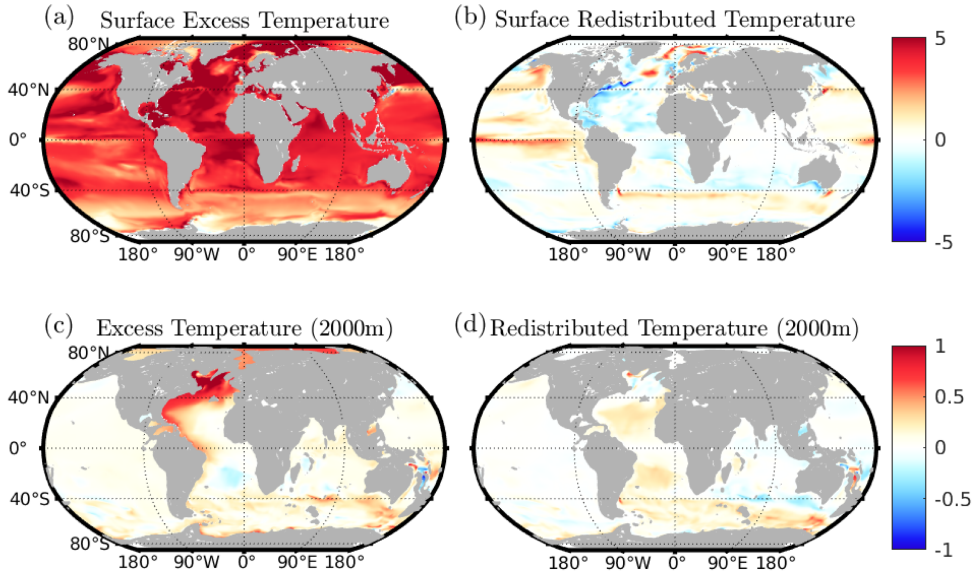

FIGURE 3.14: Maps of excess and redistributed temperature on two depth surfaces: the surface and at 2000m. Values given are the decadal mean for the decade 2090-2099. Colour axes are shared between each component at both depths.

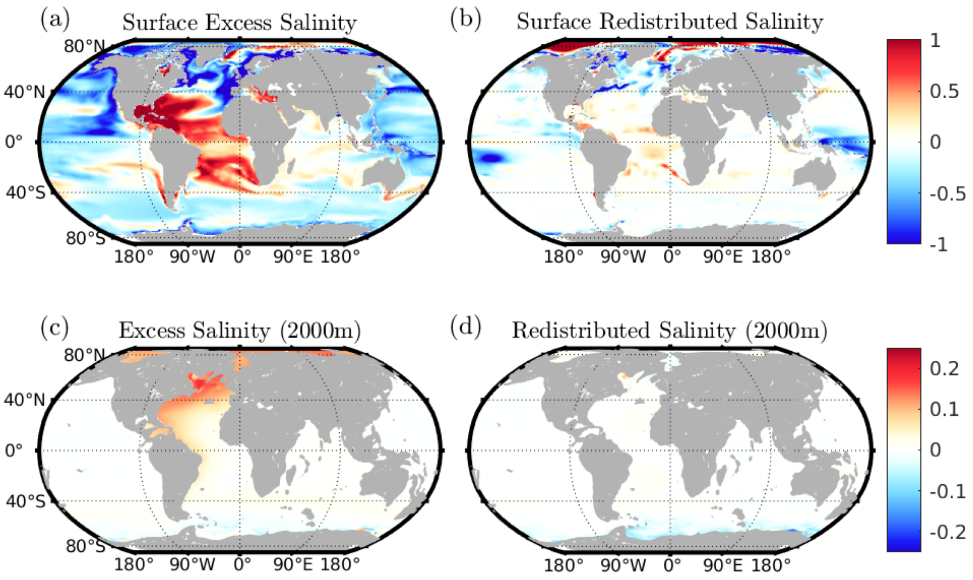

FIGURE 3.15: Maps of excess and redistributed salinity on two depth surfaces: the surface and at 2000m. Values given are the decadal mean for the decade 2090-2099. Colour axes are shared between each component at both depths.

Whilst surface warming is unsurprisingly dominated by excess temperature, at 2000m the contributions of excess and redistributed temperature to total temperature change are of comparable magnitude, with the exception of the North Atlantic. In contrast, the majority of salinity change at depth is accounted for by the excess component, though appreciable changes are generally only found in the North Atlantic. This salinity increase at depth is despite surface freshening in the Subpolar North Atlantic (Figure 3.15a, 3.15c), resulting from the propagation of surface salinification here in the 20th century. Patterns of excess and redistributed surface salinity are consistent with the results of Sathyanarayanan et al. (2021) and Levang and Schmitt (2015).

The strong surface redistributed salinity signal in the Arctic appears to result from reduced sea ice freshwater transport from the marginal seas of the Arctic inwards. Previous studies using the NEMO GCM coupled to the LIM2 sea ice model have found that Arctic sea ice tends to grow along the coastal shelves of the Arctic Ocean, before being transported by the Beaufort Gyre circulation and transpolar drift (Moreau et al., 2016). The net result of this is to transport both freshwater and DIC from the coastal shelves to the centre of the Arctic Ocean: changes in this transport will therefore act to cause large and tightly correlated changes in DIC and salinity in the surface Arctic Ocean. Our decomposition therefore partitions salinity change resulting from changes in this transport to redistribution. Similar changes in sea ice transport also act to cause redistributed freshening in the coastal Southern Ocean.

The total inventory change in heat, salt, and density by the last decade of the simulations, as well as the storage of the excess and redistributed components are shown in

Figure 3.16, for the upper 2000m of the ocean. We present these as contributions to steric sea level change, allowing for both normalisation and a comparison of contributions to steric sea level rise. On this timescale, excess (Figure 3.16c) and total (Figure 3.16a) heat inventory changes are positive nearly everywhere, with the exception of the Weddell and Ross Gyres. Redistributed heat inventories are negative generally in the North Atlantic (Figure 3.16c), with the largest values seen in the Labrador and Norwegian Seas, as well as the Subtropical Gyre. In the Pacific and Indian Oceans, redistributed heat inventories are most negative at around 30-35°S.

Salinity inventory changes show a different geographical distribution: excess salinity increases uniformly only in the Atlantic and Arctic oceans (Figure 3.16e). Total salinity change is again dominated by the excess here. As with heat, the fingerprint of AMOC slowdown can be seen in the redistributed salinity signal: there is redistribution driven cooling and freshening in the North Atlantic and redistribution driven warming and salinification in the Equatorial and South Atlantic, resulting from a weakening in the northward transport of heat and southward transport of fresh water. This redistribution driven cooling and freshening acts to oppose the warming and salinification associated with increased surface heating and concurrent increases in evaporation - precipitation (E-P).

Density inventory changes (Figure 3.16a) are relatively globally uniform compared to the individual contributions: a decrease is seen in the total change and excess inventory nearly everywhere, with the exception of the Weddell and Ross Seas, as well as the central Arctic Ocean. The Arctic Ocean decrease is dominated by the changes in freshwater transport, whereas the Weddell and Ross Sea decrease result from upwelling cool water. In the Atlantic, large changes in steric sea level resulting from excess temperature are significantly reduced by the accumulation of excess salinity, and a similar cancellation is seen in the redistributed components.

3.4 Discussion and Conclusions

In this Chapter, a new technique for estimating the redistribution of heat and salinity by the ocean in response to anthropogenic climate change has been presented, allowing for the identification of the excess signal and produce estimates consistent with other reconstructions. This method can be thought of as sitting within a family of techniques which aim to understand ocean circulation changes through the relationship between ocean temperature and DIC, along with the methods of Bronselaer and Zanna (2020) and Williams et al. (2021). It produces results which are consistent with the assumptions of both methods, without constraints to enforce this. Instead, it assumes that on decadal and subdecadal timescales, local ocean heat and carbon content are dominated

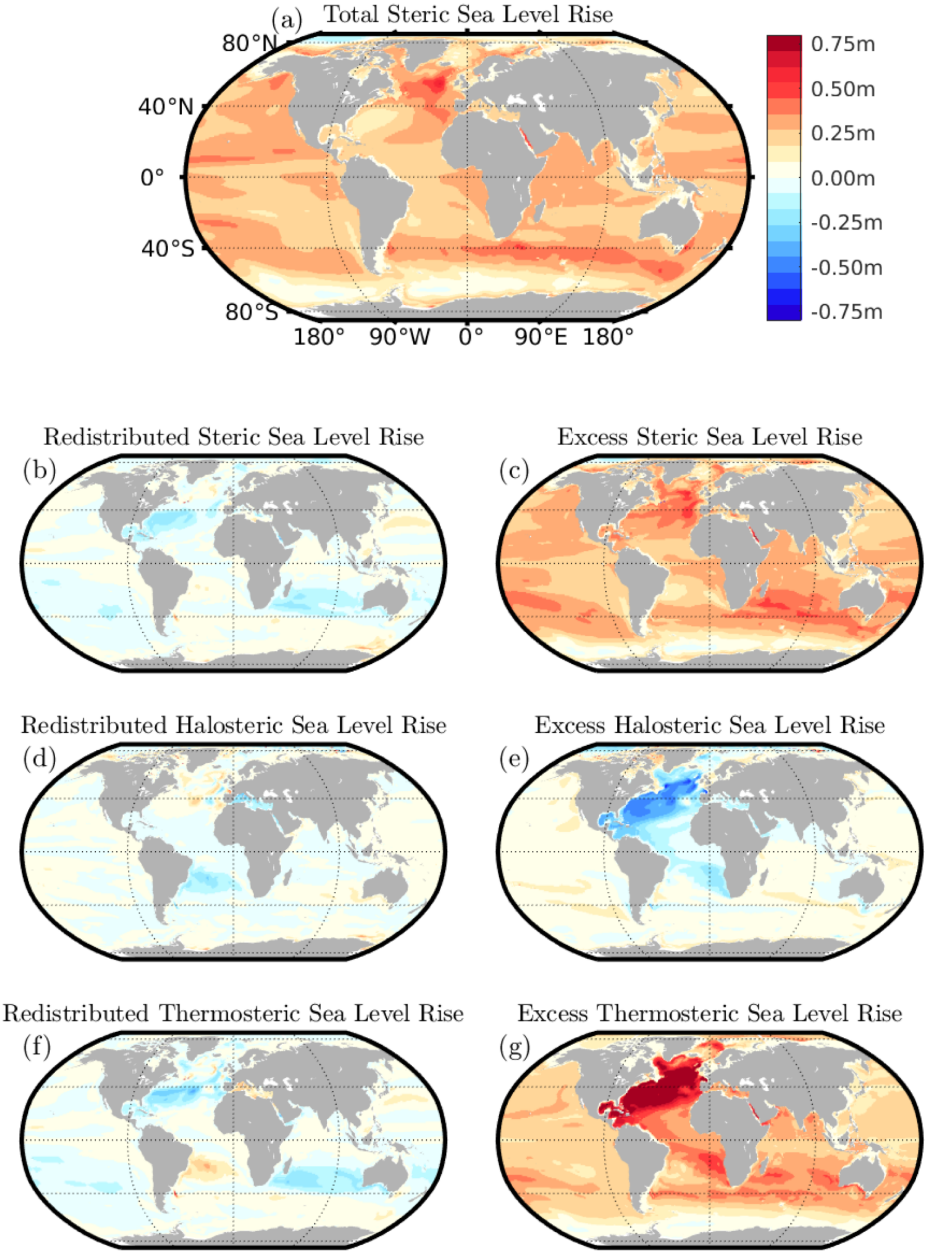

FIGURE 3.16: 2090's mean steric, halosteric and thermosteric contributions to sea level rise, as well as the total, from the upper 2000m of the ocean.

by redistribution, and that on longer (multidecadal to centennial) timescales, circulation variability dominates over biological changes in natural carbon. This first assumption is consistent with the results of Thomas et al. (2018), who investigated the relationship between ocean heat and carbon content, finding the two to be anticorrelated on decadal timescales. The results of Williams et al. (2021) suggest the assumption of circulation variability dominating over biologically driven changes is also reasonable. A key strength of this new technique is that it also allows us to estimate not only the redistribution of heat, but also salinity, and there are no obvious theoretical reasons

why it may not be extended to other tracers whose distributions evolve in response to anthropogenic climate change. Furthermore, its implementation is such that in order to identify circulation driven changes in a given tracer requires only a timeseries of the tracer in question, and of a secondary tracer which we may assume to change distribution only through redistribution, for example C_{nat}. It should therefore also be applicable to observational timeseries with little modification.

The globally integrated estimates it produces indicate that the magnitude of the excess and redistributed temperature signals are currently of a similar size, with the magnitude of excess temperature signals expected to exceed those of redistributed temperature signals towards the end of the 2020's. This is in keeping with previous studies which find excess heat beginning to dominate over redistributed heat in the period 2011-2060 (Bronselaer and Zanna, 2020). Of course, as this is only one climate change run from a single model, there is a large uncertainty associated with this and we recognise that it does not account for the spread of model responses to imposed climate change under an RCP8.5 scenario. However, these results are internally consistent, demonstrating a number of phenomena thought to occur under a changing climate explicitly in terms of the accumulation of excess heat and redistribution of preindustrial heat.

It has also allowed the production of the first modelled estimates of the redistribution of the preindustrial salinity field by the ocean and so the excess salinity field: that is, the changes in salinity due to changes in the balance of surface freshwater transport, directly excluding changes in ocean freshwater transport. By extension, it therefore also produces estimates of excess and redistributed density, and so the contributions to steric sea level rise of temperature and salinity changes. In these simulations, the penetration to depth of the redistributed salinity signal is far weaker than that of temperature, which, with the exception of the North Atlantic, accounts for a similar fraction of deep temperature change as the excess. However, there are several signals in surface excess and redistributed salinity changes consistent with hydrological amplification, as well as a salinity signal in the South Atlantic as a previously identified 'salinity pile up' in the South Atlantic consistent with AMOC slowdown (Zhu and Liu, 2020). By the 2090's, the Southern and Subtropical North Atlantic show increasing redistributed surface salinity as a result of AMOC slowdown, with a decreasing redistributed salinity in the Subpolar North Atlantic. At the surface, the majority of salinity change results from changes in E-P (excess), rather than circulation changes (redistributed), and that these patterns in excess salinity are consistent with both historical observations globally (Durack and Wijffels, 2010), and, in the Atlantic, with the salinity response to an idealised surface heat flux (Zika et al., 2018). The decrease in global mean excess salinity occurs earlier than the increase in globally integrated excess heat, consistent with previous studies which find significant sea ice loss even in the early 20^{th} century, before appreciable global warming (Wadhams and Munk (2004), Hetzinger et al. (2019)). These

results suggest that historical observations of temperature changes are dominated by redistribution, with excess temperature likely to dominate in the coming decades. Historical changes in salinity however may instead be predominately the result of excess salinity, rather than redistribution. This holds at both global and local scales, with the patterns of local excess salinity appearing to be dominated by amplification of the hydrological cycle, and is in agreement with the findings of Stott et al. (2008), Terray et al. (2012), Pierce et al. (2012) and Skliris et al. (2014), who suggest the salinification of the subtropical North Atlantic and freshening of the Western Pacific Warm Pool may constitute an early fingerprint of anthropogenic forcings.

In applying this technique to the Atlantic, it has been possible to show explicitly the redistribution of heat associated with changes to the overturning circulation, in addition to the aforementioned salinity signal. There are also fingerprints of AMOC change in both the redistributed temperature and salinity inventories of the North and South Atlantic: a large and rapid accumulation of negative redistributed heat in the North Atlantic over the period 2037-2099, as well as the accumulation of a large inventory of redistributed salinity in the South Atlantic over the same period. Over the period 2023-2099, for which the accumulation of excess heat in the North Atlantic is distinct from noise, $25 \pm 2\%$ of global excess heat accumulation is in the North Atlantic. This is remarkably similar to observational estimates of anthropogenic carbon uptake (Sabine et al., 2004) (albeit over different timescales), again indicating the close relationship between excess heat and anthropogenic carbon.

By the end of the 21^{st} century heat storage is dominated by excess heat. In addition, excess salinity storage is also largely spatially uniform, though the contributions of redistributed and excess salinity to halosteric SLR are of similar scales in most locations. The only exception to this is a large increase in excess salinity in the Atlantic, where excess salinity inventories are much larger than redistributed salinity inventories. The similar contributions of excess and redistributed salinity storage is despite patterns of regional change in sea surface salinity and salinity inventory changes being dominated by the excess component, both historically and by the end of the 21^{st} century.

By combining the estimates of excess temperature and salinity, we can directly compute the excess density change, and the redistribution of density. In the North Atlantic, there is warming and salinification in the excess components, and cooling and freshening in the redistributed components. In both cases, these changes are in a density compensating fashion. Previous studies have noted that whilst density compensated water mass changes may be a general property of the ocean, the behaviour is particularly marked in the Atlantic (Lowe and Gregory, 2006), as well as important for contemporary Atlantic deep ocean heat uptake (Mauritzen et al., 2012), though it is uncertain how this will evolve. These results suggest that in the Atlantic, even by the last decade of these simulations, changes in excess temperature and salinity act in a density compensating fashion. A consequence of this is that changes in surface freshwater fluxes associated

with climate change oppose the reduction of overturning circulation associated with increased surface warming, opposing the reduction in the North Atlantic's capacity to sequester excess heat. This suggests that the excess contribution to themosteric SLR in the Atlantic will continue to grow on centennial timescales, assuming continued CO_2 emissions, though the thermosteric SLR is greatly ameliorated by halosteric sea level fall. This is in agreement with historical observations (Antonov et al., 2002). However, the much smaller redistribution contribution to density indicates that changes to ocean circulation will have little effect on steric SLR in the North Atlantic by the end of the 21st century, although redistributed density compensation in the North Atlantic begins to break down in approximately 2050, as the redistribution of heat out of the North Atlantic significantly exceeds that of salinity by this time.

Finally, although only being applied within a single model, the patterns of excess and redistributed heat storage identified here are consistent with previous studies (Winton et al. (2013), Bronselaer and Zanna (2020), Williams et al. (2021)), despite differing assumptions used in the calculation of the redistribution of heat from carbon. A key benefit of the method introduced here compared to prior carbon based estimates of circulation change is that it is applicable both to models and observations, and across multiple tracers. In combination with other techniques, this method may therefore be a powerful tool for understanding causes of observed and projected ocean temperature, salinity and density change across multiple spatial and temporal scales.

Chapter 4

Robustness of the κ_r Decomposition: Investigating Time Evolution

The arguments of Section 2.3 strongly imply that the redistribution coefficient is time independent, both as a result of cancellation of time dependence in the full 3D formulation, and the simplification to a ratio of vertical derivatives also lacks time dependence. However, in this one dimensional simplification (Chapter 2, Subsection 2.3.1), higher order terms in the Taylor expansion about an initial state were ignored. It is now shown that consideration of these terms may in fact lead to time dependence of κ_r. Intuitively, the fact that κ_r is not time independent ought to be obvious: if the background temperature-carbon slope changes as a result of significant perturbations to ocean circulation, it is necessary to update the estimate of this slope in order to continue to 'walk along' the curve reliably. The time independence is therefore the result of approximations made in the definition of a redistribution coefficient. In this Chapter, the limitations of these approximations are explored, and how we might better improve our understanding by attempting to relax these assumptions. However, relaxing these assumptions would require a significant time investment and so has not been performed: instead, the limitations are simply categorised in order to suggest where future research might be directed.

Recall Equation 2.59:

$$\Delta\theta_r = \kappa_r \Delta C_{nat}. \tag{2.63}$$

This formulation assumes that, at any point in the ocean, changes in the background temperature and C_{nat} are related by a simple constant, κ_r. A corollary of this is that the preindustrial temperature and C_{nat} fields, if plotted against one another, would coalesce onto a straight line (this would remain true even if the spatial and temporal dependence of $\delta\theta_r$ and δC_{nat} were explicitly included). However, this cannot be the case: if it were, we could use a single, global value of κ_r in the results of the previous chapter.

That it is necessary to spatially resolve κ_r is therefore a direct result of the limitations of this assumption: we need to locally approximate the slope of the preindustrial temperature-carbon curve in order to accurately understand redistribution. This is one of the limitations of the 'correlated excess, anticorrelated redistribution' argument used by Williams et al. (2021): as noted in the previous Chapter, their arguments are not applicable in the North Pacific, where the spatial anticorrelation is reversed, and we in fact see a positive spatial correlation between the background temperature and carbon fields (Figure 3.4). It is therefore necessary to understand how to account for these spatial variations in κ_r in order to improve this technique.

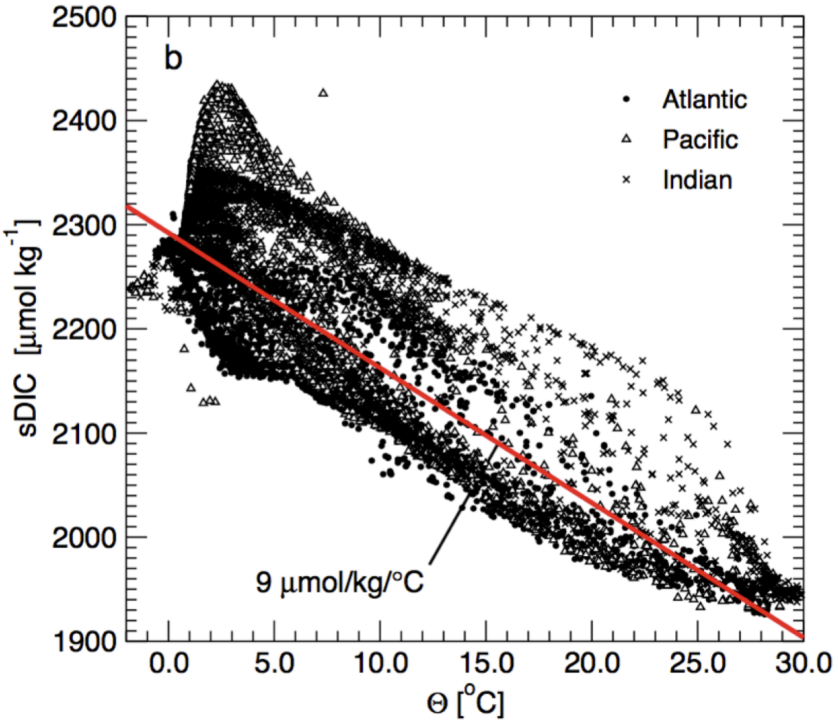

FIGURE 4.1: sDIC (DIC normalised to a salinity of 35) concentrations against potential temperature, for all observations from 3 latitudinal hydrographic sections: A16, P16, and I9N-I8S. The red line represents the expected trend if sDIC changed only as a result of CO_2 solubility due to temperature. Reproduced from Sarmiento and Gruber (2006).

Consider Figures 4.1 and 4.2 (note the x and y axes are reversed between the two figures). Whilst the straight line assumption is clearly imperfect for the more global distribution shown in Figure 4.1, it could be considered a reasonably accurate assumption. However, in Figure 4.2, there is a clear kink in the curve at around 8 degrees. Therefore, if circulation changes sufficiently to replace water at a given point which was previously 10 degrees with water which is 5 degrees, the estimation of the redistributed temperature change using the κ_r approach and the actual redistributed temperature change will be in the opposite directions: that is, what should be cooling driven by circulation change will instead be estimated to be warming. Solving this problem requires

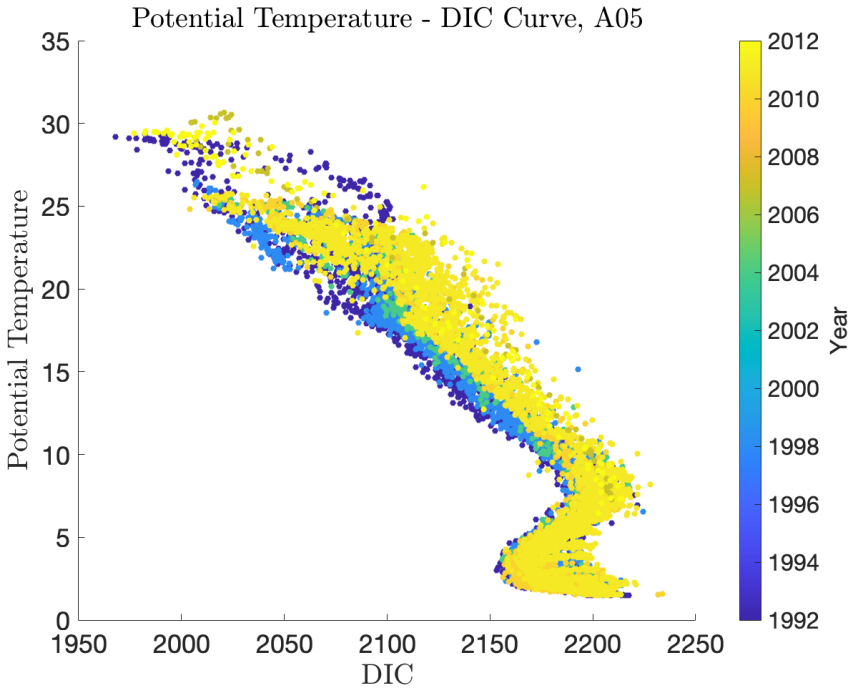

FIGURE 4.2: The distribution of potential temperature and DIC for all bottle data from the 1992, 1998, 2004, 2010 and 2011 hydrographic occupations of A05.

extending the κ_r framework.

The approximation of Equation 2.59 may be written more properly as follows:

$$\Delta\theta_r(\vec{r},t) = \kappa_r(\vec{r})\Delta C_{\text{nat}}(\vec{r},t), \tag{4.1}$$

where the space and time dependence have now been explicitly included: κ_r varies with position, but does not vary with time. All time dependence in redistributed temperature is contained in the time evolution of the natural carbon field. Noting now that κ_r does in fact evolve in time as circulation changes, we may instead write the κ_r value at a given point as a function of time:

$$\kappa_r(\vec{r},t) = \left.\frac{\vec{v'}\cdot\nabla\theta}{\vec{v'}\cdot\nabla C_{\text{nat}}}\right|_t. \tag{4.2}$$

This means that we may write our previous approximation for κ_r as

$$\kappa_r(\vec{r},t) = \left.\frac{\vec{v'}\cdot\nabla\theta}{\vec{v'}\cdot\nabla C_{\text{nat}}}\right|_{t=t_{PI}}, \tag{4.3}$$

where t_{PI} is a preindustrial time. Thus, continuing in the framework where changes are calculated at a fixed point at space, we could generalise Equation 2.59 to account for the evolution of the temperature-carbon curve slope at a fixed spatial location as

follows:

$$\Delta\theta_r(\vec{r},t) = \int_{t_{PI}}^{t} \kappa_r(\vec{r},\tau)\frac{dC_{nat}}{d\tau}d\tau. \tag{4.4}$$

This formulation may be interpreted as follows: instead of prescribing a preindustrial κ_r distribution and using this to estimate how circulation changes redistribute temperature from changes in C_{nat}, we instead compute the spatial correlation between temperature and C_{nat} at every time step. Integrating this relationship forward in time therefore ought to yield a more accurate redistributed temperature, as it allows the slope of the temperature-carbon curve to change over time. This is a straightforward extension of the previous technique, and addresses the inaccuracy in the Taylor Expansion implicitly performed in the κ_r decomposition by repeatedly Taylor expanding at each time step. This extension is demonstrated visually in Figure 4.3.

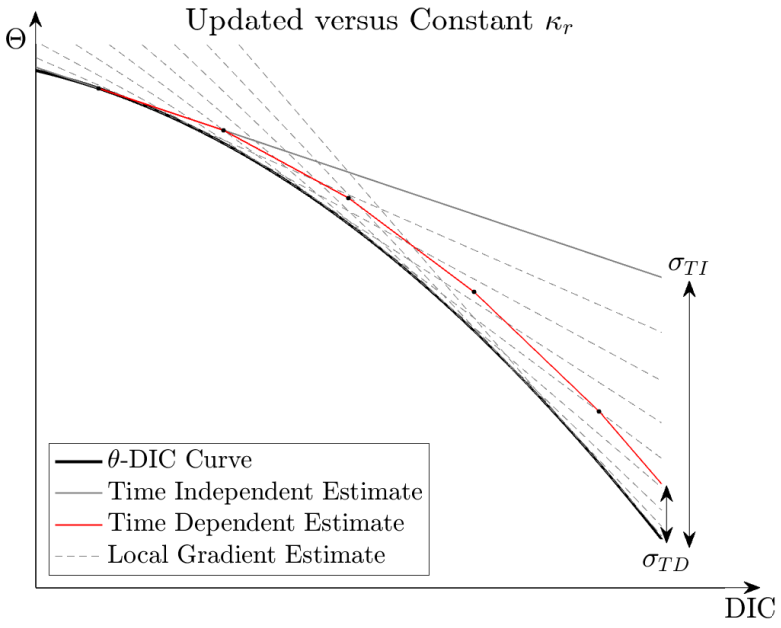

FIGURE 4.3: The effect of repeatedly recalculating the slope to the temperature-carbon curve in order to estimate redistribution, versus using a time independent estimate: total error is significantly reduced. σ_{TD} represents the error in the time dependent κ_r estimate, and σ_{TI} the error in the time independent κ_r estimate.

Now, alternative methods with which this decomposition's range of validity might be expanded are considered. After that, the range of validity, as applied in the previous chapter, is explored. However, the application of these methods is not explored: they are simply noted as likely avenues which ought to build upon the principles discussed here.

4.1 Machine Learning

Artificial Neural Networks (ANN's) are a relatively simple and straightforward machine learning technique that can used to predict a variable (sometimes called the 'label') from a number of other variables (sometimes called the 'features'). Common uses include, for example, predicting house prices from post codes, bedroom number and bathroom number. Typically, developing an ANN follows a procedure similar to the one depicted below:

1. A dataset is procured. For oceanographic purposes, this could be, for example, all bottle data from a single hydrographic occupation.
2. The dataset is split into a training dataset (typically 80% of the available data), and a validation dataset (the remaining 20%).
3. The model is trained using the training dataset: a cost function is minimised, enabling the ANN to learn an arbitrary nonlinear function which predicts the variable of interest from the available data.
4. The model is then validated on the validation dataset, to ensure that the model is not 'overfitting' the training data.
5. Once validated, the model can then be applied to new data, predicting the 'label' from 'features'.

An example of such a workflow can be found at the Tensorflow documentation: `https://www.tensorflow.org/tutorials/keras/regression`.

It is relatively straightforward to see how such an approach could be applied to the data such as those in Figure 4.2. In the current implementation of the κ_r decomposition, the approach is to estimate the local gradient of the temperature-carbon curve, and Taylor expand this gradient to estimate the change in temperature due to redistribution, effectively 'walking along' the curve to a new part of the original parameter space. An ANN approach could instead be applied to perform a similar parametrisation. For example, bottle data from a hydrographic occupation of a section could be taken, and the C_{nat}, longitude, latitude and depth used to predict temperature, for example. Other tracers could also be included, if the excess component of these tracers can be neglected. If these predictions are accurate, then this model could be applied to later occupations of the same section to predict the temperature from these variables, or indeed to nearby sections. As the model is trained on data which are assumed to be simultaneous, it therefore ought to learn how the spatial distributions covary. The temperature this model will then predict, when applied to a later occupation of the same section, will therefore represent the redistributed temperature, as the prediction

is trained on data which vary only spatially. The difference between the predicted temperature and the total change in temperature therefore represents the excess temperature. An initial exploration of this approach to predicting excess and redistributed temperature can be found at https://github.com/charles-turner-1/BioGeoChem_TempRedistribution.

Whilst the implementation of this approach differs significantly from the implementation used in Chapter 3, it is conceptually very similar: the approach used in Chapter 3 attempts to 'walk along' the $\theta - C_{nat}$ curve via a Taylor expansion, whilst the ANN approach would instead attempt to fit the $\theta - C_{nat}$ curve using an arbitrary nonlinear function, allowing a location on this curve (ie. a temperature change due to redistribution) to be predicted from other variables.

4.2 Water Mass Space

This approach is somewhat more complicated than those illustrated in the previous section, and can be considered to be an extension to the principles outlined in Chapter 2, Section 2.4, in which the full temperature-carbon space formalism is developed. In this full temperature-carbon space framework, it can be thought of as analagous to the approach of Zika et al. (2021), where total transformation between water masses in temperature-salinity space is minimised to produce a material heating and material salinification. However, as this approach has already been applied in temperature-salinity space, for which observations are more plentiful than those for carbon, repeating a similar analysis in temperature-carbon space would likely improve understanding of the carbon-heat coupling, but not the causes of oceanic temperature and salinity change - thus, due to the complexity of this approach, it is not applied here.

Alternatively, it can be thought of more simply as a water mass lookup approach. I will first describe how and why this water mass lookup approach might have utility for extending the range of validity of the κ_r based decomposition, before illustrating how it might be applied.

3 hydrographic sections in the South Atlantic are included in GO-SHIP (discussed in the following Chapter): 24S (A09.5), 30S (A10) and 34S (A10.5) (though 34S will not be further considered here). Occupations of 24S were performed in 2009 and 2018, and 30S in 1993, 2003 and 2011. As these two hydrographic sections are separated by only 6 degrees and contain essentially the same source waters, it is therefore instructive to ask the following question: can we grid observations from all 5 occupations onto a nominal reference section and compute estimates of redistribution on this nominal grid? It is clear from the maps shown in Figure 4.4 that the distributions of DIC at A09.5 (top right panel) and A10 (bottom right) are substantially different, and likewise for temperature (though this is less visible due to the colour scale). Gridding the observations from

2009 onto the same grid as those of 1993 and applying the matrix inversion calculation described in Chapter 2, Section 2.4, we obtain the excess temperature field shown in Figure 4.5. For comparison, excess temperature estimates obtained by computing the excess temperature over 1993-2003 at 30S is also shown in Figure 4.5. Though the 30S excess temperature field suffers from some noise (it has not been subjected to secondary data control which will be described in Chapter 5), it is not a completely unrealistic field. However, the excess temperature field shown on the right hand side of Figure 4.5 is.

From the previous discussion, it is clear why this excess temperature field is so unrealistic: the assumptions made about the slope of the temperature-carbon curve and changes in DIC concentration are completely violated as a result of the vastly different DIC distributions: this causes nonsensical excess temperature field. However, as previously noted, the distributions of water mass properties are similar in temperature-carbon space - see Figure 4.6. Though κ_r distributions in temperature-carbon space do differ, patterns are broadly similar. Identification of a clear tracer space for which a map of κ_r values obtained at one location may be projected onto another location therefore would also represent the identification of a more robust framework for the application of the redistribution coefficient decomposition.

Finally, it might be noted that in Figure 4.3, the time dependence shown could instead be considered to be either temperature or DIC dependence without any modification to the approach: thus, the water mass approach, and the time evolving κ_r approach which will be discussed in the following section can be considered to be equivalent.

4.3 κ_r time evolution in the NEMO OGCM

Two potential techniques which may be useful for extending the redistribution coefficient framework have now been described. Here, the simpler time evolving κ_r approach is explored, with the aim of exploring the question of how sensitive κ_r is to the period during which we estimate it. This can be broadly broken into two questions.

Firstly, how well defined is a preindustrial κ_r distribution? The formalism of Chapter 2, Sections 1.2.1, 2.3, as well as Chapter 3 all assume that the preindustrial ocean is a steady state system, with well defined covariability between temperature and carbon content. This has not been explicitly tested. Understanding the range of validity of the redistribution coefficient approach therefore requires an understanding of the extent to which this assumption is accurate.

Secondly, it is not possible to directly observe the preindustrial ocean. Therefore, it is necessary to answer a related but separate question: does estimating a redistribution

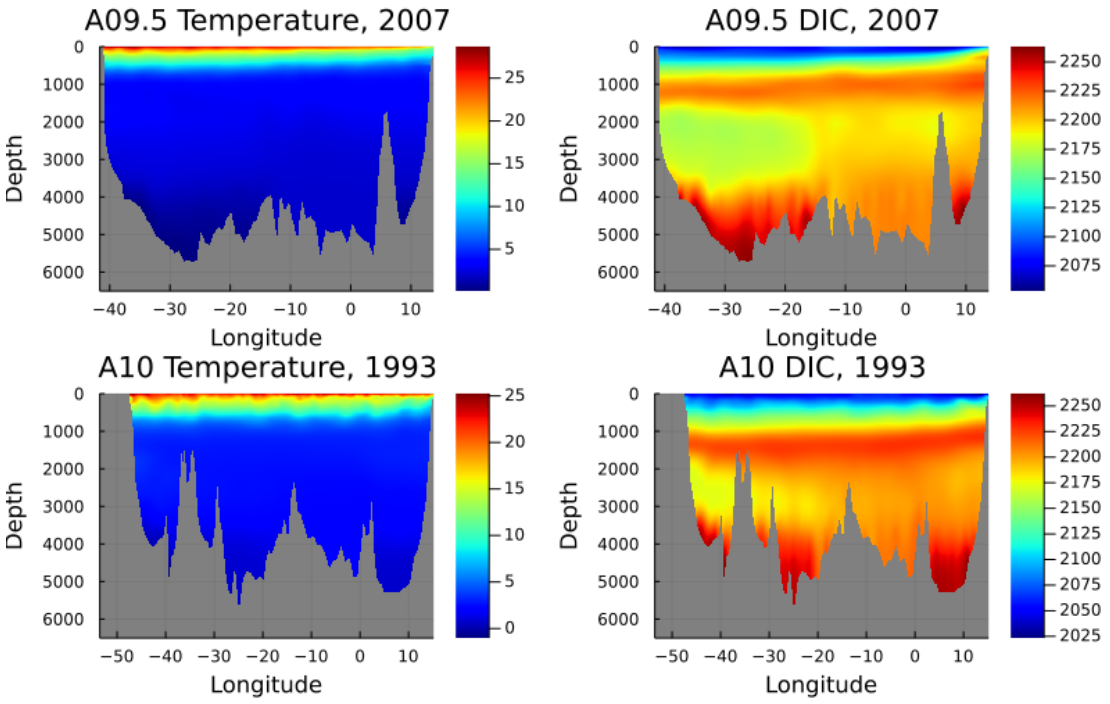

FIGURE 4.4: Maps of the initial temperature and DIC field at A09.5, and A10. Clearly, these show quite different distributions, even though they may both be gridded onto the same nominal section.

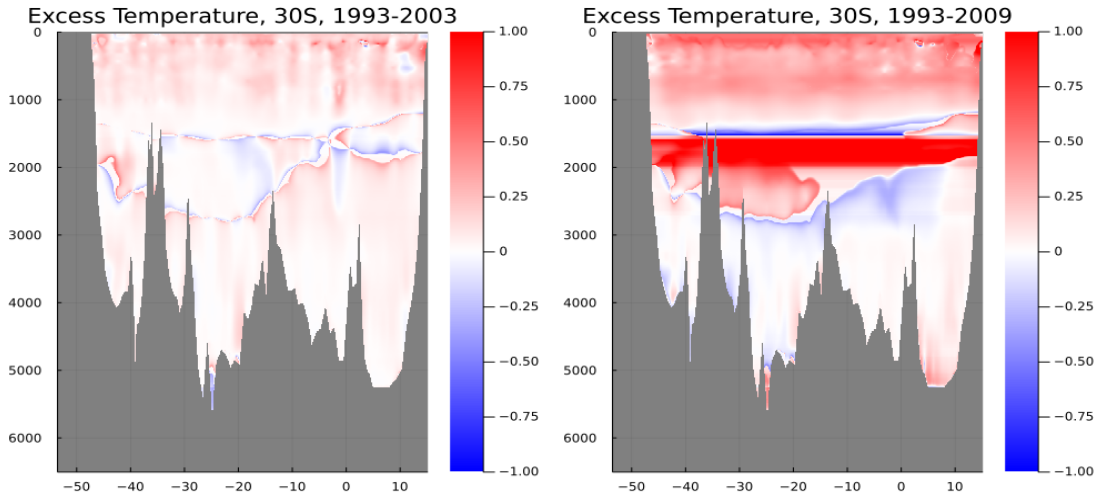

FIGURE 4.5: The excess temperature calculated using Equation 2.77, as applied to the differences between the 1993 A10 cruise and subsequent 2003 A10 cruise, and the 1993 A10 cruise and 2009 A09.5 cruise, shown in Figure 4.4. Trying to compute an excess temperature by gridding both sets of observations onto the same section yields results which are clearly nonsensical at depths of approximately 1500-2000m.

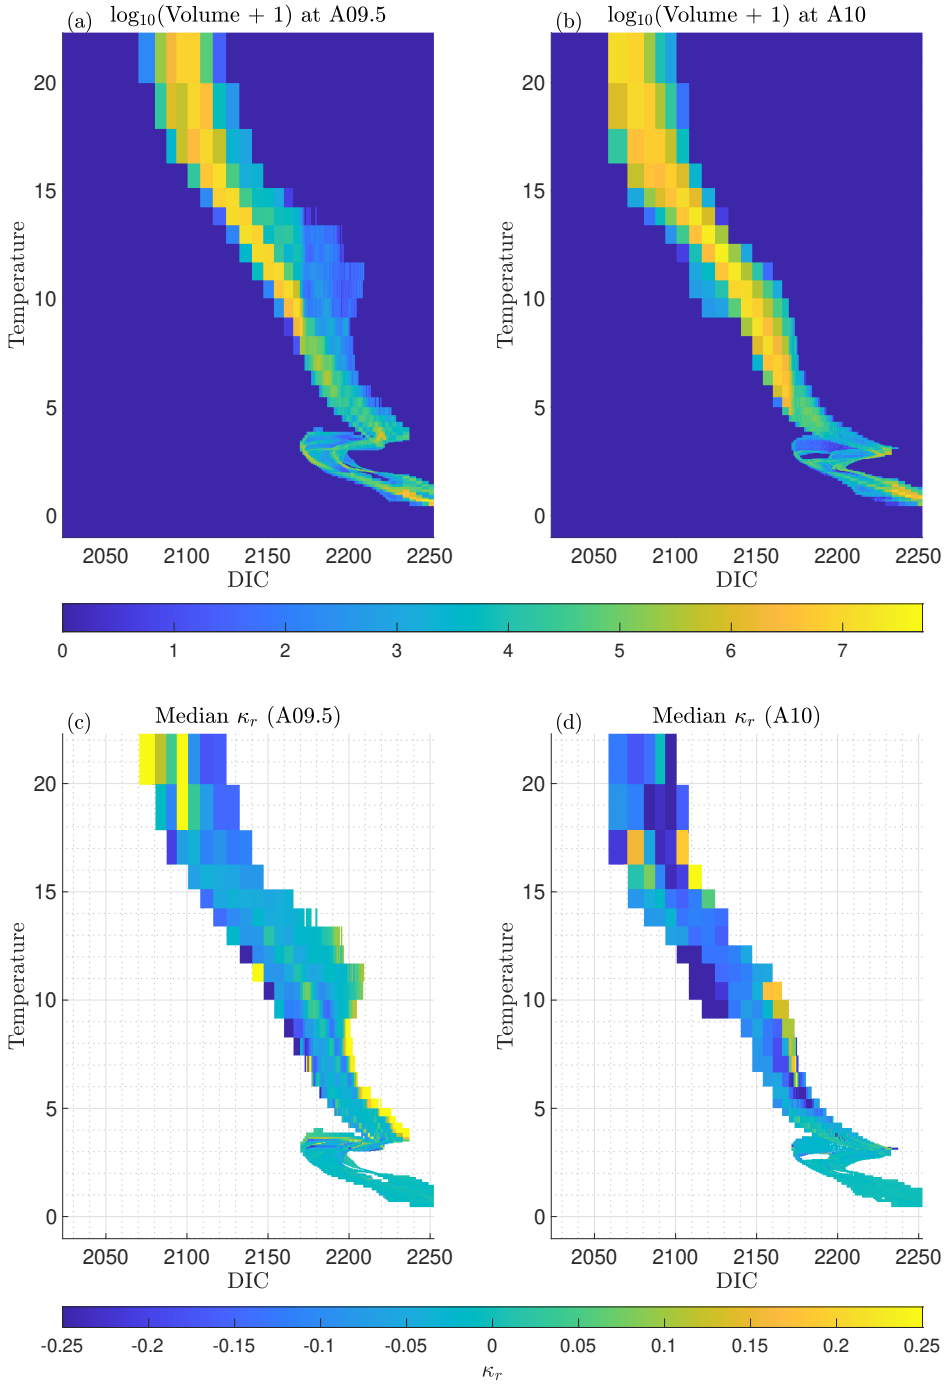

FIGURE 4.6: The volume of water found in each percentile temperature and carbon bin for A09.5 (a) and A10(b), and the median κ_r value in each of these bins ((c) and (d)). Volumes are calculated as the number of grid cells presented in Figures 4.4 and 4.5 (all grid cells are of equal area).

coefficient using contemporary distributions of temperature and carbon have a significant affect on the decomposition? In other words, when applying this technique to observations, will results be fundamentally biased due to the lack of 'true preindustrial' observations?

These questions can both be answered by recalculating κ_r fields for different time periods. To do so, the same approach detailed in Chapter 3, Section 3.2 was used, except a number of different periods were considered. As the preindustrial run uses a repeated 30 year forcing, it is not straightforward to generate a large ensemble of preindustrial covariability estimates, or the estimates produced will overlap significantly.

As a compromise between generating a large number of estimates and generating distinct estimates, 3 separate preindustrial estimates of κ_r were made: one spanning 1860-1959, one spanning 1870-1969, and one spanning 1880-1979. Though the forcing is shared between these 3 estimates, the internal variability of the model will be different, and so this allows us to estimate how well defined the preindustrial spatial covariability of temperature and carbon is.

To answer the second question, a separate, 'post-industrial' estimate of κ_r was generated, in order to compare this field with the κ_r field used in Chapter 3 and shown in Figures 3.4, 3.5. For the post industrial field, the same approach detailed in Chapter 3 was also used, except instead of comparing the high frequency variability in the DIC and temperature over the full control run, the high frequency variability over 2040-2099 in the RAD run was used.

4.3.1 Preindustrial κ_r Variability

Figure 4.7 shows maps of the zonal mean preindustrial κ_r value, calculated as the mean of the three estimates described above. Unsurprisingly, distributions are near identical to those shown in Figures 3.4 and 3.5. The standard deviations of the three estimates (thus representing a time deviation in the mean state) are extremely small relative to the mean κ_r values, indicating that the notion of a 'preindustrial state' is well defined (note that the scale of the standard deviation panels are a quarter that of the means). However, some variability is seen, particularly in the high latitude North Atlantic (panel (b)) and the Pacific sector of the Southern Ocean (panel (d)). This indicates there may be some nontrivial variability in the preindustrial covariability of temperature and carbon in these locations. However, these are also locations where large mean values are seen, and so they do not represent a significantly larger fractional error than in other ocean regions: in ocean regions where the κ_r value is not approximately zero, the ratio of the standard deviation in κ_r to the mean value is typically 15%. This indicates that the uncertainty associated with the preindustrial state estimate is on the order of 15%: within the 5-20% range of uncertainty estimated in Chapter 3.

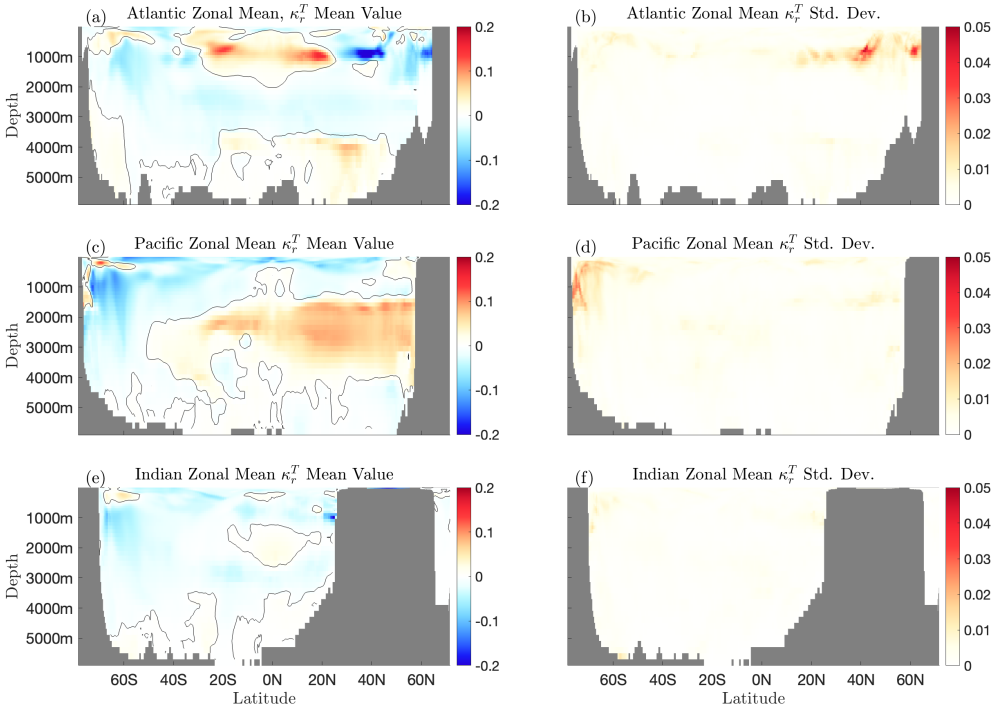

FIGURE 4.7: The mean of the three preindustrial zonal mean κ_r estimates for the Atlantic (a), Pacific (c), and the Indian (e) oceans, and the standard deviation of each (b), (d) and (f), respectively. Note that the colour axes for standard deviations are a quarter of that for mean values. All values are given in units of K/(μmol/kg).

4.3.2 κ_r Time Evolution

Figures 4.8, 4.10 and 4.9 show preindustrial and 'postindustrial' distributions of κ_r, for the Atlantic, Pacific, and Indian ocean, respectively. The bottom panel of each figure shows the difference between the preindustrial and post industrial estimates. Though patterns do differ in their details, in particular in the Atlantic, in general the patterns seen are broadly similar.

4.4 Discussion

In this short chapter, a number of the limitations of the redistribution coefficient decomposition have been explored, and several potential avenues to overcome these limitations explored. However, these extensions to the redistribution coefficient approach would require significant work: thus, only the potential to extend the range of validity of this decomposition has been noted. Subsequent chapters will instead focus on the

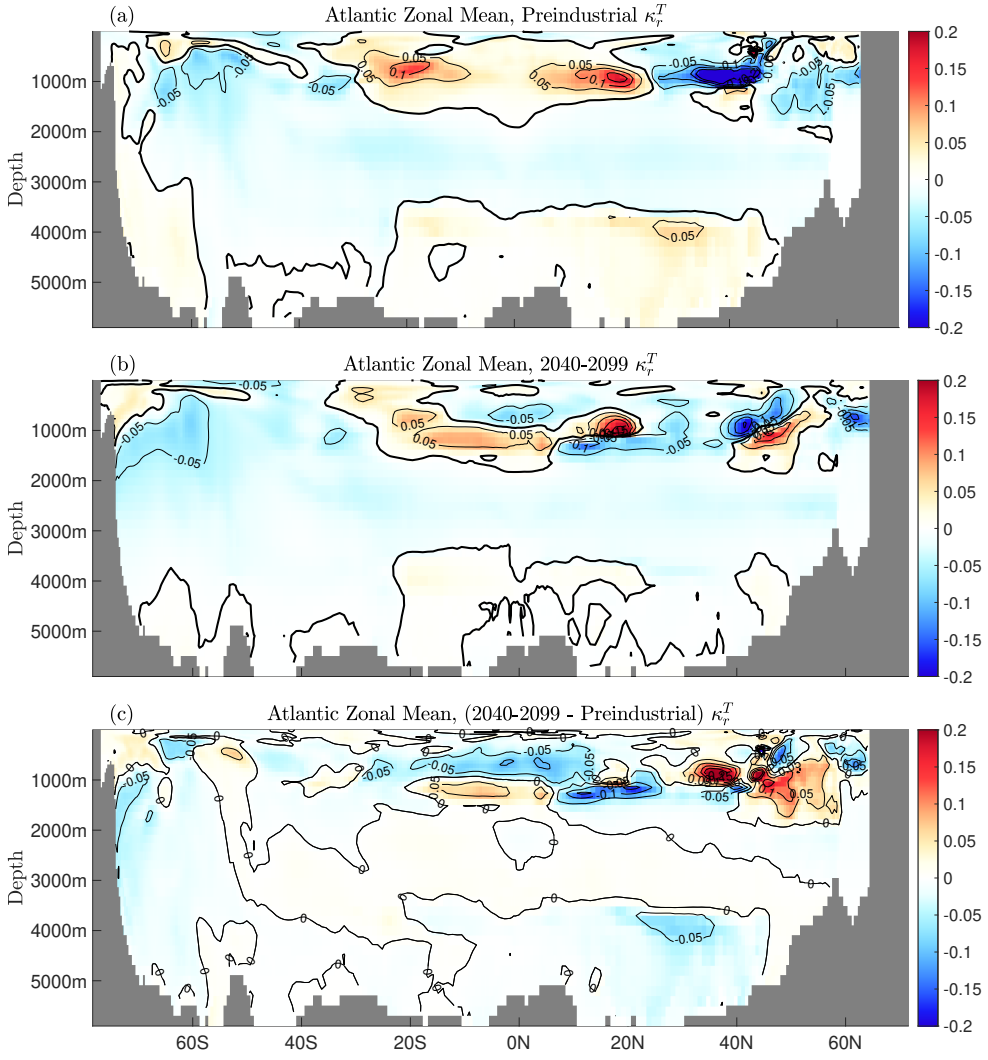

FIGURE 4.8: The preindustrial zonal mean κ_r value for the Atlantic (a), 'postindustrial' (b), and the difference between the two (c). All values are given in units of K/(μmol/kg).

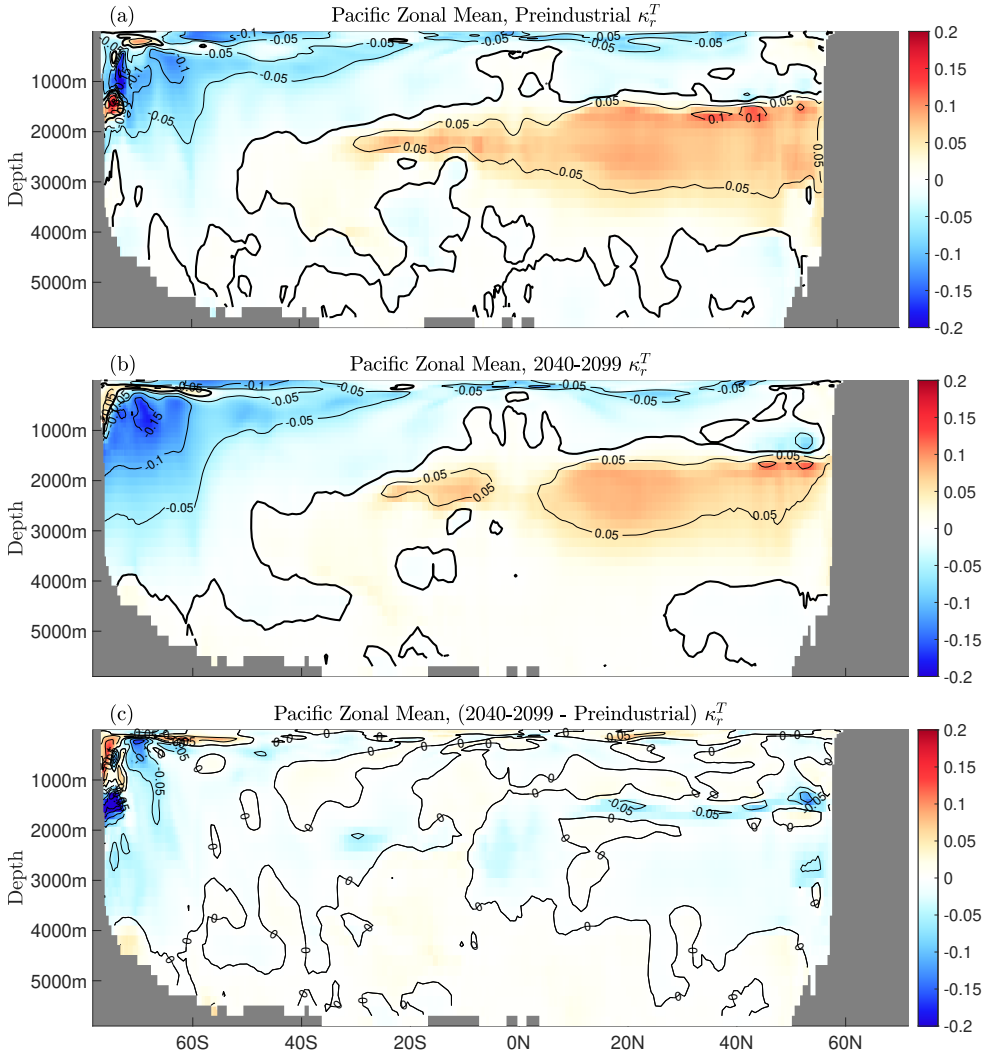

FIGURE 4.9: The preindustrial zonal mean κ_r value for the Pacific (a), 'postindustrial' (b), and the difference between the two (c). All values are given in units of K/(μmol/kg).

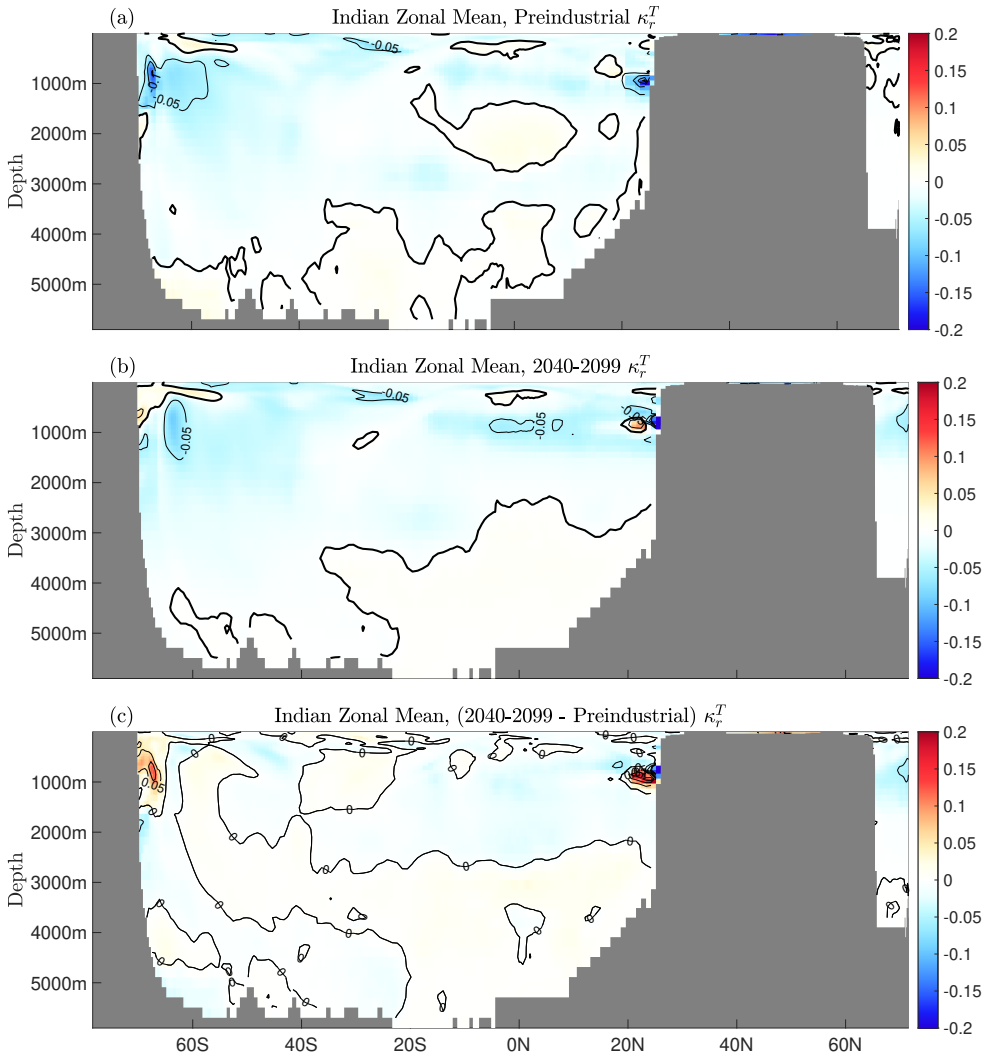

FIGURE 4.10: The preindustrial zonal mean κ_r value for the Indian ocean (a), 'postindustrial' (b), and the difference between the two (c). All values are given in units of K/(μmol/kg).

application of the matrix inversion decomposition to hydrographic data, which is simpler to implement and requires more readily available and less uncertain data: that is, DIC, rather than C_{anth} and C_{nat}.

Chapter 5

Application of carbon based temperature and salinity decomposition to repeat hydrography in the Subtropical North Atlantic

5.1 Introduction

The A05 hydrographic section is a section of ocean in the Subtropical North Atlantic, spanning Florida to the Canary Islands at a nominal latitude of 24.5 North: its location is illustrated in Figure 5.1. In addition, the bathymetry of the A05 hydrographic section, as well as the approximate locations of the main water masses found within it are shown in Figure 5.2, reproduced from Guallart et al. (2015). It was designated A05 during the World Ocean Circulation Experiment (WOCE) in the 1990's (Koltermann et al., 2011), and crosses the North Atlantic through the middle of the subtropical gyre. As a result of this location, it is an excellent laboratory to study the ocean's sequestration of heat and carbon, due to the North Atlantic's disproportionate influence on both ocean heat and carbon uptake: for example, comprising 25% of ocean C_{anth} sequestration despite only representing 13% of its surface area (Sabine et al., 2004). It is also important for the transport of heat and freshwater by the North Atlantic, being approximately the latitude at which the northward transport of heat by the Atlantic Meridional Overturning Circulation (AMOC) peaks (Srokosz and Bryden, 2015). The AMOC also transports freshwater southwards with 91% of the variance in freshwater flux into the Atlantic north of 26.5 North controlled by the strength of the AMOC at this latitude (McDonagh et al., 2015).

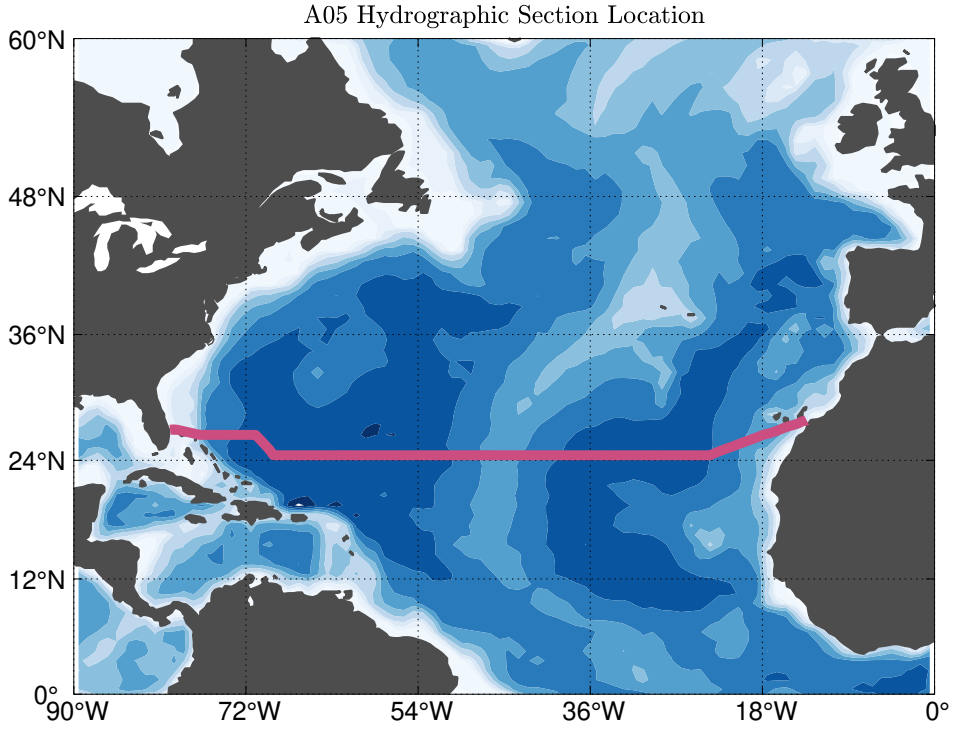

FIGURE 5.1: The location of the A05 hydrographic section (indicated by the pink line) within the North Atlantic. Bathymetry is indicated with shading.

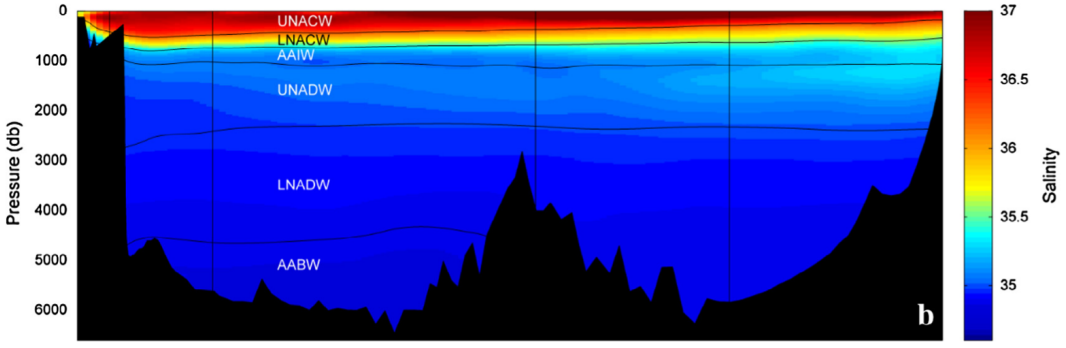

FIGURE 5.2: The bathymetry of the A05 hydrographic section. The main water masses found are also indicated, and the salinity is also shown. Abbreviations are as follows: UNACW (Upper North Atlantic Central Water), LNACW (Lower North Atlantic Central Water), AAIW (Antarctic Intermediate Water), UNADW (Upper North Atlantic Deep Water), LNADW (Lower North Atlantic Deep Water) and AABW (Antarctic Bottom Water). Vertical lines delineate regions used for data aggregation in Guallart et al. (2015). Reproduced from Guallart et al. (2015).

Due to its importance for global ocean heat and carbon storage, the A05 section has been occupied numerous times. After an initial occupation during the international geophysical year (1957), during which temperature and salinity measurements were taken, it has subsequently been occupied in 1992, 1998, 2004, 2010, 2011, 2015 and 2020, with high quality observations of temperature, salinity and DIC (in addition to various other observations). Additionally, an occupation was made in 1981, but no carbon data was collected then. Collection of data at A05 in 2020 (GLODAP expocode `74EQ20220209`) formed part of this work: however, due to COVID-19 related delays, carbon date were unable to be calibrated in time for inclusion in the analysis used here.

In this Chapter, the Temperature-Carbon space formalism, introduced in Chapter 2, Section 2.4, which allows for the estimation of excess and redistributed temperature, carbon and salinity is further developed. It is then applied to the repeat hydrographic occupations of A05. The choice of parametrisations: whether to use a single α_T value, or allow it to vary with temperature, and the sensitivity of results to this choice is explored. Finally, excess heat, salinity, and carbon accumulation is explored, the rates of accumulation of excess temperature and salinity are estimated.

5.2 Theory

The theory underlying the decomposition of temperature and carbon used here was introduced in Chapter 2, Section 2.4.

A brief recap is now provided, as well as a more detailed investigation of mathematical issues that must be considered when applying this decomposition. We may decompose temperature and carbon changes as follows:

$$\Delta\theta = \Delta\theta_e + \Delta\theta_r \tag{5.1}$$

$$\Delta\text{DIC} = \Delta\text{DIC}_e + \Delta\text{DIC}_r, \tag{5.2}$$

where the subscript e refers to changes due to excess temperature or carbon, and the subscript r refers to changes due to the redistribution of background temperature or carbon. Changes in temperature at a fixed spatial location are therefore due either to the introduction of excess heat, or the redistribution of the existing heat field. Changes in the excess components are then related by a constant, α_T:

$$\Delta\theta_e = \alpha_T \Delta\text{DIC}_e \tag{5.3}$$

and changes in the redistributed components by β_T (note that henceforth, the redistribution coefficient for temperature, previously called κ_r^T will be referred to as β_T to

improve compatibility with the nomenculature of Bronselaer and Zanna (2020)):

$$\Delta\theta_r = \beta_T \Delta \mathrm{DIC}_r \tag{5.4}$$

Both are given the subscript T to identify that they relate changes in temperature and DIC: later this equation will be expanded to also include salinity. Also note at this point that α_T is referred to as a constant in space, following the convention of Bronselaer and Zanna (2020). It is additionally assumed here that the timescales of interest are sufficiently short that we may ignore the time dependence of this parameter. However, the assumption of spatial uniformity will be explicitly tested: it is straightforward to extend α_T to make it a function of initial parameters, spatially resolving it. This allows the inclusion of the temperature dependence of α_T discussed in Chapter 2. Rearranging the above equations yields expressions for the excess and redistributed DIC in terms of the change in the temperature and DIC:

$$\Delta \mathrm{DIC}_e = \frac{\Delta\theta - \beta_T \Delta \mathrm{DIC}}{\alpha_T - \beta_T} \tag{5.5}$$

and

$$\Delta \mathrm{DIC}_r = \frac{\alpha_T \Delta \mathrm{DIC} - \Delta\theta}{\alpha_T - \beta_T}, \tag{5.6}$$

which may be combined in a vector equation as follows:

$$\begin{pmatrix} \Delta \mathrm{DIC}_e \\ \Delta \mathrm{DIC}_r \end{pmatrix} = \frac{1}{\alpha_T - \beta_T} \begin{pmatrix} 1 & -\beta_T \\ -1 & \alpha_T \end{pmatrix} \cdot \begin{pmatrix} \Delta\theta \\ \Delta \mathrm{DIC} \end{pmatrix} \tag{5.7}$$

Here, approximating $\mathrm{C_{anth}}$ with excess DIC allows us to use this decomposition to identify anthropogenic carbon. This will later be used to examine the validity of this decomposition. However, as the primary concern of this decomposition is identifying excess and redistributed temperature and salinity, the relationships of Equation 5.3 and 5.4, combined into a vector equation as

$$\begin{pmatrix} \Delta\theta_e \\ \Delta\theta_r \end{pmatrix} = \begin{pmatrix} \alpha_T & 0 \\ 0 & \beta_T \end{pmatrix} \cdot \begin{pmatrix} \Delta \mathrm{DIC}_e \\ \Delta \mathrm{DIC}_r \end{pmatrix} \tag{5.8}$$

will be used to calculate excess and redistributed temperature:

$$\begin{pmatrix} \Delta\theta_e \\ \Delta\theta_r \end{pmatrix} = \frac{1}{\alpha_T - \beta_T} \begin{pmatrix} \alpha_T & -\alpha_T\beta_T \\ -\beta_T & \alpha_T\beta_T \end{pmatrix} \cdot \begin{pmatrix} \Delta\theta \\ \Delta \mathrm{DIC} \end{pmatrix} \tag{5.9}$$

We may estimate excess and redistributed temperature from Equation 5.9, provided the values of α_T and β_T are known. The matrix in Equation 5.9 will henceforth be referred to, where appropriate, as the Decomposition Matrix, D:

$$D = \frac{1}{\alpha_T - \beta_T} \begin{pmatrix} \alpha_T & -\alpha_T\beta_T \\ -\beta_T & \alpha_T\beta_T \end{pmatrix} \tag{5.10}$$

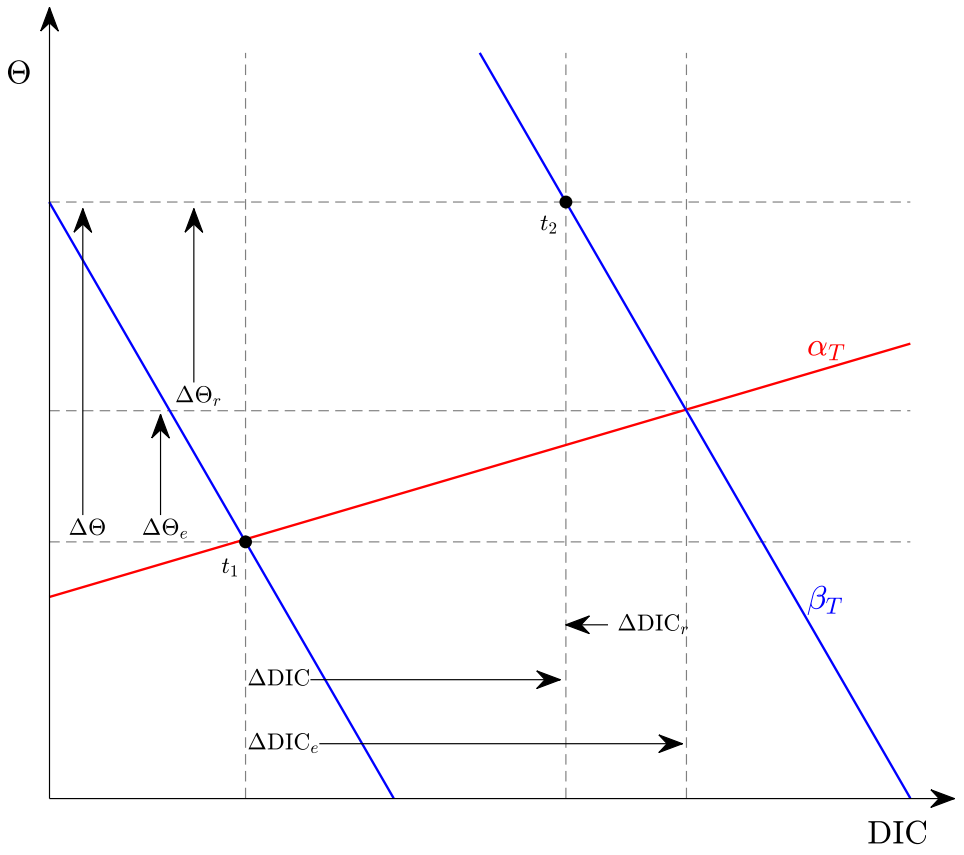

FIGURE 5.3: A diagram illustrating the decomposition of excess and redistributed temperature in temperature-carbon space. Changes in temperature and DIC components are indicated with arrows, and the transient and redistribution axes in red (transient axis, α_T) and blue (redistribution axis, β_T) respectively. For a change in temperature and DIC, $(\Delta\theta, \Delta\mathrm{DIC})$, between times t_1 and t_2, the excess components are linked by the transient axis α_T and the redistributed components by the redistribution axis β_T.

The decomposition encoded in Equation 5.9 is shown diagrammatically in Figure 5.3. α_T relates increases in ocean heat and carbon content, and therefore specifies a curve in temperature-carbon space linking excess heat and carbon. A global mean value can be derived theoretically by considering the radiative forcing of CO_2 and the oceans ability to sequester carbon: this was the focus of Section 2.1.2, where a contemporary value of approximately $0.0195 \pm 0.006 \mathrm{K}/\mu\mathrm{mol/kg}$ was obtained. However, as the focus here is on the A05 hydrographic section, a value directly estimated for this hydrographic section is used, as Chapter 2, Section 2.1.2 demonstrated it may be necessary to allow for variability in α_T locally. This assumption of spatial variability in α_T will be explicitly tested.

The redistribution coefficient β_T relates changes in the temperature and carbon content of water at a given geographical location due to redistribution, and therefore specifies a curve in temperature-carbon space linking redistributed heat and carbon. This is the same assumption that motivates the use of T-S space to remove the effects of circulation variability when studying physical changes in water properties (Evans et al., 2014). Mathematically, β_T is defined by

$$\beta_T = \frac{\vec{v}' \cdot \nabla\theta}{\vec{v}' \cdot \nabla \mathrm{DIC}}, \tag{5.11}$$

thus representing a weighted sum of the gradient in θ-DIC space by a general perturbation to the mean velocity at that point. In practice, however, it can generally be approximated as the ratio of the vertical stratification of the two properties:

$$\beta_T \approx \frac{\partial\theta}{\partial z} \Big/ \frac{\partial \mathrm{DIC}}{\partial z}. \tag{5.12}$$

This was demonstrated in Section 2.3. This approximation is used here, instead of statistically estimating β_T (the approach taken by Turner et al. (2022) and in Chapter 3), as hydrographic occupations provide a temporal snapshot of the ocean, rather than high frequency timeseries of temperature and DIC.

Now, we must consider a potential issue with this decomposition: the curve α_T is fixed, but β_T is not. Therefore, in the limit $\beta_T \to \alpha_T$, the estimates of excess and redistributed temperature will asymptotically approach $\pm\infty$. To see this, consider the eigenvalues of the Decomposition Matrix, D (Equation 5.10). Solving the characteristic equation, we obtain

$$\lambda_{1,2} = \frac{\alpha_T(\beta_T + 1) \pm \sqrt{\alpha\left(\beta_T^2(\alpha_T + 4) - \alpha_T(\beta_T + 4)\right)}}{2(\alpha_T - \beta_T)}. \tag{5.13}$$

From this, it is clear that the eigenvalues of this equation approach infinity as $\beta_T \to \alpha_T$. We may therefore remove unreliable estimates of excess and redistributed temperature based on the eigenvalues of this matrix. To do so, consider a masking factor, M, defined as

$$M = \min(1, \max(1 - \log_{10}|\lambda_1|), 0) \tag{5.14}$$

Values of excess temperature and salinity are excluded where $M < 1$: this removes virtually all unreliable estimates of excess and redistributed temperature, whilst leaving reliable estimates unaltered. Intuitively, this approach may be expressed with the following heuristic: "if excess temperature changes are of a scale which approaches or is larger than that of DIC changes, estimates are likely to be unreliable and are so discarded".

The inclusion of salinity in this decomposition is now considered. As previously noted, redistribution does not change distributions in tracer-tracer space, and so the θ-S curve

is unaltered by redistribution (this insight also forms the basis of the approach used by Zika et al. (2021) to estimate material warming). We may therefore estimate redistributed salinity, ΔS_r, as either of two equivalent forms:

$$\Delta S_r \approx \left(\frac{\partial S}{\partial z} \middle/ \frac{\partial DIC}{\partial z} \right) \Delta DIC_r = \beta_S \Delta DIC_r \tag{5.15}$$

or

$$\Delta S_r \approx \left(\frac{\partial S}{\partial z} \middle/ \frac{\partial \theta}{\partial z} \right) \Delta \theta_r = \tau^{-1} \Delta \theta_r. \tag{5.16}$$

Salinity can therefore be included in this decomposition as follows:

$$\begin{pmatrix} \Delta\theta_e \\ \Delta\theta_r \\ \Delta S_r \end{pmatrix} = \frac{1}{\alpha_T - \beta_T} \begin{pmatrix} \alpha & -\alpha_T\beta_T \\ -\beta_T & \alpha_T\beta_T \\ -\beta_S & \alpha_T\beta_S \end{pmatrix} \cdot \begin{pmatrix} \Delta\theta \\ \Delta DIC \end{pmatrix} \tag{5.17}$$

Note that the 2nd and 3rd rows of the matrix in Equation 5.17 are linearly dependent, expressing the linear relationship between redistributed temperature and salinity at a point. Though it is not further pursued here, again this is in principle extensible to any tracer (even those with uniform vertical distributions may be redistributed, though redistribution of a uniform field has no effect). This may be expressed as

$$\begin{pmatrix} \Delta\theta_e \\ \Delta\theta_r \\ \Delta S_r \\ \Delta Q^1_r \\ \vdots \\ \Delta Q^i_r \end{pmatrix} = \frac{1}{\alpha_T - \beta_T} \begin{pmatrix} \alpha & -\alpha_T\beta_T \\ -\beta_T & \alpha_T\beta_T \\ -\beta_S & \alpha_T\beta_S \\ -\beta_{Q^1} & \alpha_T\beta_{Q^1} \\ \vdots & \vdots \\ -\beta_{Q^i} & \alpha_T\beta_{Q^i} \end{pmatrix} \cdot \begin{pmatrix} \Delta\theta \\ \Delta DIC \end{pmatrix} \tag{5.18}$$

for an arbitrary set of tracers Q^i. As with the previous equation, no additional information used to decompose changes is introduced by the introduction of additional tracers Q^i. However, if additional relationships between changes in excess temperature and additional tracers could be found, for example, it would in principle be possible to include them in this decomposition, improving its reliability.

5.3 Data and Methods

5.3.1 Data and Interpolation

This technique is applied to the 6 repeat hydrographic occupations of the A05 section in the Subtropical North Atlantic contained in the GLODAP dataset Lauvset et al. (2021).

This comprises 6 cruises: 1992, 1998, 2004, 2010, 2011, and 2015. 2020 data, collected as part of this work, have not been included. For the 1992 cruise, two teams (Spanish and American) worked on carbon, with the results of the American team being submitted to GLODAP. However, the Spanish team collected more samples with a greater coverage (DIC measurements were made at every station, unlike the American team). As such, the DIC data used by the Spanish team is considered in this work: this is implemented as an exception in the lookup routines of the Julia toolbox described in Appendix A.

These data are used to estimate excess and redistributed temperature and salinity relative to the year 1992 for the 5 hydrographic occupations since. Bottle data from each cruise are gridded onto a 651 vertical level by 670 longitude grid, using DIVAnd gridding (Barth et al., 2014). Each vertical level spans 10db, and each longitudinal bin spans approximately one tenth of a degree of latitude. It was decided to use temperature and salinity bottle data, rather than CTD data, as only bottle data are available for DIC. This ensures all data are gridded in an identical manner, reducing the likelihood of mapping techniques influencing results. As cruise data are taken from the GLODAP dataset, they are automatically adjusted with GLODAP recommended adjustments: this process is described in Appendix A.

Gridding produces 6 sections of temperature, salinity and DIC: one for each occupation, and all on the same grid. β_T and β_S are then estimated using the 1992 occupation by computing vertical derivatives for each field as central differences for each grid cell. Changes in temperature and salinity are then decomposed by applying the Decomposition Matrix D independently at each point. Unreliable values of temperature and salinity are masked out using the masking factor M, and missing data points are reinterpolated, again using DIVAnd, to produce final fields of excess temperature and salinity.

Observational uncertainties in DIC measurements are approximately 2μmol/kg, temperature measurement uncertainty approximately $2\times10^{-3}\,^{\circ}$C, and salinity 2×10^{-3}PSU. These generally result from undetectable systematic measurement uncertainty, and so are uniform within each cruise. As such, nearly all uncertainty in the estimates of excess and redistributed temperature, salinity and carbon result from uncertainty in the parameters α_T, β_T and β_S.

In addition, four C_{anth} reconstructions (TTD, ϕC_T^0, TrOCA and ΔC^*) are available for the 1992-2011 cruises. These are used to validate the estimates of excess DIC, as well as to produce an initial estimate of α_T. Validation of excess DIC using C_{anth} requires C_{anth} data on the same grid as excess DIC: thus, C_{anth} data were also gridded using DIVAnd.

5.3.2 Decomposition Parameters and Uncertainty

Bronselaer and Zanna (2020) approximated the relationship between excess heat and

anthropogenic carbon, α, as being globally uniform in order to estimate excess heat: this is a necessary approximation in their method. However, the results of Chapter 3 were able to relax this assumption, and explicitly calculate a local value of α_T. These suggest, in agreement with the CO2SYS calculation presented in Chapter 2, that α_T is better described as a function of temperature. Here, whether it is necessary to account for variability in the value of α_T in this new method is directly quantified.

To do so, the methods of both Bronselaer and Zanna (2020) and Turner et al. (2022) are applied independently on the same dataset. If a systematic correction for variability in α_T were necessary, this should show that the method of Bronselaer and Zanna (2020) overestimates excess temperature in warmer waters, relative to that of Turner et al. (2022): the findings of Turner et al. (2022) and Chapter 2 indicate systematic variability in α_T, with larger α_T values in cooler waters.

However, it is important to note at this point this analysis is cannot be extrapolated to conclude whether it is necessary to account for the temperature dependence of α_T globally. The results of Chapter 2, Section 2.2 suggest a theoretical temperature dependence for α_T which gives the same value as the section mean α_T at approximately 16 degrees Celsius. Waters of this temperature are found in the gyre, where the majority of excess heat accumulation is expected to occur (Zanna et al., 2019). Thus, agreement in excess heat accumulation from the constant α_T and the temperature dependent formulation here may result from simple coincidence. At abyssal depths (3000 metres and deeper), waters are both much cooler and expected to have negligible excess heat accumulation: thus, differences in estimates of excess heat between the two methods will be undetectable due to the lack of excess heat accumulation here. It is therefore not obvious that this analysis is also applicable to waters in polar regions, where cooler waters are found at the surface and are thus expected to contain greater quantities of excess heat due to their more recent ventilation.

In a manner analogous to Equation 5.9, the approach of Bronselaer and Zanna (2020) may be expressed as

$$\begin{pmatrix} \Delta\theta_e \\ \Delta\theta_r \end{pmatrix} = \begin{pmatrix} \alpha_T & 0 & 0 \\ -\alpha_T & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} \Delta C_{\text{anth}} \\ \Delta DIC \\ \Delta\theta \end{pmatrix}, \tag{5.19}$$

where $\Delta\theta_e$ is the excess temperature change between two observations, $\Delta\theta_r$ is the redistributed temperature change, α_T the global mean carbon-temperature coupling, and $\Delta\theta$ the total temperature change. It is important to note that the decomposition expressed by Equation 5.19 is not the same as presented in Section 5.2: thus, the terms presented are not mathematically consistent with that decomposition.

In a similar fashion to the above, the redistribution coefficient decomposition may be written as

$$\begin{pmatrix} \Delta\theta_e \\ \Delta\theta_r \end{pmatrix} = \begin{pmatrix} \beta_T & -\beta_T & 1 \\ -\beta_T & \beta_T & 0 \end{pmatrix} \cdot \begin{pmatrix} \Delta C_{anth} \\ \Delta DIC \\ \Delta\theta \end{pmatrix}, \tag{5.20}$$

where β_T is the redistribution coefficient (note $\Delta C_{nat} = \Delta DIC - \Delta C_{anth}$). Again, this is a separate decomposition to that presented in Section 5.2, it instead represents the decomposition presented in Chapter 3. Thus, it does not include a factor of $1/(\alpha_T - \beta_T)$, for example.

As the two estimation techniques expressed in Equations 5.19 and 5.20 enforce different assumptions about the nature of the relationship between temperature and carbon changes, there is no guarantee that they will produce the same estimates of excess and redistributed temperature change. This can therefore be used to estimate how reliable an estimate of excess temperature is obtained by using a constant α_T value. We may add uncertainty terms to our two expressions for excess and redistributed temperature:

$$\begin{pmatrix} \Delta\theta_e \\ \Delta\theta_r \end{pmatrix} = \begin{pmatrix} \alpha_T + \sigma_\alpha & 0 & 0 \\ \alpha_T + \sigma_\alpha & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} \Delta C_{anth} \\ \Delta DIC \\ \Delta\theta \end{pmatrix} \tag{5.21}$$

$$\begin{pmatrix} \Delta\theta_e \\ \Delta\theta_r \end{pmatrix} = \begin{pmatrix} \beta_T + \sigma_\beta & -\beta_T + \sigma_\beta & 1 \\ -\beta_T + \sigma_\beta & \beta_T + \sigma_\beta & 0 \end{pmatrix} \cdot \begin{pmatrix} \Delta C_{anth} \\ \Delta DIC \\ \Delta\theta \end{pmatrix}, \tag{5.22}$$

where σ_α and σ_β are uncertainties in the value of α_T and β_T, respectively. This allows both estimates to vary within observational uncertainty in order to produce results consistent with one another. To do so, we will aim to minimise the quantity $\delta\Delta\theta_e$, defined as

$$\begin{pmatrix} \delta\Delta\theta_e \\ \delta\Delta\theta_r \end{pmatrix} = \left[\begin{pmatrix} \alpha_T + \sigma_\alpha & 0 & 0 \\ \alpha_T + \sigma_\alpha & 0 & 1 \end{pmatrix} - \begin{pmatrix} \beta_T + \sigma_\beta & -\beta_T + \sigma_\beta & 1 \\ -\beta_T + \sigma_\beta & \beta_T + \sigma_\beta & 0 \end{pmatrix} \right] \cdot \begin{pmatrix} \Delta C_{anth} \\ \Delta DIC \\ \Delta\theta \end{pmatrix}, \tag{5.23}$$

whilst simultaneously minimising the total fractional error in both estimates:

$$\Sigma^2 = \left(\frac{\sigma_\alpha}{\alpha_T}\right)^2 + \left(\frac{\sigma_\beta}{\beta_T}\right)^2, \tag{5.24}$$

$\delta\Delta\theta_e$ therefore represents the difference between the excess temperature as estimated by the approach of Bronselaer and Zanna (2020) and the approach of Turner et al. (2022). As this is an underdetermined problem and the aim is to simultaneously minimise $\delta\Delta\theta_e$ and Σ^2, fractional errors are limited to be less than half their 'initialisation' values.

When $\delta\Delta\theta_e = 0$, both estimation techniques give the exact same decomposition of temperature into excess and redistributed components. I define the local value of α_T, $l\alpha_T$ as the quantity $\alpha_T + \sigma_\alpha$ in each grid cell where $\delta\Delta\theta_e = 0$. Similarly, I define the 'best' estimate of β_T as the quantity $\beta_T + \sigma_\beta$ in each grid cell where $\delta\Delta\theta_e = 0$.

For the A05 section, we may approximate the global mean α_T value with the section mean α_T value: over our study period, this is 0.0171Kμmol/kg for the ΔC^* decomposition, 0.0173 for the TTD and φC_T^0 decompositions and 0.0135 for TrOCA. These values are calculated using the ratio of section mean C_{anth} to section mean temperature change. As the estimates for all but the TrOCA decompositions are remarkably close, a value of 0.017K/μmol/kg is used as an initial estimate of α_T.

To calculate β_T (avoiding divide by zero errors at inflection points) a 5 point linear polynomial fit is used, from 2 vertical levels above the grid cell of interest to two below. This is not a problem in the full matrix inversion decomposition, as infinities due to inversions in the carbon profile analytically cancel in this case.

Only the results from the ΔC^* decomposition are shown here: results from the TrOCA, and ϕC_T^0 decompositions are similar. The differences between the technique of Bronselaer and Zanna (2020) with no variability in α_T allowed) are shown in Figure 5.4a, and with a best estimate of a local α_T value, labelled $l\alpha_T$, are shown in Figure 5.4b. In addition, the difference between the two is shown in Figure 5.4c, and the C_{anth} change used in this analysis is shown in Figure 5.4d. The results of the analysis are presented in Figure 5.5.

To obtain values of $l\alpha_T$, the optimal value of α_T is computed for each pair of cruises in the set 1992, 1998, 2004, 2010 and 2011 (ie. 1992-1998, 1992-2004, 1998-2011 etc). In each grid cell, the mean value of $l\alpha_T$ is then computed using all available data points, before these mean values are smoothed horizontally and then vertically to reduce spurious variability as a result of the averaging process: this final distribution is shown in Figure 5.5a.

Whilst no clear spatial patterns of $l\alpha_T$ are discernible from Figure 5.5a, the zonal mean values presented in Figure 5.5b show there is a clear increase in mean $l\alpha_T$ values with depth in both the Western and Eastern Basin. However, 5.5c shows that when presented as a function of temperature without horizontal averaging or vertical smoothing, variability is far less clear, with noise far exceeding any clear relationship between depth and $l\alpha_T$ values.

As a result, although there does appear to be a relationship between temperature and $l\alpha_T$ (with cooler waters exhibiting larger values as indicated by the results of Chapter 2, Section 2.2.1 and Chapter 3, Section 3.3) the strength of this relationship does not appear large enough to warrant using a spatially resolved α_T value, at least not for the A05 section. In fact, it will be shown in Section 5.4.2 that for the A05 section, the

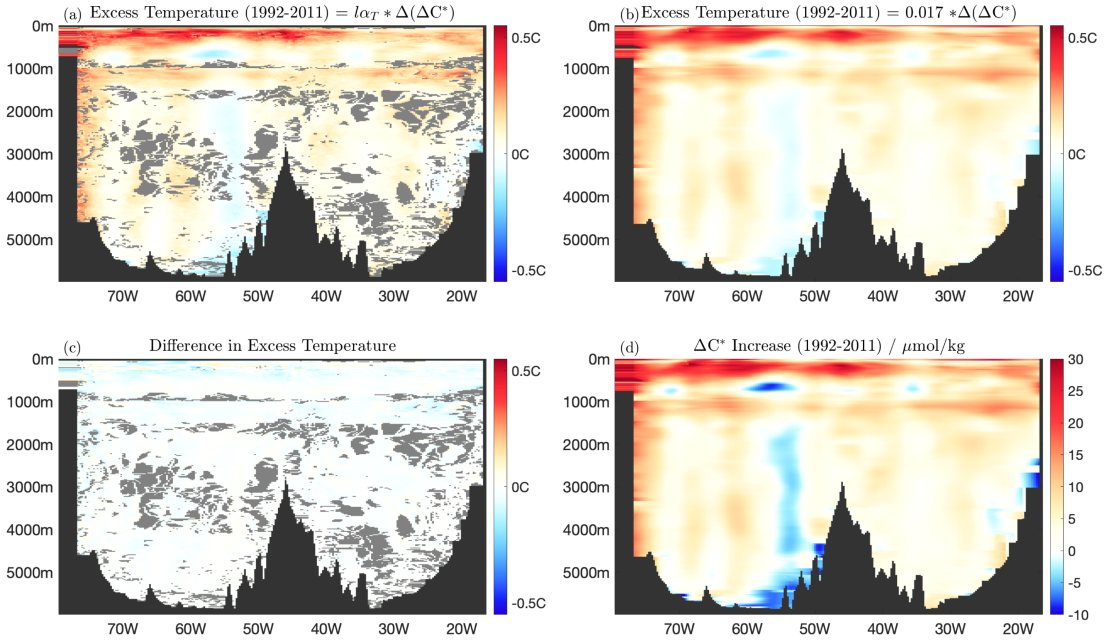

FIGURE 5.4: An excess temperature derived using an optimal local value of α_T: $l\alpha_T$ (Panel a), an excess temperature derived using a section mean α_T of 0.017 K/(μmol/kg (Panel b), the difference between the two estimates (Panel c), and the C_{anth} change used to calculate the optimal $l\alpha_T$ values (Panel d).

choice of uniform or temperature dependent α_T does not significantly affect estimates of excess temperature.

Figure 5.5d shows the distribution of $l\alpha_T$ values obtained from this analysis, presented as a histogram. The distribution obtained well approximates a Gaussian with mean 0.0157K/(μmol/kg) and standard deviation 0.0013K/(μmol/kg). As the initialisation value of 0.0171K/(μmol/kg) is approximately 1σ from the distribution mean, the distribution peak value is not used as a mean α_T value: instead, a rounded initialisation value of 0.017K/(μmol/kg) is used, reflecting both imperfections in this analysis and that the distribution obtained indicates 3 significant figures is an unrealistic level of confidence in a value of α_T.

Now consider the error in β_T. The estimate of the optimal value of $l\alpha_T$ also yields an optimal value of β_T: however, unlike α_T, initial estimates of β_T are spatially variable. Therefore, the error in β_T is presented as a fractional change from the initial to the optimal value: this is shown in Figure 5.6. The distribution of changes is well approximated by the sum of two Gaussian distributions, with nearly all changes from initial to optimal values less than 5%. For simplicity, it will be instead approximated as a single Gaussian. This distribution has mean and standard deviation of 2%, and, unlike that of α_T, does not appear to contain any spurious spatial variability: therefore, it has not been smoothed.

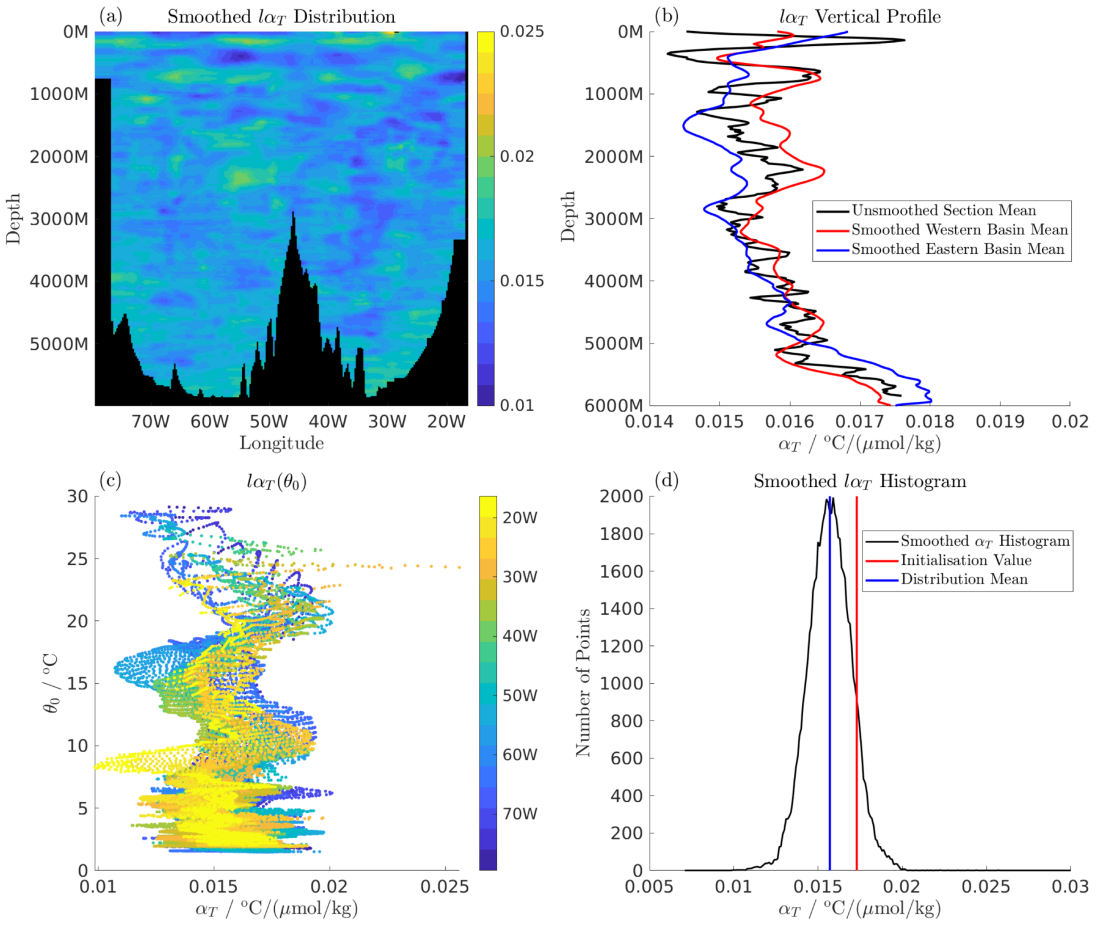

FIGURE 5.5: The distribution of $l\alpha_T$ values at A05: a spatially smoothed distribution (a), the mean vertical profiles (b), smoothed $l\alpha_T$ values as a function of initial temperature (c) and the histogram of values obtained (d).

From this analysis, it is apparent that for the A05 section, we may consider α_T and β_T as having uncertainties which are well approximated by Gaussian distributions centred on the calculated values. Therefore, in order to compute uncertainties in excess and redistributed temperature and salinity in the matrix inversion decomposition, we may perturb the coefficients α_T and β_T randomly within uncertainty in order to bootstrap confidence intervals. For the case of constant α_T, values are randomly sampled from a Gaussian distribution, with a mean value of $0.017^oC/\mu mol/kg$, and a standard deviation of $0.0013^oC/\mu mol/kg$. For β_T, I use the locally computed value (from Equation 5.12) as the mean value, and perturb the values by applying a normally distributed scale factor, with mean 1 and standard deviation 0.02.

In the case of temperature dependent α_T, it is necessary to give an *a priori* estimate of the temperature dependence: it cannot be estimated from observational data. To do so, the CO2sys calculations performed in Chapter 1, Section 2.2.1 are used. For the range of

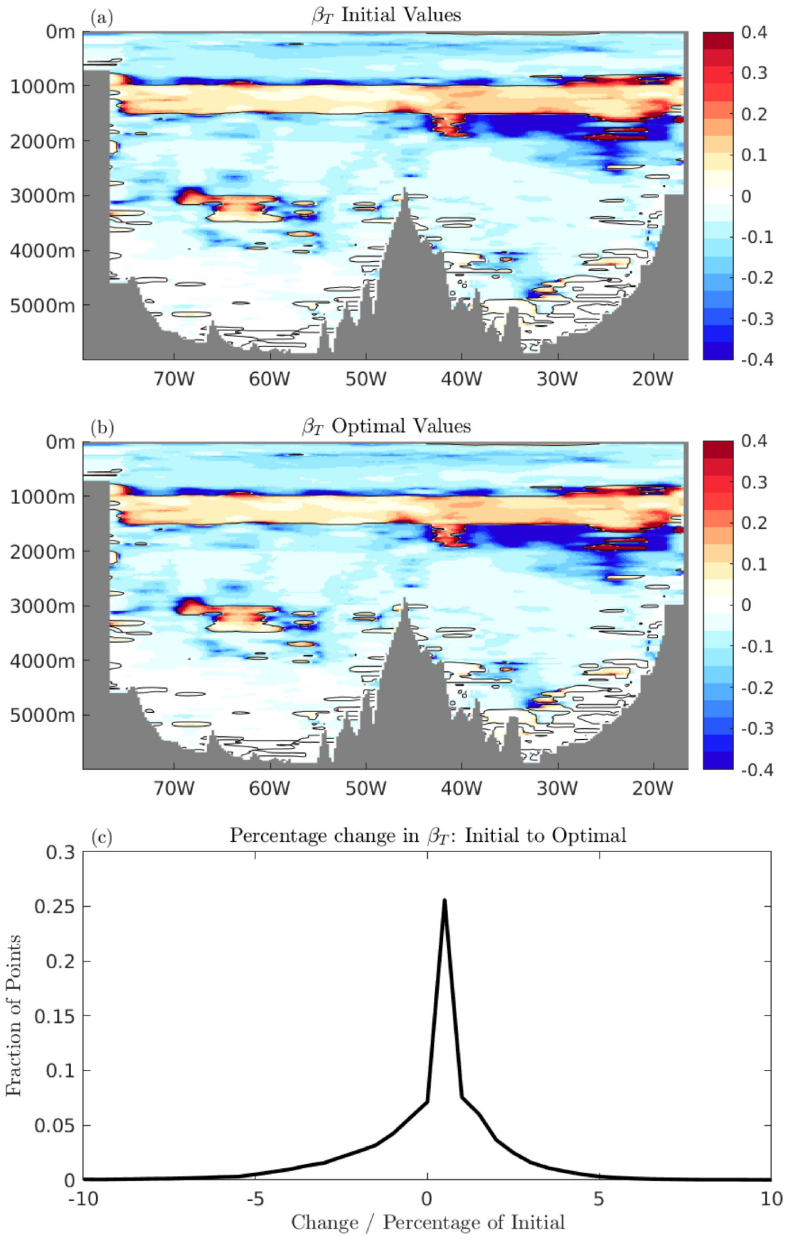

FIGURE 5.6: The uncertainty in estimates of β_T. Initial values are shown in Panel (a), the optimal values in Panel (b), and the percentage change in Panel (c). As with α_T, all values of β_T given are in $^{o}C/\mu mol/kg$)

initial temperatures shown in Figure 2.6, the mean temperature and DIC change were computed at each pCO_2 value, using global mean warming estimates from CMIP6 data. A linear fit of the coefficient $\Delta\theta/\Delta\text{DIC}$ was then computed. This yields

$$\alpha_T(\theta) = 0.025 - 0.0005\theta, \tag{5.25}$$

with an R^2 value of 0.98. This linear dependence on background temperature is thus an excellent approximation. Varying the global mean warming within the range of uncertainty of warming given by CMIP6 models does not affect the coefficient of the temperature dependent term in Equation 5.25, but does change the constant significantly. This variation, from a lower limit of 0.018 to an upper limit of 0.032, is unrealistically large: it implies an uncertainty almost 4 times the size of the uncertainty ($0.0068^oC/\mu mol/kg$) found using observational data to calculate a section mean value for α_T. As uncertainty in warming in CMIP6 models are largely the result of uncertainty in modelling the response of climate to cloud and aerosol radiative forcing (Myhre et al., 2013), the same distribution shape is instead used for this constant as was for the temperature independent case: a Gaussian, with width $0.0013^oC/\mu mol/kg$ and mean value 0.0025. The coefficient of the temperature dependent term was not varied: allowing it to vary by 10% had a negligible impact on results, and so it is ignored.

5.3.3 Implementation

Implementation of this technique is straightforward, being a simple matrix multiplication. However, care must be taken to avoid numerical infinities during implementation (analytically, these infinities are not problematic). It is therefore implemented as

$$\begin{pmatrix} \Delta\theta_e \\ \Delta\theta_r \\ \Delta S_r \end{pmatrix} = \left[g(\zeta,0,\sigma) \begin{pmatrix} 0 & \alpha_T \\ 1 & -\alpha_T \\ \tau^{-1} & -\alpha_T\tau^{-1} \end{pmatrix} + \frac{1-g(\zeta,0,\sigma)}{\alpha_T-\beta_T} \begin{pmatrix} \alpha_T & -\alpha_T\beta_T \\ -\beta_T & \alpha_T\beta_T \\ -\beta_S & \alpha_T\beta_S \end{pmatrix} \right] \cdot \begin{pmatrix} \Delta\theta \\ \Delta\text{DIC} \end{pmatrix} \tag{5.26}$$

where $g(\zeta,0,\sigma)$ is a Gaussian with zero mean and width σ, and $\zeta = \partial_z\text{DIC}$. When $\zeta \to 0$, $\beta_T \to \infty$. Thus, problematic infinities caused by inversions in the DIC profile are taken care of by using the small ζ limit of the Decomposition Matrix when β_T is large enough to cause numerical issues. It was not necessary to implement a similar approach for inversions in salinity. A width of 0.005 is used for σ.

To estimate the uncertainty in the reconstructions of excess and redistributed temperature and salinity, the paramaters α_T and β_T are perturbed via bootstrapping with an ensemble of 1000 members, as the error covariance of these parameters is unclear. To do so, estimates of α_T and β_T are perturbed at random as described above: according to Gaussian distributions centred on their mean values ($0.017K/\mu mol/kg$ for α_T), with

standard deviations of 0.0013K/(μmol/kg) for α_T and 2% of the mean β_T value for β_T. Additionally, uniform perturbations (that is, for each grid cell in a given cruise the same perturbation is applied) within the range of observational uncertainty to DIC, temperature and salinity measurements ($\pm 2\mu$mol/kg, $\pm$0.002K, $\pm$0.002PSU, respectively).

After excess temperature and salinity are estimated with the perturbed parameters, the field were then masked using the masking factor M (Equation 5.14), to remove spurious large values. However, unlike the unperturbed fields and due to computational expense, no reinterpolation was performed: the fields were instead left blank.

Uncertainty on the excess and redistributed temperature and salinity fields are then calculated as twice the standard deviation of the bootstrap ensemble at each point: this produces a 95% confidence interval on the local estimate of these quantities, assuming the ensembles are normally distributed: this is typically an excellent approximation.

5.4 Results

5.4.1 Constant α

Figure 5.7 shows the mean vertical profiles of excess and redistributed temperature and salinities. Excess temperature are appreciable generally only in the top 1000m, and with the exception of 2011, increases monotonically with time in this region. In general, both excess and redistributed temperature decrease monotonically with depth, with the redistributed temperature exhibiting a much larger scale. However, unlike excess temperature, redistributed temperature does not increase monotonically with time: profiles are similar in 1998, 2004 and 2011, with the 2010 and 2015 profiles much smaller. This is as expected: redistribution of the preindustrial temperature field is thought to comprise the majority of contemporary temperature change, as well as variability (Bronselaer and Zanna (2020), Zika et al. (2021)).

Excess salinity does not share the same structure as excess temperature, nor does it exhibit a smaller scale than redistributed salinity. However, like excess temperature, excess salinity accumulation is negligible below 1500m at all times.

Excess salinity at depths of 250-750m evolves following a similar pattern to temperature, increasing monotonically with time except in 2011. This pattern of correlated change between excess salinity and excess temperature appears to be a consistent feature (Turner et al., 2022) and appears to be a manifestation of density compensated changes in water mass properties. This density compensation is thought to be a particularly strong effect in the Atlantic (Lowe and Gregory, 2006), (Mauritzen et al., 2012).

It is clear from Figure 5.7a that the uncertainty in excess temperature increases with the value of excess temperature. In fact, the uncertainty in excess temperature, σ_{θ_e} is well

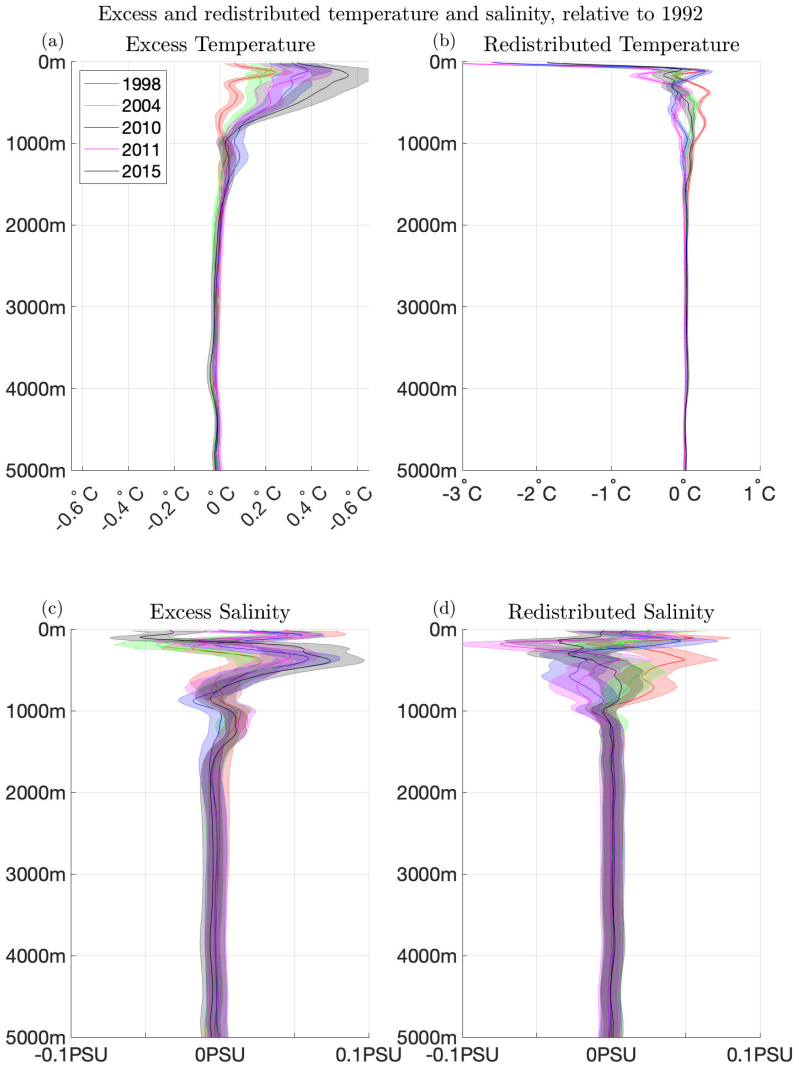

FIGURE 5.7: The mean profiles of excess (a) and redistributed (b) temperature, and excess (c) and redistributed (d) salinity, for the 5 occupations since 1992. 95% confidence intervals are represented by shading, and scales are not shared for excess and redistributed components.

described by

$$\sigma_{\theta_e} = 0.001 + 0.14|\theta_e|, \tag{5.27}$$

with this linear approximation yielding an R^2 value of 0.95. The uncertainty in excess temperature is therefore approximately a seventh the excess temperature. This uncertainty is of a similar magnitude to that of the material heat uncertainty calculated by Zika et al. (2021). However, such a simple parametrisation is not possible for redistributed temperature, nor excess or redistributed salinity: the R^2 values obtained are too low to make such parametrisations useful (for example 0.27 for excess salinity).

The spatial patterns of excess, redistributed and total change in temperature is shown in Figure 5.8, and the patterns of excess, redistributed and total salinity change in Figure 5.9. As expected, the excess temperature signal increases relatively smoothly with time, with the greatest increases in the upper 1000m of the ocean. The temperature change is dominated again by the redistribution signal. For salinity, however, the dominance of redistribution is less obvious: by 2015, salinity changes are contributed to equally by both excess and redistributed salinity.

Unlike the mean profiles and salinity, excess and redistributed temperatures are generally distinct from zero throughout the section. This likely reflects overconfidence in our technique: from the horizontal mean, it is clear that the decomposition is assigning balanced positive and negative redistributed temperatures. This most probably results from our vertical only approximation. However, the effect of this on our estimates is minimal: though the excess warming and cooling at depths of several thousand metres is nonzero with 95% confidence, these warming and cooling trends are extremely weak, and cancel in the horizontal mean.

Conversely, the majority of the excess and redistributed salinity signals at depths greater than 1000m are calculated to be insignificant at 95% confidence. The exception to this is the excess salinity at depth in the Western Basin in 2015 (Figure 5.9m). However, the DIC measurements in this location at this cruise are somewhat suspect (King (2015), Page 61 & personal communications with Ute Schuster), and this likely results from this. It also appears to be the case that the statistically significant cooling signals seen in the Western Basin in 2015 (Figure 5.10 are also the result of issues with these anomalously low DIC measurements. Conversely, the penetration to depth of a warming signal in the western boundary (most visible in 2004 and 2011, Figure 5.10 (d) and (j)) does appear to be a genuine signal. The excess temperature distribution in 2011 (Figure 5.10j) shows a number of features consistent with the change in ΔC^* from 1992-2011 (Figure 5.4d): an increase in excess temperature and anthropogenic carbon in the top 1000m and the deep western boundary current region, and a decrease in both at depths below 2000m just west of the Mid Atlantic Ridge.

By considering whether trends are distinct from zero, it is possible to estimate the error in our excess temperature estimates associated with the gridding process and vertical

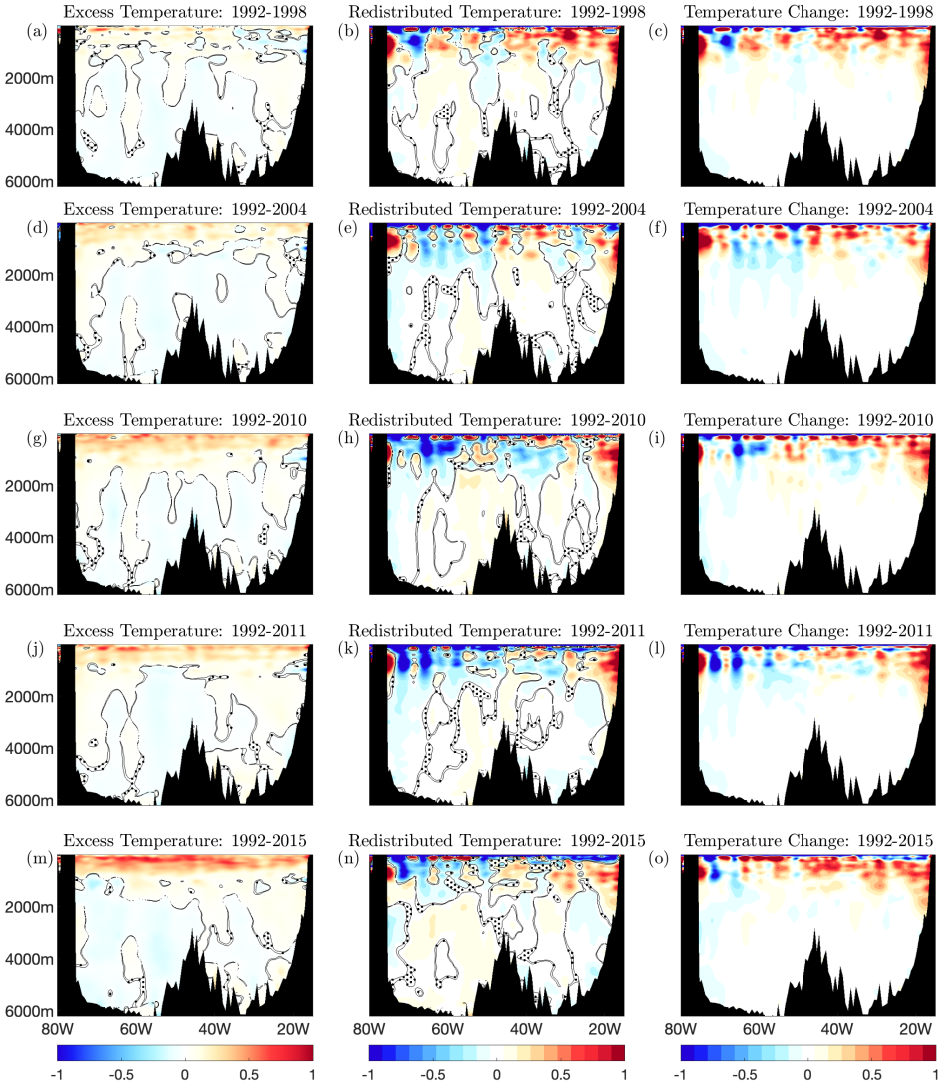

FIGURE 5.8: Excess temperature, redistributed temperature, and total temperature change for each of the 5 occupations subsequent to 1992. Regions for which estimates of excess or redistributed temperature are not significant at 2σ are shown by stippling.

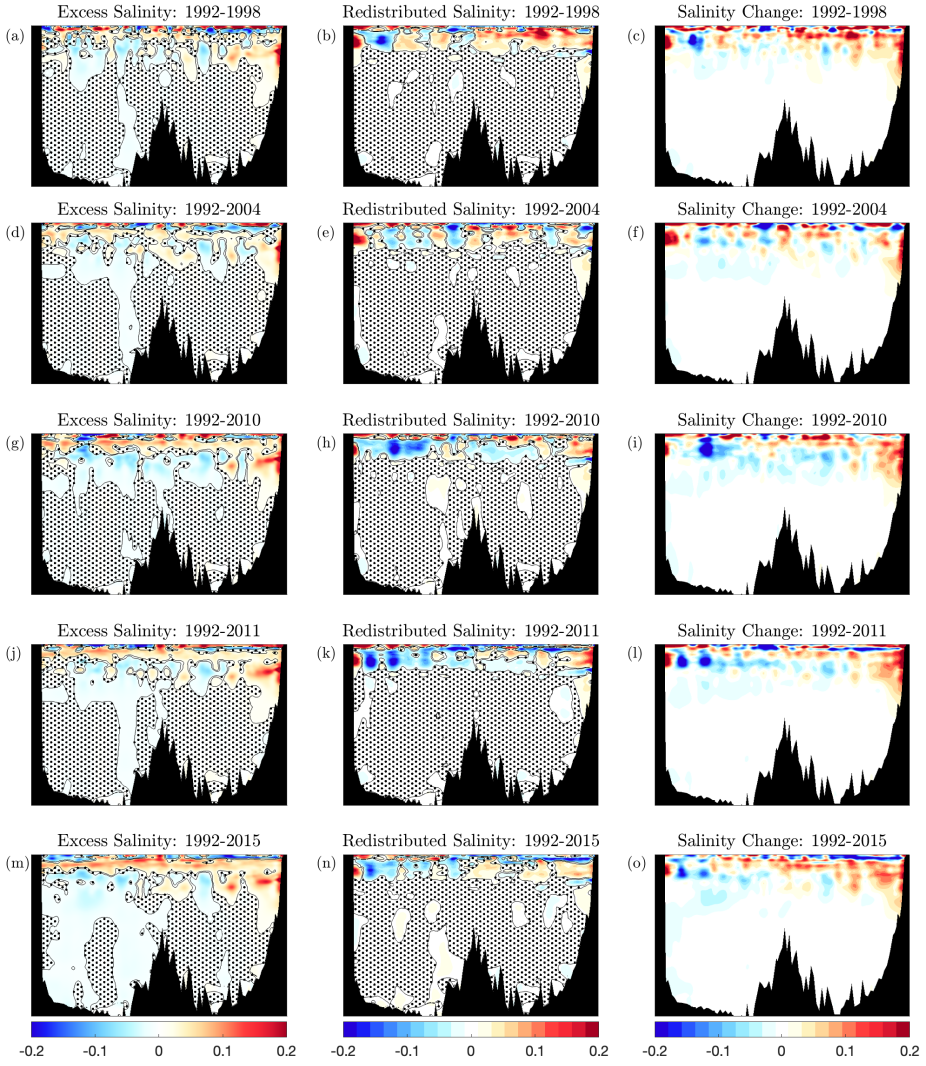

FIGURE 5.9: Excess salinity, redistributed salinity, and total salinity change for each of the 5 occupations subsequent to 1992. Regions for which estimates of excess or redistributed salinity are not significant at 2σ are shown by stippling.

only approximation. Figure 5.7 indicates that at depths below approximately 1000m (1500m in 2010), we do not expect to be able to detect any meaningful trends in excess temperature. By explicitly requiring this, we may estimate this error. To do so, a uniform error is added to each grid cell in addition to the bootstrapped errors shown in Figure 5.8, until excess temperature estimates below approximately 1000m are no longer distinct from noise. This is shown in Figure 5.10, and indicates that the uncertainty introduced to our estimates of excess and redistributed temperature through the combination of gridding error and our vertical only approximation is 0.05K. Alternatively, we could explicitly calculate the error fields associated with the DIVAnd gridding process in order to progagate these errors into our error calculation. However, doing so would not account for the interaction of these errors with the uncertainty introduced by the assumption that redistribution can be described entirely described by vertical motions. Thus, the empirical approach used here is more stringent.

Thus, adjusting to include for this, we may parametrise the error in our excess temperature as

$$\sigma_{\theta_e} = 0.051 + 0.14|\theta_e|, \tag{5.28}$$

where the R^2 value of 0.95 remains unchanged by the addition of a constant factor.

5.4.2 Constant versus variable α

Figure 5.5 suggests that there may be some temperature dependence in the local value of α_T, but that it is not straightforward to spatially resolve this. One way to test explicitly whether it is necessary to account for this temperature dependence is to decompose temperature under the assumption of constant α_T and under variable α_T and to inspect the differences in the two quantities. This is shown in Figure 5.11.

From these plots, it is clear that there are systematic differences in the two estimates: generally, the constant α_T estimate gives higher excess temperatures directly in the subsurface and western basin, and lower excess temperatures at mid depths and in the eastern basin. However, the two estimates are statistically indistinguishable at 2σ virtually everywhere. Inspecting the section mean accumulation of excess temperature and salinity, accounting for the temperature dependence of α_T, we again see similar patterns of excess and redistributed temperature and salinity. It appears that this estimate produces excess temperature accumulation in the upper 200m of the ocean (Figure 5.12), relative to the constant α_T estimate. However, these differences are minor (see Figure 5.13), with differences that are visibly nonzero only found in the upper 250m or so. Whilst these differences are statistically indistinct, the difference in mean value is noticeable.

A priori, it does not seem sensible that the accumulation of excess temperature in the upper 200m of the ocean should be less than that at depths of 500-1000m, as is the

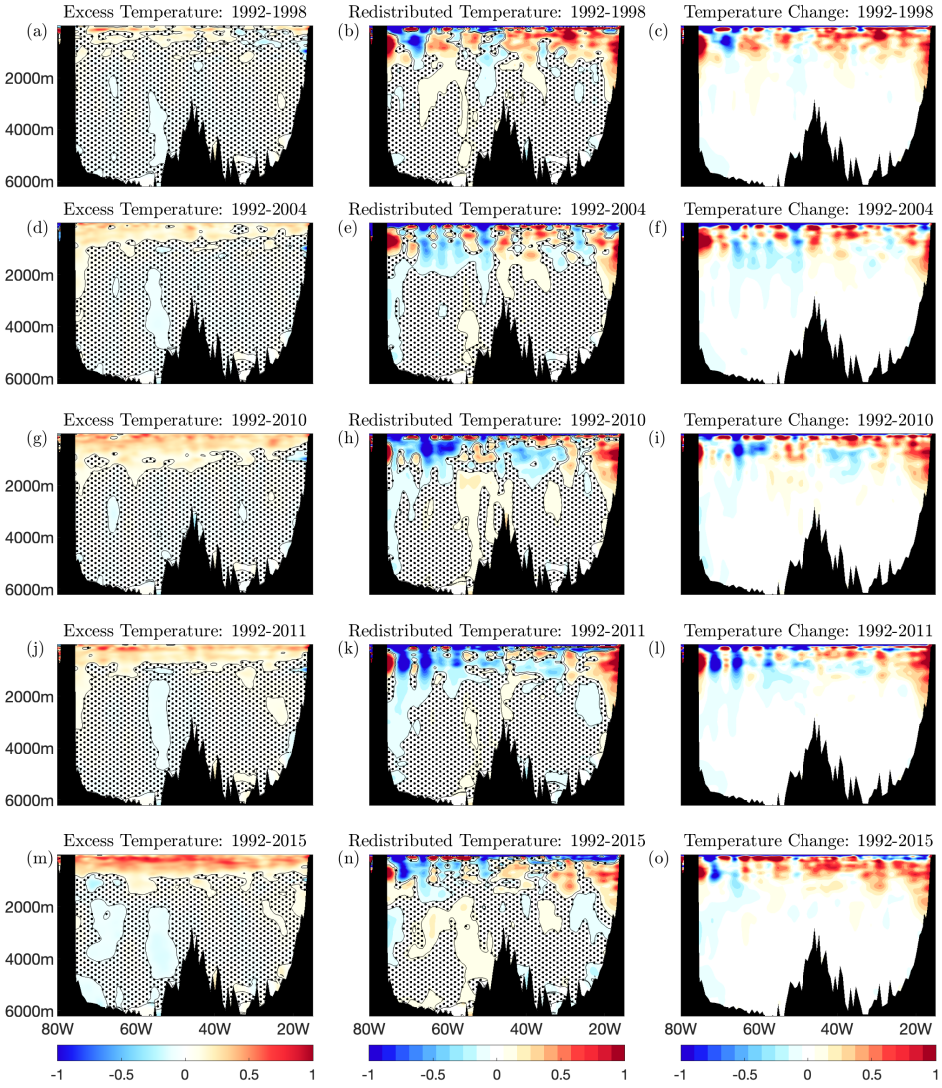

FIGURE 5.10: Excess temperature, redistributed temperature, and total temperature change for each of the 5 occupations subsequent to 1992, as in Figure 5.8. A uniform error of 0.05 has been added to each point. Regions for which estimates of excess or redistributed temperature are not significant at 2σ are shown by stippling.

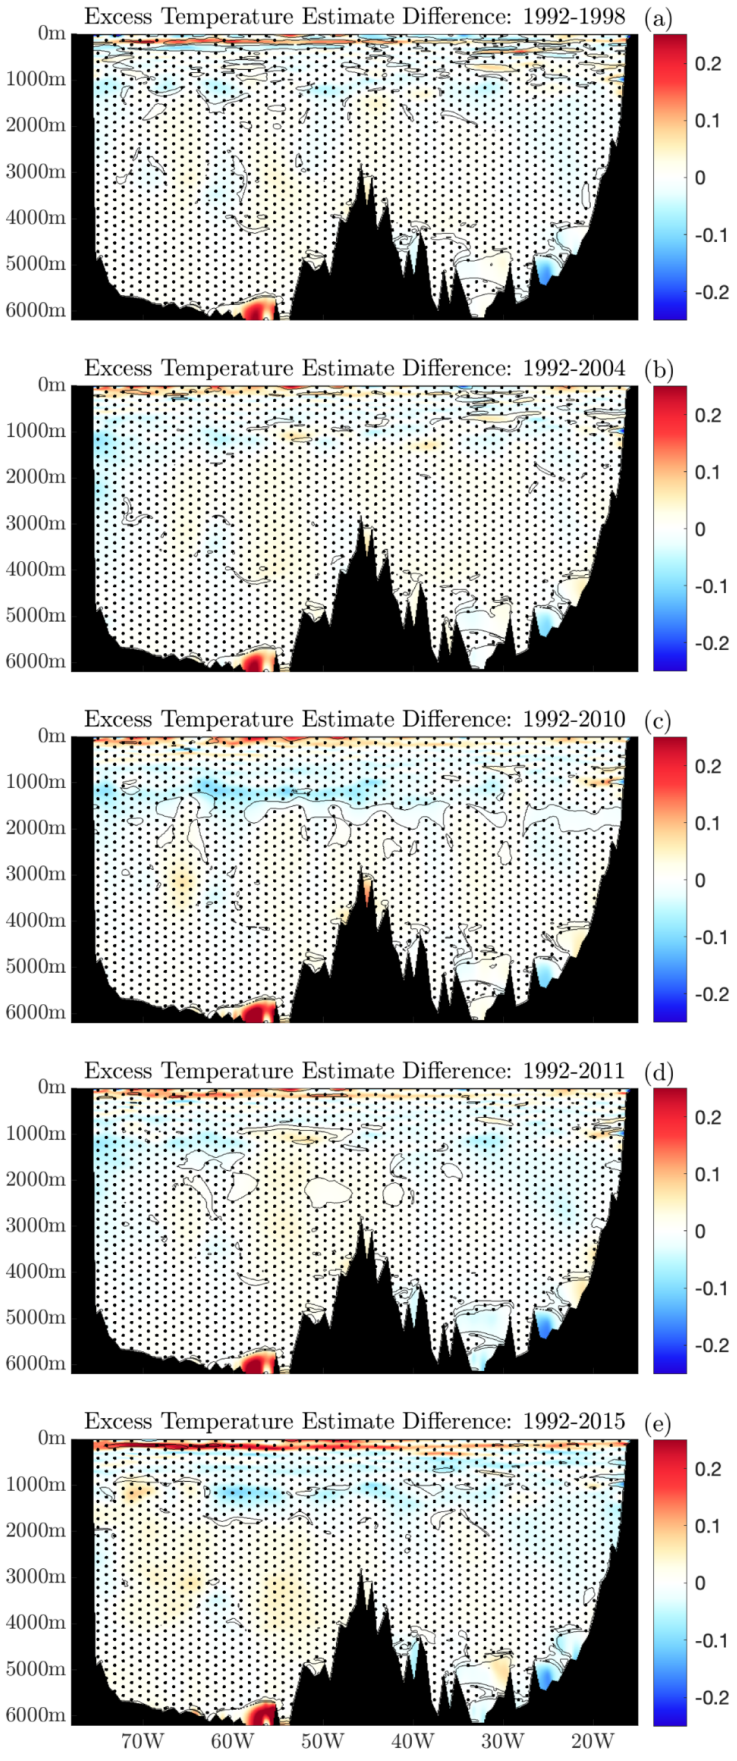

FIGURE 5.11: The difference in excess temperature estimates between the two methods of calculation: the assumption of constant α_T (=0.017K/μmol/kg) and temperature dependent α_T (= 0.025 - 0.0005θ K/μmol/kg)). Regions where the two estimates are indistinguishable at 2σ are indicated with stippling. Values presented are calculated as constant estimate - variable estimate.

case for both the constant and variable α_T estimates. In the case of the variable α_T estimate, it could be argued that this is the result of the temperature dependence of α_T: as this parameter is lower in warmer waters, we expect a reduction in excess temperature accumulation relative to excess DIC ($\approx C_{anth}$) accumulation in the warmer surface waters. However, as this behaviour is also apparent with the fixed α_T approximation, it does not appear to be the result of the temperature dependence of α_T, although the issue is slightly exacerbated by this parametrisation. Section 5.4.3 will show that this behaviour is in fact due to the 1992 cruise being performed in summer, rather than during the winter as for all the other cruises. The issue thus is the result of the reduced solubility of DIC in warmer water (summer), and so this behaviour is also replicated in C_{anth} reconstructions.

From the conceptual framework of excess and redistributed temperature, it is clear that seasonal variability in surface and near subsurface temperatures are properly described as excess temperature, rather than redistributed temperature. We might therefore expect to see a degree of cooling in the near subsurface excess temperature for all cruises subsequent to 1992, as the scale of seasonal variability here exceeds excess temperature accumulation. However, this is not the case: in fact, most subsurface cooling is instead captured by the redistributed temperature (see for example Figure 5.8). This implies that this decomposition in fact describes the majority of seasonal variability as redistribution, rather than excess. Though this is not consistent with the precise definition of excess and redistributed temperature, it is convenient: that this decomposition appears to act to deseasonalise the accumulation of excess temperature by default is a convenient property, as it makes identifying trends in excess temperature more straightforward. However, this deseasonalisation is not perfect, and so we see reduced excess temperature accumulation in the near subsurface here.

The parametrisation of error for excess temperature accumulation in the case of temperature dependent α_T, and its correspondence to errors for constant α_T is now considered. A linear correlation between total excess temperature accumulated in 2015 and the error associated with these values yields

$$\sigma_{\theta_e} = 0.05 + 0.12|\theta_e|, \tag{5.29}$$

with this linear approximation yielding an R^2 value of 0.87, and having added the error associated with gridding and the vertical only approximation from Equation 5.28 (the intercept of this linear approximation is negligible, approximately 5×10^{-4}). This linear approximation implies that the error in excess temperature, calculated using the assumption of temperature dependent α_T, is slightly smaller than that of the error as calculated under the assumption of constant α_T.

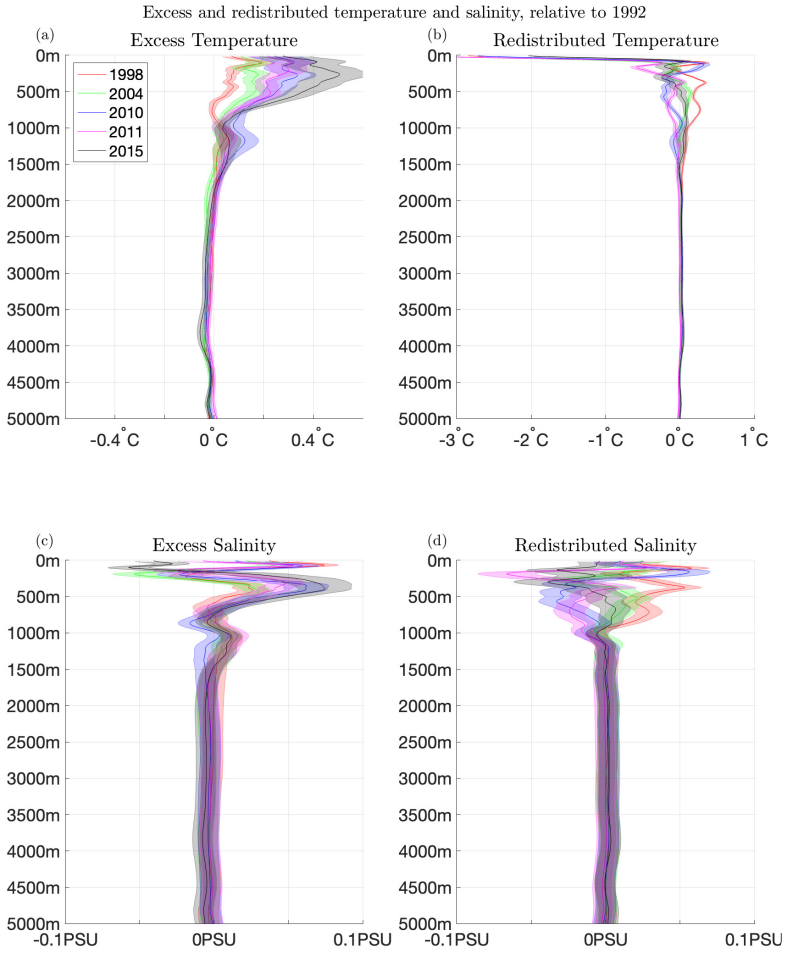

FIGURE 5.12: The mean profiles of excess (a) and redistributed (b) temperature, and excess (c) and redistributed (d) salinity, for the 5 occupations since 1992. 95% confidence intervals are represented by shading, and scales are not shared for excess and redistributed components.

The smaller coefficient relating the uncertainty in excess temperature value to their absolute value can be explained as follows: in the case of a temperature dependent α_T, uncertainty in excess temperature accumulation depends on the initial temperature, and cooler waters exhibit larger errors due to their larger α_T values. As such, the error in the warmer and younger surface waters is smaller than under the assumption of constant α_T, and vice versa in the cooler, less recently ventilated deep waters. This leads to reduced total error in warmer waters and therefore a weaker relationship between the accumulated excess temperature and it's error. Whilst this result appears appropriate here as a parametrisation, it therefore may underestimate the uncertainty in excess heat uptake at higher latitudes, where cooler waters are more recently ventilated: explicit testing will be required in such locations in order to better understand the appropriateness of this parametrisation.

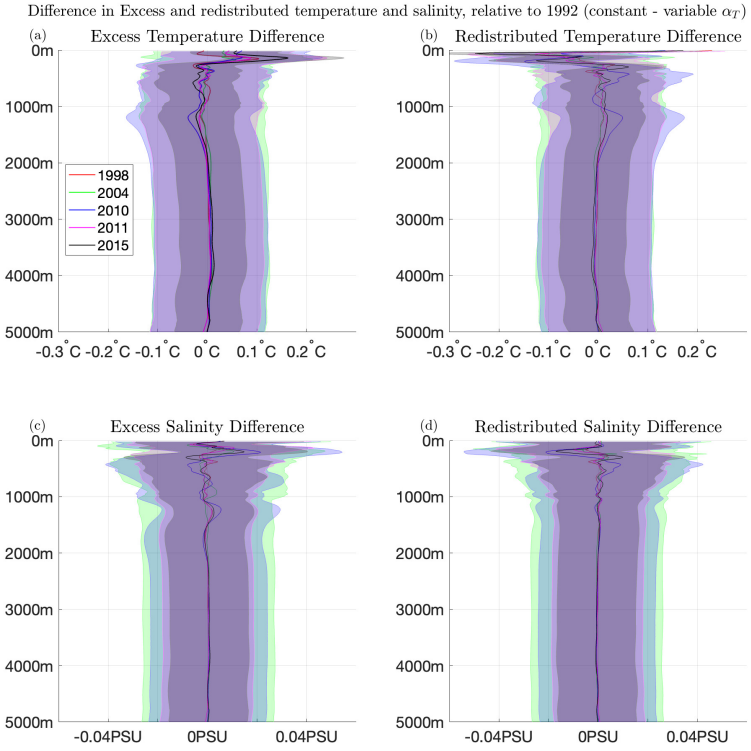

FIGURE 5.13: The difference in mean profiles of excess (a) and redistributed (b) temperature, and excess (c) and redistributed (d) salinity, for the 5 occupations since 1992. 95% confidence intervals are represented by shading, and scales are not shared for excess and redistributed components. Differences are calculated as constant α_T - variable α_T estimates.

5.4.3 Excess DIC and Anthropogenic Carbon

As alluded to in Section 5.3, we may compare our estimates of excess DIC accumulation and C_{anth} accumulation in order to empirically validate this decomposition. This may be performed by comparison of the rates of accumulation of excess DIC and C_{anth} accumulation, as well as total DIC accumulation, in order to understand how and why these quantities differ. Recall Equation 5.5:

$$\Delta \text{DIC}_e = \frac{\Delta\theta - \beta_T \Delta \text{DIC}}{\alpha_T - \beta_T} \tag{6.28}$$

Williams et al. (2021) notes that on short enough timescales, we can consider the excess and anthropogenic carbon fields to be almost identical, and this is supported by Figure 3.2, at least on a global basis: the γ factor used in Chapter 3 generally does not change significantly enough on timescales of 20 to 30 years that we would expect excess carbon and anthropogenic carbon to diverge. This is also discussed in Section 5.4.3.

Summarising these arguments, on short timescales, we expect the divergence between anthropogenic and excess carbon to be small enough that we can consider the two to be essentially equal (see Figure 1.8). As there are four C_{anth} reconstructions available for the repeat occupations of A05 (TTD, TrOCA, ΔC^* and ϕC_T^0), we may therefore compare the accumulation of excess DIC and C_{anth} in order to empirically validate this reconstruction. Though the period is slightly shorter to excess DIC (excess DIC spans 1992-2015, the back calculation methods span 1992-2011, the TTD method spans 1992-2010), these periods are broadly comparable. Provided that excess DIC well approximates C_{anth}, the reconstruction of excess temperature should therefore be accurate: Section 5.4.2 showed that the excess temperature fields obtained were not significantly different when calculated under the approximation of constant or temperature dependent α_T, and so a close approximation of C_{anth} by excess DIC implies the accurate reconstruction of excess temperature. For simplicity, only the excess DIC calculated under the assumption of constant α_T is investigated here, though the results are similar for the excess DIC under the assumption of temperature dependent α_T: unlike excess temperature, excess DIC does not depend strongly on the value of α_T.

Figure 5.14 shows the longitudinally averaged rates of C_{anth}, DIC, and excess DIC accumulation, along with 95% confidence intervals. As with the uncertainties in excess temperature and salinity, the errors in Figure 5.14 are calculated by bootstrapping: observations are uniformly perturbed according to a Gaussian distribution for each cruise, and the trend in these perturbed values calculated. For DIC, observational error is approximately $\pm 2\mu$mol/kg, and so perturbations are sampled from a Gaussian distribution with mean 0 and standard deviation 1μmol/kg. For C_{anth} reconstructions, observational error is approximately $\pm 6\mu$mol/kg, and so perturbations are sampled from a Gaussian distribution with mean 0 and standard deviation 3μmol/kg. For excess DIC, a uniform perturbation within the previously calculated uncertainty, divided by α_T was applied (it is trivial to show $\sigma_{\theta_e} = \alpha_T \sigma_{DIC_E}$). Linear fits were then performed at every point, and the standard deviation of these linear fits at each point in an 1000 member ensemble used to calculate uncertainties. Note that whilst perturbations are uniform within each cruise, a unique perturbation is applied to each cruise (or the error in trends would collapse to zero).

In general, excess DIC (red line) closely follows total DIC change (blue line). All back calculation methods (ΔC^*, ϕC_T^0 and TrOCA) give approximately equal estimates of the rate of C_{anth} accumulation, with the TTD method showing a profile which is noticeably different to all others, though still in agreement to back calculation methods to within observational uncertainty. Though excess DIC more closely agrees with the trend in total DIC than the C_{anth} reconstruction methods, the difference between the excess DIC and the back calculation methods is generally far less than that between the back calculation methods and the TTD reconstruction. Finally, as previously noted, for all back calculation reconstructions, as well as the excess and total DIC, the trend is largest at

approximately 200 metres, and reduced above this. This is due to the temperature dependence of CO_2 solubility: see Guallart et al. (2015).

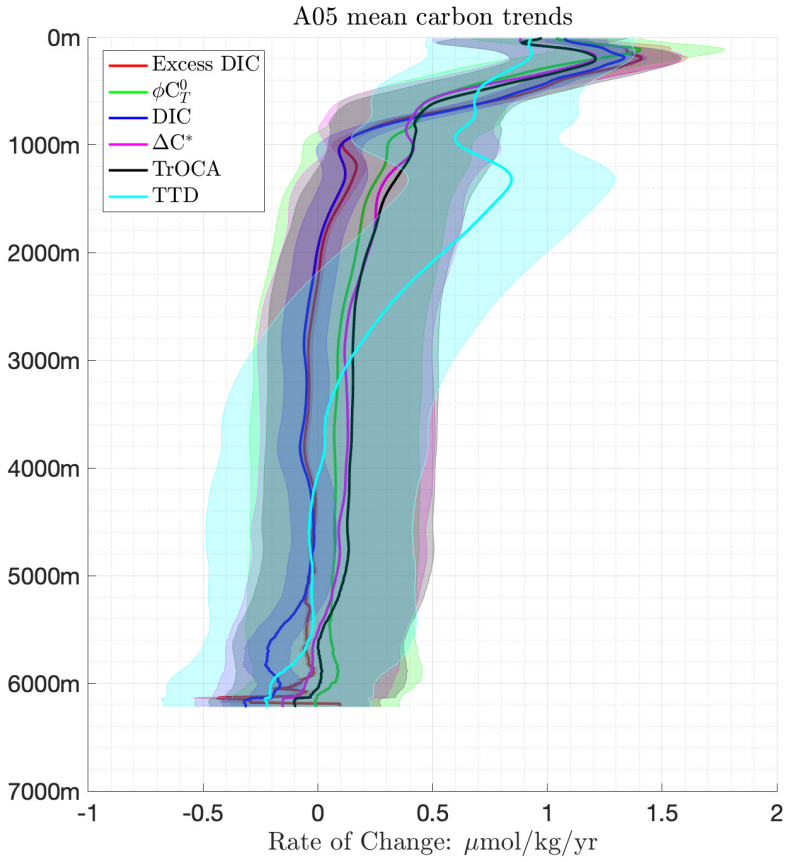

FIGURE 5.14: The longitudinal mean rate of accumulation of each of the four C_{anth} reconstructions, Excess DIC, and total DIC. 95% confidence intervals are indicated with shading.

Worth noting is the significantly lower error in excess DIC than in any of the C_{anth} reconstruction methods or the total DIC. This appears to be the result of the geometric nature of the estimation technique.

Finally, it may appear that this test is somewhat circular: α_T is derived by considering increases in temperature and C_{anth}, with the excess DIC obtained using this α_T compared to the C_{anth} fields used to generate them. Thus, it would appear that the good agreement between the fields is merely a test of internal consistency. However, it will now be shown that excess DIC accumulation is in fact relatively insensitive to the precise value of α_T.

To do so, estimates of the section mean accumulation of excess temperature and carbon, over the top 1000m, were recalculated, using seven different values of α_T, ranging from 0.001 to 0.04K/(μmol/kg): a range approximately equal to the 'tuned' value of 0.017K/(μmol/kg), and far greater than the error in this value. Thus, these represent an

extremely pessimistic uncertainty (and unrealistically so) associated with excess temperature and DIC uptake due to the uncertainty specifically in the chosen value of α_T. This is shown in Figure 5.15. Whilst excess temperature estimates (b) are highly sensitive to the value of α_T, with those calculated using using $\alpha_T = 0.001$Kμmol/kg showing virtually no excess heat accumulation and those calculated using an unrealistically large value showing over half a degree, the sensitivity of excess DIC is much lower. By 2011, we find a spread of approximately 5μmol/kg, compared to a mean value of approximately 17μmol/kg. Thus, though the precise value of excess DIC obtained is dependent on the value of α_T chosen, the dependence is weak, and thus it appears that the excess DIC field does contain novel information, rather than just acting as a check on self consistency. Again, note that this is an extreme range of α_T values: over ten times the calculated uncertainty.

5.4.4 Rates of Accumulation of Excess Temperature and Salinity

Whilst this chapter so far has been primarily concerned with the total accumulation of excess heat and salinity relative to 1992, Chapter 6 will be concerned primarily with the rates of accumulation of excess temperature and salinity. It would therefore be preferable to derive an error parametrisation for rates of change of excess temperature and salinity, rather than simply an error as a function of the total excess temperature and salinity. Such a parametrisation is now investigated.

The section mean rate of accumulation of excess temperature and salinity, for both the temperature dependent and independent parametrisations of α_T, are shown in Figure 5.16. The peak rate of excess temperature accumulation is found at approximately 200 metres for the temperature independent parametrisation, and approximately 300 metres for the temperature dependent parametrisation. The temperature independent parametrisation has a slightly larger peak value (0.021 degrees/year versus 0.020 degrees/year), but both peaks are well within the 95% confidence intervals of one another (confidence intervals are calculated the same way as for excess DIC accumulation). In addition, the secondary peak in excess temperature accumulation at approximately 1200 metres is slightly enhanced by the temperature dependent parametrisation, though again both values fall within the uncertainty of one another. The 95% confidence intervals of the two trends overlap everywhere, and it is only between 100 and 200 metres that the mean trends are not within the confidence intervals of the other parametrisation.

The rates of accumulation of excess salinity (Figure 5.16b) are less affected by the choice of temperature dependent or independent parametrisation, with no clear pattern of differences in the two values. As with excess temperature, mean rates of accumulation do not significantly differ anywhere, and the agreement between the two parametrisations is better than is the case for temperature: both parametrisations show a peak in excess

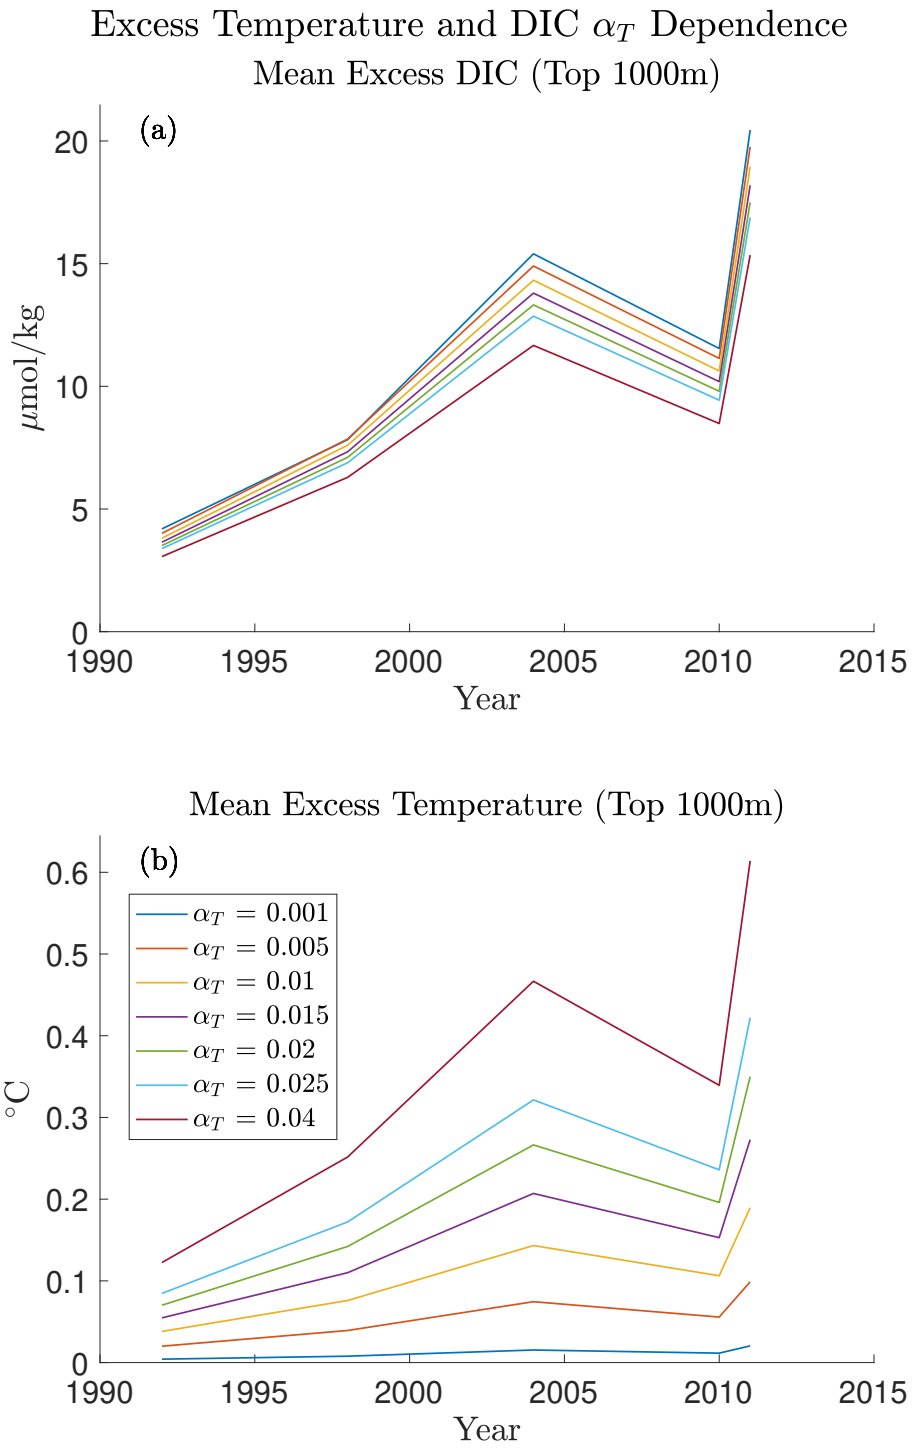

FIGURE 5.15: Section longitudinal and top 1000m mean values of excess DIC accumulated since 1992 (a), and excess temperature (b), for a range of α_T values much larger than the calculated uncertainty in Figure 5.5. Both plots are scales such that the y axes run from 0 to 1.05 times the largest value obtained.

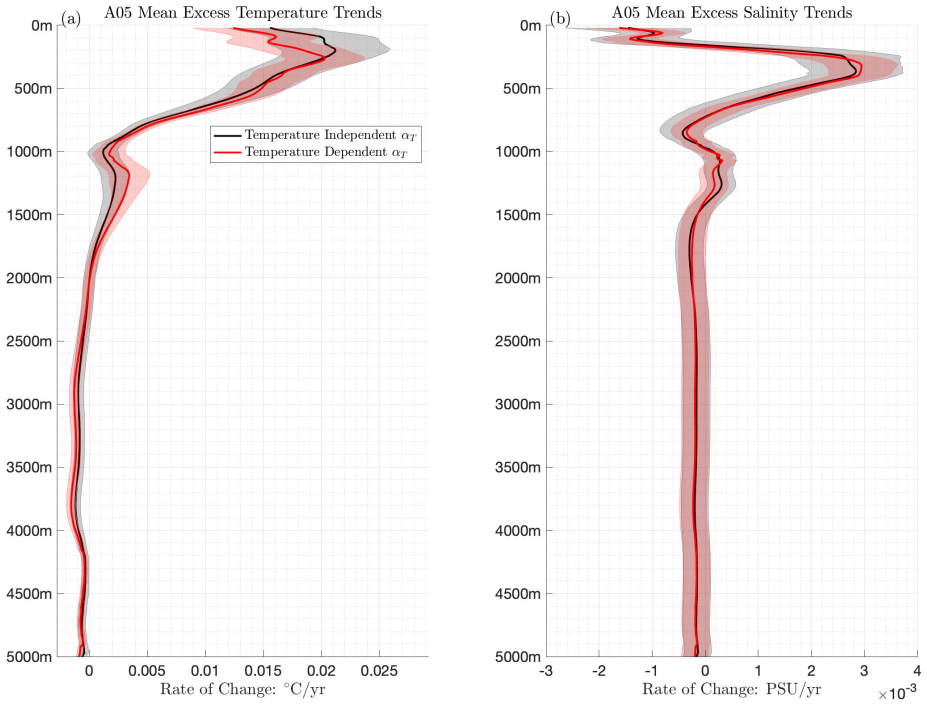

FIGURE 5.16: The section mean trends in excess temperature (a) and salinity (b) as a function of depth. Both the temperature dependent (red) and independent (black) parametrisations of α_T are shown. 95% confidence intervals are indicated with the shaded regions.

salinity accumulation of approximately 1.8×10^{-3}PSU/year at deths between 300 and 400 metres.

Interestingly, despite the statistically insignificant accumulation of excess temperature at depths below a thousand metres, both parametrisations show a statistically significant cooling at depths between approximately 2500 and 4000 metres. This corresponds approximately to the depth of lower North Atlantic Deep Water, formed in the Nordic seas, and may therefore indicate changes in the properties of water masses formed in this region.

In the case of temperature, it is visually clear that there is a relationship between the mean value and the size of the uncertainty, however, the uncertainty on salinity trends appears to be relatively constant. Figure 5.17 shows that this is in fact the case. For the constant α_T parametrisation, the uncertainty in the trend may be parametrised as

$$\sigma_{\dot{\theta}_e} = 0.010|\dot{\theta}_e|, \tag{5.30}$$

where $\dot{x} = dx/dt$. This error parametrisation yields an R^2 value of 0.88. For the case of temperature dependent α_T, the relationship

$$\sigma_{\dot{\theta}_e} = 0.0084|\dot{\theta}_e|, \tag{5.31}$$

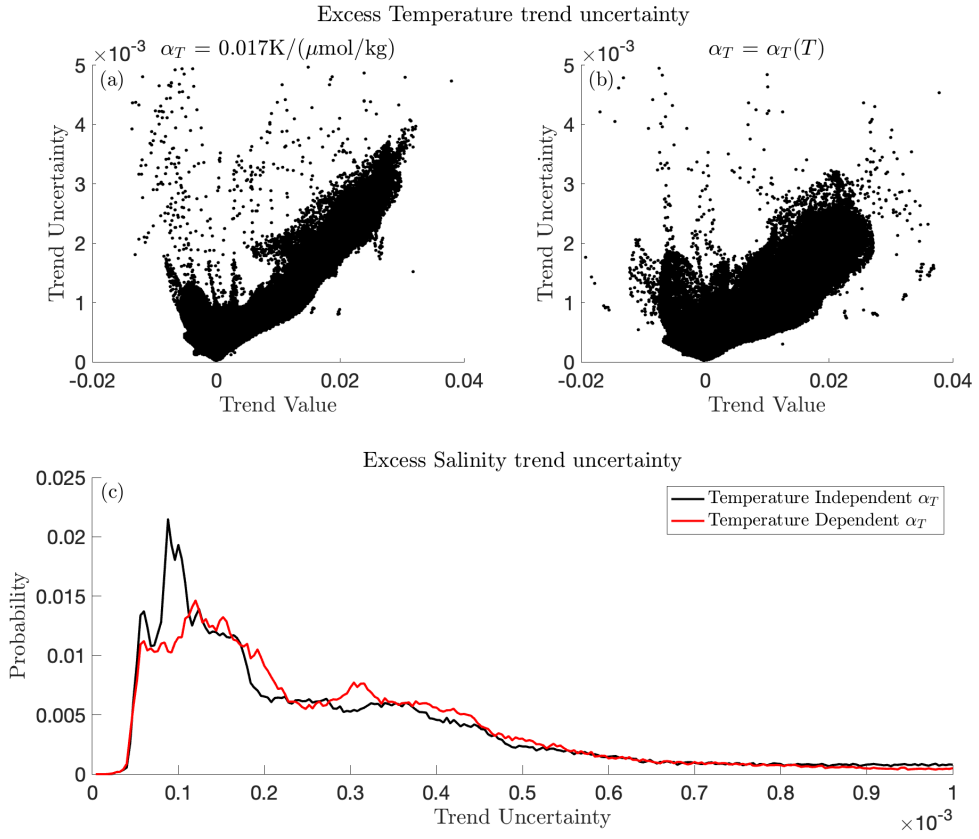

FIGURE 5.17: Uncertainties in trends in excess temperature as a function of the trend value for the fixed temperature parametrisation (a), the temperature dependent parametrisation (b), and the distribution of uncertainties on trends in excess salinity for both parametrisations (c).

with an R^2 value of 0.77. In both cases, the intercept is negligible: on the order of 10^{-4}. As the R^2 value is large in both cases, and the size of the relationship is similar (10% of the trend, 8.4% of the trend), this result appears robust: provided it can be replicated in other sections, it should be readily applicable to global fields. Additionally, the relationship between the total accumulated excess temperature under the temperature dependent α_T parametrisation is also approximately 10% (Equation 5.29). This suggests an error of approximately 10% is a relatively robust feature of this estimate of excess temperature accumulation.

In the case of salinity, the R^2 values are small for both the temperature dependent and independent parametrisations of α_T: 0.27 and 0.26, respectively. Thus, a trend dependent error parametrisation is not further pursued for salinity. However, it can be seen from Figure 5.17, panel c, that both the temperature dependent and independent parametrisations have relatively similar distributions in their uncertainty, with both peaking at approximately 10^{-4} PSU/yr (this peak is significantly stronger and at a slightly lower value for the temperature independent parametrisation). Virtually no trends have uncertainties of 7×10^{-4} PSU/yr or larger. The distribution of

errors in the temperature independent parametrisation has mean 4.97×10^{-4} and median 2.84×10^{-4}, and the temperature dependent parametrisation has mean 4.51×10^{-4} and median 2.60×10^{-4}. As a result, a characteristic uncertainty of 5×10^{-4}PSU/yr in excess salinity trends ought to well capture the associated uncertainty, provided these results can be replicated in other sections.

5.5 Conclusions

In this Chapter, a novel decomposition of temperature and salinity changes into their excess and redistributed components has been demonstrated. This decomposition can be thought of as combining the carbon based methodologies of Bronselaer and Zanna (2020), Williams et al. (2021) and Turner et al. (2022) (Chapter 3) in order to remove the necessity for the decomposition of DIC before using changes in DIC to decompose changes in temperature and salinity. It also can be considered the analogue to the decomposition used by Clément et al. (2022) in temperature-carbon space, instead of temperature-salinity space.

By applying this technique to repeat hydrography at A05, we are able to identify excess temperature in the upper 1000m on timescales of only six years, despite the signal of redistributed temperature far exceeding the excess temperature signal here. Though impressive, this is not unprecedented: Zika et al. (2021) found a similar result when applying their minimum transformation method to estimate material heat content change globally. This suggests that the application of techniques which aim to separate changes in temperature due to excess heat content and due to circulation variability may in fact be able to identify the thermodynamic warming of the ocean on such short timescales: a remarkably powerful result. Both Zika et al. (2021) and Bronselaer and Zanna (2020) note that a fundamental issue with historical simulations of heat content change is that changes in local ocean heat content have been dominated by circulation variability, rather than the accumulation of excess heat content. Continued progress in understanding how the excess and redistributed heat content of the ocean have and continued to evolve therefore represents an enormous opportunity for improved understanding of the physical response of the ocean to anthropogenically forced climate change. A key benefit to the method presented here is that its assumptions are orthogonal to those of Zika et al. (2021), and therefore represents a completely independent method for observationally estimating excess heat content change in the ocean. Thus, its application to global ocean heat uptake represents an opportunity for independent verification of the estimates of Zika et al. (2021) and Zanna et al. (2019).

Interestingly, Parrilla et al. (1994) found warming at A05 over the period 1957-1992, peaking at approximately half a degree of warming at depths of approximately 1000m,

with a consistent warming signal across the entire hydrographic section at depths between 700 and 2000m. This signal is remarkably horizontally uniform, and thus plausibly represents the addition of excess heat to the North Atlantic at this location between 1957 and 1992. A subsequent spice/heave analysis by Bryden et al. (1996) found that over this period, both warming and salinification are found at these depths between 1981 and 1992: a signal consistent with the addition of excess heat. However, prior to 1981, most of the changes in temperature were explained by the downward displacement of isopycnals, though a spice heave analysis may be unreliable due to potential issues with the 1957 salinity measurements. Whilst the warming and salinification found by Bryden et al. (1996) is on isopycnal surfaces and thus is not directly comparable, there are other indications that changes in water mass properties in the Atlantic tend to be density compensating (for example Lowe and Gregory (2006), Mauritzen et al. (2012)). Thus, it is plausible that there was significant addition of excess heat to the Subtropical North Atlantic prior to 1992, with the structure of warming found between 1981 and 1992 (Parrilla et al. (1994) Figure 1c) strongly indicating the presence of excess temperature. Unfortunately, due to the lack of DIC measurements prior to 1992, this cannot be explicitly verified within the framework developed here. In addition, a reconstruction of excess temperature including this prior work, were DIC observations available, would likely require an extension of the framework developed here: as timescales become sufficiently large, it is likely that it would become necessary to account for the time dependence of α_T.

In the same way that Turner et al. (2022) was able to extend the decomposition to salinity thanks to the generality of the relationship utilised, this decomposition may also be used to decompose changes in salinity into excess and redistributed salinity. As with those findings, changes in salinity appear to be less dominated by the redistributed component of salinity than is the case for temperature. This appears to be an emerging pattern: whilst changes in temperature are dominated by the redistribution of the background temperature field, changes in salinity are often contributed equally to by both excess and redistributed salinity, or indeed predominately excess salinity. It is currently unclear whether this is the case due to the smaller dynamical range of background salinity values, or because excess salinity is fundamentally an atmospheric, rather than oceanic phenomenon (as is the case for excess temperature), and thus the energy requirement for substantial changes is lower. Whether the former or latter is the case, understanding excess and redistributed salinity represents an opportunity to identify the propagation of atmospheric anomalies into the ocean in a way that may be less straightforward for temperature.

With regards to the choice of whether to parametrise α_T as a function of the initial temperature, as the results of Chapter 1, Section 2.2.1 suggest, or a single value in the manner of Bronselaer and Zanna (2020), the results presented here suggest that the decomposition is relatively insensitive to the choice of whether to use a constant α_T

or not. However, as previously noted, cooler waters at this location are only found at depth, and are therefore relatively poorly ventilated. At higher latitudes, where cooler waters are better ventilated, this result will need to be explicitly tested.

The trend in excess temperature peaks at approximately 0.02°/year at depths of 200 to 300 metres, with a secondary peak at depths between 1000 and 1500 metres. At depths between 2500 and 4000 metres, a small but statistically significant negative trend in excess temperature is seen, likely reflecting changes in properties of water masses formed in the Nordic seas. Excess salinity reveals a slight freshening trend in the top 100 metres, with a strong salinification, peaking at approximately 2×10^{-3} PSU/year at depths of 300 to 400 metres. No statistically significant trends in excess salinity are seen at depth.

Comparison of the excess DIC field with 3 back calculation C_{anth} estimates, total DIC change and the TTD reconstruction of anthropogenic carbon reveals that patterns of excess DIC accumulation are broadly consistent with the accumulation of anthropogenic carbon at this location. This is unsurprising, as we expect the two changes to be approximately equivalent on short timescales (Williams et al. (2021), Turner et al. (2022)). However, it is reassuring, providing evidence for the validity of this approach. Additionally, the simplicity of this approach, orthogonality of the assumptions employed in the decomposition of DIC changes by it, and its close resemblance to the back calculation C_{anth} fields suggest utility as a 'quick and easy' approximation of anthropogenic carbon: a potential new way to assess changes in anthropogenic carbon inventories.

Though the definitions of excess and redistributed temperature mean that seasonal variability in upper ocean temperature are best described as excess, not redistribution driven changes, this is not the case in this decomposition: the results presented in this chapter suggest instead that this reconstruction will class seasonal variability predominately as redistribution. This is convenient, as the implication is that this decomposition will therefore automatically deseasonalise changes in excess temperature. Such a property means that trends in excess heat content ought to be more readily identifiable than they might otherwise be.

In the case of excess temperature, we can parametrise errors as a simple linear function of the excess temperature value. Provided it can be shown that this parametrisation holds on other hydrographic sections, this therefore represents an opportunity to estimate the error associated with a global field of excess heat uptake, without needing to resort to the computationally expensive bootstrapping procedure. This parametrisation appears to result from the geometric nature of the decomposition, and thus there is no reason to suspect it should not be the case, however, it will be necessary to validate this explicitly. Similarly, we can parametrise trends in excess temperature in this fashion: the error in trends is approximately 10% of the trend, regardless of whether α_T is prescribed to be constant or allowed to vary with temperature. The uncertainty

in excess salinity trends is typically on the order of 5×10^{-4}PSU/year. However, these error parametrisations on trends will also require explicit testing.

Finally, by considering the error in our parameters by independently applying and comparing the results from the techniques of Bronselaer and Zanna (2020) and Turner et al. (2022), it has been possible to obtain robust error estimates for the parameters associated with this method, and thus the total error inherent in the decomposition applied here. Interestingly, it is the case that the error in excess DIC at depth is generally much smaller than the error in total DIC change at depth: this reflects the power of this technique in removing noise from signals. It also implies that the uncertainty must instead be contained within the redistribution signal. Whilst the error calculation performed here does not explicitly account for the gridding error, by assuming *a priori* that signals of excess temperature at abyssal depths should not be detectable, this may be empirically quantified. These results suggest an error in excess temperature of 0.05 degrees, plus 12% of the total accumulated excess temperature. However, again, it will be necessary to validate this on other hydrographic sections before applying this approach to global fields of excess temperature.

Chapter 6

Understanding global ocean excess heat uptake and salinity changes using the GLODAP dataset

6.1 Introduction: Previous studies and context

In the previous chapter, a technique for decomposing changes in temperature, salinity and carbon into changes due to their excess and redistributed components was developed, and its implementation, limitations and the associated uncertainties were investigated in detail at the A05 hydrographic section. However, the primary motivation for developing such a technique is not simply to understand the accumulation of excess heat in the subtropical North Atlantic, but to improve understanding of global ocean heat uptake: this requires applying this technique globally. As such, this chapter will be concerned with the global application of this technique. However, such an extension is not straightforward: as such, the majority of this Chapter will be concerned with reliably extending the analysis of Chapter 5 in order to generate global gridded fields of excess heat and salinity accumulation.

There are three previous studies which the results presented in this chapter will be primarily comparable to: Zanna et al. (2019), Gruber et al. (2019) and Zika et al. (2021). Zanna et al. (2019) and Zika et al. (2021) are concerned with ocean heat uptake, and Gruber et al. (2019) with the storage of C_{anth}. Zanna et al. (2019) investigated the change in the global ocean heat inventory over the full historical period from preindustrial, whereas Zika et al. (2021) investigated global ocean heat content change only over the period 2006-2017. Gruber et al. (2019) investigated the change in the global ocean C_{anth} inventory over the period 1994-2017. Some salient results from these studies are discussed below.

Zanna et al. (2019) estimate a total heat content change of the global ocean of 436 ± 91ZJ over the period 1871-2017. They also estimate a change of 153 ± 44ZJ over 1990-2015, a period more comparable to that of Gruber et al. (2019) (1994-2017) and Zika et al. (2021) (2006-2017). In addition to the globally integrated heat content change, Zanna et al. (2019) use a Green's Function based reconstruction to estimate 'passive' warming over the period 1955-2017. This passive warming is essentially the heat content change due to the addition of a Passive Anomalous Tracer as described in Chapter 1, Section 1.2.1, and so is closely related to the excess heat content of the global ocean: the two can be thought of as essentially equivalent.

Zika et al. (2021) instead report 'material heat content change', total heat content change, and calculate redistributed heat as the difference between the two of these. Material heat content change can be also be thought of as essentially the same as excess heat content change, but is described as a material change as it is calculated using water mass space principles. In practical terms, this will result in the homogenisation of material heat content throughout each water mass. However, it is demonstrated that their results are in general relatively insensitive to the resolution of their water mass bins, and thus this ought not to constitute a source of anomalous smoothing of their material heat content change estimates. They report their results in terms of a heating rate, rather than heat content changes, finding a global ocean heating rate of 398 ± 81TW over their time period. In the units of Zika et al. (2021), the warming of 153 ± 44ZJ calculated by Zanna et al. (2019) over 1990-2015 is equivalent to a global ocean heating rate of 184 ± 78ZJ.

6.2 Methods

In this chapter, data from the full GLODAPv2 (Lauvset et al., 2021) is used. It is applied in two complementary ways:

1. Trends in excess temperature and salinity are computed on the repeat hydrographic sections defined by GO-SHIP and included in their toolbox, as well as the Julia toolbox described in Appendix A. This allows for the direct examination of trends in a number of ocean regions. It also allows for the explicit testing of the parametrisation of error produced in Chapter 5 in other regions of the ocean.

2. Gridded global fields of trends in excess temperature and salinity are computed, using a standard 1 degree resolution grid on 33 vertical levels. This is the same grid as used in Sabine et al. (2004), for GLODAP climatological data (Lauvset et al., 2016) and by Gruber et al. (2019).

As the procedure used to calculate excess temperature and salinity on other hydrographic sections is identical to that used for the A05 section, discussed in detail in the

previous chapter, it is not further discussed here. However, the process used to generate the gridded 3D fields is now described in detail, as it differs significantly.

6.2.1 Generating global gridded fields for excess temperature and salinity

As with the hydrographic sections, the gridded global field is produced using the DIVA software package (Barth et al., 2014). Two gridding procedures were followed: one inspired by the procedure of Lauvset et al. (2016), but differing in several key ways as a result of the different nature of the estimate. This will be referred to as a pointwise estimate. Alternatively, a global estimate may be generated by interpolating the rates of change obtained from the repeat hydrographic occupations of ocean sections contained within the GO-SHIP Easy Ocean toolbox: this is similar to the method of Sabine et al. (2004). This will be referred to as the section interpolation.

6.2.1.1 Pointwise Estimate

The pointwise gridding is performed as follows:

1. An empty 1 degree by 1 degree grid is initialised. This grid has 245 vertical levels: one grid cell every ten metres from the surface to 2000 metres, followed by one grid cell every 100 metres from 2100 metres to a maximum possible depth of 6500 metres. This oversampling helps to ensure that once estimates of excess temperature and salinity where $\alpha_T \approx \beta_T$ are removed, sufficient data remain to reinterpolate over these gaps smoothly.

2. In each grid cell occupied by ocean, the subset of GLODAPv2 data within 1 degree of latitude and longitude the midpoint of the grid cell are extracted. The variables extracted are `G2tco2`, `G2tco2f`, `G2theta`, `G2salinity`, `G2salinityf`, `G2pressure`, `G2maxsampdepth`, `G2cruise`.

3. Flagged data (DIC and Salinity) are removed. This is performed using a function contained within the GLODAP Easy Ocean toolbox.

4. Each unique cruise for which this latitude, longitude box contains observations is then determined. For each of these unique cruises, the maximum sample depth is checked. If at this point it is found that less than 3 unique cruises occupied at least 80% of the full cruise profile, as calculated using the maximum sample depth of all cruises combined, then the data from this latitude, longitude box are discarded.

5. If at least 3 cruises have sampled at least 80% of the full water column, the data from all cruises in this box are combined, and these combined data used to generate a mean profile of temperature, salinity and DIC within this grid cell using a

one dimensional DIVA interpolation. Unlike the gridded hydrographic sections, the vertical correlation length is prescribed rather than calculated here.

6. Once a background profile for a grid cell has been generated, each cruise is interpolated in turn. To do so, GLODAPv2 recommended adjustments for DIC values for each cruise are first applied, using functions contained within the GLODAP Easy Ocean toolbox. The background profiles are then subtracted from the observations for each cruise to obtain a profile of anomalies, which are then interpolated, before the means are added back to the anomalies to produce gridded profiles of temperature, salinity and DIC for each cruise.

7. The mean date for each cruise is calculated. If fewer than 3 occupations containing DIC observations are found in this cell, or the time separation between the first and final cruises is less than ten years, then no further calculations are performed.

8. If all conditions are met, the excess and redistributed temperature and salinity for this column is then calculated at each of the 245 depth levels, using Equation 5.17. The eigenvalues are then calculated, and the excess and redistributed temperature and salinity fields are then masked using the same criteria as used for the A05 section.

9. For the vertical levels which are not masked, a linear fit is then performed to calculate the mean rate of accumulation of excess and redistributed temperature and salinity at each depth. This yields a semi-regular grid of rates of change of excess and redistributed temperature and salinity: all values are found on a regular grid, but it is sparsely populated.

10. A mean profile for the rate of change of excess temperature on each of the 245 depth surfaces of the initial was then calculated, in order to be used as a first guess for the DIVA gridding algorithm. This mean profile was linearly interpolated onto the 33 depths surfaces used in the final grid. As all but two of these 33 depths surfaces were also included on the 245 depth grid, linear interpolation was preferred over more sophisticated techniques (note these two points are not neighbours and so loss of accuracy should be minimal). The DIVA gridding algorithm will relax its final gridded field towards this first guess in the absence of data, and so it is necessary to provide a realistic initial guess. At depths greater than 1000m, the mean profile was artificially relaxed to zero: this was achieved using a depth dependent scaling factor. This is described in detail below. As we expect the global integrals of redistributed temperature and salinity to be zero, and excess salinity approximately zero, no mean value was calculated for these fields: thus, these fields will be relaxed towards zero by the DIVA gridding algorithm.

11. In the case of excess temperature accumulation rates, the mean profile was subtracted, and the sparsely populated array of anomalies interpolated to produce a 3D gridded field. The mean profile was then added back in. In the case of excess salinity and redistributed temperature and salinity, the sparsely populated array of rates of change are directly interpolated to produce 3D gridded fields describing their rates of change.

Further relevant details about the gridding process are now provided: parameters, data cleaning, and so forth.

Unlike the analysis performed on the A05 hydrographic section, for which correlation lengths were calculated explicitly for the data used, for the global estimates of excess heat uptake, correlation lengths were prescribed, following the procedure used in Lauvset et al. (2016). They state that the data distribution in the mapped and merged GLODAPv2.2016 data product leads to an optimised correlation length of order 25°. However, they choose instead to use a globally uniform correlation length of 7° in the latitudinal direction and 14° in the longitudinal direction: this shorter correlation length is chosen as a balance between the better resolution of fronts and other features associated with shorter correlation lengths, and the smoother data fields and smaller mapping errors associated with longer correlation lengths. The correlation length is doubled in the longitudinal correlation to account for the fact that the oceans generally mix more easily in the zonal, rather than meridional direction.

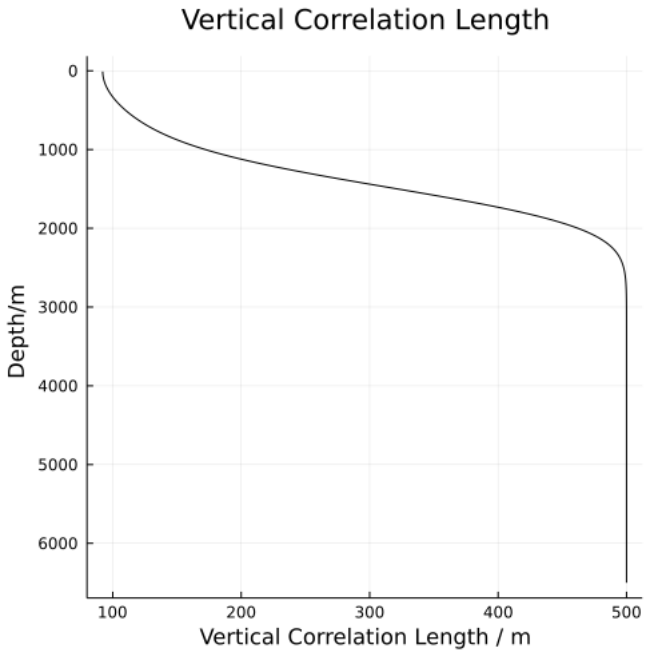

FIGURE 6.1: The vertical correlation length used in the calculation of gridded excess temperatue and salinity fields.

As excess heat is known to be very spatially uniform from both modelling experiments and the observational estimate of Zika et al. (2021), it is anticipated excess heat fields

have extremely long correlation lengths. As such, a longitudinal correlation length of 50° and a latitudinal correlation length of 25° are used here. For the vertical correlation length, a length of 70 metres was prescribed in the top 200m, 100m between 200m and 750m, 150 metres between 760m and 1500m, and 500 metres below this. These lengths were found to be a reasonable compromise between smooth fields and vertical resolution, as discussed above for horizontal correlation lengths, and were determined from manual inspection of the correlation lengths of excess temperature and salinity fields obtained from hydrographic sections. To avoid issues associated with sudden change in correlation lengths, the vertical correlation lengths were then smoothed using the `DIVAnd.smoothfilter` function with a diffusion length of 400m. The final smoothed vertical correlation lengths are shown in Figure 6.1: the smoothing leads to a final minimum length of approximately 92 metres at the surface.

For the parameter ϵ_2, a value of 3 was used. This is substantially higher than the values used in the gridding of hydrographic sections (0.1 for temperature and salinity, 0.2 for DIC and the reinterpolated excess temperature fields). This reflects the noise found in these estimates.

In the absence of data, the DIVA algorithm relaxes the gridded field values towards a background estimate. It is therefore important that a valid first guess is given. Notably, a nonzero background estimate of excess temperature at abyssal depth will result in grid cells far from data relaxing towards this value, substantially and artificially increasing estimates of excess heat content accumulation.

To generate a background field, a simple mean of the profiles of excess temperature and salinity accumulation at each depth was calculated. As substantial accumulation of excess heat content at abyssal depths is not expected, the background field was then relaxed towards zero to ensure this artificial inflation did not occur and bias estimates: this was necessary as the horizontal mean trend did not approach zero at depths as it should. To do so, a sigmoid scaling function, $\sigma(z)$, was applied:

$$\sigma(z) = \frac{1}{1 + e^{(x-\bar{x})/w}}, \tag{6.1}$$

where $\bar{x}$ represents a depth offset at which $\sigma(z) = 0$ and w is a width scale. A value of 1000 metres was chosen for $\bar{x}$, and a value of 300 metres for w: the results at A05 suggest that globally, the majority of excess heat ought to accumulate in the upper 1000m or so, approximately the depth of the subtropical gyres. The mean profile of the rate of change of excess temperature, before and after this scaling is applied, is shown in Figure 6.2. It is currently unclear whether this mean profile not approaching zero at depth is a deficiency of the method or a result of sampling bias. However, it is likely that sampling bias plays an important role, as areas of oceanographic interest tend to be sampled more frequently, with these areas being ones in which changes are expected

to be seen. Notably, the North Atlantic, a region of global importance for deep ocean heat uptake, is far better sampled than other ocean regions.

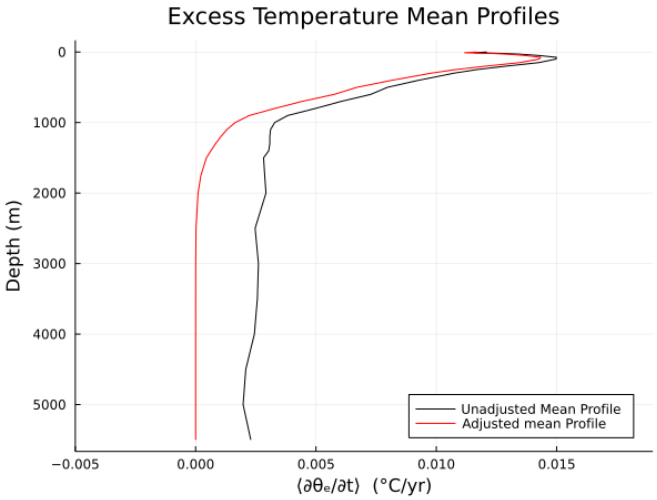

FIGURE 6.2: The mean profile of the time derivative of excess temperature, before scaling (black), and after scaling (red).

In addition to forcing the estimate of the mean profile of the rate of change of excess temperature at depth to zero, some additional data processing was performed: namely, adding artificial profiles to 'pin' the field to a global mean value in certain regions, and the removal of certain profiles which were manually inspected and deemed to be unreliable. This is shown in Figure 6.3. Though the mapping uncertainty in these regions which were 'pinned' to the mean profile is large, as expected, these estimates are unrealistic and the values large enough to notably bias results. As the addition of artificial profiles does not reduce the mapping uncertainty to unrealistically low values, this procedure was considered to produce the most robust results. Additionally, some profiles were removed by manual inspection: these were typically regions in which a single repeated vertical profile, far from other observations, contributed strongly to the accumulation of excess temperature and salinity, and did not appear to be realistic. If manual inspection of the profiles of temperature, salinity and DIC revealed anything unusual, they were removed.

Finally, before investigating the results of this decomposition of temperature and salinity, it is also important to understand the period over which these results are calculated: as previously noted, due to the asynchronous nature of hydrographic occupations, these results cannot be considered to be representative of a single period of time. Instead, they represent the amalgamation of a number of slightly different time periods: the results obtained are therefore likely to be affected by the differences in these time periods. To visualise how these differences in time periods may affect the results, the dates of the initial and final cruises used to compute each trend were interpolated, according to the same procedure and with the same parameters (correlation lengths, signal to noise ratio) as the excess temperature and salinity fields: the only difference

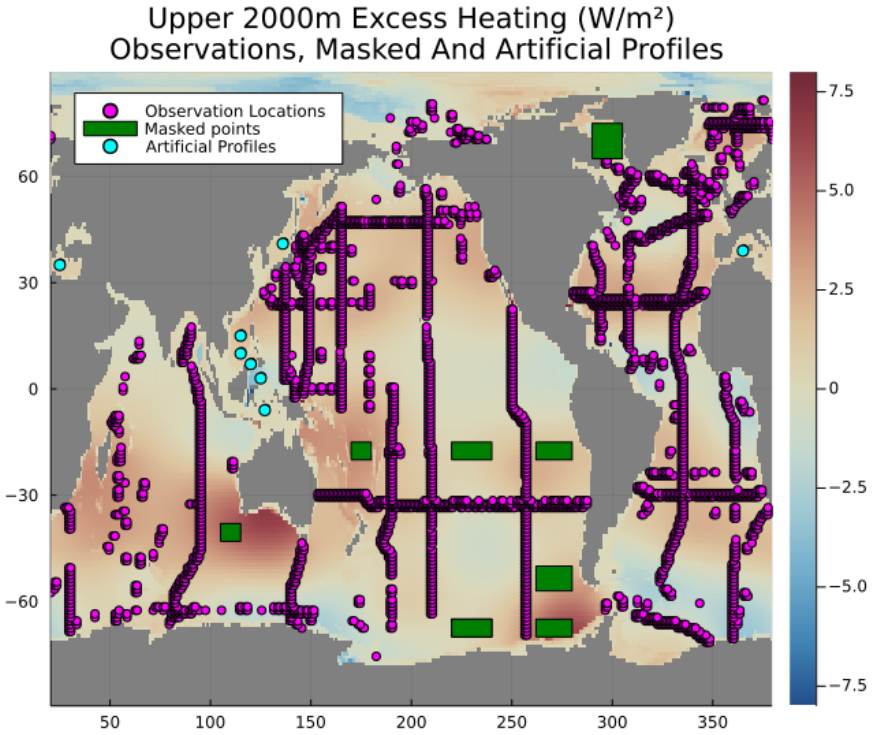

FIGURE 6.3: The upper 2000m excess heating rate in W/m^2, with the location of observations used shown in magenta, the location of artificial profiles shown in cyan, and regions for which profiles have been deemed to be unreliable and been removed in green.

being a two dimensional interpolation was used. A background field of 2000 was prescribed (ie. the DIVA interpolation will relax dates towards the year 2000 in the absence of data). The result of this interpolation is shown in Figure 6.4. In general, most locations in the global ocean span a period of at least 20 years, with periods of less than 15 years in only a few locations: typically the Arctic and some marginal seas (Mediterranean, Baltic, Black, and South China seas). Interpolation periods should be less than 10 years nowhere, but in some of these marginal seas periods of zero are seen. Essentially, this reflects the complete lack of data here, with both the initial and final field being strongly relaxed towards the background value (the year 2000). These regions can therefore also be thought of as also being regions where estimates of the rate of change of excess temperature and salinity are unreliable.

However, there are some features which will act to confound comparisons with prior estimates. Notably, in the South Atlantic, at approximately 30°S, there is a sharp gradient in time periods. This is also a region in which strong gradients in material heat content change are found by Zika et al. (2021). As a result of this feature in the time over which the interpolation presented here is performed, this feature is unlikely to be reproduced.

Additionally, though throughout the tropical and subtropical north Atlantic the time

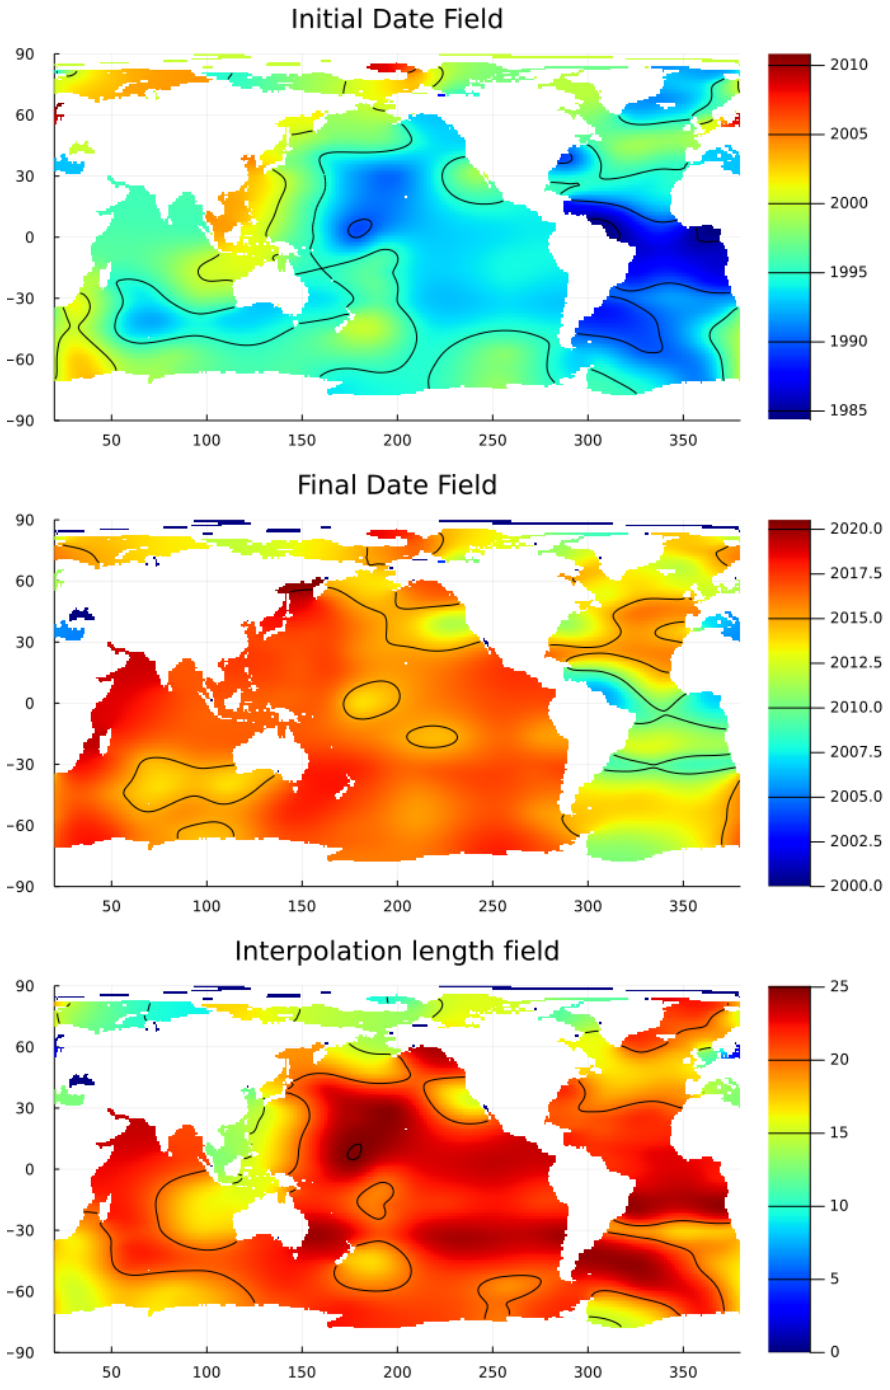

FIGURE 6.4: The interpolated dates of the initial cruise (top), final cruise (middle), and the period over which trends are calculated (bottom).

period over which our interpolation is performed is relatively constant, the initial and final dates show far more variability. Though the smoothly varying nature of excess fields ought to make this much less of an issue than might be the case for a simple interpolation of the rate of change of temperature or salinity to a global field, it is still plausible that this will cause an issue: care must therefore be taken to examine the signals in these regions carefully.

To produce an approximate time period over which our rates of change accumulation of excess temperature and salinity are calculated, a volume weighted average of these date fields over the global ocean is performed. This yields a start date of 1994.7 ± 3.65 years and an end date of 2015.2 ± 2.56 years, roughly corresponding to August 1994 and February 2015, respectively. Uncertainties are standard deviations of the start and end date, rather than standard deviations of their means.

6.2.1.2 Section Interpolation

In addition to the pointwise interpolation described above, an additional interpolation was performed in order to generate complementary estimates of excess temperature and salinity storage. This interpolation was generated from the repeat hydrographic occupations of oceanographic sections included in the GO-SHIP Easy Ocean toolbox: the results discussed from A05 in the previous chapter are an example of this. This interpolation is therefore more similar to the methodology used to generate C_{anth} estimates by Sabine et al. (2004). As with the pointwise interpolation, the condition that three repeat occupations of a section with temperature, salinity and DIC data, with a minimum separation between the first and last cruise of at least ten years was enforced. This reduced the number of sections for which sufficient data was available to twelve: A05, A16-A23, A22, A10, I06S, I09N-I08S, P02, P07, P10, P15, P18, SR03. They are shown in Figure 6.5. These data are then gridded onto the same 1 degree latitude, 1 degree longitude, 33 depth bin grid as the pointwise interpolation.

Due to the much sparser data coverage of the section interpolation as compared to the pointwise interpolation, it is expected that smoother fields will be obtained: the DIVAnd algorithm relaxes towards a mean value in the absence of data. As with the pointwise interpolation, the section interpolation uses a mean excess temperature estimate for each of the 33 depth bins. This estimate was calculated by taking horizontal means of the rates of excess temperature accumulation for each of the 12 hydrographic sections included. The median profile was then calculated, and this was scaled with the same sigmoid scaling function $\sigma(z)$ as the as for the pointwise interpolation, enforcing a background estimate of no warming below 2000m.

As with the pointwise estimate, horizontal correlation lengths of 25 degrees in the latitudinal direction, and 50 degrees in the longitudinal direction were used. However,

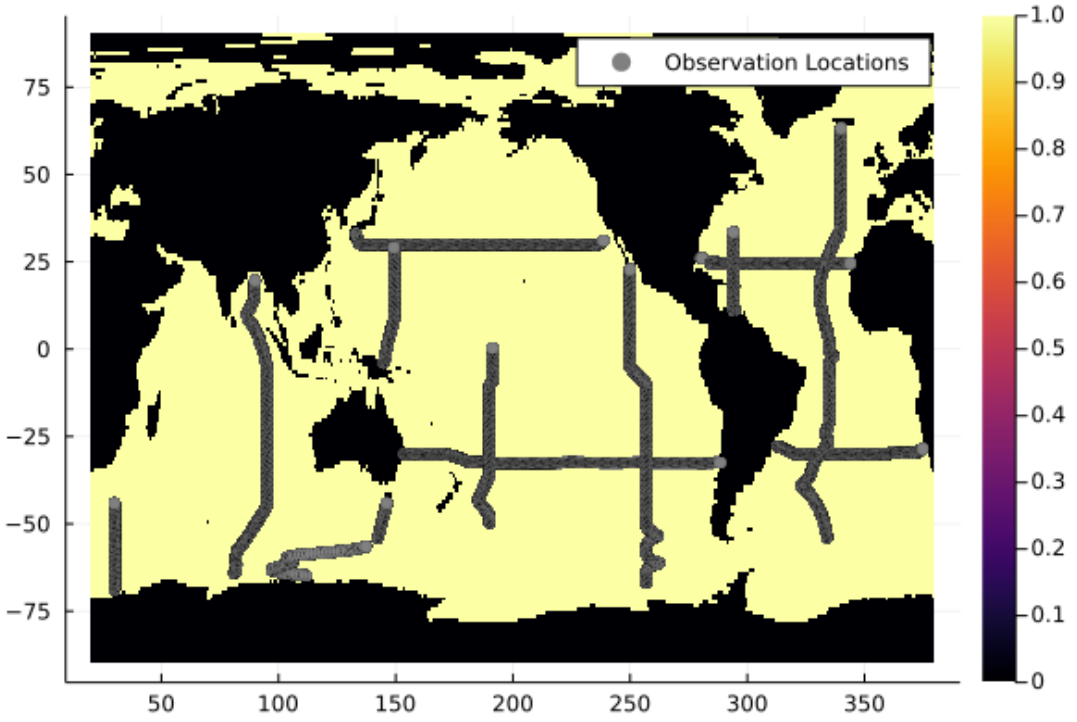

FIGURE 6.5: The locations of the hydrographic sections included in the section interpolation estimate of excess heat and salinity accumulation.

due to the greater vertical data coverage (hydrographic sections being on a 10db grid), a vertical correlation length of 250m was used.

6.3 Results

6.3.1 Testing error parametrisation

In Chapter 5, it was established that the uncertainties in trends in temperature could be parametrised as approximately 10% of the trend value, and that the uncertainty in salinity trends were better described by a constant, with a value of approximately 5×10^{-4} PSU/yr. However, this was only tested for one section, A05, located in the Subtropical North Atlantic. These parametrisations therefore cannot be assumed to be representative of the uncertainties in excess temperature and salinity throughout the global ocean: this parametrisation must be tested on other hydrographic sections.

The I08S hydrographic section is located in the Indian sector of the Southern Ocean, spanning approximately 78-95 degrees east and 28-66 degrees south. It has been occupied 3 times: 1994-1995 (316N19941201), 2007 (33RR20070204) and 2016 (33RR20160208): the cruise tracks from these cruises are shown in Figure 6.6.

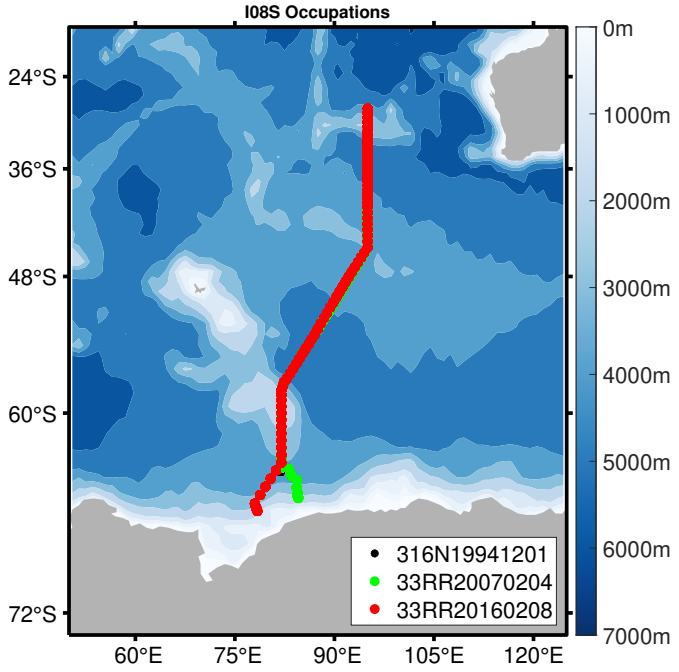

FIGURE 6.6: The 3 occupations of I08S. Cruise tracks have not been offset for visibility, instead highlighting the excellent colocation of repeated occupations.

Like the subtropical North Atlantic, the Southern Ocean is also a key contributor to the ocean sequestration of anthropogenic carbon and heat. However, it exists in a vastly different parameter space, with surface waters ranging from approximately zero to ten degrees. This makes it a suitable location to test the parametrisations established in Chapter 5. Due to its vastly different surface temperature and being highly ventilated, it is also a suitable section to test whether the assumption of a globally uniform α_T value, as opposed to a temperature dependent one, significantly affects the estimates of excess temperature and salinity accumulation obtained, in a region with substantially different background temperature. In addition, the vertical stratification of both temperature and carbon is substantially different in this region to the subtropical North Atlantic, and so testing the parametrisation on this hydrographic section also allows the exploration of a different regions of the parameter space of β_T.

To test these parametrisations, the same process was performed on I08S as on A05: bootstrapping was performed (following the same procedure) to perturb the observations, as well as the parameters used (α_T, β_T, β_S). Trends were calculated on these perturbed estimates, and 95% confidence intervals for the trends were calculated, in the same way as for A05.

The results obtained for the mean vertical profile of excess heat and salinity accumulation is shown in Figure 6.7. Several results from this are notable. Firstly, the depths to

which trends in excess temperature can be robustly detected to is much deeper: warming is seen until depths of approximately 2000m. The warming below 4500 metres appears to be driven by warming in the Antarctic Bottom Water.

Trends in excess salinity are less clear: a strong and significant freshening is seen in the top 200m, and between 300 and 600 metres an oscillating signal is seen. However, this oscillating appears to be spurious, driven by the lack of spatial coherence in excess temperature trends. This results in reduction in the spatial coherence of the redistributed temperature signal, which in turn leads to a less spatially coherent redistributed salinity signal: this finally imprints into the excess salinity, causing spurious oscillations.

It is unclear whether the lack of spatial coherence in excess temperature trends at high latitude, as compared to lower latitude, is real or an artifact of the estimation technique. It is the case that the processes by which excess heat enters the ocean at high latitude, for example wintertime deep convection, exhibit greater spatial and temporal variability than the processes by which excess heat enters the ocean at low latitude: this is similarly the case for C_{anth} (Bopp et al., 2015). However, it also remains plausible that the reduced homogeneity of excess heat and salinity accumulation as estimated here is simply due to the weak vertical stratification of temperature and salinity at high latitude, as compared to at A05. However, as the trends are clear in the section mean and interpolations onto a global grid will necessarily require correlation lengths long enough to effectively smooth out these inhomogeneities, this issue is not further considered.

Figure 6.8 shows the results for the parametrisation of errors at I08S, as performed for A05. Again, there is, for the temperature independent parametrisation, a clear relationship between the absolute value of the excess temperature trend and its error (Figure 6.8a). This relationship is less obvious in the case of the temperature dependent parametrisation (Figure 6.8b), but still visible.

In the case of a temperature independent α_T, errors in excess temperature trends may be parametrised as 12% of the trend, with a negligible intercept and an R^2 value of 0.78. Whilst the R^2 value is lower in the temperature dependent case, 0.60, the error is also slightly lower: 9% of the trend. The intercept is again negligible.

In the case of excess salinity, the distribution of errors is much tighter at I08S than at A05, with both the temperature dependent and independent parametrisations having sharp peaks at approximately 0.15×10^{-3} PSU/yr.

Though the correlations between the excess temperature accumulation rates and their associated errors are weaker at I08S than A05, the correlations remain strong. In addition, the intercepts of these parametrisations remain negligible, and the slopes are similar: approximately 10%. Thus, this parametrisation is also valid at I08S, and will

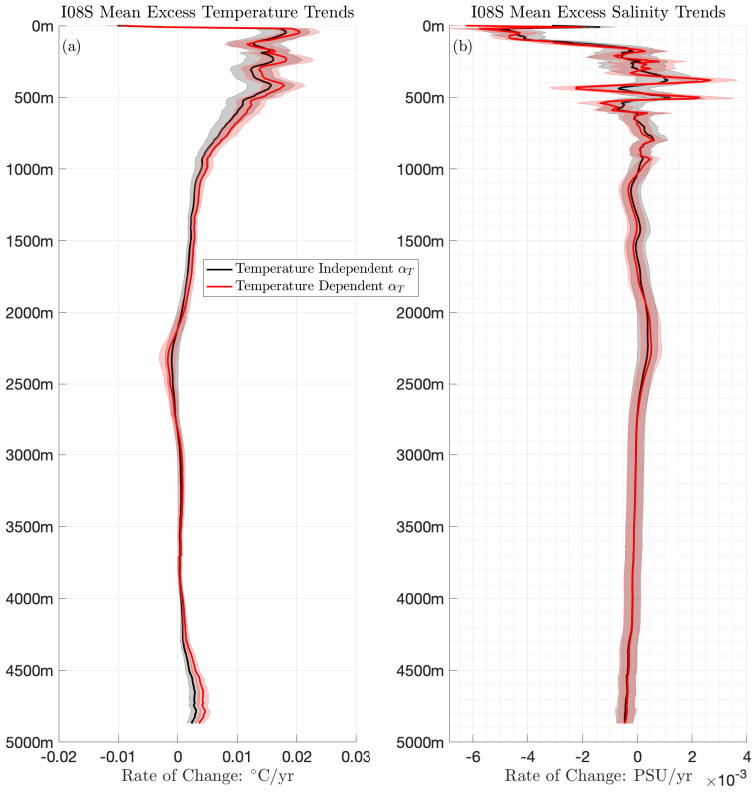

FIGURE 6.7: The mean profiles of rates of accumulation of excess temperature (panel (a)) and excess salinity (panel (b)), for I08S. Mean profile for temperature independent α_T are shown in black, and temperature dependent α_T in red.

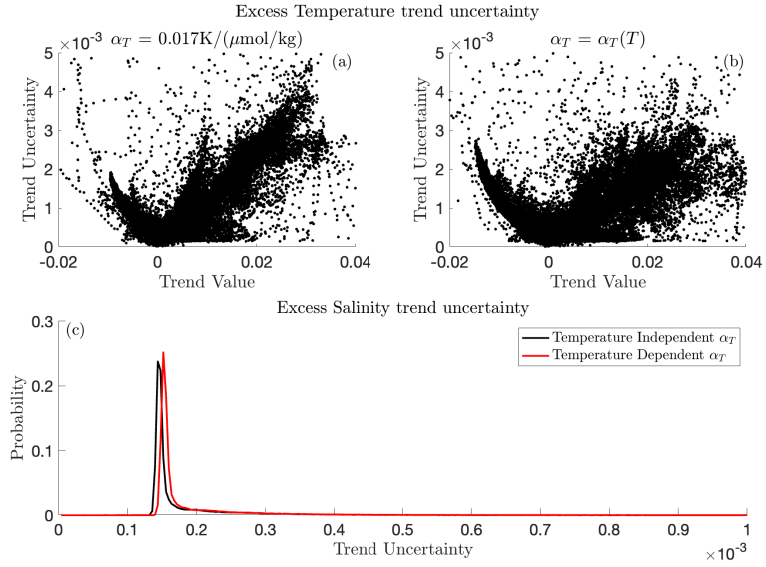

FIGURE 6.8: The errors in the calculated excess temperature trends for constant α_T (panel (a)), for temperature dependent α_T (panel (b)), and a probability distribution of the excess salinity trend error (panel (c)). Temperature trends are in °C/year, and salinity trends in PSU/yr.

so be used for estimates of global trend uncertainties: the uncertainty in the excess temperature trend is 10% of the trend.

Given the weaker correlations at I08S than A05 between the excess temperature and its error, it is somewhat surprising that the errors for excess salinity trends are smaller. As a result, the larger errors obtained at A05 will be used to estimate uncertainty in global trends of excess salinity: a characteristic uncertainty of $2.5 \times 10^{-}4$ PSU/year will be used.

6.3.2 Global Accumulation of Excess Heat and Salinity

6.3.2.1 Vertically Integrated Heating

Figure 6.9 shows the estimated globally integrated heating rate as a function of depth, integrated from the surface downwards. For the pointwise interpolation, a full depth globally integrated heat content accumulation rate of 387±82 TW is obtained. This is remarkably similar to the estimate of Zika et al. (2021), who found a globally integrated heating rate of 398±81TW over the period of their study, 2006-2017. Their estimate and its associated uncertainty is shown by the red line and shading in Figure 6.9. However, a notably higher estimate is found for the cruise interpolation: 632 ± 140TW.

Zanna et al. (2019), who investigated global ocean heat uptake over the full historical period, found a globally integrated heat content accumulation of 153±44ZJ over the period 1990-2015, a period more comparable to the one investigated here. This corresponds to a mean heating rate of 184±78TW (shown in green in Figure 6.9), which is approximately half to a third the heating rates found here. This is surprising, given that the period is more directly comparable. Zanna et al. (2019) also report total warming for 3 depth horizons: the top 700 metres, the top 2000 metres, and the full ocean depth. Over these depth horizons and the time period 1955-2017, they find approximately 90% of the ocean warming occurs in the upper 2000m, with 70 ± 15% of the warming in the upper 2000m or 63 ± 13.5% of the total in the upper 700m.

For our pointwise interpolation, approximately 80% of the top 2000m warming is contained in the top 800m: 267 of 334 terawatts. For the cruise interpolation, approximately 60% of the top 2000m warming is contained in the top 700m: 303 out of 511 terawatts. Despite the different period over which these estimates are calculated, these are both within the uncertainty of the estimate of the distribution of excess heat of Zanna et al. (2019). For the pointwise interpolation, approximately 14% of the total warming is found below 2000 metres, whereas for the cruise interpolation, 20% is found below 2000m.

It is unclear if these warming signals at these abyssal depths are real: it will be shown later that some deep warming signals appear to be spurious and the result of the weak

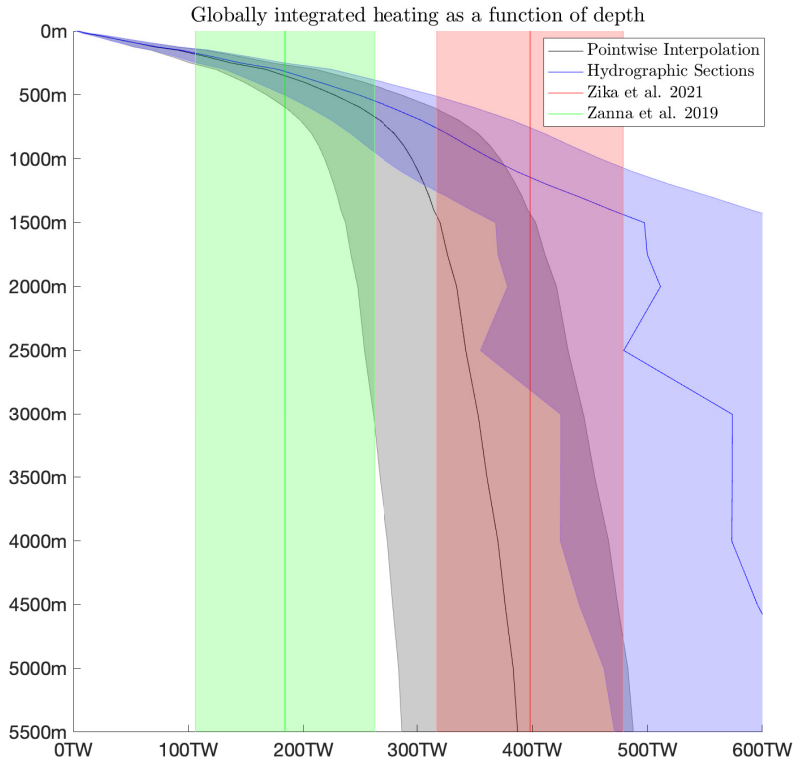

FIGURE 6.9: The globally integrated heating rate obtained, integrated as a function of depth, for the pointwise interpolation (black) and the cruise interpolation (blue). The shading represents a 95% confidence interval. Estimates of global ocean heat uptake from Zanna et al. (2019) and Zika et al. (2021) are shown in green and red, respectively, along with their uncertainties.

vertical stratification of DIC at depth. Conversely, the method agrees well with both the vertical distribution of warming from Zanna et al. (2019), indicating that spurious deep warming signals likely cancel out in the horizontal mean. Both the pointwise and cruise interpolation agree to within uncertainty with globally integrated heating rate of Zika et al. (2021), and the agreement between the total heating rate found by the pointwise estimate (387 ± 82TW), and that of Zika et al. (2021) (398 ± 81TW), is remarkable.

However, as will be shown later, it is likely that the cruise interpolation is an overestimate of the global ocean excess heat content uptake, due to the lack of spatial coverage. Thus, it is likely that the 'true' ocean excess heat content change is instead closer to that of Zika et al. (2021) and Zanna et al. (2019). Further investigations ought to reveal which estimate the excess heat content accumulation presented here should better resemble: *a priori*, we expect better agreement with the methods of Zanna et al. (2019) due to the time horizons of the two studies being closer. It is clear though that the estimate of ocean excess heat content accumulation presented here, in particular deep ocean heat content, is biased high. Accounting for this uncertainty will require a more careful interpolation and treatment of errors.

6.3.2.2 Column Inventories

Figure 6.10 shows the spatially resolved excess heating rate, for the full depth range of the ocean (a, b), upper 2000m (c, d) and deeper than 2000m (e, f), for the pointwise interpolation (left column), and cruise interpolation (right column). These are presented as heating rates over the given depth ranges. Several results are notable.

Firstly, both the cruise and pointwise interpolation show significant warming below 2000m. This is surprising, given previous studies (eg. Zanna et al. (2019)) indicate around 80-90% of excess heat uptake is expected to be in the top 2000m. Additionally, below 2000m, both the pointwise and cruise interpolation show regions of heating and cooling. This is not entirely unexpected: for example, Gebbie and Huybers (2019) investigated the response of the global ocean to the propagation of surface temperature anomalies from the 'Little Ice Age' of the 18^{th} century. They found that this could explain cooling in the deep North Pacific Ocean. Whilst this does not explain the warming signal found here by both interpolations, it provides a potential explanation for strong and statistically significant excess warming and cooling in regions of the ocean which are not thought to be well ventilated. However, the differences between both interpolations strongly suggest that below 2000m, warming signals are unreliable. This will later be explored in more detail.

In the top 2000m, both estimates generally show warming: however, the pointwise estimate is far less homogeneous, with substantial regions of cooling, and warming generally greatest in the subtropical gyres. The cruise interpolation shows warming everywhere, except for in a region of the Southern Ocean starting with the Weddell Sea and extending eastwards: cooling is also found in this region in the pointwise interpolation: this cooling signal therefore appears robust. However, cooling signals in other regions in the pointwise estimate are less clear.

Noticeably, the region of cooling in excess temperature storage seen in the upper 2000m in the North Atlantic Drift region is not seen in the deep ocean heat storage. This region of cooling therefore likely corresponds to the influence of redistribution feedback: a northward shift in the position of the North Atlantic Drift will redistribute warm water northwards. This will alter the radiative balance, reducing the total downward radiative flux and leading to a cooling which is diagnosed as excess temperature, in a process referred to as redistribution feedback. Reduced data coverage in the cruise interpolation means this feature is absent in that estimate.

However, for the pointwise interpolation, this pattern is not the case in other regions of cooling: for example, in the equatorial Atlantic, statistically significant cooling is seen in both the upper and deep ocean. This is also the case in the Atlantic and parts of the Indian sectors of the Southern Ocean. Thus, an alternative explanation is required for this cooling.

GLODAP does not recommend adjustment for DIC data unless crossover analysis finds that DIC values are differ by at least $\pm 4\mu$mol/kg. Therefore, within the boundaries set by GLODAP quality control, artificial DIC differences of up to 4 μmol/kg can result from separate cruises occupying the same latitude/longitude box. The effect of a uniform DIC offset to excess temperature accumulation is demonstrated in Figure 6.11. Below 1400 metres, all estimated excess temperature is within the range of uncertainty introduced by a permutation of $\pm 4\mu$mol/kg. Thus, it is highly likely that pointwise estimates are strongly influenced by systematic offsets between cruises, in latitude/longitude boxes where cruises cross over. This contributes additional noise to the pointwise estimate in a way that will not affect the cruise interpolation. Calibration means that errors due to these systematic offsets should sum to (approximately) zero: thus, the global excess heat accumulation estimate from the pointwise interpolation is likely more accurate, due to its greater data coverage. However, spatially resolved estimates from the pointwise interpolation are therefore likely less reliable than those from the cruise interpolation.

Figure 6.12 shows the spatially resolved excess salinification rate, for the full depth range of the ocean (a, b), upper 2000m (c, d) and deeper than 2000m (e, f), for the pointwise interpolation (left column), and cruise interpolation (right column). These are presented as freshwater fluxes in metres per year over the given depth ranges. As for excess heating, estimates below 2000m appear highly unreliable, and above 2000m, the pointwise interpolation shows far more spatial variability due to its increased data coverage and noise.

Due to the noise in the pointwise interpolation, results from it will no longer be considered. However, there are several results from the cruise interpolation which are in agreement with previous studies and appear robust.

In the Atlantic, freshening is seen at the Equator, and salinification elsewhere. This is in line with the results of Durack and Wijffels (2010), who found a decrease in sea surface salinity over 1950-2000 in the Equatorial Atlantic, and salinification elsewhere (their Figure 5b). Additionally, they found freshening in the North Pacific, freshening in the tropical and subtropical South Pacific, and salinification throughout much of the Indian Ocean. Though their trends are better spatially resolved than those presented here and thus show greater spatial variability, the results are qualitatively similar. Thus, it appears that the cruise interpolation is reliably identifying excess salinity, and that the trends in freshwater fluxes over 1950-2000 have continued over the GLODAP period. This is in agreement with the findings of Chapter 3, which showed changes in excess salinity generally precede those in excess temperature.

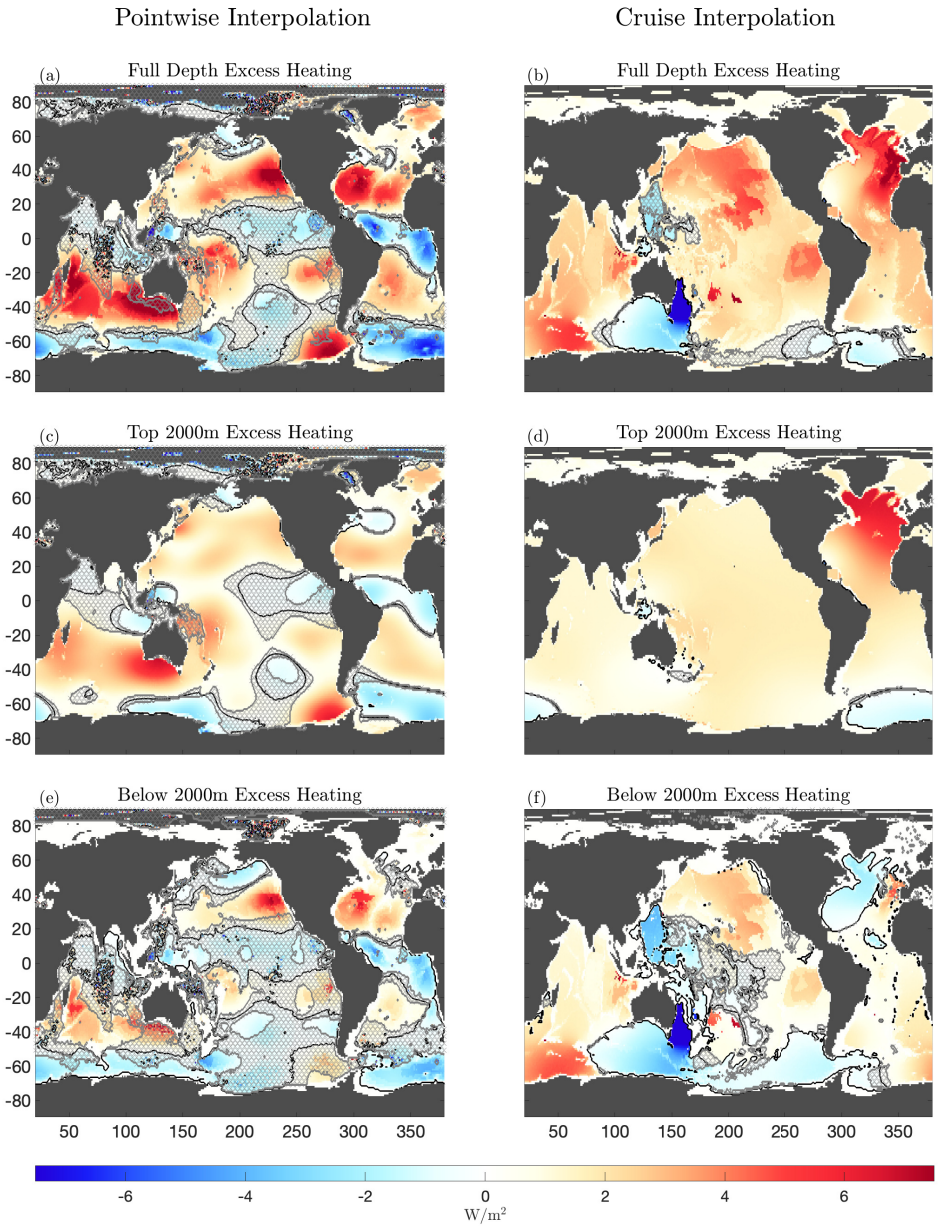

FIGURE 6.10: Column inventories of excess heating in units of Watts per square meter. Panels (a) and (b) shows the full depth heating rate, panels (c) and (d) the upper 2000m, and panels (e) and (f) depths below 2000m. The zero contour is indicated in black, areas where estimates are not statistically significant are indicated with stippling and surrounded by grey contours.

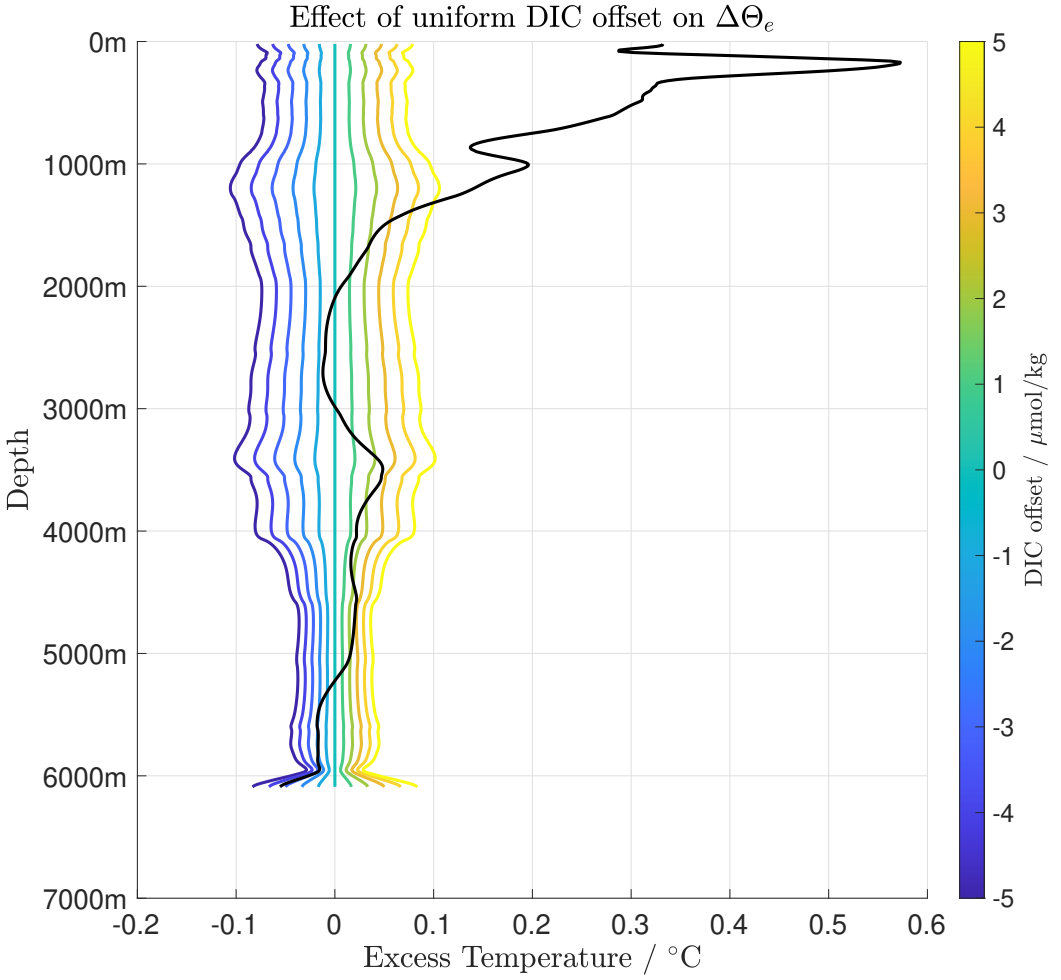

FIGURE 6.11: The effect of applying a uniform DIC offset of up to $\pm 5\mu$mol/kg to a sample column of water. The black line shows the accumulated mean excess temperature for a sample column of water at A05, between 1990 and 2015. The coloured lines indicate how much excess heat is over or underestimated due to the uniform addition of up to $\pm 5\mu$mol/kg.

6.3.2.3 Trends at Depth

In the previous section, it was noted that at depths below 2000m, trends did not appear reliable. Whilst for the pointwise interpolation this may be explained in terms of noise introduced through systematic offsets which are not caught through crossover analyses in GLODAP, this explanation cannot be applied to the trends from the cruise interpolation, which also appear to be artificial. As trends at depth appear reliable for the hydrographic sections A05 and I08s, the most likely explanation is therefore not that the decomposition itself is at fault, but that these unreliable trend estimates are the result of the interpolation from cruises (or the pointwise estimate) to a 3D gridded field.

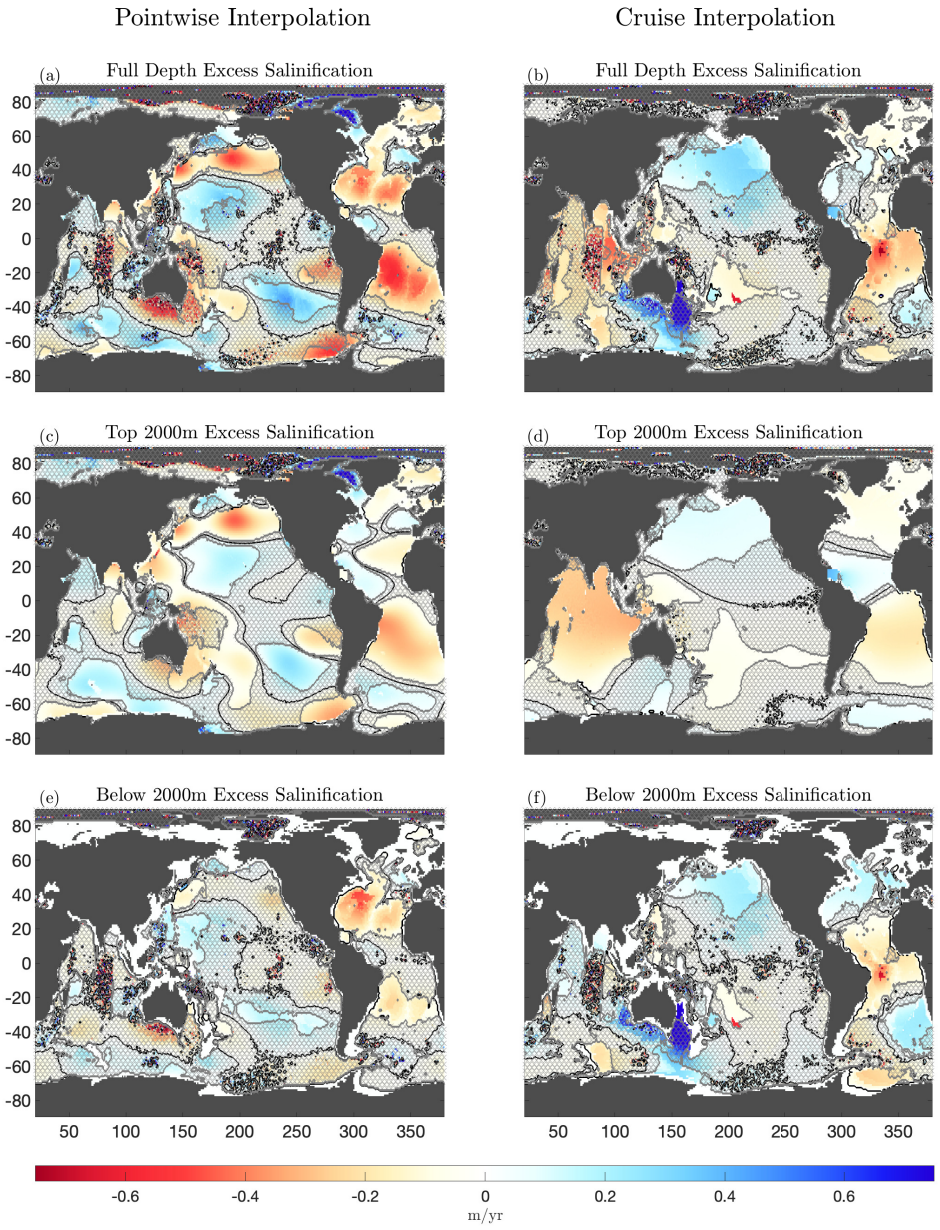

FIGURE 6.12: Column inventories of excess salinification in units of metres of freshwater per year. Panels (a) and (b) shows the full depth salinification rate, panels (c) and (d) the upper 2000m, and panels (e) and (f) depths below 2000m. The zero contour is indicated in black, areas where estimates are not statistically significant are indicated with stippling and surrounded by grey contours.

To test whether these large trends at depth are the result of the gridding algorithm, a simple interpolation test was performed using MATLAB. In each ocean basin, trends from each cruise, and their location, were pooled within the range of each depth box of the 3D grid. These trends were then interpolated onto the same one degree grid, using Natural Neighbour interpolation. The results are shown in Figure 6.13.

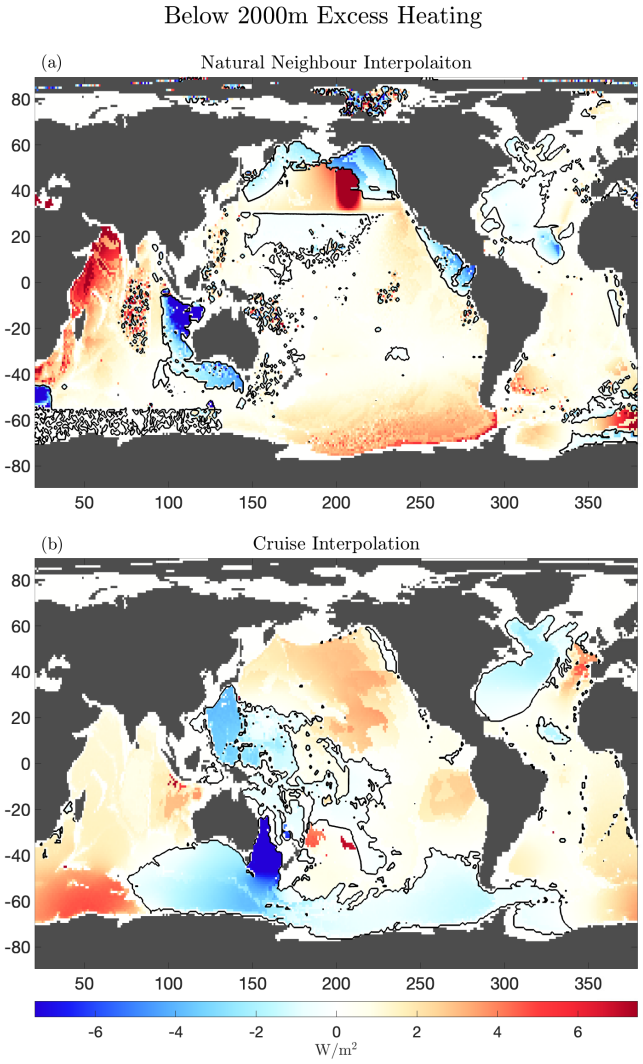

FIGURE 6.13: Column inventories of excess heating below 2000m in units of Watts per square meter. Panel (a) shows the results from the simple Natural Neighbour interpolation, and (b) shows the results from the cruise interpolation. The zero contour is indicated in black.

Though the patterns of heating below 2000m are significantly different, regions of strong warming and cooling are still seen. This indicates that the likely cause of this potentially spurious deep warming is not the interpolation, but simply that additional

refinements to the decomposition are necessary in order to accurately decompose excess and redistributed temperature below 2000m. This is most likely the result of the fact that the Decomposition Matrix D enforces an exact solution due to being an analytic, rather than numerical inversion, and so does not allow for uncertainty. Thus, some locations display spurious excess warming, which is propagated into the global estimates. Further improvement of this technique will therefore require improvements of the decomposition technique in order to rectify this: subsequent investigations have revealed that the masking factor M is not sufficient to remove unreliable estimates in the deep ocean, and that improved estimates of deep ocean heat content will require more aggressive masking. This should likely bring the results of ocean heat content change into better agreement with those of Zanna et al. (2019). However, for the rest of this work, I shall instead focus on the top 2000m, where results appear to be robust.

6.3.2.4 Zonal Means

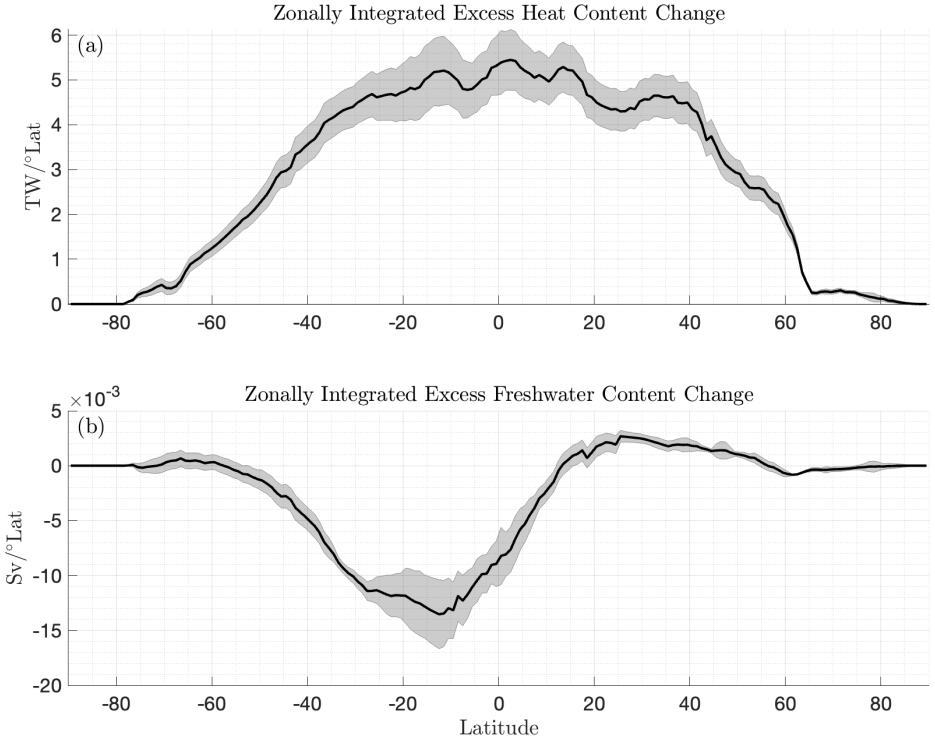

FIGURE 6.14: Global zonal mean excess heat (a) and freshwater (b) accumulation rates for the top 2000m. Shading represents a 95% confidence interval.

Figures 6.14 and 6.15 show the zonal mean accumulation of excess heat and freshwater, in Terawatts and Sverdrups, respectively, for the Global Ocean, and broken down by basin. Like both Zika et al. (2021) and Zanna et al. (2019), there is substantial accumulation of excess heat in the subtropics, corresponding to the accumulation of excess heat in the subtropical gyres. This also mirrors the accumulation of anthropogenic carbon

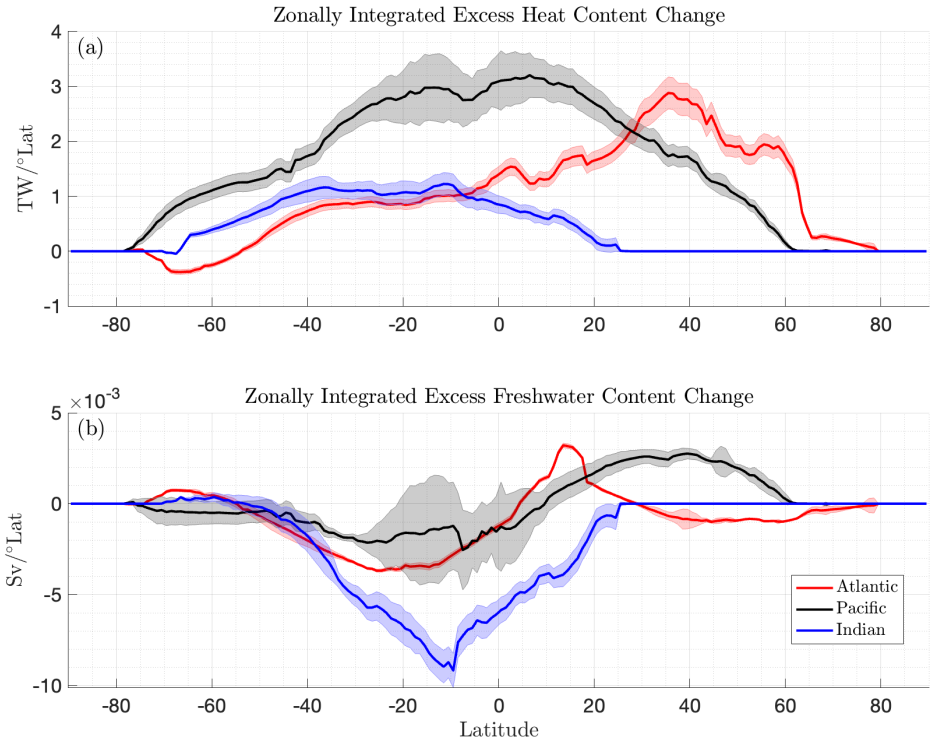

FIGURE 6.15: Basin zonal mean excess heat (a) and freshwater (b) accumulation rates for the top 2000m, for the Atlantic, Pacific and Indian oceans. Shading represents a 95% confidence interval.

over a similar time period found by Gruber et al. (2019). These are all patterns which would be expected from previous modelling studies and the close analogy between excess heat and anthropogenic carbon.

In the global mean (Figure 6.14), excess heat uptake appears almost entirely passive, with small peaks at the Equator, approximately ± 15 degrees, and approximately 35 degrees north. The peak at 35°N is also found by Zanna et al. (2019) and Zika et al. (2021). However, they also found a peak at 35°S: this is likely not shared by our dataset due to the lack of spatial resolution. This is likely to be the case in general: the lack of spatial resolution in the cruise interpolation will lead to generally smoother fields than those of Zika et al. (2021).

Despite the lack of spatial resolution, we also find slight peaks in warming at low latitude, in agreement with Zika et al. (2021) but not with Zanna et al. (2019). Zika et al. (2021) suggest that the warming they find here may be related to decadal variability. However, as the period over which the excess heat trends computed here is approximately 1990-2015, as compared to 2006-2017 for Zika et al. (2021), the results presented here suggest it is unlikely that this is the case. Instead, differences between the distributions found by Zika et al. (2021) and Zanna et al. (2019) likely result from methodological differences.

Patterns of excess salinity accumulation are less clear. Generally, in a warming climate, we would expect to see increases in salinity in the subtropics, and freshening in the tropics and at high latitude due to water cycle amplification (Durack and Wijffels (2010), Zika et al. (2018)). However, in the global zonal mean, the results presented here instead suggest a salinification between 60 south and 15 North, and freshening elsewhere. It is not straightforward to interpret this pattern, given the strong salinification in the Indian Ocean seen in Figure 6.12.

To better understand the patterns of excess salinification, it is necessary to split up changes by ocean basin. This is shown in Figure 6.15. From the patterns seen here, it is clear that the salinification signal seen in the global zonal mean is dominated by salinification in the Indian ocean. In addition, a similar pattern to the global mean, though of smaller amplitude, is seen in the Pacific ocean.

In the Atlantic, a pattern more closely resembling what might be expected from water cycle amplification is seen. Though the peak is shifted northwards of what might be expected, a strong freshening signal is observed at approximately 13 degrees north, with salinification on either side. In the Southern Ocean, where there are observations, freshening is then also seen south of approximately 55 degrees. However, observations contributing to the gridded 3D fields in the Subpolar North Atlantic are sparse, and so we do not see the expected freshening here due to a lack of observations: instead, the gridding 'drags' the salinification northwards.

Patterns of excess heat storage on a basin by basin basis are much more easily interpreted. Like Zika et al. (2021), we find a peak in warming in the Indian Ocean at approximately 40 degrees south, though again it is less pronounced than theirs. Similarly, the largest peak in the Pacific Ocean excess heat content change is found at and just north of the equator, with significantly less warming elsewhere. However, there is not a secondary peak at approximately 35 degrees south, unlike Zika et al's fields.

In the Atlantic, the pattern of excess heat storage strongly resembles that of Zika et al. (2021): relatively uniform warming of approximately 1TW/°Lat in the Southern Hemisphere, and a peak on the order of 3TW/°Lat at approximately 40 degrees north. Their excess heating declines precipitiously north of this, unlike ours: however, this is again likely due to the differences in data coverage. In this region where excess heating peaks, we also see excess salinification: this may indicate density compensation in the changes to water mass properties in the North Atlantic. This density compensating behaviour is thought to be a pronounced and robust feature in the North Atlantic (Lowe and Gregory (2006), Mauritzen et al. (2012)), and so is now investigated in further detail.

6.3.3 Contributions to excess density

Results in Subsection 6.3.2.4 suggest that trends in excess temperature and salinity may occur in a density compensating fashion in the North Atlantic, with warming and salinification found here. To investigate this, changes in excess density are now directly computed. To do so, the trend in excess temperature and salinity was added to climatological temperature and salinity fields from the GLODAP dataset. This allows the computation of an excess density trend. In addition, the 'excess temperature only' and 'excess salinity only' trends were computed, by computing only the density change due to the addition of either component to the climatological fields: whilst the total excess density trend will not be exactly the sum of the two due to the nonlinear equation of state of seawater, it does allow for the examination of the extent to which changes to heat content and the freshwater cycle control changes in excess density. Confidence intervals were calculated via the same bootstrapping procedure used to compute zonally integrated trends in excess heat and freshwater content.

Figure 6.16 shows the results of this calculation for the Atlantic, Pacific and Indian oceans. Unsurprisingly, patterns of total excess density change more strongly resemble the salinity driven excess density change: heat driven excess density change is generally a relatively uniform, negative perturbation. However, it is surprising how in both the South Atlantic and Indian Oceans, the total excess density change is much closer to the salinity driven excess density change. This implies that the change in excess density here is more strongly controlled by changes in patterns of evaporation and precipitation than changes in surface heat fluxes. In contrast, in the North Atlantic, excess density change is most strongly controlled by excess heat change, as a result of the large increase in excess heat storage here.

Throughout much of the Atlantic, and the entire Indian Ocean, changes in excess heat and salinity content act in a density compensating fashion. However, good density compensation is not seen in any locations, implying that the storage of excess heat and freshwater by the ocean is currently affecting ocean dynamics.

6.4 Discussion & Conclusions

In this Chapter, two gridded excess heat and salinity accumulation techniques have been presented: a product generated from the interpolation of hydrographic cruises, and a pointwise interpolation of DIC, temperature, and salinity profiles. In addition, it has been shown that the choice of a temperature dependent formulation of α_T, or the simple choice of a global mean α_T does not strongly affect estimates of the accumulation of excess temperature and salinity. However, whilst accumulation rates on the hydrographic sections for which high quality repeat observations are available are

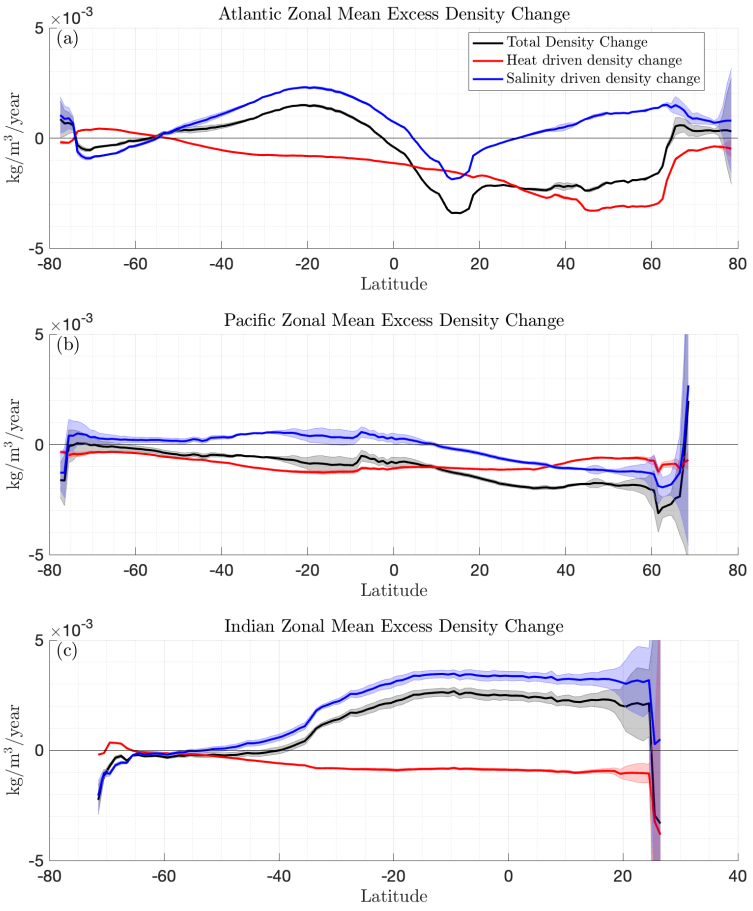

FIGURE 6.16: Zonal mean rates of accumulation of excess density, for the Atlantic (a), Pacific (b), and Indian oceans (c). Total excess density changes are shown in black, the contribution of heat driven changes in red, and the contribution of salinity driven changes in blue. 95% confidence intervals are indicated with shading.

reliable at all depths, interpolation of these sections to a global field at depths greater than about 2000m will require additional care, likely comprising careful manual data cleaning.

Both interpolations have different strengths and weaknesses: the pointwise interpolation greatly overestimates spatial variability in excess heat, due to the way in which multiple cruises occupying the same latitude/longitude box can affect results. However, due to its greater observational coverage, it is possible that it better estimates the globally integrated excess heat uptake. It implies, over the GLODAP period (spanning approximately 1994-2015), a globally integrated ocean heat uptake of 387 ± 82TW: remarkably similar to the heating rate of 398 ± 81TW found by Zika et al. (2021). However, this is substantially higher than the rate of ocean heat uptake found by Zanna et al. (2019) of 184 ± 78TW over the more comparable period of 1990-2015. For this global estimate, approximately 14% of the total heat uptake is found below 2000m, and

80% of the top 2000m warming is found in the top 700m: again in agreement with the results of Zanna et al. (2019).

The cruise interpolation, which likely underestimates spatial variability but is generated using far more reliable data, produces a larger estimate of globally integrated ocean heat uptake: 632 ± 140TW. This is substantially larger than the results of Zika et al. (2021) and Zanna et al. (2019), suggesting that it may overestimate ocean heat uptake. The zonal mean heat accumulation results in the Pacific ocean, showing no clear peak at approximately 35S and appearing to be due to the gridded field being interpolated from a cruise at this location, suggest that this overestimate may be due to interpolation of signal peaks across larger regions of ocean than is realistic. However, the vertical distribution of excess heat accumulation is again realistic: approximately 80% of the total warming is found in the top 2000m, and 60% of the top 2000m warming is found in the top 700m.

In addition, it is highly likely that the Subtropical Pacific is not the only region where the trend in ocean excess heat content is exaggerated due to interpolation over a larger region than is realistic. For example, the pointwise interpolation shows reduced excess heating in the North Atlantic, relative to the cruise interpolation (Figure 6.10). In most ocean regions it is highly likely that the cruise interpolation is more reliable: as shown in Figure 6.11, a small DIC offset can have a substantial result on the excess heat content and thus pointwise estimates are likely less reliable than the cruise interpolation. However, the North Atlantic benefits from plentiful observations (Figure 6.3, and thus it is less likely that excess heat content changes here from the pointwise interpolation are highly unreliable. As a result, it seems plausible that in this location, where the pointwise interpolation shows substantially lower and spatially smoother excess heat uptake, the cruise interpolation also overestimates the trend in excess heat content, with the true trend less than the calculated trend. Plausibly, this effect is repeated elsewhere in the global ocean: again, it is surprising that our globally integrated heat content is more similar to that of Zika et al. (2021) than Zanna et al. (2019). A positive bias in regions over which we have interpolated the trend in excess heat content is a plausible candidate for this.

Zika et al. (2021) finds a maximum in their material heating at the Equator, and Zanna et al. (2019) find warming at the Equator to be similar to other regions. The results presented here suggest a small peak in excess heat accumulation at the Equator, but not substantially larger than in other regions. Together, these results suggest that the heat content change of the Equatorial ocean is not substantially more variable than other ocean regions as suggested by Zika et al. (2021), but that methodological differences between the approach of Zanna et al. (2019) and the approach used here and by Zika et al. (2021) instead accounts for these differences.

Unlike Zika et al. (2021), the results presented here show regions of both excess heating and cooling, though the excess heat estimate here does not show more spatial structure than theirs. The presence of both warming and cooling suggests that the excess heat content changes we find are not driven entirely by ocean heat uptake in response to a warming climate. As cooling is only found in the high latitude southern ocean, it seems highly likely that this cooling is due to the redistribution feedback: that is, surface fluxes responding to changes in ocean circulation, driving radiative cooling in regions where redistribution has warmed the ocean. It also appears plausible that the cooling in the North Atlantic current regions, as estimated by the pointwise interpolation, is another example of redistribution feedback driving negative excess heating. Despite these differences in distribution, that the magnitude of our globally integrated heating rate (from the pointwise estimate) and that estimated by Zika et al. (2021) are so close is remarkable, as both techniques employ very different assumptions and methodologies. This suggests that their net change in global ocean heat content is strongly controlled by the thermodynamic response of the ocean to global warming, as they suggest, and not due to internal variability in ocean heat content: it would be expected that this rate would not fluctuate greatly, as it is essentially set by top of atmosphere radiative fluxes.

The finding that excess density changes are in many locations dominated by changes to excess salinity, rather than temperature, implies that a number of assumptions about the response to the ocean to anthropogenic climate change may need to be rethought. For example, the FAFMIP experiments (Gregory et al., 2016) generally consider changes to freshwater to be the least influential surface flux (see `http://www.fafmip.org/CMIP6_FAFMIP.pdf`). The results presented here suggest this assumption may lead to an underestimate of the importance of changes to the freshwater cycle in understanding ocean circulation change over the anthropogenic period. However, again, these results still require validation through an external estimate: due to the lack of other estimates of ocean heat uptake, it is difficult to say whether these remarkable results are correct, or simply due to deficiencies in the estimation technique. Though there no direct evidence to suggest these results are inaccurate, and indeed the patterns of excess salinity change in the Atlantic are commensurate with other estimates (Durack and Wijffels (2010), Skliris et al. (2014), Stott et al. (2008) for example), explicit validation of these results through an alternative method should be performed before they can be considered truly reliable.

Finally, despite the far lower spatial resolution of observations going into the cruise interpolation of excess heat accumulation, the basin resolved zonal mean estimates show remarkable agreement with those of Zika et al. (2021), particularly in the Atlantic. This not only provides independent confirmation of the spatial structure of excess heating found by Zika et al. (2021), but again underscores the importance of the North Atlantic for global ocean uptake of heat and carbon, and of continued observational monitoring of both physical and chemical properties in the ocean for the understanding of the

future ocean response to anthropogenic climate change.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

- Two new methods for diagnosing the excess temperature and salinity of the ocean have been designed. The first (described in Chapter 3 and referred to as the redistribution coefficient method) relies on an understanding of how changes in circulation at to couple perturbations to the background temperature (or salinity) and carbon field in the ocean. The second (described in Chapter 5 and referred to as the matrix inversion method) combines this technique with the method of Bronselaer and Zanna (2020) to diagnose excess temperature and salinity in the ocean relative to an initial profile through a simple matrix inversion. This second technique can be thought of as a derivative form of the framework introduced by Williams et al. (2021) to identify regions of the ocean dominated by excess or redistributed heat changes. These two techniques have then been applied to a single ocean model as proof of concept, and to repeat hydrography on the A05 hydrographic line and the entire GLODAP dataset, respectively. Notably, both techniques employ assumptions which make them in principle extensible to other tracers.

- In the development of the second technique, it has been shown that the global mean approximation of Bronselaer and Zanna (2020) represents an oversimplification of the transient response link between changes in ocean temperature and carbon content. Instead, it appears this transient response link, which they refer to as the carbon-heat coupling, is better parametrised as a linear function of background temperature due to the correlations between background carbonate chemistry and preindustrial ocean temperature.

- By applying the redistribution coefficient method to the NEMO OGCM, a number of expected features of anthropogenic climate change were demonstrated explicitly: a redistribution of heat out of the North Atlantic, a redistribution of salinity

into the South Atlantic, and changes to the water cycle that significantly precede changes to ocean heat content. It was explicitly shown to well capture heat redistribution into regions of the ocean which can be considered unventilated, and is thus thought to be accurate. It has also allowed for the first modelled estimates of excess and redistributed salinity to be produced, which produce results consistent with a number of previous studies.

- By applying the matrix inversion method to the A05 and I08S hydrographic sections, it has been shown that the decision of whether to follow the approach of Bronselaer and Zanna (2020) and use a single value for the carbon-heat coupling, α, or to make it temperature dependent does not strongly influence estimates of excess heat uptake. However, some statistically significant differences between estimates generated under the assumption of global uniformity or temperature dependence are seen. This suggests that whilst the validity of the estimates generated by Bronselaer and Zanna (2020) are not strongly affected by their use of a globally uniform constant, estimates would be improved by the inclusion of temperature dependence on the carbon-heat coupling. Results at A05 suggest statistically significant trends in excess temperature and salinity are seen typically only in the upper 1500m, whereas results at I08S suggest at higher latitudes, statistically significant trends in excess temperature can be seen at depths of up to 2000m. Indeed, it appears that some excess heat content can even be seen in bottom water at this location. Error analyses on these locations indicate that the error on excess temperature trends may accurately be parametrised as approximately 10% of the trend, and the error on excess salinity trends may be parametrised as approximately 5×10^{-4} PSU/year.

- By applying the matrix inversion method to the full GLODAP dataset, two global estimates of the rates of accumulation of excess heat and salinity over a period spanning approximately 1994-2015 have been produced. The pointwise interpolation suggests a globally integrated heat content accumulation rate of 387 ± 82 terawatts: remarkably close to the 398 ± 81 terawatt estimate of Zika et al. (2021), over the period 2006-2017. The cruise interpolation, which uses more reliable data but may erroneously interpolate peaks in the excess warming signal, instead finds a globally integrated heat content accumulation rate of 632 ± 140 terawatts. Whilst significantly larger, the uncertainty range in both estimates overlaps with the uncertainty range of Zika et al. (2021). However, both are significantly higher than the passive heating estimate of 184 ± 78 terawatts over the period 1990-2015 found by Zanna et al. (2019). Zika et al. (2021) suggests the difference between their estimate and that of Zanna et al. (2019) is due to the different period over which the estimates are calculated, however, the results presented here suggest that this may instead be due to metholodogical differences. This is also the case for the latitudinal distribution of heat accumulation.

- Calculating trends in excess density due to excess temperature and salinity change, as well as just excess temperature or salinity reveals that the impact of salinity changes on excess density change is in many cases larger than that of temperature. This a remarkable result, despite some indications that it is not necessarily surprising, and will require explicit validation against other methods of estimating excess salinity. However, if true, it suggests that a good understanding of changing ocean dynamics depends more heavily on an understanding of perturbations to the freshwater cycle than the effects of heating.

7.2 Future Work

In this work, two novel and related techniques which leverage changes in the local carbon content of the ocean to diagnose changes in temperature and salinity have been developed, and applied to a number of different datasets, both observational and model based. Whilst it has been possible to compare the estimates of excess temperature and DIC with prior estimates in a number of ways, this is only the case due to the significant efforts of others in developing techniques to separate out changes in temperature and carbon into excess and redistributed temperature, and into natural and anthropogenic carbon. Unfortunately, such a comparison has not been possible for excess salinity, due to the lack of prior studies investigating it. Whilst there have been numerous studies investigating perturbations to the freshwater cycle due to anthropogenic climate change, these typically focus on sea surface salinity: to my knowledge, the only other comparable estimate of excess salinity is that alluded to by Zika et al. (2021), which is not further discussed in the paper and remains unpublished. Thus, whilst the estimates of excess salinity presented here are of interest, they currently remain unvalidated independently. Though the data presented in this work regarding excess salinity are both plausible and consistent in many ways with a number of other studies, it will be necessary to explicitly validate the results presented here regarding excess salinity against other techniques in order to understand their accuracy.

Furthermore, both the redistribution coefficient and matrix inversion techniques are in essence simply first order Taylor expansions, which will therefore become less valid as the ocean state further diverges from the preindustrial conditions. Though there is no indication that this validity will expire in the coming decades, extension of the techniques via the use of Water Mass principles, Machine Learning techniques, or perhaps other methods, as alluded to in Chapter 4 remains a promising avenue for improving our understanding of ocean temperature and salinity change through carbon based methods.

More broadly, scope exists to extend the understanding of many of the techniques developed in this work, and the tools used in it. Though Chapters 5 and 6 found that the

globally uniform α approximation used by Bronselaer and Zanna (2020) does not produce significantly different results from those obtained using a temperature dependent formulation, there is now strong evidence that this temperature dependent formulation is more accurate. It is therefore likely that better understanding spatial variability in α would lead to a better understanding of how the ocean sequesters heat and carbon, and precisely how these two intimately related processes are linked. For example, Davila et al. (2022) used a transport matrix inversion (TMI) to investigate how anthropogenic carbon is propagated into the ocean. By repeating this analysis for temperature, and then relating the results of the two, it ought to be possible to deconvolve the temporal and spatial variability in α: this would likely lead to a far richer understanding of the evolution of temperature and carbon covariability in the ocean.

In addition, the redistribution coefficient approach, developed in Chapters 2 and 3, makes use of the spatial covariability of temperature and carbon in order to diagnose the redistribution of temperature and salinity from changes in natural carbon. However, as noted here, whilst carbon is a convenient tracer for this purpose as it can be readily separated into a 'natural' and 'anthropogenic' component, it is not necessarily unique as a choice of tracer to use in order to diagnose redistribution. There is therefore scope for the extension of this technique, both in using carbon to diagnose the redistribution of tracers other than temperature and salinity, but also in using tracers other than carbon to diagnose redistribution.

Finally, the matrix inversion technique, which specifies two curves linking excess and redistributed temperature and carbon, is currently implemented as an exact, analytical, matrix inversion. As a result, it does not allow for uncertainty in the model, nor does it consider additional information about how local carbon concentrations may have changed, for example due to biological activity. Current work at NORCE is focusing on extending this matrix inversion technique to include the additional information contained by the link between temperature and oxygen changes: changes in DIC and oxygen due to biological activity are similarly linked. It therefore seems likely that this work will lead to both an improved understanding of excess and redistributed temperature, DIC and salinity, in addition to better understanding of ocean deoxygenation. However, this represents a single avenue for extending the techniques developed in this work. Most fundamentally, by identifying reliable links between changes in temperature and carbon in the ocean, this work has showed that it is possible to link biogeochemistry and physical quantities in ways that allows the accurate local decomposition of changes into circulation driven variability and and the 'excess' tracer. Though it would require the identification or use of other relationships, this allows it to serve as a template for the use of other biogeochemical measurements to inform understanding of ocean physics.

7.3 Final Remarks

The results presented in this work have contributed to the continual improvement of our understanding of how changes in temperature and salinity are propagated into the interior ocean, and how ocean circulation responds to changing climate. However, a major theme connecting these findings is that changes in ocean carbon content can be leveraged to understand other changes to the ocean state: carbon should not be considered simply an isolated, passive tracer which is used to estimate to estimate how the ocean ameliorates increases in atmospheric CO_2. Instead, due to its strong mechanistic covariability with temperature changes in the ocean, it enables us to uncover key information about changes in the ocean state, both anthropogenically forced and due to natural variability. Increased monitoring capabilities of oceanic carbon content therefore represent an opportunity not only to improve understanding of the oceanic storage of carbon in response to climate change, but also to understand how the ocean circulation is changing. Thus, oceanic carbon is important not just for its biogeochemical impacts, but can be utilised to improve understanding of physical changes in the ocean. The results presented in this work therefore most fundamentally underscore the need for continued, and hopefully improved, monitoring and understanding of carbon in the ocean as a method for better understanding the response of the global ocean to anthropogenic climate change.

Appendix A

GLODAP Easy Ocean Toolbox: A Julia wrapper package to easily grid hydrographic sections

In this short chapter, I describe a toolbox produced as part of this project in the Julia programming language, which may be used to generate gridded 2D (horizontal coordinate, depth) or 3D (horizontal coordinate, depth, time) fields from any variable contained within the GLODAP dataset. Gridded fields produced using this toolbox are discussed in the following two chapters, and therefore the key principles of this toolbox and a top level overview are provided here for simplicity. Work on this toolbox is ongoing, however, functionality is sufficiently mature that it may be of use to others. The first four sections in this Chapter describe the various components integrated into this toolbox, and the 5th how the toolbox top levelAPI works.

A.1 GLODAP

GLODAP (Global Ocean Data Analysis Project) (Lauvset et al., 2021) is a project which aims to synthesise data for ocean biogeochemical observations. It comprises over 1.3 million samples taken from 989 cruises, and contains data for 12 core variables: salinity, oxygen, phosphate, nitrate, silicate, dissolved inorganic carbon, total alkalinity, pH, CFC-11, CFC-12, CFC-113 and CCl4, which are subjected to primary and secondary quality control, with the aim of identifying outliers and correcting for measurement biases. It also contains data for potential temperature, neutral density and a number more variables. The cruises included within GLODAP span 1870-2020, although the density of data is substantially lower pre 1990 and post 2016 (due to data paucity and the work required to integrate measurements into the dataset, respectively). In total,

over 100 variables are included, including more esoteric variables such as Neon and Helium isotopes. It therefore represents an extremely valuable resource for ocean biogeochemical research.

However, GLODAP is limited in its scope: while a 3D gridded climatology is provided, due to the challenges associated with gridding observations onto standard grids, particularly when these observations are not simulaneous, no gridded fields are provided. Instead, all variables are provided as a vector of observations. As such, whilst is is straightforward to use observations contained in GLODAP in order to visualise distributions of data from hydrography, comparison of, for example, temperature, from one cruise to the next is not necessarily straightforward, as observations are not necessarily collocated. It is therefore down to the user to grid GLODAP data in the fashion that best suits their requirements.

A.1.1 GO-SHIP

Similar to GLODAP, GO-SHIP (The Global Ocean Ship-Based Hydrographic Investigations Program) is an effort to coordinate ship-based repeat hydrographic occupations into a coordinated observational project. However, unlike GLODAP, which aims to compile, cross validate and synthesise existing measurements, GO-SHIP instead aims to organise the collection of data in order to maximise the scientific returns of these data collection efforts. By coordinating hydrographic surveys amongst nations and research institutes, GO-SHIP provides repeat hydrographic occupations of sections of interest of the global ocean at approximately decadal resolution, and at higher temporal resolution for key data. Further detail can be found at `https://www.go-ship.org/index.html` and `http://cchdo.ucsd.edu/`.

Additionally, GO-SHIP provide a toolbox, GO-SHIP Easy Ocean (Katsumata et al., 2022), which contains grid data, masks, and Temperature, Salinity, and Oxygen data for a number of cruises contained within the GO-SHIP database. The spatial data contained within this toolbox form the basis for the gridding used in the Julia toolbox described in this work.

A.2 Julia

Julia is a modern, high level, high performance, dynamic programming language, developed primarily for use in computational science (Bezanson et al., 2017). Code written in Julia is superficially similar to both Python and MATLAB code, with features such as list comprehensions, broadcasting, and efficient built in libraries for linear algebra. However, unlike Python and MATLAB, it is designed with performance as a

primary consideration. As such, Julia uses a 'just-ahead-of-time' compiler, rather than a 'just-in-time' interpreter, as is the case for MATLAB and Python.

In Julia, upon the first call to a function in a given REPL session, code is compiled to bytecode. All subsequent calls to this function then call the bytecode implementation, avoiding the overhead associated with an interpreted language. This allows the user to write code which performs comparably to a compiled language such as C or FORTRAN. This makes it an ideal language for the implementation of expensive computations: more details can be found at `https://julialang.org/`

A.2.1 DIVAnd

DIVAnd (Data Interpolating Variational Analysis in n dimensions) is a gridding algorithm and associated software package which, similarly to optimal interpolation and Kriging, may be used to grid ocean observations (Barth et al., 2014). Application of DIVAnd to observational data will create an interpolated field by minimising a cost function, which penalises deviations from observations, deviation from an initial guess, and abrupt variation in the field based on a correlation length. It can also include additional constraints, such as ocean circulation. However, unlike optimal interpolation, it naturally decouples observations in topologically disconnected regions. This is useful as we do not expect topologically disconnected water masses to have similar properties necessarily. For example, the Western Pacific sea surface is typically much fresher than the Gulf of Mexico. Due to the narrow separation between these regions, standard optimal interpolation will tend to artificially salinify the Western Pacific and freshen the Gulf of Mexico, unless care is taken. DIVAnd does not suffer from this issue.

The DIVAnd software package (accessible at `https://github.com/gher-ulg/DIVAnd.jl` or through the Julia package manager (`using Pkg; Pkg.add("DIVAnd")`) contains a number of functions for gridding, calculation of correlation lengths (horizontal and vertical), and to estimate the error associated with the gridding process. Full documentation can be found at `https://gher-ulg.github.io/DIVAnd.jl/latest/index.html`.

Finally, detailed theory describing how the DIVAnd gridding algorithm works can be found in Barth et al. (2014), or in Troupin et al. (2019).

A.3 GLODAP Easy Ocean Toolbox

By combining these resources, we may interpolate data from GLODAP onto hydrographic sections. At present, this toolbox relies on sections defined by the GO-SHIP Easy Ocean Toolbox Katsumata et al. (2022), as well as a user defined section for 24 South in the Atlantic (Nominal WOCE code A09.5), see discussion at `https://www.`

go-ship.org/Docs/A9.5_A10_A10.5__final.pdf. Whilst is possible for a user to define sections as above, this functionality is not currently built into the toolbox as an intended use case: this will be added at a later date, and discussion herein limited to only the predefined sctions from the GO-SHIP Easy Ocean toolbox.

At the top level, the toolbox provides two pipeline functions, which allow a user to generate a section or repeat of gridded data with a single line of code: `gridCruisePipeline` and `gridSectionPipeline`. Function calls in Julia may take positional and/or keyword arguments, and defaults are readily specified. A typical call to `gridCruisePipeline` is now described, and how the toolbox operates in turn. Note that some functionality still requires refining.

`gridCruisePipeline` takes three data arguments: `glodapDir`, `goshipDir` and `maskMatFile`, all of which must be strings. They specify the directory in which the pipeline will look for the GLODAP database file (the current implementation uses the .mat file version), the top level directory in which the GO-SHIP Easy Ocean toolbox is saved, and the file containing all of the masks which specify topography for the available sections, respectively, and all have defaults set. Four arguments without defaults are then required: `sectionName`, `horzCoordinate`, `variableName` and `expocode`, again all of which must be strings. `sectionName` is a WOCE code, which specifies the section which we wish to grid, for example "A05". `horzCoordinate` specifies whether to use "latitude" or "longitude" as the horizontal coordinate for this section (future versions will automatically compute the horizontal coordinate from the section name). `variableName` specifies which variable we wish to grid: it must be contained within the GLODAP dataset, for example "G2theta" for temperature.

All further arguments have defaults set, and need not be specified. However, they are described here for completeness. The argument `gridding`, which must be a string, specifies how to grid data. By default, it is set to "isobaric", but may also be set to "isopycnic" to perform isopycnic gridding. Isopycnic gridding is currently still experimental, and occasionally gives nonsensical results, but usually performs well.

`meanValue` specifies the how the background field value is calculated. The DIVA gridding algorithm penalises deviations from a background field value, and so it can be necessary to correctly set the background value to avoid spurious results. By default, `meanValue` takes an unweighted mean value of all observations for a given cruise and uses this as a mean background value for the field: this tends to work well in most cases. However, `meanValue` can also be set to either "horzMean" or "climatology". In the case of "horzMean", the horizontal mean value of the variable to be gridded will be calculated and used as a background field. In the case of "climatology", the climatological mean value of the variable to be gridded will be used as a background field. However, "climatology" relies on the existence of a climatology file from GLODAP, and it's use is currently limited to temperature, salinity, and DIC.

`epsilonVal`, default value 0.1, specifies the error variance of the observations. This value is passed directly to the DIVAndrun (the function used to call the DIVA gridding algorithm in the Julia package) argument `epsilon2`, and so is documented there. The smaller `epsilonVal`, the more tightly constrained the output field is to conform to the observations. Larger values therefore smooth the gridded field.

`plotResults` is a boolean with default value false. If specified to be true, after gridding, the gridded field will be plotted as a colourmap, with the section name, variable name, and expocode of the cruise which has been gridded.

`autoTruncateMask` is also a boolean with default value false. If set to true, in the case of a partial section occupation, regions of the section which were not occupied will be automatically masked based on recorded longitude values.

Finally, there are two variables which relate to automatic cross validation: `crossValidate`, which is a boolean, default false, and `crossValidationNum`, which is an integer, default 5. If `crossValidate` is set to true, then a cross validation will be performed in order to optimise correlation lengths and the epsilon value supplied to the final gridding. If this cross validation is performed, `crossValidationNum` will be used to determine the number of cross validations which are performed.

`gridCruisePipeline` therefore works as follows:

1. Inputs are checked to ensure they are compatible.
2. The variable and expocode are checked for exceptions. If an exception is found, the exception data is loaded from a separate .mat file, rather than from the main GLODAP file.
3. If DIC is to be intepolated, GLODAP recommended adjustments are read and automatically applied.
4. If the mask is to be automatically truncated, whether the cruise occupied the full section is checked, and an if it is found not to, the mask is truncated.
5. Horizontal and vertical correlation lengths are then calculated. The user may supply a list of correlation length multiplicative factors (contained in a .csv file), which can then be applied to the correlation lengths if necessary for a given cruise and variable. This can be necessary when data are only collected for a given variable every 5th station, for example.
6. The DIVA gridding routine is called, and the results returned. If the user specifies, the final gridded field is also plotted as a heatmap.

The `gridSectionPipeline` function works similarly, although here, expocodes are not specified manually (through the argument `expocode`): instead, a section name is passed (as with `gridCruisePipeline` using the argument `sectionName`). This function instead will find all expocodes associated with a given section, either automatically if the argument `convDir` is left blank, or from a manually supplied list if `convDir` is specified. It then performs the same analysis as `gridCruisePipeline` for each expocode found in the supplied list of expocodes. However, there is no functionality to perform a cross validation, nor plot results.

As only the top level functionality is necessary to understand how this toolbox may be used to generate gridded sections of variables from the GLODAP dataset, no further details about code structure will be supplied here. However, the full code (with extensive documentation and some examples) are available at https://github.com/charles-turner-1/GLODAP_Section_Gridder.jl

Glossary

The glossary presented here is based upon that from the paper Turner et al. (2022), in which I explore how we can use the principles illustrated in Section 2.3 in order to understand changes in global ocean temperature and carbon.

For an arbitrary tracer Q, transported by a velocity field $\mathbf{v}$, we may write

$$\vec{v}Q = (\vec{v}_0 + \vec{v}')(Q_0 + Q') = \underbrace{\vec{v}_0 Q_0}_{Preindustrial} + \underbrace{\vec{v}' Q_0}_{Redistributed} + \underbrace{\vec{v}_o Q' + \vec{v}' Q'}_{Excess}, \tag{A.1}$$

where $\vec{v}_o$ and Q_0 refer to the preindustrial components of the velocity field, $\vec{v}$, and the tracer field, Q, and $\vec{v}'$ and Q' the perturbations. The excess and redistributed changes in Q, denoted Q_e and Q_r respectively, are therefore given by

$$Q_e(t) = \int_{t_0}^{t} \left(F'_Q - (\vec{v}_0 + \vec{v}') \cdot \nabla Q' \right) dt, \tag{A.2}$$

and

$$Q_r(t) = -\int_{t_0}^{t} \left(\vec{v}' \cdot \nabla Q_0 \right) dt, \tag{A.3}$$

where F'_Q is the anomalous surface flux in Q, t_0 is a preindustrial time, and t is a generic time. These definitions are described in further detail in Williams et al. (2021) and below.

- **Excess Q**: Changes in the local ocean Q field value due to the imposition of changes in the surface forcing of the Q field. Excess Q may be positive or negative, depending on changes in surface forcing.
- **Excess Temperature**: Change in local ocean temperature due to changing surface heat fluxes, for example warming due to increased radiative forcing at the sea surface.
- **Excess Salinity**: Change in local ocean salinity due to changing ocean freshwater fluxes, for example salinification as a result of increased evaporation and/or reduced precipitation at the sea surface.

- **Redistributed Q**: Changes in the local ocean Q field value due to changes in ocean transport, either imposed in response to climate change or as the result of natural variability. As redistribution can only rearrange the inventory of Q within the global ocean, the global ocean inventory of redistributed Q must always sum to zero, as positive redistributed Q in one location must be compensate for by negative redistributed Q in another location.

- **Redistributed Temperature**: Changes in local ocean temperature due to circulation change, for example cooling in the North Atlantic due to the reduction of northward heat transport associated with AMOC decline.

- **Redistributed Salinity**: Changes in local ocean salinity due to circulation change, for example salinification in the South Atlantic due to the reduction of northward freshwater transport associated with AMOC decline.

- **DIC**: Dissolved Inorganic Carbon, also known as tCO_2. This is the total local inorganic carbon content. It may be decomposed as

$$\mathrm{DIC} = \mathrm{DIC_{sat}} + \mathrm{DIC_{carb}} + \mathrm{DIC_{soft}} + \mathrm{DIC_{diseq}} + \mathrm{C_{anth}} = \mathrm{C_{nat}} + \mathrm{C_{anth}} \tag{A.4}$$

where

$$\mathrm{C_{nat}} = \mathrm{DIC_{sat}} + \mathrm{DIC_{carb}} + \mathrm{DIC_{soft}} + \mathrm{DIC_{diseq}} \tag{A.5}$$

- **$\mathrm{DIC_{sat}}$**: Saturation carbon, the DIC content which a parcel of water would have, if allowed to equilibrate with the preindustrial atmosphere at its potential temperature and salinity. It accounts for the bulk of DIC concentrations, around 2000μmol/kg.

- **$\mathrm{DIC_{carb}}$**: Carbonate carbon, DIC content due to the remineralisation of calcium carbonate. Concentrations increase with age, with concentrations up to approximately 60 μmol/kg in the oldest waters.

- **$\mathrm{DIC_{soft}}$**: Soft tissue carbon, DIC content due to the remineralisation of soft tissue. As with DIC_{carb}, its' concentration increases with the age of water, up to approximately 200μmol/kg in the oldest waters.

- **$\mathrm{DIC_{diseq}}$**: Disequilibrium carbon, the DIC content due to the disequilibrium of a parcel of water with the overlying atmosphere, when subducted away from the surface. It may be either positive or negative.

- **$\mathrm{C_{anth}}$**: Anthropogenic carbon, the DIC content of a parcel of water due to equilibration with the increased atmospheric CO_2 content of the atmosphere, relative to preindustrial. It is defined as having a preindustrial concentration of zero, and hence is closely related to excess DIC (to see this, let $Q_0 = 0$ in Equation F1.

- **C_{nat}**: Natural carbon, the DIC content of a parcel of water with the contribution from increased atmospheric CO_2 concentrations removed. It is the sum of the saturation, soft tissue, carbonate and disequilibrium pools of DIC. As excess DIC and C_{anth} are closely related, C_{nat} therefore approximates redistributed DIC. However, in response to a warming ocean, the global C_{nat} inventory will decline, causing it to systematically differ from redistributed DIC.

- **C_{nat}^{adj}**: Adjusted Natural Carbon, calculated as $C_{nat}^{adj} = C_{nat} + \gamma C_{anth}$, where γ is a factor between 0 and 1. This aims to correct for the outgassing of Saturation Carbon in response to ocean warming, in order to adjust for the systematic reduction in natural carbon leading to inconsistency in the definition of changes in natural carbon and the redistribution of DIC.

- **Excess DIC**: Changes in local DIC content driven by changes in surface conditions: these include changes to surface wind forcing, SST driven change in CO_2 solubility in surface, but predominately those due increases in atmospheric CO_2 concentrations.

- **Redistributed DIC**: Changes in local ocean DIC content due to circulation change, for example an increase in the DIC concentration of the deep North Atlantic due to reduced formation of North Atlantic Deep Water.

- **Carbon-Heat Coupling**: An emergent near-linear relationship between changes in both local and global heat and carbon content. Typically referred to as α when presented in units of heat and carbon inventories, or α_T when presented in units of local temperature and carbon concentration change.

- **Redistribution Coefficient**: A coefficient that relates the change in two tracers due to redistribution. Typically referred to as either κ_r when used in isolation, or as β when used in conjuction with the carbon heat coupling.

References

Myles R. Allen, David J. Frame, Chris Huntingford, Chris D. Jones, Jason A. Lowe, Malte Meinshausen, and Nicolai Meinshausen. Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature*, 458(7242):1163–1166, 2009. ISSN 00280836. .

Richard B. Alley and Peter U. Clark. The deglaciation of the northern hemisphere: A global perspective. *Annual Review of Earth and Planetary Sciences*, 27(1):149–182, 1999. . URL https://doi.org/10.1146/annurev.earth.27.1.149.

John I. Antonov, Sydney Levitus, and Timothy P. Boyer. Steric sea level variations during 1957–1994: Importance of salinity. *Journal of Geophysical Research: Oceans*, 107 (C12):SRF 14–1–SRF 14–8, 2002. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2001JC000964.

David Archer, Haroon Kheshgi, and Ernst Maier-Reimer. Multiple timescales for neutralization of fossil fuel co2. *Geophysical Research Letters*, 24(4):405–408, 1997. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/97GL00168.

Helene T. Banks and Jonathan M. Gregory. Mechanisms of ocean heat uptake in a coupled climate model and the implications for tracer based predictions of ocean heat uptake. *Geophysical Research Letters*, 33(7):3–6, 2006. ISSN 00948276. .

A. Barth, J.-M. Beckers, C. Troupin, A. Alvera-Azcárate, and L. Vandenbulcke. DIVAnd-1.0: n-dimensional variational data analysis for ocean observations. *Geoscientific Model Development*, 7:225–241, 2014. . URL http://www.geosci-model-dev.net/7/225/2014/gmd-7-225-2014.html.

Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017. URL https://doi.org/10.1137/141000671.

Nathaniel L. Bindoff and Trevor J. Mcdougall. Diagnosing climate change and ocean ventilation using hydrographic data. *Journal of Physical Oceanography*, 24(6):1137 – 1152, 1994. . URL https://journals.ametsoc.org/view/journals/phoc/24/6/1520-0485_1994_024_1137_dccaov_2_0_co_2.xml.

L. Bopp, M. Lévy, L. Resplandy, and J. B. Sallée. Pathways of anthropogenic carbon subduction in the global ocean. *Geophysical Research Letters*, 42(15):6416–6423, 2015. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2015GL065073.

J. Boé, A. Hall, and X. Qu. Deep ocean heat uptake as a major source of spread in transient climate change simulations. 36(22):L22701, 2009. ISSN 0094-8276. . URL http://doi.wiley.com/10.1029/2009GL040845.

Wallace S. Broecker. *The Great Ocean Conveyor*. 1991. URL https://doi.org/10.5670/oceanog.1991.07.

Ben Bronselaer and Laure Zanna. Heat and carbon coupling reveals ocean warming due to circulation changes. *Nature*, 584(7820):227–233, 2020. ISSN 0028-0836, 1476-4687. . URL http://www.nature.com/articles/s41586-020-2573-5.

Harry L. Bryden, Michael J. Griffiths, Alicia M. Lavin, Robert C. Millard, Gregorio Parrilla, and William M. Smethie. Decadal changes in water mass characteristics at 24°n in the subtropical north atlantic ocean. *Journal of Climate*, 9(12):3162–3186, 1996. ISSN 08948755, 15200442. URL http://www.jstor.org/stable/26201281.

Martha W. Buckley and John Marshall. Observations, inferences, and mechanisms of the atlantic meridional overturning circulation: A review. 54(1):5–63. ISSN 8755-1209, 1944-9208. . URL https://onlinelibrary.wiley.com/doi/abs/10.1002/2015RG000493.

L. Caesar, S. Rahmstorf, A. Robinson, G. Feulner, and V. Saba. Observed fingerprint of a weakening atlantic ocean overturning circulation. *Nature*, 556(7700):191–196, Apr 2018. ISSN 1476-4687. . URL https://doi.org/10.1038/s41586-018-0006-5.

J.A Church, N.J White, L.F Konikow, C.M Domingues, J.G Cogley, E Rignot, J.M. Gregory, M.R van der Broeke, A.J Monaghan, and I Velicogna. Revisiting the Earth's sea-level and energy budgets from 1961 to 2008. *Geophysical Research Letters*, 38, 2011. ISSN 0370629X. .

J.A. Church, P.U. Clark, A. Cazenave, J.M. Gregory, S. Jevrejeva, A. Levermann, M.A. Merrifield, G.A. Milne, R.S. Nerem, P.D. Nunn, A.J. Payne, W.T. Pfeffer, D. Stammer, and A.S. Unnikrishnan. Sea level change. In T.F. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P.M. Midgley, editors, *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, pages 1137–1216. Cambridge University Press, 2013. ISBN ISBN 978-1-107-66182-0. . URL www.climatechange2013.org. Section: 13 Type: Book Section.

John A. Church, J. Stuart Godfrey, David R. Jackett, and Trevor J. McDougall. A model of sea level rise caused by ocean thermal expansion. *Journal of Climate*, 4(4):438

– 456, 1991. . URL https://journals.ametsoc.org/view/journals/clim/4/4/1520-0442_1991_004_0438_amoslr_2_0_co_2.xml.

Dominic Clement and Nicolas Gruber. The emlr(c*) method to determine decadal changes in the global ocean storage of anthropogenic co2. *Global Biogeochemical Cycles*, 32(4):654–679, 2018. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2017GB005819.

Louis Clément, E. L. McDonagh, J. M. Gregory, Q. Wu, A. Marzocchi, J. D. Zika, and A. J. G. Nurser. Mechanisms of ocean heat uptake along and across isopycnals. *Journal of Climate*, 35(15):4885 – 4904, 2022. . URL https://journals.ametsoc.org/view/journals/clim/35/15/JCLI-D-21-0793.1.xml.

W. J. Collins, N. Bellouin, M. Doutriaux-Boucher, N. Gedney, P. Halloran, T. Hinton, J. Hughes, C. D. Jones, M. Joshi, S. Liddicoat, G. Martin, F. O'Connor, J. Rae, C. Senior, S. Sitch, I. Totterdell, A. Wiltshire, and S. Woodward. Development and evaluation of an earth-system model – hadgem2. *Geoscientific Model Development*, 4(4):1051–1075, 2011. . URL https://gmd.copernicus.org/articles/4/1051/2011/.

Matthew Couldrey. *Mechanisms of ocean carbon cycle variability in the 21st Century*. PhD thesis, University of Southampton, 2018.

Matthew P. Couldrey, Kevin I. C. Oliver, Andrew Yool, Paul R. Halloran, and Eric P. Achterberg. On which timescales do gas transfer velocities control north atlantic co2 flux variability? *Global Biogeochemical Cycles*, 30(5):787–802, 2016. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2015GB005267.

Matthew P. Couldrey, Kevin I. C. Oliver, Andrew Yool, Paul R. Halloran, and Eric P. Achterberg. Drivers of 21st Century carbon cycle variability in the North Atlantic Ocean. *Biogeosciences Discussions*, (January):1–33, 2019. ISSN 1810-6285. .

U Cubasch, G Meehl, G J Boer, Ron Stouffer, M Dix, A Noda, C A Senior, S Raper, K S Yap, A Abe-Ouchi, S Brinkop, M Claussen, M Collins, J Evans, I Fischer-Bruns, G Flato, J C Fyfe, A Ganopolski, J M Gregory, Z Z Hu, Fortunat Joos, T Knutson, R Knutti, C Landsea, L O Mearns, C Milly, J F Mitchell, T Nozawa, H Paeth, J Raisanen, R Sausen, Steven J Smith, T Stocker, A Timmermann, U Ulbrich, A Weaver, J Wegner, P Whetton, T M Wigley, M Winton, F Zwiers, J W Kim, and J Stone. Projections of future climate change. Technical report. URL https://www.osti.gov/biblio/901481.

H. Damon Matthews, Katarzyna B. Tokarska, Joeri Rogelj, Christopher J. Smith, Andrew H. MacDougall, Karsten Haustein, Nadine Mengis, Sebastian Sippel, Piers M. Forster, and Reto Knutti. An integrated approach to quantifying uncertainties in the remaining carbon budget. *Communications Earth & Environment*, 2(1):7, 2021. . URL https://doi.org/10.1038/s43247-020-00064-9.

Xabier Davila, Geoffrey Gebbie, Ailin Brakstad, Siv K. Lauvset, Elaine L. McDonagh, Jörg Schwinger, and Are Olsen. How is the ocean anthropogenic carbon reservoir filled? *Global Biogeochemical Cycles*, 36(5):e2021GB007055, 2022. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2021GB007055. e2021GB007055 2021GB007055.

Russ E. Davis, Lynne D. Talley, Dean Roemmich, W. Brechner Owens, Daniel L. Rudnick, John Toole, Robert Weller, Michael J. McPhaden, and John A. Barth. 100 years of progress in ocean observing systems. *Meteorological Monographs*, 59:3.1 – 3.46, 2019. . URL https://journals.ametsoc.org/view/journals/amsm/59/1/amsmonographs-d-18-0014.1.xml.

Damien Desbruyères, Elaine L. McDonagh, Brian A. King, and Virginie Thierry. Global and full-depth ocean temperature trends during the early twenty-first century from argo and repeat hydrography. *Journal of Climate*, 30(6):1985 – 1997, 2017. . URL https://journals.ametsoc.org/view/journals/clim/30/6/jcli-d-16-0396.1.xml.

Tim DeVries. The oceanic anthropogenic co2 sink: Storage, air-sea fluxes, and transports over the industrial era. *Global Biogeochemical Cycles*, 28(7):631–647, 2014. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2013GB004739.

Tim DeVries, Mark Holzer, and Francois Primeau. Recent increase in oceanic carbon uptake driven by weaker upper-ocean overturning. *Nature*, 542(7640):215–218, 2017. ISSN 14764687. . URL http://dx.doi.org/10.1038/nature21068.

Catia M Domingues, John A Church, Neil J White, Peter J Gleckler, Susan E Wijffels, Paul M Barker, and Jeff R Dunn. Improved estimates of upper-ocean warming and multi-decadal sea-level rise. *Nature*, 453(7198):1090—1093, June 2008. ISSN 0028-0836. . URL https://doi.org/10.1038/nature07080.

Paul J. Durack and Susan E. Wijffels. Fifty-year trends in global ocean salinities and their relationship to broad-scale warming. *Journal of Climate*, 23(16):4342 – 4362, 2010. . URL https://journals.ametsoc.org/view/journals/clim/23/16/2010jcli3377.1.xml.

Eric S. Egleston, Christopher L. Sabine, and François M. M. Morel. Revelle revisited: Buffer factors that quantify the response of ocean chemistry to changes in DIC and alkalinity: REVELLE REVISITED. 24(1):n/a–n/a, 2010. ISSN 08866236. . URL http://doi.wiley.com/10.1029/2008GB003407.

Dana Ehlert, Kirsten Zickfeld, Michael Eby, and Nathan Gillett. The sensitivity of the proportionality between temperature change and cumulative co emissions to ocean mixing. *Journal of Climate*, 30(8):2921–2935, 2017. ISSN 08948755, 15200442. URL https://www.jstor.org/stable/26387983.

Captain Henry Ellis and Stephen Hales. Henry Ellis. *Philosophical Transactions of the Royal Society*, 47:211–216, 1751. .

Dafydd Gwyn Evans, Jan D. Zika, Alberto C. Naveira Garabato, and A. J. George Nurser. The imprint of southern ocean overturning on seasonal water mass variability in drake passage. *Journal of Geophysical Research: Oceans*, 119(11):7987–8010, 2014. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2014JC010097.

Raffaele Ferrari and David Ferreira. What processes drive the ocean heat transport? *Ocean Modelling*, 38(3):171–186, 2011. ISSN 1463-5003. . URL https://www.sciencedirect.com/science/article/pii/S1463500311000485.

Yvonne L. Firing, Elaine L. McDonagh, Brian A. King, and Damien G. Desbruyères. Deep temperature variability in drake passage. *Journal of Geophysical Research: Oceans*, 122(1):713–725, 2017. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2016JC012452.

G. Gebbie and P. Huybers. The little ice age and 20th-century deep pacific cooling. *Science*, 363(6422):70–74, 2019. . URL https://www.science.org/doi/abs/10.1126/science.aar8413.

Nathan P. Gillett, Vivek K. Arora, Damon Matthews, and Myles R. Allen. Constraining the ratio of global warming to cumulative CO2 emissions using CMIP5 simulations*. 26(18):6844–6858, 2013. ISSN 0894-8755, 1520-0442. . URL https://journals.ametsoc.org/jcli/article/26/18/6844/34422/Constraining-the-Ratio-of-Global-Warming-to.

Philip Goodwin, Richard G. Williams, Andy Ridgwell, and Michael J. Follows. Climate sensitivity to the carbon cycle modulated by past and future changes in ocean chemistry. *Nature Geoscience*, 2(2):145–150, Feb 2009. ISSN 1752-0908. . URL https://doi.org/10.1038/ngeo416.

Philip Goodwin, Richard G. Williams, and Andy Ridgwell. Sensitivity of climate to cumulative carbon emissions due to compensation of ocean heat and carbon uptake. *Nature Geoscience*, 8(1):29–34, Jan 2015. ISSN 1752-0908. . URL https://doi.org/10.1038/ngeo2304.

W. John Gould and Stuart A. Cunningham. Global-scale patterns of observed sea surface salinity intensified since the 1870s. *Communications Earth & Environment*, 2(1):76, Apr 2021. . URL https://doi.org/10.1038/s43247-021-00161-3.

H. D. Graven, R. F. Keeling, S. C. Piper, P. K. Patra, B. B. Stephens, S. C. Wofsy, L. R. Welp, C. Sweeney, P. P. Tans, J. J. Kelley, B. C. Daube, E. A. Kort, G. W. Santoni, and J. D. Bent. Enhanced seasonal exchange of co2 by northern ecosystems since 1960.

Science, 341(6150):1085–1089, 2013. . URL https://www.science.org/doi/abs/10.1126/science.1239207.

Heather D. Graven. The carbon cycle in a changing climate. *Physics Today*, 69(11):48–54, 2016. . URL https://doi.org/10.1063/PT.3.3365.

J. M. Gregory, N. Bouttes, S. M. Griffies, H. Haak, W. J. Hurlin, J. Jungclaus, M. Kelley, W. G. Lee, J. Marshall, A. Romanou, O. A. Saenko, D. Stammer, and M. Winton. The flux-anomaly-forced model intercomparison project (fafmip) contribution to cmip6: investigation of sea-level and ocean climate change in response to co_2 forcing. *Geoscientific Model Development*, 9(11):3993–4017, 2016. . URL https://gmd.copernicus.org/articles/9/3993/2016/.

Nicolas Gruber. Warming up, turning sour, losing breath: ocean biogeochemistry under global change. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 369(1943):1980–1996, 2011. . URL https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2011.0003.

Nicolas Gruber, Jorge L. Sarmiento, and Thomas F. Stocker. An improved method for detecting anthropogenic CO2 in the oceans. *Global Biogeochemical Cycles*, 10:809–837, 1996. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/96GB01608.

Nicolas Gruber, Dominic Clement, Brendan R Carter, Richard A Feely, Steven Van Heuven, Mario Hoppema, Masao Ishii, Robert M Key, Alex Kozyr, Siv K Lauvset, Claire Lo Monaco, Jeremy T Mathis, Akihiko Murata, Are Olsen, Fiz F Perez, Christopher L Sabine, Toste Tanhua, and Rik Wanninkhof. The oceanic sink for anthropogenic CO2 from 1994 to 2007. *Science*, 363(March):1193, 2019. . URL www.sciencemag.org/content/363/6432/1193/suppl/DC1.

Elisa F. Guallart, Ute Schuster, Noelia M. Fajar, Oliver Legge, Peter Brown, Carles Pelejero, Marie Jose Messias, Eva Calvo, Andrew Watson, Aida F. Ríos, and Fiz F. Pérez. Trends in anthropogenic CO2 in water masses of the Subtropical North Atlantic Ocean. *Progress in Oceanography*, 131:21–32, 2015. ISSN 00796611. . URL http://dx.doi.org/10.1016/j.pocean.2014.11.006.

Timothy M Hall, Thomas W N Haine, and Darryn W Waugh. Inferring the concentration of anthropogenic carbon in the ocean from tracers. *Global Biogeochemical Cycles*, 16(4):78–1–78–15, 2002. ISSN 08866236. . URL http://doi.wiley.com/10.1029/2001GB001835.

Steffen Hetzinger, Jochen Halfar, Zoltán Zajacz, and Max Wisshak. Early start of 20th-century Arctic sea-ice decline recorded in Svalbard coralline algae. *Geology*, 47(10):963–967, 08 2019. ISSN 0091-7613. . URL https://doi.org/10.1130/G46507.1.

Jochen Hinkel, Daniel Lincke, Athanasios T. Vafeidis, Mahé Perrette, Robert James Nicholls, Richard S. J. Tol, Ben Marzeion, Xavier Fettweis, Cezar Ionescu, and Anders Levermann. Coastal flood damage and adaptation costs under 21st century sea-level rise. *Proceedings of the National Academy of Sciences*, 111(9):3292–3297, 2014. . URL https://www.pnas.org/doi/abs/10.1073/pnas.1222469111.

Nagio Hirota, Yukari N. Takayabu, Masahiro Watanabe, and Masahide Kimoto. Precipitation reproducibility over tropical oceans and its relationship to the double itcz problem in cmip3 and miroc5 climate models. *Journal of Climate*, 24(18):4859 – 4873, 2011. . URL https://journals.ametsoc.org/view/journals/clim/24/18/2011jcli4156.1.xml.

Sirpa Häkkinen, Peter B. Rhines, and Denise L. Worthen. Heat content variability in the north atlantic ocean in ocean reanalyses. *Geophysical Research Letters*, 42(8):2901–2909, 2015. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2015GL063299.

J. Imbrie, A. Berger, E. A. Boyle, S. C. Clemens, A. Duffy, W. R. Howard, G. Kukla, J. Kutzbach, D. G. Martinson, A. McIntyre, A. C. Mix, B. Molfino, J. J. Morley, L. C. Peterson, N. G. Pisias, W. L. Prell, M. E. Raymo, N. J. Shackleton, and J. R. Toggweiler. On the structure and origin of major glaciation cycles 2. the 100,000-year cycle. *Paleoceanography*, 8(6):699–735, 1993. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/93PA02751.

D. R. Jackett, T. J. McDougall, M. H. England, and A. C. Hirst. Thermal expansion in ocean and coupled general circulation models. *Journal of Climate*, 13(8):1384 – 1405, 2000. . URL https://journals.ametsoc.org/view/journals/clim/13/8/1520-0442_2000_013_1384_teioac_2.0.co_2.xml.

W. E. Johns, M. O. Baringer, L. M. Beal, S. A. Cunningham, T. Kanzow, H. L. Bryden, J. J. M. Hirschi, J. Marotzke, C. S. Meinen, B. Shaw, and R. Curry. Continuous, array-based estimates of atlantic ocean heat transport at 26.5°n. *Journal of Climate*, 24(10):2429 – 2449, 2011. . URL https://journals.ametsoc.org/view/journals/clim/24/10/2010jcli3997.1.xml.

A. R. Karspeck, D. Stammer, A. Köhl, G. Danabasoglu, M. Balmaseda, D. M. Smith, Y. Fujii, S. Zhang, B. Giese, H. Tsujino, and A. Rosati. Comparison of the atlantic meridional overturning circulation between 1960 and 2007 in six ocean reanalysis products. *Climate Dynamics*, 49(3):957–982, Aug 2017. ISSN 1432-0894. . URL https://doi.org/10.1007/s00382-015-2787-7.

Anna Katavouta, Richard G. Williams, Philip Goodwin, and Vassil Roussenov. Reconciling atmospheric and oceanic views of the transient climate response to emissions. *Geophysical Research Letters*, 45(12):6205–6214, 2018. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018GL077849.

Anna Katavouta, Richard G. Williams, and Philip Goodwin. The effect of ocean ventilation on the transient climate response to emissions. *Journal of Climate*, 32(16): 5085–5105, 2019. ISSN 08948755, 15200442. URL https://www.jstor.org/stable/26754331.

Katsuro Katsumata, Sarah G. Purkey, Rebecca Cowley, Bernadette M. Sloyan, Stephen C. Diggs, Thomas S. Moore, Lynne D. Talley, and James H. Swift. Go-ship easy ocean: Gridded ship-based hydrographic section of temperature, salinity, and dissolved oxygen. *Scientific Data*, 9(1):103, Mar 2022. ISSN 2052-4463. . URL https://doi.org/10.1038/s41597-022-01212-w.

S Khatiwala, F Primeau, and M Holzer. Ventilation of the deep ocean constrained with tracer observations and implications for radiocarbon estimates of ideal mean age. *Earth and Planetary Science Letters*, 325-326:116–125, 2012. ISSN 0012821X. . URL http://dx.doi.org/10.1016/j.epsl.2012.01.038.

S. Khatiwala, T. Tanhua, S. Mikaloff Fletcher, M. Gerber, S. C. Doney, H. D. Graven, N. Gruber, G. A. McKinley, A. Murata, A. F. Ríos, and C. L. Sabine. Global ocean storage of anthropogenic carbon. *Biogeosciences*, 10(4):2169–2191, 2013. . URL https://bg.copernicus.org/articles/10/2169/2013/.

Samar Khatiwala, Martin Visbeck, and Mark A. Cane. Accelerated simulation of passive tracers in ocean circulation models. *Ocean Modelling*, 9(1):51–69, 2005. ISSN 1463-5003. . URL https://www.sciencedirect.com/science/article/pii/S1463500304000307.

Brian A. King. RRS Discovery Cruise DY040 (RAGNARoCC, 6 December 2015 - 22 January 2016, The 2015 transatlantic hydrography section at 24.5N. Technical report, National Oceanography Centre, 2015. URL https://cchdo.ucsd.edu/data/14741/74EQ20151206_do.pdf.

Klaus Peter Koltermann, Viktor Gouretski, and Karin Jancke. Hydrographic atlas of the world ocean circulation experiment (woce). volume 3: Atlantic ocean. 2011.

S. K. Lauvset, R. M. Key, A. Olsen, S. van Heuven, A. Velo, X. Lin, C. Schirnick, A. Kozyr, T. Tanhua, M. Hoppema, S. Jutterström, R. Steinfeldt, E. Jeansson, M. Ishii, F. F. Perez, T. Suzuki, and S. Watelet. A new global interior ocean mapped climatology: the $1^\circ \times 1^\circ$ glodap version 2. *Earth System Science Data*, 8(2):325–340, 2016. . URL https://essd.copernicus.org/articles/8/325/2016/.

S. K. Lauvset, N. Lange, T. Tanhua, H. C. Bittig, A. Olsen, A. Kozyr, M. Álvarez, S. Becker, P. J. Brown, B. R. Carter, L. Cotrim da Cunha, R. A. Feely, S. van Heuven, M. Hoppema, M. Ishii, E. Jeansson, S. Jutterström, S. D. Jones, M. K. Karlsen, C. Lo Monaco, P. Michaelis, A. Murata, F. F. Pérez, B. Pfeil, C. Schirnick, R. Steinfeldt, T. Suzuki, B. Tilbrook, A. Velo, R. Wanninkhof, R. J. Woosley, and R. M.

Key. An updated version of the global interior ocean biogeochemical data product, glodapv2.2021. *Earth System Science Data*, 13(12):5565–5589, 2021. . URL https://essd.copernicus.org/articles/13/5565/2021/.

John Lazier, Ross Hendry, Allyn Clarke, Igor Yashayaev, and Peter Rhines. Convection and restratification in the labrador sea, 1990–2000. *Deep Sea Research Part I: Oceanographic Research Papers*, 49(10):1819–1835, 2002. ISSN 0967-0637. . URL https://www.sciencedirect.com/science/article/pii/S096706370200064X.

C. Le Quéré, R. M. Andrew, P. Friedlingstein, S. Sitch, J. Hauck, J. Pongratz, P. A. Pickers, J. I. Korsbakken, G. P. Peters, J. G. Canadell, A. Arneth, V. K. Arora, L. Barbero, A. Bastos, L. Bopp, F. Chevallier, L. P. Chini, P. Ciais, S. C. Doney, T. Gkritzalis, D. S. Goll, I. Harris, V. Haverd, F. M. Hoffman, M. Hoppema, R. A. Houghton, G. Hurtt, T. Ilyina, A. K. Jain, T. Johannessen, C. D. Jones, E. Kato, R. F. Keeling, K. K. Goldewijk, P. Landschützer, N. Lefèvre, S. Lienert, Z. Liu, D. Lombardozzi, N. Metzl, D. R. Munro, J. E. M. S. Nabel, S. Nakaoka, C. Neill, A. Olsen, T. Ono, P. Patra, A. Peregon, W. Peters, P. Peylin, B. Pfeil, D. Pierrot, B. Poulter, G. Rehder, L. Resplandy, E. Robertson, M. Rocher, C. Rödenbeck, U. Schuster, J. Schwinger, R. Séférian, I. Skjelvan, T. Steinhoff, A. Sutton, P. P. Tans, H. Tian, B. Tilbrook, F. N. Tubiello, I. T. van der Laan-Luijkx, G. R. van der Werf, N. Viovy, A. P. Walker, A. J. Wiltshire, R. Wright, S. Zaehle, and B. Zheng. Global carbon budget 2018. *Earth System Science Data*, 10(4): 2141–2194, 2018. . URL https://essd.copernicus.org/articles/10/2141/2018/.

Martin Leduc, H. Damon Matthews, and Ramón de Elía. Regional estimates of the transient climate response to cumulative co2 emissions. *Nature Climate Change*, 6(5):474–478, May 2016. ISSN 1758-6798. . URL https://doi.org/10.1038/nclimate2913.

Samuel J. Levang and Raymond W. Schmitt. Centennial changes of the global water cycle in cmip5 models. *Journal of Climate*, 28(16):6489 – 6502, 2015. . URL https://journals.ametsoc.org/view/journals/clim/28/16/jcli-d-15-0143.1.xml.

Qingxiang Li, Wenbin Sun, Boyin Huang, Wenjie Dong, Xiaolan Wang, Panmao Zhai, and Phil Jones. Consistency of global warming trends strengthened since 1880s. *Science Bulletin*, 65(20):1709–1712, 2020. ISSN 2095-9273. . URL https://www.sciencedirect.com/science/article/pii/S2095927320303881.

Jason A. Lowe and Jonathan M. Gregory. Understanding projections of sea level rise in a hadley centre coupled climate model. *Journal of Geophysical Research: Oceans*, 111(C11), 2006. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2005JC003421.

Andrew H. MacDougall and Pierre Friedlingstein. The origin and limits of the near proportionality between climate warming and cumulative co2 emissions. *Journal of Climate*, 28(10):4217 – 4230, 2015. . URL https://journals.ametsoc.org/view/journals/clim/28/10/jcli-d-14-00036.1.xml.

Gurvan Madec and Maurice Imbard. A global ocean mesh to overcome the north pole singularity. *Climate Dynamics*, 12(6):381–388, May 1996. ISSN 1432-0894. . URL https://doi.org/10.1007/BF00211684.

Gurvan Madec, Bourdallé-Badie Romain, Bouttier Pierre-Antoine, Bricaud Clément, Bruciaferri Diego, Calvert Daley, Chanut Jérôme, Clementi Emanuela, Coward Andrew, Delrosso Damiano, Ethé Christian, Flavoni Simona, Graham Tim, Harle James, Iovino Doroteaciro, Dan Lea, Lévy Claire, Lovato Tomas, Martin Nicolas, Masson Sébastien, Mocavero Silvia, Paul Julien, Rousset Clément, Storkey Dave, Storto Andrea, and Vancoppenolle Martin. Nemo ocean engine, October 2017. URL https://doi.org/10.5281/zenodo.3248739. Fix broken cross-references, still revision 8625 from SVN repository.

The HadGEM2 Development Team: G. M. Martin, N. Bellouin, W. J. Collins, I. D. Culverwell, P. R. Halloran, S. C. Hardiman, T. J. Hinton, C. D. Jones, R. E. McDonald, A. J. McLaren, F. M. O'Connor, M. J. Roberts, J. M. Rodriguez, S. Woodward, M. J. Best, M. E. Brooks, A. R. Brown, N. Butchart, C. Dearden, S. H. Derbyshire, I. Dharssi, M. Doutriaux-Boucher, J. M. Edwards, P. D. Falloon, N. Gedney, L. J. Gray, H. T. Hewitt, M. Hobson, M. R. Huddleston, J. Hughes, S. Ineson, W. J. Ingram, P. M. James, T. C. Johns, C. E. Johnson, A. Jones, C. P. Jones, M. M. Joshi, A. B. Keen, S. Liddicoat, A. P. Lock, A. V. Maidens, J. C. Manners, S. F. Milton, J. G. L. Rae, J. K. Ridley, A. Sellar, C. A. Senior, I. J. Totterdell, A. Verhoef, P. L. Vidale, and A. Wiltshire. The hadgem2 family of met office unified model climate configurations. *Geoscientific Model Development*, 4(3):723–757, 2011. . URL https://gmd.copernicus.org/articles/4/723/2011/.

H. Damon Matthews, Nathan P. Gillett, Peter A. Stott, and Kirsten Zickfeld. The proportionality of global warming to cumulative carbon emissions. *Nature*, 459(7248): 829–832, 2009. ISSN 00280836. . URL http://dx.doi.org/10.1038/nature08047.

C. Mauritzen, A. Melsom, and R. T. Sutton. Importance of density-compensated temperature change for deep North Atlantic Ocean heat uptake. *Nature Geoscience*, 5(12): 905–910, December 2012. ISSN 1752-0894, 1752-0908. . URL http://www.nature.com/articles/ngeo1639.

Elaine L. McDonagh, Brian A. King, Harry L. Bryden, Peggy Courtois, Zoltan Szuts, Molly Baringer, Stuart A. Cunningham, Chris Atkinson, and Gerard McCarthy. Continuous estimate of atlantic oceanic freshwater flux at 26.5°n. *Journal of Climate*, 28, 2015.

T. J. McDougall and P. M. Barker. Getting started with teos-10 and the gibbs seawater (gsw) oceanographic toolbox. *SCOR/IAPSO WG127*, page 28, 2011.

Gordon McGranahan, Deborah Balk, and Bridget Anderson. The rising tide: assessing the risks of climate change and human settlements in low elevation coastal zones.

Environment and Urbanization, 19(1):17–37, 2007. . URL https://doi.org/10.1177/0956247807076960.

M. Meinshausen, E. Vogel, A. Nauels, K. Lorbacher, N. Meinshausen, D. M. Etheridge, P. J. Fraser, S. A. Montzka, P. J. Rayner, C. M. Trudinger, P. B. Krummel, U. Beyerle, J. G. Canadell, J. S. Daniel, I. G. Enting, R. M. Law, C. R. Lunder, S. O'Doherty, R. G. Prinn, S. Reimann, M. Rubino, G. J. M. Velders, M. K. Vollmer, R. H. J. Wang, and R. Weiss. Historical greenhouse gas concentrations for climate modelling (cmip6). *Geoscientific Model Development*, 10(5):2057–2116, 2017. . URL https://gmd.copernicus.org/articles/10/2057/2017/.

Sébastien Moreau, Martin Vancoppenolle, Laurent Bopp, Oliver Aumont, Gurvan Madec, Bruno Delille, Jean-Louis Tison, Pierre-Yves Barriat, and Hugues Goosse. Assessment of the sea-ice carbon pump: Insights from a three-dimensional ocean-sea-ice biogeochemical model (NEMO-LIM-PISCES). *Elementa: Science of the Anthropocene*, 4, 08 2016. ISSN 2325-1026. . URL https://doi.org/10.12952/journal.elementa.000122. 000122.

D. M. Murphy, S. Solomon, R. W. Portmann, K. H. Rosenlof, P. M. Forster, and T. Wong. An observationally based energy balance for the earth since 1950. *Journal of Geophysical Research: Atmospheres*, 114(D17), 2009. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2009JD012105.

G. Myhre, D. Shindell, F.-M. Bréon, W. Collins, J. Fuglestvedt, J. Huang, D. Koch, J.-F. Lamarque, D. Lee, B. Mendoza, T. Nakajima, A. Robock, G. Stephens, T. Takemura, and H. Zhang. *Anthropogenic and natural radiative forcing*, pages 659–740. Cambridge University Press, Cambridge, UK, 2013. .

Gunnar Myhre, Eleanor J. Highwood, Keith P. Shine, and Frode Stordal. New estimates of radiative forcing due to well mixed greenhouse gases. *Geophysical Research Letters*, 25(14):2715–2718, 1998. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/98GL01908.

Emily Newsom, Laure Zanna, and Samar Khatiwala. Relating patterns of added and redistributed ocean warming. *Journal of Climate*, pages 1 – 41, 2022. . URL https://journals.ametsoc.org/view/journals/clim/aop/JCLI-D-21-0827.1/JCLI-D-21-0827.1.xml.

Andreas Oschlies, Peter Brandt, Lothar Stramma, and Sunke Schmidtko. Drivers and mechanisms of ocean deoxygenation. *Nature Geoscience*, 11(7):467–473, Jul 2018. ISSN 1752-0908. . URL https://doi.org/10.1038/s41561-018-0152-2.

A. K. Pardaens, J. A. Lowe, S. Brown, R. J. Nicholls, and D. de Gusmão. Sea-level rise and impacts projections under a future scenario with large greenhouse gas emission reductions. *Geophysical Research Letters*, 38(12), 2011. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2011GL047678.

Gregorio Parrilla, Alicia Lavín, Harry Bryden, Maria García, and Robert Millard. Rising temperatures in the subtropical north atlantic ocean over the past 35 years. *Nature*, 369(6475):48–51, May 1994. ISSN 1476-4687. . URL https://doi.org/10.1038/369048a0.

M. Perrette, F. Landerer, R. Riva, K. Frieler, and M. Meinshausen. A scaling approach to project regional sea level rise and its uncertainties. *Earth System Dynamics*, 4(1): 11–29, 2013. . URL https://esd.copernicus.org/articles/4/11/2013/.

David W. Pierce, Peter J. Gleckler, Tim P. Barnett, Benjamin D. Santer, and Paul J. Durack. The fingerprint of human-induced changes in the ocean's salinity and temperature fields. *Geophysical Research Letters*, 39(21), 2012. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2012GL053389.

Sarah G. Purkey and Gregory C. Johnson. Global Contraction of Antarctic Bottom Water between the 1980s and 2000s*. *Journal of Climate*, 25(17):5830–5844, September 2012. ISSN 0894-8755, 1520-0442. . URL http://journals.ametsoc.org/doi/10.1175/JCLI-D-11-00612.1.

Keywan Riahi, Shilpa Rao, Volker Krey, Cheolhung Cho, Vadim Chirkov, Guenther Fischer, Georg Kindermann, Nebojsa Nakicenovic, and Peter Rafaj. Rcp 8.5—a scenario of comparatively high greenhouse gas emissions. *Climatic Change*, 109(1):33, Aug 2011. ISSN 1573-1480. . URL https://doi.org/10.1007/s10584-011-0149-y.

K. B. Rodgers, M. Ishii, T. L. Frölicher, S. Schlunegger, O. Aumont, K. Toyama, and R. D. Slater. Coupling of surface ocean heat and carbon perturbations over the subtropical cells under twenty-first century climate change. *Journal of Climate*, 33(23):10321 – 10338, 2020. . URL https://journals.ametsoc.org/view/journals/clim/33/23/jcliD191022.xml.

Barry Ruddick. A practical indicator of the stability of the water column to double-diffusive activity. *Deep Sea Research Part A. Oceanographic Research Papers*, 30(10):1105–1107, 1983. ISSN 0198-0149. . URL https://www.sciencedirect.com/science/article/pii/0198014983900638.

Joellen L. Russell, Keith W. Dixon, Anand Gnanadesikan, Ronald J. Stouffer, and J. R. Toggweiler. The southern hemisphere westerlies in a warming world: Propping open the door to the deep ocean. *Journal of Climate*, 19(24):6382 – 6390, 2006. . URL https://journals.ametsoc.org/view/journals/clim/19/24/jcli3984.1.xml.

Christopher L. Sabine, Richard A. Feely, Nicolas Gruber, Robert M. Key, Kitack Lee, John L. Bullister, Rik Wanninkhof, C. S. Wong, Douglas W. R. Wallace, Bronte Tilbrook, Frank J. Millero, Tsung-Hung Peng, Alexander Kozyr, Tsueno Ono, and Aida F. Rios. The oceanic sink for anthropogenic co2. *Science*, 305(5682):367–371, 2004. . URL https://www.science.org/doi/abs/10.1126/science.1097403.

Jorge L. Sarmiento and Nicolas Gruber. *Ocean Biogeochemical Dynamics*. Princeton University Press, 2006. ISBN 978-0-691-01707-5. . URL www.jstor.org/stable/j.ctt3fgxqx.

Jorge L. Sarmiento, James C. Orr, and Ulrich Siegenthaler. A perturbation simulation of co2 uptake in an ocean general circulation model. *Journal of Geophysical Research: Oceans*, 97(C3):3621–3645, 1992. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/91JC02849.

Anju Sathyanarayanan, Armin Kohl, and Detlef Stammer. Ocean salinity changes in the global ocean under global warming conditions part 1: Mechanisms in a strong warming scenario. *Journal of Climate*, pages 1 – 56, 2021. . URL https://journals.ametsoc.org/view/journals/clim/aop/JCLI-D-20-0865.1/JCLI-D-20-0865.1.xml.

Raymond W. Schmitt. The ocean's role in climate. *Oceanography*, 31(2):32–40, 2018. ISSN 10428275, 2377617X. URL https://www.jstor.org/stable/26542649.

Christopher R. Schwalm, Spencer Glendon, and Philip B. Duffy. Rcp8.5 tracks cumulative co(2) emissions. *Proceedings of the National Academy of Sciences of the United States of America*, 117(33):19656–19657, Aug 2020. ISSN 1091-6490. . URL https://pubmed.ncbi.nlm.nih.gov/32747549.

Jörg Schwinger, Jerry F. Tjiputra, Christoph Heinze, Laurent Bopp, James R. Christian, Marion Gehlen, Tatiana Ilyina, Chris D. Jones, David Salas-Mélia, Joachim Segschneider, Roland Séférian, and Ian Totterdell. Nonlinearity of ocean carbon cycle feedbacks in cmip5 earth system models. *Journal of Climate*, 27(11):3869 – 3888, 2014. . URL https://journals.ametsoc.org/view/journals/clim/27/11/jcli-d-13-00452.1.xml.

Giovanni Sgubin, Didier Swingedouw, Sybren Drijfhout, Stefan Hagemann, and Eddy Robertson. Multimodel analysis on the response of the amoc under an increase of radiative forcing and its symmetrical reversal. *Climate Dynamics*, 45(5):1429–1450, Sep 2015. ISSN 1432-0894. . URL https://doi.org/10.1007/s00382-014-2391-2.

Nicholas J. Shackleton. The 100,000-year ice-age cycle identified and found to lag temperature, carbon dioxide, and orbital eccentricity. *Science*, 289(5486):1897–1902, 2000. . URL https://www.science.org/doi/abs/10.1126/science.289.5486.1897.

Jeremy D. Shakun, Peter U. Clark, Feng He, Shaun A. Marcott, Alan C. Mix, Zhengyu Liu, Bette Otto-Bliesner, Andreas Schmittner, and Edouard Bard. Global warming preceded by increasing carbon dioxide concentrations during the last deglaciation. *Nature*, 484(7392):49–54, Apr 2012. ISSN 1476-4687. . URL https://doi.org/10.1038/nature10915.

Nikolaos Skliris, Robert Marsh, Simon A. Josey, Simon A. Good, Chunlei Liu, and Richard P. Allan. Salinity changes in the world ocean since 1950 in relation to changing surface freshwater fluxes. *Climate Dynamics*, 43(3):709–736, Aug 2014. ISSN 1432-0894. . URL https://doi.org/10.1007/s00382-014-2131-7.

M. A. Srokosz and H. L. Bryden. Observing the atlantic meridional overturning circulation yields a decade of inevitable surprises. *Science*, 348(6241):1255575, 2015. . URL https://www.science.org/doi/abs/10.1126/science.1255575.

M. Steinacher and F. Joos. Transient earth system responses to cumulative carbon dioxide emissions: linearities, uncertainties, and probabilities in an observation-constrained model ensemble. *Biogeosciences*, 13(4):1071–1103, 2016. . URL https://bg.copernicus.org/articles/13/1071/2016/.

Thomas F. Stocker, Dahe Qin, Gian Kasper Plattner, Melinda M.B. Tignor, Simon K. Allen, Judith Boschung, Alexander Nauels, Yu Xia, Vincent Bex, and Pauline M. Midgley. Climate change 2013 the physical science basis: Working Group I contribution to the fifth assessment report of the intergovernmental panel on climate change. *Climate Change 2013 the Physical Science Basis: Working Group I Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, 9781107057:1–1535, 2013. .

Peter A. Stott, Rowan T. Sutton, and Doug M. Smith. Detection and attribution of atlantic salinity changes. *Geophysical Research Letters*, 35(21), 2008. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2008GL035874.

Taro Takahashi, Jon Olafsson, John G. Goddard, David W. Chipman, and S. C. Sutherland. Seasonal variation of co2 and nutrients in the high-latitude surface oceans: A comparative study. *Global Biogeochemical Cycles*, 7(4):843–878, 1993. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/93GB02263.

L. D. Talley and M. S. McCartney. Distribution and circulation of labrador sea water. *Journal of Physical Oceanography*, 12(11):1189 – 1205, 1982. . URL https://journals.ametsoc.org/view/journals/phoc/12/11/1520-0485_1982_012_1189_dacols_2_0_co_2.xml.

Karl E. Taylor. Summarizing multiple aspects of model performance in a single diagram. *Journal of Geophysical Research: Atmospheres*, 106(D7):7183–7192, 2001. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2000JD900719.

Karl E. Taylor, Ronald J. Stouffer, and Gerald A. Meehl. An overview of cmip5 and the experiment design. *Bulletin of the American Meteorological Society*, 93(4):485 – 498, 2012. . URL https://journals.ametsoc.org/view/journals/bams/93/4/bams-d-11-00094.1.xml.

Laurent Terray, Lola Corre, Sophie Cravatte, Thierry Delcroix, Gilles Reverdin, and Aurélien Ribes. Near-surface salinity as nature's rain gauge to detect human influence on the tropical water cycle. *Journal of Climate*, 25(3):958 – 977, 2012. . URL https://journals.ametsoc.org/view/journals/clim/25/3/jcli-d-10-05025.1.xml.

Jordan Thomas, Darryn Waugh, and Anand Gnanadesikan. Relationship between ocean carbon and heat multidecadal variability. *Journal of Climate*, 31(4):1467 – 1482, 2018. . URL https://journals.ametsoc.org/view/journals/clim/31/4/jcli-d-17-0134.1.xml.

B.J. Thompson, J. Crease, and John Gould. *The Origins, Development and Conduct of WOCE*. Elsevier Ltd, 2001. URL https://www.nodc.noaa.gov/woce/woce{_}v3/wocedata{_}1/wocedocs/wocebook/Chapter{_}1.3.pdf.

Ralph Timmermann, Hugues Goosse, Gurvan Madec, Thierry Fichefet, Christian Ethe, and Valérie Dulière. On the representation of high latitude processes in the orca-lim global coupled sea ice–ocean model. *Ocean Modelling*, 8(1):175–201, 2005. ISSN 1463-5003. . URL https://www.sciencedirect.com/science/article/pii/S146350030300074X.

J. R. Toggweiler and David W. Lea. Temperature differences between the hemispheres and ice age climate variability. *Paleoceanography*, 25(2), 2010. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2009PA001758.

F. Touratier, C. Goyet, C. Coatanoan, and C. Andrié. Assessments of anthropogenic co2 distribution in the tropical atlantic ocean. *Deep Sea Research Part I: Oceanographic Research Papers*, 52(12):2275–2284, 2005. ISSN 0967-0637. . URL https://www.sciencedirect.com/science/article/pii/S096706370500227X.

Franck Touratier and Catherine Goyet. Definition, properties, and atlantic ocean distribution of the new tracer troca. *Journal of Marine Systems*, 46(1):169–179, 2004. ISSN 0924-7963. . URL https://www.sciencedirect.com/science/article/pii/S0924796303001842.

C Troupin, M Ouberdous, D Sirjacobs, A Alvera-Azcarte, A Barth, M.-E. Toussaint, S Watelet, and J.-M. Beckers. Diva user guide. 2019. . URL https://github.com/gher-ulg/Diva-User-Guide/raw/master/DivaUserGuide.pdf.

C. E. Turner, P. J. Brown, K. I. C. Oliver, and E. L. McDonagh. Decomposing oceanic temperature and salinity change using ocean carbon change. *Ocean Science*, 18(2):523–548, 2022. . URL https://os.copernicus.org/articles/18/523/2022/.

Steven Van Heuven, Denis Pierrot, J.W.B Rae, E. Lewis, and Douglas W.R. Wallace. CO2SYS, 2011.

M. Vázquez-Rodríguez, X. A. Padin, A. F. Ríos, R. G. J. Bellerby, and F. F. Pérez. An upgraded carbon-based method to estimate the anthropogenic fraction of dissolved CO2 in the Atlantic Ocean. *Biogeosciences Discussions*, 6(2):4527–4571, 2009. ISSN 1810-6285. . URL https://bg.copernicus.org/preprints/6/4527/2009/.

Peter Wadhams and Walter Munk. Ocean freshening, sea level rising, sea ice melting. *Geophysical Research Letters*, 31(11), 2004. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2004GL020039.

S. Wang, J. K. Moore, F. W. Primeau, and S. Khatiwala. Simulation of anthropogenic co_2 uptake in the ccsm3.1 ocean circulation-biogeochemical model: comparison with data-based estimates. *Biogeosciences*, 9(4):1321–1336, 2012. . URL https://bg.copernicus.org/articles/9/1321/2012/.

Rik Wanninkhof. Relationship between wind speed and gas exchange over the ocean revisited. *Limnology and Oceanography: Methods*, 12(6):351–362, jun 2014. ISSN 15415856. . URL http://doi.wiley.com/10.4319/lom.2014.12.351.

Spencer R. Weart. Global Warming, Cold War, and the Evolution of Research Plans. *Historical Studies in the Physical and Biological Sciences*, 27(2):319–356, 01 1997. ISSN 0890-9997. . URL https://doi.org/10.2307/27757782.

Andrew J. Weaver, Jan Sedláček, Michael Eby, Kaitlin Alexander, Elisabeth Crespin, Thierry Fichefet, Gwenaëlle Philippon-Berthier, Fortunat Joos, Michio Kawamiya, Katsumi Matsumoto, Marco Steinacher, Kaoru Tachiiri, Kathy Tokos, Masakazu Yoshimori, and Kirsten Zickfeld. Stability of the atlantic meridional overturning circulation: A model intercomparison. *Geophysical Research Letters*, 39(20), 2012. . URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2012GL053763.

R.G. Williams and M.J. Follows. *Ocean Dynamics and the Carbon Cycle - Principles and Mechanisms*. Camridge University Press, New York, 2011. ISBN 978-0-521-84369-0.

Richard G. Williams, Anna Katavouta, and Vassil Roussenov. Regional asymmetries in ocean heat and carbon storage due to dynamic redistribution in climate model projections. *Journal of Climate*, 34(10):3907 – 3925, 2021. . URL https://journals.ametsoc.org/view/journals/clim/34/10/JCLI-D-20-0519.1.xml.

Michael Winton, Stephen M. Griffies, Bonita L. Samuels, Jorge L. Sarmiento, and Thomas L. Frölicher. Connecting changing ocean circulation with changing climate. 26(7):2268–2278. ISSN 0894-8755, 1520-0442. . URL https://journals.ametsoc.org/jcli/article/26/7/2268/33101/Connecting-Changing-Ocean-Circulation-with.

Michael Winton, Ken Takahashi, and Isaac M. Held. Importance of ocean heat uptake efficacy to transient climate change. *Journal of Climate*, 23(9):2333 – 2344, 2010. .

URL https://journals.ametsoc.org/view/journals/clim/23/9/2009jcli3139.1.xml.

Michael Winton, Stephen M. Griffies, Bonita L. Samuels, Jorge L. Sarmiento, and Thomas L. Frölicher. Connecting changing ocean circulation with changing climate. *Journal of Climate*, 26(7):2268 – 2278, 2013. . URL https://journals.ametsoc.org/view/journals/clim/26/7/jcli-d-12-00296.1.xml.

Peng Xie and Geoffrey K. Vallis. The passive and active nature of ocean heat uptake in idealized climate change experiments. *Climate Dynamics*, 38(3):667–684, Feb 2012. ISSN 1432-0894. . URL https://doi.org/10.1007/s00382-011-1063-8.

Haijun Yang, Qing Li, Kun Wang, Yu Sun, and Daoxun Sun. Decomposing the meridional heat transport in the climate system. *Climate Dynamics*, 44(9):2751–2768, May 2015. ISSN 1432-0894. . URL https://doi.org/10.1007/s00382-014-2380-5.

A. Yool, E. E. Popova, and T. R. Anderson. MEDUSA-2.0: an intermediate complexity biogeochemical model of the marine carbon cycle for climate change and ocean acidification studies. 6(5):1767–1811, 2013. ISSN 1991-9603. . URL https://gmd.copernicus.org/articles/6/1767/2013/.

Laure Zanna, Samar Khatiwala, Jonathan M. Gregory, Jonathan Ison, and Patrick Heimbach. Global reconstruction of historical ocean heat storage and transport. *Proceedings of the National Academy of Sciences*, 116(4):1126–1131, 2019. . URL https://www.pnas.org/doi/abs/10.1073/pnas.1808838115.

Chenyu Zhu and Zhengyu Liu. Weakening atlantic overturning circulation causes south atlantic salinity pile-up. *Nature Climate Change*, 10(11):998–1003, Nov 2020. ISSN 1758-6798. . URL https://doi.org/10.1038/s41558-020-0897-7.

Kirsten Zickfeld, Michael Eby, H. Damon Matthews, and Andrew J. Weaver. Setting cumulative emissions targets to reduce the risk of dangerous climate change. *Proceedings of the National Academy of Sciences*, 106(38):16129–16134, 2009. . URL https://www.pnas.org/doi/abs/10.1073/pnas.0805800106.

Jan D Zika, Nikolaos Skliris, Adam T Blaker, Robert Marsh, A J George Nurser, and Simon A Josey. Improved estimates of water cycle change from ocean salinity: the key role of ocean warming. *Environmental Research Letters*, 13(7):074036, jul 2018. . URL https://dx.doi.org/10.1088/1748-9326/aace42.

Jan D. Zika, Jonathan M. Gregory, Elaine L. McDonagh, Alice Marzocchi, and Louis Clément. Recent water mass changes reveal mechanisms of ocean warming. 34(9):3461–3479, 2021. ISSN 0894-8755, 1520-0442. . URL https://journals.ametsoc.org/view/journals/clim/34/9/JCLI-D-20-0355.1.xml.