UNIVERSITY OF LIVERPOOL

Impacts of thermal and carbon processes on climate reconstructions and projections

Thesis submitted in accordance with the requirements of the University of Liverpool for the degree of

Doctor in Philosophy

by

Katherine Elise Turner

March 2022

in the Faculty of Science and Engineering
School of Environmental Sciences
Department of Earth, Ocean and Ecological Sciences

Declaration of Authorship

I declare that this thesis titled "Impacts of thermal and carbon processes on climate reconstructions and projections" and the work presented within are my own work. The material contained in the thesis has not been presented, nor is currently being presented, either wholly or in part, for any other degree or qualification.

Signed:

Katherine Elise Turner
Date: 10 March 2022

Abstract

The unprecedented release of CO_2 into the atmosphere through anthropogenic activities has led to a mass transformation of the Earth system. The Earth system response to emissions involves a carbon response and a thermal response, which respectively act to redistribute the extra carbon and heat added by carbon emissions. To first order, the carbon response and heat response compensate each other, which leads to a proportional relationship between cumulative emissions and surface warming, but this compensation relies on a balance between regional differences in the carbon and thermal responses and can be impacted by different regional processes, climate feedbacks, or nonlinearities related to the sign of emissions. In this thesis, the thermal and carbon responses are investigated, with an emphasis on the ocean's role in setting these responses, and the impact of these processes are examined for climate reconstructions and future projections.

The thesis begins by exploring how ocean observations can be used to characterise historical ocean carbon variability. A method is presented through which observational information is propagated to data-sparse regions using CMIP6 model covariances. The method is tested separately with ocean carbon time series sites locations and Argo temperature and salinity profile locations. The ocean carbon time series can be used to describe local ocean carbon variability, but are limited in representing larger-scale variability due to the distance between time series sites being orders of magnitude larger than the ocean carbon correlation length scale. Reconstructions using temperature and salinity measurements are found to represent about half of the interannual and decadal ocean carbon variability in global ocean model tests, due to the strong correlations between ocean carbon content and temperature and salinity content coupled with the wide-spread sampling within the Argo programme.

The thesis continues by expanding theory on the relationship between cumulative carbon emissions and surface warming to consider regional thermal and carbon responses. The expanded theory is applied to an intermediate-complexity Earth System Model integrated under two emissions scenarios to understand the hemispheric responses. The compensation seen on the global scale between the thermal and carbon responses is not reflected on hemispheric scales. The hemispheric responses to emissions show a distinct asymmetry, set mainly by the thermal response. Asymmetries in the radiative forcing induce cross-equatorial heat transport and opposing tendencies in the climate feedback parameter. To first order, the hemispheric carbon responses are symmetric when accounting for the hemispheres' different land/ocean partition, as the carbon response is set mainly by

ocean thermocline ventilation.

Finally, the response to carbon dioxide removal is tested through a large perturbed parameter ensemble of an idealised ocean-atmosphere box model. The model is forced with a set of scenarios that have equivalent carbon drawdown by year 2200 to examine how the response to carbon dioxide removal is constrained by the integrated drawdown. Temperature reductions from carbon dioxide removal are found to continue after drawdown ceases, leading to earlier carbon dioxide removal implementations being more effective at mitigating surface warming. The added benefit of early carbon dioxide removal is enhanced under a low climate feedback parameter. Carbon dioxide removal schemes that add alkalinity to the ocean have the co-benefit of raising surface ocean pH. The surface ocean pH reaches its peak at the year carbon dioxide removal ceases, resulting in a slower carbon dioxide removal implementation having a higher surface pH when the drawdown timelines converge. Ocean overturning acts to impact the responses of temperature and pH changes in similar ways as it works to sequester altered tracers away from the atmosphere and into the ocean thermocline.

The work within this thesis indicates that the asymmetry between the thermal and carbon responses to emissions is a useful feature to understand the response to emissions and its uncertainties. The asymmetric thermal and carbon responses to emissions are present throughout different time periods and forcing scenarios, making them a ubiquitous characteristic of the climate system. The presence of these asymmetries allows for reconstructions of carbon in spite of historically poor observational coverage, quantification of uncertainties in the transient climate response to emissions through regional processes, and carbon dioxide removal policy recommendations based on which climate benefits are most desired.

Acknowledgements

A huge thanks to my supervisor Ric, for offering me the PhD position and being a guiding force in the messy world of research and academia, particularly in the last 18 months of various lockdowns. I would also like to thank my mentor and unofficial supervisor Anna for her hours helping me with my model code and helping me make my work more understandable.

This work would not have been able to be completed without the help of many outside collaborators. A huge thanks to Andreas Oschlies, David Beerling, and Doug Smith for always being able to meet up, even across time zones, and for providing your valuable insights into the three very disparate sections of my PhD.

I never would have been able to even come to the UK if it were not for Andy Cooper and his generous offer to study as part of the Leverhulme Research Centre for Functional Materials Design. Many thanks additionally to Adele, Lucy, and Ben for all of their help and counsel over the years.

A massive thanks as well to Ole and Dave as well as Ric and Anna for taking the time to read through my thesis and providing very helpful feedback. Thank you also to my examiners Jonathan Sharples and Thomas Frölicher for reading my thesis and providing useful constructive criticism for the final submission.

And lastly, a big thanks to my family for supporting my stints in Germany and the UK for my studies, and to my partner Dave for his support and encouragement through the entire PhD.

Contents

Abstract ii

Acknowledgements iv

List of Figures ix

List of Tables xvii

Glossary of Acronyms xviii

1 Introduction 1

1.1 Carbon emissions and warming in the Anthropocene 1

1.2 The ocean's role in understanding the Earth system response to emissions . 3

1.2.1 Regional behaviour in heat and carbon uptake 4

1.3 Utilising the response to emissions to calculate carbon budgets 5

1.4 Carbon Dioxide Removal as a forcing on the Earth system 6

1.4.1 Impacts of CDR on Earth's thermal and carbon responses 8

1.5 The role of climate models in understanding the response to emissions . . . 9

1.5.1 Model complexity and uses 10

1.5.2 Uncertainties and errors within climate models 12

1.6 Thesis Scope and Overview 14

2 Representativeness of ocean carbon time series for regional variability 17

2.1 Recent trends and variability in ocean carbon 18

2.1.1 Availability of ocean interior carbon measurements 19

2.1.2 Expanding sparse observations using optimal interpolation 20

2.1.3 Driving questions and objectives 22

2.2 The optimal interpolation method 23

2.2.1 The CMIP6 model ensemble 24
2.2.2 Choosing a background field 26
2.3 Representativeness of carbon timeseries 30
2.3.1 Correlation fields for carbon at BATS and HOT 31
2.3.2 Error reduction from expanding BATS and HOT time series data . . 33
2.4 Extrapolation limitations of super-sparse datasets 34
2.4.1 Lack of patterns of climate variability 35
2.4.2 Length scale analysis 39
2.5 Discussion and Conclusions 42
2.5.1 Suitability of the Ensemble Optimal Interpolation method for ocean carbon 42
2.5.2 Extrapolating observations using CMIP6 carbon variability 43
2.5.3 Potential and limits of extrapolating time series data 44

3 Using temperature and salinity fields to reconstruct historical ocean carbon changes 48
3.1 Introduction 49
3.1.1 Oceanographic inputs for reconstructing carbon 49
3.1.2 First principles relationships between ocean carbon, temperature, and salinity 50
3.1.3 Driving questions and objectives 52
3.2 Ensemble Optimal Interpolation with temperature and salinity 54
3.2.1 Spatial coverage of Argo floats 54
3.3 Local reconstructions of ocean carbon variability 55
3.3.1 Correlations between upper-ocean temperature, salinity, and carbon 55
3.3.2 Reconstructing carbon with pointwise temperature and salinity . . . 57
3.3.3 Model tests of pointwise carbon reconstructions 59
3.4 Impacts of Argo distribution 64
3.5 Discussion and Conclusions 69
3.5.1 Reconstructing ocean carbon using temperature and salinity correlations 69
3.5.2 Ability to reconstruct global carbon variability using Argo-like inputs 71

4 Hemispheric contrasts in the relationship between surface warming and cumulative carbon emissions 76

4.1 The transient climate response to emissions 77
4.1.1 Theory connecting carbon emissions to surface warming 77
4.1.2 Driving questions and objectives 80
4.2 Methods 81
4.2.1 Expansion of relationships to hemispheric behaviours 81
4.2.2 The University of Victoria Earth System Climate Model 82
4.3 Global and hemispheric UVic ESCM responses to emissions 84
4.3.1 Global UVic ESCM response to RCP forcings 84
4.3.2 Proportion of surface warming to cumulative emissions 87
4.4 Contributions to the hemispheric proportion of surface warming to cumulative emissions 88
4.5 Hemispheric contributions to the thermal response 90
4.5.1 Hemispheric impacts on the climate feedback parameter 91
4.5.2 Diagnosing the UVic ESCM cross-equatorial heat flux 94
4.5.3 Hemispheric energy balances with asymmetric and symmetric radiative forcing 96
4.6 Hemispheric contributions to the global carbon response 98
4.6.1 Hemispheric terrestrial and oceanic carbon responses 98
4.6.2 Structure of ocean carbon pools 99
4.7 Discussion and Conclusions 102
4.7.1 Hemispheric thermal responses to local radiative forcing set by heat transport 104
4.7.2 Hemispheric carbon responses dominated by ocean ventilation symmetry 105
4.7.3 Limitations of the UVic ESCM diagnostics 106
4.7.4 Potential to extend theory to more complex regions and models . . . 107

5 Controls on climate change mitigation through Carbon Dioxide Removal 109
5.1 Carbon dioxide removal as a potential climate change solution 110
5.1.1 Thermal and carbon responses associated with carbon dioxide removal 111
5.1.2 Challenges for climate model studies of negative emissions 112
5.1.3 Driving questions and objectives 113
5.2 The Gnanadesikan model ensemble 115
5.2.1 Model parameter space 116
5.2.2 Model setups with CO_2 emissions and carbon dioxide removal 117

5.3 Model responses to forcings 119
5.3.1 The climate response to carbon dioxide removal 120
5.3.2 Relationship between post-emissions warming and post-drawdown cooling 121
5.4 Physical controls on the path-dependency of temperature mitigation signals 123
5.4.1 Thermal and carbon controls on peak temperature mitigation 124
5.5 Integral constraints of carbon dioxide removal co-benefits 128
5.5.1 The role of overturning for thermal and carbon benefits 130
5.6 Discussion and Conclusions 133
5.6.1 Path dependency of temperature mitigation encourages early implementation 133
5.6.2 Temperature benefits and carbon co-benefits can have opposing path dependencies 135
5.6.3 Implications of results for carbon dioxide removal policy 136

6 Final Discussion and Concluding Remarks 138
6.1 A variety of methods to understand thermal and carbon processes 138
6.1.1 High-complexity CMIP6 models to extrapolate observational datasets 139
6.1.2 Intermediate-complexity models to understand regional thermal and carbon contributions to the Earth system response 140
6.1.3 Idealised box model ensembles to understand the controlling processes behind the timing of CDR benefits 141
6.2 Future pathways for this work 142
6.3 Broader implications of the thesis 143

A Construction of the Gnanadesikan ocean model 145
A.0.1 Mass transformation and overturning balances 147
A.0.2 Forcings of carbon and heat into the ocean-atmosphere system . . . 148

Bibliography 150

List of Figures

1.1 Relationship between cumulative carbon emissions (in PgC and $GtCO_2$) and surface average warming relative to 1850-1900 (in °C) for the historical period (grey line) and CMIP6 climate projections from 2015 on (colour lines). The spread in the idealised TCRE metric is included as grey shading for comparison. Figure taken from Canadell et al. (2021). 2

1.2 Schematic of processes involved in converting carbon emissions ΔI to surface warming ΔT. Arrows denote additions/removals of carbon and heat from the atmosphere, with blue arrows denoting those areas in which the ocean plays a substantial role. Red dashed arrows denote the TCRE, a commonly-used metric to describe the relationship between surface warming and carbon emissions that builds upon the Transient Climate Response (TCR). The effective TCRE includes radiative forcing caused by greenhouse gases other than CO_2. Figure adapted from Williams et al. (2016). 3

1.3 Selection of different nature-based carbon dioxide removal schemes that may impact the ocean's role in the thermal and carbon response to emissions. Figure adapted from Rita Erven (https://www.spp-climate-engineering.de). 7

1.4 Schematic of different climate model complexities in terms of detail/resolution (x-axis), number of processes (y-axis), and timescales that can be covered by integration. Ellipsoids illustrate the domains of conceptual models, Earth Models of Intermediate Complexity (EMICs), and high-complexity atmosphere-ocean global climate models (AOGCMs) and Earth System Models. Schematic reconstructed from Claussen et al. (2002). 10

2.1 Location of current ocean interior carbon sampling. (a) Spatial coverage of repeat transects used in the GLODAP project up to 2021, with recent cruises in dark red. Figure from www.glodap.info, Lauvset et al. (2016). (b) Location of carbon timeseries. BATS and HOT are the locations used for the study in this chapter. Figure courtesy of Bates et al. (2014). 20

2.2 (a-c) Comparison of BATS upper ocean (0-500m) carbon behaviour for a UKESM model realisation when removing various first guesses: (a) the climatological mean from 1955-2014, (b) the linear trend from 1955-2014, and (c) a regression against atmospheric CO_2 concentrations. (d-f) Pointwise correlations with carbon at BATS for residuals from (d) the global climatology, (e) global detrended carbon data, and (f) global carbon residuals calculated from a regression against atmospheric CO_2 concentrations. 29

2.3 (a) Model ensemble regression coefficients f_i for 0-500m integrated DIC against atmospheric pCO_2 following (2.9), in units mol m^{-2} (μmol CO_2)$^{-1}$ and (b) standard deviation of the residual DIC, in units mol m^{-2}. 30

2.4 Integrated upper 500m DIC anomalies relative to 1955, in units mol C m^{-2}, at (a) BATS and (b) HOT. Time series shown are representive runs for each model within the ensemble. Inset is a map showing the locations of BATS and HOT. 32

2.5 Ensemble correlations against integrated DIC at BATS and HOT for different depth horizons: 0-100m (a,b) and 0-500m (c,d). 33

2.6 (a) Standard deviation of the NorESM residual, the "model truth", (b) the root mean squared error (RMSE) for the analysis using NorESM 0-500m DIC data at BATS and HOT, and (c) the relative improvement of the analysis over the first-guess field, $(RMSE(T_i - B_i) - RMSE(A_i))/RMSE(T_i - B_i)$. 34

2.7 Ensemble correlations for 0-10m integrated DIC against integrated DIC at (a) BATS and (b) HOT. 36

2.8 Correlations against BATS 0-10m integrated DIC for each model within our model ensemble. 37

2.9 Correlations against HOT 0-10m integrated DIC for each model within our model ensemble. 38

2.10 Characteristic zonal and meridional length scales, b_x ad b_y respectively, in units km, for 0-100m integrated DIC residuals (a,b) and 0-500m integrated DIC residuals (c,d). 41

3.1 GLODAP gridded 1986-2013 climatology surface DIC, in units μmol kg^{-1}, against a) surface temperatures, in units °C, and b) surface salinity, in units psu. Data are plotted between the latitudes of 45° S and 45° N, with colour of markers indicating the latitude of the surface variables. Data from (Lauvset et al., 2016). 51

3.2 Binned Argo profiles for the years (left) 2002 and (right) 2015. Profiles were collected into $1° \times 1°$ bins for 01 January - 31 December of each year. . . . 53

3.3 Coverage of binning procedure for Argo reconstructions using observation locations at the year 2015 (dark blue), with overage of experiments with search radii of 1° (bright medium blue) or 5° (light blue). Binning procedure and bin selection conducted from an original count of $n = 334881$ profile locations. 56

3.4 Pointwise correlations for upper 100m residuals of a) DIC and temperature, b) DIC and salinity, and c) temperature and salinity taken from the CMIP6 ensemble. Ensemble consists of 6 models and contains 5 historical runs for each model. Contours occur at correlation increments of 0.2. 58

3.5 CMIP6 ensemble weights for reconstructing 0-100m integrated carbon variability: a) ensemble temperature weights, in units 10 mol m^{-2} (°C m)$^{-1}$ and b) the ensemble spread divided by the ensemble temperature weights; c) ensemble salinity weights, in units 10 mol m^{-2} (psu m)$^{-1}$, and d) the ensemble spread divided by the ensemble salinity weights. 60

3.6 Diagnostics for testing the carbon reconstruction using pointwise temperature and salinity with the NorESM model: (a) standard deviation of the NorESM top 100m integrated carbon residual, in units mol m^{-2}, (b) RMSE of the top 100m integrated carbon reconstruction, in units mol m^{-2}, and (c) relative improvement of the reconstruction to the prior guess $(RMSE_{prior} - RMSE_{reconstruct})/RMSE_{prior}$, in which blue areas indicate an improvement, where an improvement of 1 indicates a perfect reconstruction, and red areas indicate that errors are increased under the reconstruction method. 61

3.7 Sensitivity tests using the model ensemble reconstructions of 0-100m carbon content using pointwise temperature and salinity. RMSE improvement is relative to a first-guess of carbon content being proportional to atmospheric CO_2 concentrations. a) Pointwise minimum improvement seen in the sensitivity tests, b) Average RMSE improvement seen across the model ensemble, and c) Pointwise maximum RMSE improvement seen in the sensitivity tests. 62

3.8 Sensitivity tests using the model ensemble reconstructions of 0-500m carbon content using pointwise temperature and salinity. RMSE improvement is relative to a first-guess of carbon content being proportional to atmospheric CO_2 concentrations. a) Pointwise minimum improvement seen in the sensitivity tests, b) Average RMSE improvement seen across the model ensemble, and c) Pointwise maximum RMSE improvement seen in the sensitivity tests. 63

3.9 NorESM carbon reconstruction tests using synthetic Argo profiles of temperature and salinity. (a-d) Root mean squared errors for reconstructions a) assuming a linear response to atmospheric CO_2 (the first-guess field), and using Argo profile information within a radius of b)1°, c)2°, and d) 5°. (e-g) RMSE improvement relative to the first-guess RMSE for reconstructions using search radii of e)1°, f)2°, and g) 5°. 65

3.10 Sensitivity tests for reconstructing 0-100m carbon content variability using Argo profiles within a 5° search radius. Error improvement is relative to the RMSE when using the first-guess that carbon is proportional to atmospheric CO_2 concentrations. Blue areas indicate regions where the Argo-style reconstruction improves upon the first guess (up to 1 = 100% of variability is accounted for in the reconstruction), whereas red areas indicate regions where the Argo-style reconstruction is worse than the first guess. 67

3.11 Snapshots of integrated 0-100m residual carbon DIC' for one realisation of each model, taken at year 1950. Note that the patterns of DIC' should not line up as the atmospheric components of each model are not set to match up with reanalysis. 68

4.1 RCP 4.5 (blue lines) and RCP 8.5 (orange lines) forcing agents: (a) cumulative CO_2 emissions ΔI_{em}, in Pg C, (b) global average radiative forcing from non-CO_2 greenhouse gases, ΔR_{GHG}, in W m^{-2}, (c) hemispheric-averaged radiative forcings from sulphate aerosols, ΔR_{sulph}, in W m^{-2}, for the Northern Hemisphere (dashed lines) and the Southern Hemisphere (dot-dashed lines). 83

4.2 (a) Atmospheric CO_2 mixing ratios for RCP 4.5 (blue) and RCP 8.5 (orange). (b) Surface air temperature for RCP 4.5 (blue) and RCP 8.5 (orange), for the global average (solid lines), Northern Hemisphere averages (dashed lines), and Southern Hemisphere averages (dot-dashed lines). 84

4.3 21st Century global ocean response for RCP 4.5 (blue) and RCP 8.5 (orange): (a) Average ocean temperature, in °C, (b) Average sea surface temperature, in °C, and (c) Meridional overturning strength in Sv. 85

4.4 Globally-averaged radiative forcing ΔR, radiative response $\lambda\Delta T$, and planetary heat uptake $\mathcal{N}$, in units W m^{-2}, for (a) RCP 4.5 and (b) RCP 8.5. Diagnosed climate feedback parameter λ (solid lines), with 11-year average (dashed lines), for (c) RCP 4.5 and (d) RCP 8.5. 86

4.5 (a) Average global and hemispheric temperature changes versus cumulative CO_2 emissions for RCP 4.5 and RCP 8.5. (b) Global and hemispheric effective transient climate responses to emissions, with units °C (1000 Pg C)$^{-1}$, for RCP 4.5 and RCP 8.5. 87

4.6 Logarithmic changes in the terms setting the responses to cumulative carbon emissions for the global response under (a) RCP 4.5 and (b) RCP 8.5, the Northern Hemisphere response under (c) RCP 4.5 and (d) RCP 8.5, and the Southern Hemisphere response for (e) RCP 4.5 and (f) RCP 8.5. Note that the hemispheric response to emissions has two terms different from the global carbon response term $\Delta I_{atm}/\Delta I_{em}$; these are set as dashed lines. 89

4.7 Integrated thermal response terms - Radiative forcing ΔR, radiative response $\lambda\Delta T$, and heat content tendency dQ/dt, for (a) RCP 4.5 Northern Hemisphere, (b) RCP 8.5 Northern Hemisphere, (c) RCP 4.5 Southern Hemisphere, and (d) RCP 8.5 Southern Hemisphere. 92

4.8 (a) Diagnosed global and hemispheric climate feedback parameters for RCP 4.5 and 8.5, with full forcing sets including CO_2 emissions, GHG radiative forcing, and sulfate aerosol radiative forcing (b) As for (a), but using runs with RCP 4.5 CO_2 emissions and GHG radiative forcing only. 93

4.9 (a) Atmospheric cross-equatorial heat transport calculated by heat flux convergence for RCP 4.5 (blue) and RCP 8.5 (orange). Values derived from Northern Hemisphere output are solid lines, values derived from Southern Hemisphere output are dashed lines. (b) Difference in cross-equatorial heat transports between calculations using Northern Hemisphere and Southern Hemisphere values. (c) As in (a), but calculated for the ocean cross-equatorial heat transport. (d) As in (b), but for the ocean heat transport. 95

4.10 Integrated ocean heat content and top-of-the-atmosphere, air-sea, and cross-equatorial heat fluxes over the 21st Century for RCP 4.5 and 8.5 scenarios, with and without sulfate aerosol forcings (top and bottom rows, respectively). Arrows denote direction of flux (with positive values indicating downwards/northwards) and area of arrows scale with magnitude. Atmospheric heat content changes are not included as they are 1-2 orders of magnitude less than the energy transfers and ocean heat content changes. Year 2000 values are integrated for that year only, and the relative sizes of the grey arrows do not correspond with the sizes of the white arrows for the integrated RCP responses. 97

4.11 (Left hand) Total reservoir changes relative to 1860 for (top) ocean carbon, (middle) regenerated ocean carbon, and (bottom) terrestrial carbon. (Right hand) As for the left-hand column but reservoir changes as a percentage of cumulative emissions. 100

4.12 (a,b) Normalised depth-integrated DIC changes relative to the preindustrial for RCP 4.5 and 8.5. Values are normalised by the global mean and standard deviation of carbon uptake for each RCP scenario. (c,d) Zonal mean change in ocean DIC relative to 1860 for RCP 4.5 and 8.5 (e,f) Zonal mean change in regenerated carbon relative to 1860 for RCP 4.5 and 8.5. Note that the colourbars for (c,d) differ from those for (e,f). 103

5.1 Schematic of the Gnanadesikan model, as expanded upon by Katavouta et al. (2018). Boxes indicate the ocean boxes within the model, with black arrows indicating the water transformations solved for. Fluxes (arrows) and tracer concentrations (coloured text) are shown for temperature (red), carbon (yellow), and alkalinity (green). 115

5.2 (a) Control Historical (1860-2005) and RCP (2005-2100) CO_2 cumulative emissions (ΔI_{em}, in PgC) used to force the Gnanadesikan model. Forcing starts at 1860 after a spin-up to equilibrium conditions under an atmospheric CO_2 concentration of 280 ppm. (b) Cumulative carbon drawdown (ΔI_{CDR}, in PgC) for the CDR-FAST and CDR-SLOW timelines. 117

5.3 Control RCP surface temperature responses for all model setups (grey lines) and the ensemble median response (thick black lines). The number of runs conducted successfully under each forcing scenario is denoted by n. 119

5.4 Illustration of atmospheric CO_2 concentrations, surface temperature increases, and impacts of CDR for two model ensemble members forced with enhanced weathering along the CDR-FAST and CDR-SLOW timelines. 120

5.5 Scatterplots comparing the year of peak warming under the control runs with (left) the year of peak temperature mitigation from the EW-FAST scenario; (centre) the absolute temperature mitigation difference between EW-FAST and EW-SLOW at 2200 (5.3), in units °C; and (right) the relative temperature mitigation difference between EW-FAST and EW-SLOW at 2200 (5.4). Colours and markers show different emissions scenarios and climate feedback parameters, arrows indicate the direction that these relationships move in stronger overturning regimes. 122

5.6 Percentage of EW-FAST temperature mitigation achieved by EW-SLOW at year 2200 for the RCP scenarios (separate boxes), climate feedback parameter λ (y-axis), and preindustrial overturning strength (marker colours). 123

5.7 (a) Heat uptake post-emissions for a control run (solid line) and EW-FAST run (dotted line) for a reference setup of RCP 4.5, $\lambda = 1.8$ W m^{-2} K^{-1}, $\delta = 0.5$, $\tau = 0.08$ N m^{-2}.(b) Radiative forcing post-emissions for the reference control run (dashed line) and reference EW-FAST run (dashed-dotted line). (c-f) Thermal contributions (solid lines) and carbon contributions (dashed lines) towards the year of peak EW-FAST temperature impact for the reference weathering run (black lines) and runs with varying physical parameters (coloured lines). 125

5.8 Impacts of CDR on ocean carbonate chemistry using two representative Gnanadesikan model runs: (a) Changes in ocean carbon inventory from control RCP scenarios (black/grey lines) and RCP scenarios with CDR representations (coloured lines), (b) Ocean carbon inventory differences between runs with CDR and control runs, (c) Surface average pH for control RCP runs and runs with CDR, and (d) Surface average pH difference between runs with CDR and control runs. 129

5.9 Surface pH difference between EW-FAST and EW-SLOW at year 2200 for the RCP scenarios (separate boxes), climate feedback parameter λ (y-axis), and preindustrial overturning strength (marker colours). 130

5.10 Schematic of preindustrial DIC and T, and tracer differences between EW and control runs under a strong overturning. Note that in these plots we have not visualised the change in volumes of the deep ocean box and the thermocline box. 131
5.11 Schematic of preindustrial DIC and T, and tracer differences between EW and control runs under a weak overturning. Note that in these plots we have not visualised the change in volumes of the deep ocean box and the thermocline box. 132

List of Tables

2.1 CMIP6 models and realisations used to create the analysis, with citations for models' documentation. Below the doubled line is the realisation of the Norwegian model used for testing the ensemble covariances. 25

5.1 Perturbed variables and ranges used to create the Gnanadesikan ensemble. 117
5.2 Description of forcing scenarios used for the Gnanadesikan ensemble. 118

A.1 Dimensions within the Gnanadesikan atmosphere/ocean model. Note that the Southern Ocean zonal and meridional extent are used for calculating Southern Ocean intermediate water formation within the Antarctic Circumpolar Current and thus $A_{SH} > L_x \times L_y$. 146
A.2 Constants used in calculating volume, heat, and carbon transports within the Gnanadesikan model. 146

Glossary of Acronyms

Acronym	Meaning
AMO	Atlantic Multidecadal Oscillation
BATS	Bermuda Atlantic Time Series
BECCS	Bio-Energy Carbon Capture and Storage
CCS	Carbon Capture and Storage
CDR	Carbon Dioxide Removal
CMIP	Coupled Model Intercomparison Project
CO_2	Carbon Dioxide
DIC	Dissolved Inorganic Carbon
EMIC	Earth system Model of Intermediate Complexity
EnOI	Ensemble Optimal Interpolation
ENSO	El Niño-Southern Oscillation
ESM	Earth System Model
EW	Enhanced Weathering
GCM	General Circulation Model
GHG	Greenhouse Gases
HOT	Hawaii Ocean Time Series
IPCC	Intergovernmental Panel on Climate Change
NAO	North Atlantic Oscillation
PDO	Pacific Decadal Oscillation
RCP	Representative Concentration Pathway
RMSE	Root Mean Squared Error
TCR	Transient Climate Response
TCRE	Transient Climate Response to Emissions
UVic ESCM	University of Victoria Earth System Climate Model

Chapter 1

Introduction

1.1 Carbon emissions and warming in the Anthropocene

The industrialisation of the global economy over the past 160 years has led to a large-scale transformation of carbon stored in geological or terrestrial reservoirs into atmospheric carbon dioxide (CO_2) through fossil fuel burning and land use changes. Between 1876-2017, over 2000 Pg CO_2 ($\sim$ 600 PgC, 1 Pg = 10^{15} grams) has been emitted through human activities, mainly through the burning of fossil fuels (IPCC 1.5 Report Summary for Policymakers, Fig. 1.1). This release of CO_2 is associated with rising surface air temperatures, and the current cumulative emissions have likely raised surface temperatures 0.8-1.3 °C above preindustrial levels (IPCC, 2021a). This massive change in the climate system is unprecedented in the recent geological past associated with increases in extreme events that are detrimental to both natural habitats and human environs.

The climate system responds to an increase in carbon through chemical and thermal responses (Fig. 1.2). First, the addition of CO_2 into the atmosphere through anthropogenic emissions moves the climate system away from its chemical equilibrium. This disequilibrium induces an exchange of carbon between the atmosphere and the ocean and terrestrial reservoirs through physical, chemical, or biological processes. The uptake and exchange of carbon between reservoirs can be considered the Earth system's *carbon response*, which affects how emissions are translated into atmospheric radiative forcing.

The atmospheric fraction of CO_2 emissions, coupled with non-CO_2 greenhouse gases such as CH_4, NO_x, and sulfate aerosols, induce radiative forcing by altering Earth's ability to release outgoing longwave radiation into space. The radiative forcing causes more heat to be stored in the climate system, leading to a planetary heat uptake and a radiative response.

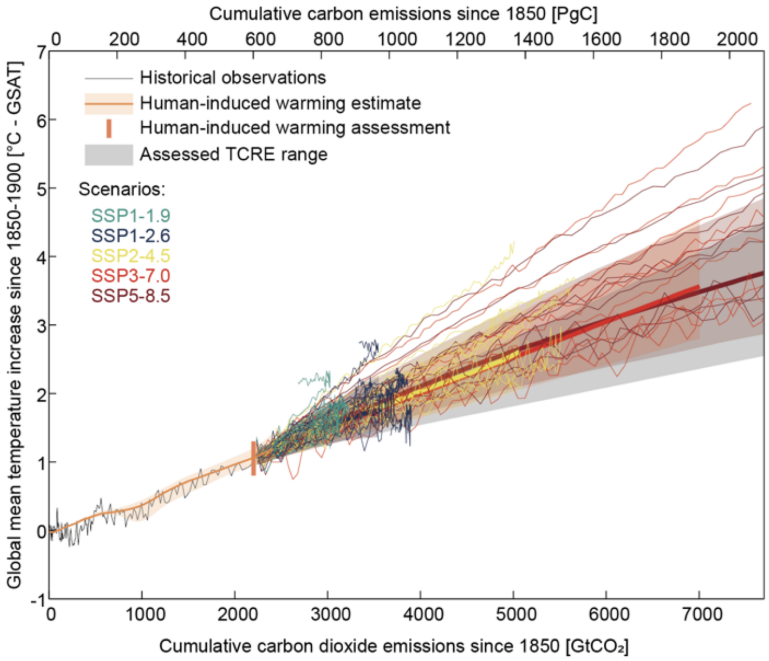

Figure 1.1: Relationship between cumulative carbon emissions (in PgC and $GtCO_2$) and surface average warming relative to 1850-1900 (in °C) for the historical period (grey line) and CMIP6 climate projections from 2015 on (colour lines). The spread in the idealised TCRE metric is included as grey shading for comparison. Figure taken from Canadell et al. (2021).

Excess planetary heat is, like carbon, partitioned among the different thermal reservoirs, and the remaining excess heat dictates the amount of global warming alongside climate feedbacks. This radiative response and uptake of heat can be considered the Earth system's *thermal response*, which affects how radiative forcing is translated to surface warming.

The combined carbon response and thermal response dictate how carbon emissions are translated to surface warming. While the carbon and thermal responses can appear linear and separate, they are linked through numerous climate feedbacks, in which temperature changes induce changes in the strength of the carbon sinks or induce changes in the Earth's energy balance and ability to take up extra heat.

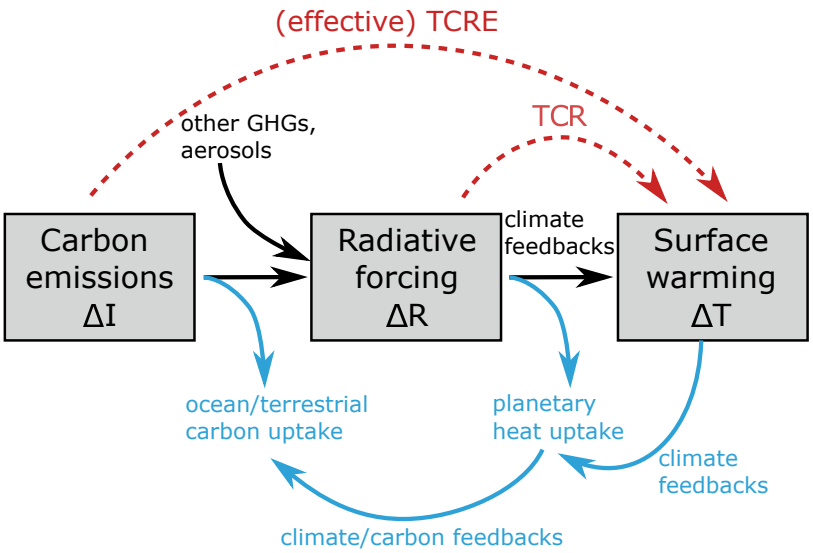

Figure 1.2: Schematic of processes involved in converting carbon emissions ΔI to surface warming ΔT. Arrows denote additions/removals of carbon and heat from the atmosphere, with blue arrows denoting those areas in which the ocean plays a substantial role. Red dashed arrows denote the TCRE, a commonly-used metric to describe the relationship between surface warming and carbon emissions that builds upon the Transient Climate Response (TCR). The effective TCRE includes radiative forcing caused by greenhouse gases other than CO_2. Figure adapted from Williams et al. (2016).

1.2 The ocean's role in understanding the Earth system response to emissions

The global ocean plays an important role in setting both the Earth system's carbon and thermal responses to emissions. The ocean takes up over 90% of anthropogenic heat (Rhein et al., 2013) and around 25% of added carbon from emissions and constitutes the only true net sink for historical CO_2 emissions (Sabine et al., 2004; Gruber et al., 2019a). This uptake of heat and carbon strongly modulates the amount of warming associated with emissions; if all emissions were to remain in the atmosphere and the only two responses to radiative forcing were rising atmospheric temperatures and increased outward longwave radiation, the amount of warming seen with current emissions would easily be orders of magnitude higher. The ocean is able to play such a large role in the uptake of heat and carbon because of the high heat capacity of water and the ocean buffer factor, where aqueous CO_2 reacts with water and transformed mainly into bicarbonate, which allows for further carbon uptake (Williams and Follows, 2011). Understanding the ocean's role in setting the thermal and

carbon responses is vital for understanding the overall connection between emissions and warming.

Additional heat and carbon from anthropogenic emissions mainly enter the ocean through atmosphere-ocean fluxes. Ocean ventilation transfers these properties in the ocean interior, where the added heat and carbon are sequestered from the atmosphere and stored on timescales ranging from decades to centuries/millennia.

1.2.1 Regional behaviour in heat and carbon uptake

The ocean displays substantial regional differences in how it takes up and stores heat and carbon. The North Atlantic and Southern Ocean play outsize roles in the uptake of heat and carbon respective to their size (Frölicher et al., 2015; Marshall et al., 2014b). The Southern Ocean accounts for over 40% of anthropogenic carbon and 75% of heat uptake while occupying 30% of the global ocean surface area. The North Atlantic, having a much smaller surface area (around 15% of the global ocean surface area), also takes up a disproportionate amount of anthropogenic carbon, resulting in the highest concentrations of anthropogenic carbon (Mikaloff Fletcher et al., 2006). The North Atlantic heat uptake is much smaller for the historical period because of the compensating effect of aerosols on the local radiative balance; under continued emissions, heat uptake in the North Atlantic may reach up to 25% of the total ocean uptake (Shi et al., 2018). These regions play important roles in the heat and carbon balances as they are regions of substantial intermediate and mode water formation (Sabine et al., 2004).

The role of ocean transport in redistributing heat and carbon becomes apparent when considering spatial differences between ocean heat and carbon uptake and storage (Frölicher et al., 2015). While the Southern Ocean remains a significant reservoir for both heat and carbon, vigorous northward transport results in the low-latitude Southern Hemisphere also playing a substantial role in heat and carbon storage. Ocean ventilation results in a vertical structure for added storage of heat and carbon that follows the thermocline (Zanna et al., 2019).

The behaviour of ocean heat and carbon uptake and storage are heavily dependent on circulation processes which may undergo changes from external forcing or natural variability. Climate modes of variability such as the North Atlantic Oscillation (NAO) or the El Niño-Southern Oscillation (ENSO) impact sea surface temperatures and air-sea heat and carbon fluxes over large regions (Jones et al., 2001; Keller et al., 2012b; Cheng et al., 2019). Under continued emissions, changes in ocean circulation have the potential to alter the pathways

for added ocean heat (Rugenstein et al., 2013; Winton et al., 2013). Changes to these regional behaviours can have substantial impacts on the behaviour of global warming; for instance, different patterns of sea surface temperature changes associated with changes in overturning can impact air-sea fluxes of heat and carbon and ultimately the ocean's efficacy in modulating the global response to emissions (Armour et al., 2013; Winton et al., 2010, 2013).

Changes in circulation impact not only the distribution of added ocean heat and carbon but also naturally-present ocean heat and carbon. Added heat and carbon have been assumed to act as passive tracers once entering the ocean (Zanna et al., 2019; Khatiwala et al., 2009). Passive anthropogenic heat and carbon are expected to have similar distributions in the ocean interior, as their pathways are set by the location at which they transfer from the atmosphere to the ocean. The redistribution of heat and carbon from changing circulation (from either climate variability or climate feedbacks) reduces the correspondence between ocean heat and carbon changes, as redistribution operates on opposite vertical gradients for natural ocean heat content and carbon content (Williams et al., 2021). Observed ocean heat and carbon content variability on multidecadal time scales are anticorrelated due to this redistribution (Thomas et al., 2018). The level to which the response reflects an added versus redistributed response is regionally dependent (Bronselaer and Zanna, 2020; Williams et al., 2021).

1.3 Utilising the response to emissions to calculate carbon budgets

On timescales of multiple decades to centuries, there is a near-constant proportionality between cumulative carbon emissions and average surface warming (Allen et al., 2009; Matthews et al., 2009; Zickfeld et al., 2009; MacDougall et al., 2017) (Fig. 1.1). The proportionality between cumulative emissions and average surface warming arises as the ocean's decreased uptake of additional heat and carbon over time increases the global thermal response and decreases the global carbon response, respectively (Goodwin et al., 2014; Williams et al., 2016; Ehlert et al., 2017; MacDougall, 2017). The resulting proportionality has led to a new climate metric, called the Transient Climate Response to Cumulative Emissions (TCRE). The TCRE intrinsically includes all the carbon and thermal processes and expands on the Transient Climate Response (TCR) to include the range of carbon cycle responses in climate models (Figure 1.2, red line). The TCRE is formally defined as

the proportion of surface warming to cumulative emissions under an idealised 140-year 1% atmospheric CO_2 increase. Current best estimates of the TCRE is between 1.0-2.3 °C (1000 PgC)$^{-1}$, with a best estimate of 1.65 °C (1000 PgC)$^{-1}$ (Canadell et al., 2021). While the TCRE provides a needed expansion of inter-model comparisons into the realm of carbon dynamics, its uncertainties are not necessarily reflective of uncertainties in its driving physical processes. Similar spreads in the TCRE can arise from various uncertainties within climate modelling, and similar TCRE behaviour may arise under different feedback behaviour or regional warming patterns (Liddicoat et al., 2016; Williams et al., 2020).

The definition of the TCRE using a 1% atmospheric CO_2 increase draws upon studies indicating a high level of path independency in the response to emissions (e.g., Herrington and Zickfeld (2014); Leduc et al. (2015)). However, these studies generally rely on intermediate-complexity model runs; using the GDFL model, Krasting et al. (2014) find that the proportion of warming to cumulative emissions is highest for high and low rates of emissions. The path dependency in the GFDL model is connected to how the thermal and carbon contributions evolve as the climate system adjusts towards its new equilibrium. It has been suggested that the path-independency found in more idealised models is related to their higher diffusive responses (MacDougall, 2017).

To first order, if a constant proportionality is assumed and non-CO_2 emissions are ignored or explicitly defined, surface warming thresholds can be translated to allowable carbon emissions, or "carbon budgets" (Matthews et al., 2017; Goodwin et al., 2018; Millar et al., 2017; Rogelj et al., 2019). At current rates of carbon emissions, the remaining budget for the Paris Agreement's 2.0 °C threshold will be surpassed within the next few decades (Goodwin et al., 2018; Tokarska et al., 2018). If international warming thresholds are to be maintained, climate engineering in the form of carbon dioxide removal is likely necessary. The most recent set of climate projections in which temperatures remain below or stabilize at 2 °C include some form of negative emissions (Rogelj et al., 2018; Gidden et al., 2019).

1.4 Carbon Dioxide Removal as a forcing on the Earth system

Carbon dioxide removal (CDR) involves creating new technologies or enhancing natural processes to reduce atmospheric CO_2 levels (Figure 1.3). Carbon may be captured from the flue gases of fossil fuel-powered plants, which prevents emitted carbon from perturbing Earth's carbon or radiative balances. Alternatively, direct air capture involves the creation

of materials to capture CO_2 from the atmosphere post-emissions. Creating these materials and extracting captured CO_2 for long-term storage can be highly energy- and resource-intensive.

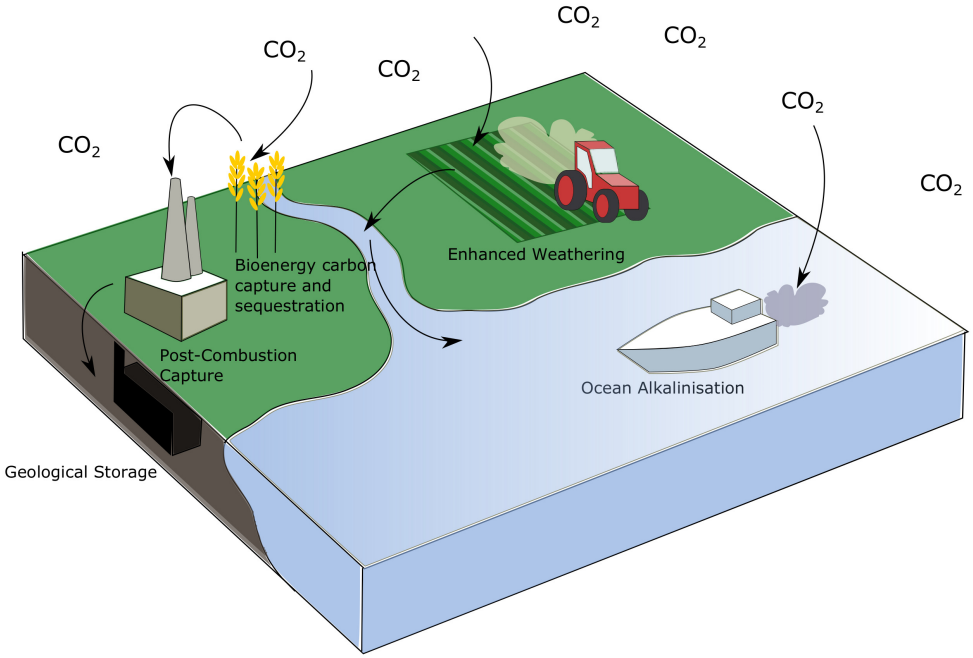

Figure 1.3: Selection of different nature-based carbon dioxide removal schemes that may impact the ocean's role in the thermal and carbon response to emissions. Figure adapted from Rita Erven (https://www.spp-climate-engineering.de).

CDR may also be formulated so that it enhances natural carbon sinks, either in magnitude or in the timescale at which they capture carbon. Carbon is sequestered temporarily away from the atmosphere, but it is not entirely removed from the Earth system. In some cases, nature-based carbon capture may be coupled with technology to produce energy and store carbon, as is the case with bioenergy and carbon capture and storage (BECCS) (Fuss et al., 2018; IPCC, 2018). CDR that relies solely on enhancing natural carbon cycle components may involve enhancing the biotic carbon cycle (e.g., through ocean fertilization or afforestation) or the chemical carbon cycle (e.g., through artificial ocean upwelling, ocean alkalinisation, or enhanced weathering) (Keller et al., 2018b). The time period in which these methods sequester carbon away from the atmosphere can vary from seasons/years (in the case of artificial ocean upwelling) to potentially several millennia (in the case of alkalinity-based methods that transfer carbon into ocean sediments).

1.4.1 Impacts of CDR on Earth's thermal and carbon responses

All CDR technologies impact the global heat and carbon balances by altering the amount of radiative forcing and partitioning of carbon, respectively. Just as the Earth system adjusts to perturbations caused by the emission of carbon dioxide, it will further adjust to any externally-forced removal. However, the response to emissions and negative emissions are not symmetrical. The artificial drawdown of carbon can weaken existing carbon sinks (Jones et al., 2016; Keller et al., 2018a; Zickfeld et al., 2021). Weakened and potentially outgassing carbon sinks under these scenarios increase the amount of drawdown necessary to reach specific atmospheric CO_2 levels. Hysteresis in the thermal response under negative emissions impacts the constant proportionality between carbon drawdown and surface temperature decreases that would exist under current TCRE theory (Tokarska and Zickfeld, 2015). This hysteresis is exacerbated in high-sensitivity climates (Jeltsch-Thömmes et al., 2020), indicating that a thorough understanding of the thermal processes that control the response to emissions is necessary not only for understanding current-day climate change but also potential climate change under drawdown scenarios. However, there can also be compensating changes in the thermal and carbon responses under moderate negative emissions scenarios, resulting in behaviour that is close to path-independent for low overshoot scenarios (with negative emissions at or below 300 PgC) (Tokarska et al., 2019). Thus, for heavily mitigated emissions scenarios that require relatively small amounts of carbon drawdown, estimates of the TCRE may still be relevant for negative emissions budgets, while for stronger overshoot scenarios the carbon budget estimates under positive emissions will likely be overestimates.

By altering natural carbon processes, carbon dioxide removal schemes can include impacts that are not directly related to the removal of atmospheric CO_2. Carbon dioxide removal involving the use of alkalinity, such as enhanced weathering or ocean alkalinisation, has the potential to mitigate some of the ocean acidification associated with climate change by transporting basic materials into the ocean (Keller et al., 2018a; Renforth and Henderson, 2017; Beerling et al., 2018) (Figure 1.3). The release of bases may additionally provide nutrients or toxins to aquatic ecosystems (Bach et al., 2019). In addition to the potential impacts on marine life, the addition of alkalinity will alter the ocean's role in the carbon response by changing the ocean carbon saturation state. Any implementation of carbon dioxide removal will need to fully consider both the direct and indirect impacts on carbon and thermal processes to avoid worsening some of the problems associated with carbon emissions and to maximise the effectiveness of carbon dioxide removal technologies

(Field and Mach, 2017).

1.5 The role of climate models in understanding the response to emissions

Observations of the climate system provide a complete view as to how the Earth responds to emissions. Observed responses are always the result of multiple competing processes, making process-based understanding of the system difficult at best. Additionally, gaining a globally-resolved view of the climate state through observations requires vast amounts of resources. Climate models are thus valuable tools in studying the Earth system response to forcing, as they include a finite number of known processes with explicit formulae and discretisations. Model experiments can be set up to explore the separate impacts of natural variability and different setups of external forcing. Climate models represent the complexities of the Earth system mathematically through differential equations and empirical formulas; when coupled together, these systems of equations can be solved for transient and equilibrium solutions under a variety of forcings.

The use of differential equations and numerical methods for geophysical prediction goes back to Richardson (1922), who suggested that human "calculators" could be arranged to create weather forecasts. Along with the development of the computer and numerical methods, climate models have evolved from 2-dimensional quasi-geostrophic models (Phillips, 1956) to coupled atmosphere-ocean global circulation models with realistic topography (Manabe et al., 1975) and, in turn, representations of non-CO_2 radiative forcing elements like sulfate aerosols (Mitchell et al., 1995) and carbon cycling (Cox et al., 2000). The complexity and resolution of today's highest complexity climate models requires the use of high-performance computing clusters and integrate parallelisation procedures.

Climate models have been instrumental in quantifying the potential climate response to carbon dioxide emissions since the early energy budget model of Arrhenius (1896). Using a 1-dimensional radiative convective model, Manabe and Wetherald (1967) estimated a warming of 2 °C if atmospheric CO_2 were doubled. Early models were instrumental to identify important aspects of the climate system such as the role of ice-albedo feedbacks in setting surface temperatures (Budyko, 1969) and amplified warming in the Northern Hemisphere (Stouffer et al., 1989). While the development of climate models has been necessary to better understand the intricacies behind the response to emissions, it should be noted that the bulk responses of simplified models of the mid-20th Century fit within

today's climate uncertainty; the equilibrium climate sensitivity estimated by Charney et al. (1979) of 2-4.5 °C for a doubling of CO_2 is similar to the equilibrium climate sensitivity found in the CMIP6 ensemble (IPCC, 2021a).

1.5.1 Model complexity and uses

Current computational power prevents the creation of a "one-size-fits-all" climate model which would be equally applicable to medium-range weather forecasting as to multi-millennial climate reconstructions. Instead, there is a wide range of climate model types available, with varying complexities; the complexities and running time of the models dictate the types of studies to which they are best suited.

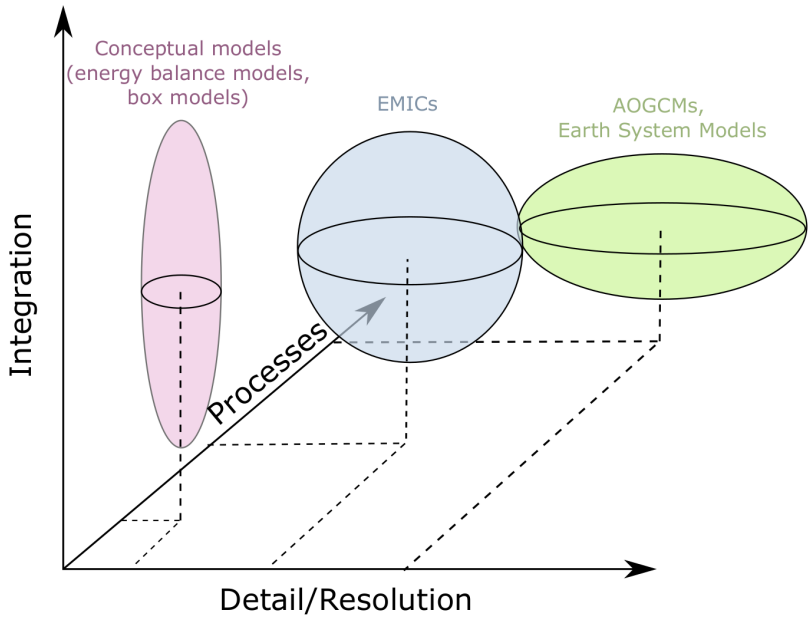

Figure 1.4: Schematic of different climate model complexities in terms of detail/resolution (x-axis), number of processes (y-axis), and timescales that can be covered by integration. Ellipsoids illustrate the domains of conceptual models, Earth Models of Intermediate Complexity (EMICs), and high-complexity atmosphere-ocean global climate models (AOGCMs) and Earth System Models. Schematic reconstructed from Claussen et al. (2002).

In the simplest extreme lie conceptual models, which include energy balance models and box models and consist of a few equations to represent energy or carbon transfers. While these models grossly oversimplify the climate system, they are flexible and computationally cheap to solve. With these aspects, conceptual models can be solved for processes at almost any timescale. Additionally, the flexibility and overt simplifications involved in setting

up these models mean that assumptions can be tested and altered to provide fundamental mechanistic insights. Even as more complex models have been developed, conceptual models are still used to justify hypotheses for more complex models (Katavouta et al., 2018), explore first principle mechanisms, and to create emulators that span large amounts of parameter space and resemble integral behaviours of more complex systems (Meinshausen et al., 2011b; Smith et al., 2018).

Earth system Models of Intermediate Complexity (EMICs) raise the complexity of at least one of these components, although for computational reasons some model components remain highly simplified. These models provide a combination of computational efficiency and realistic basic climate responses. For instance, the overturning reduction in the University of Victoria EMIC under atmospheric CO_2 increases is comparable to those for higher-complexity models, but its slab atmosphere does not allow for any changes to cloud convection, which can play an important role in setting the strength of regional and global climate feedbacks. The more efficient running requirements make the models easy to integrate forwards, and the models can be used for multi-millennial timescales or perturbed parameter ensembles.

The most complex climate models consist of Global Circulation Models (GCMs) and Earth System Models (ESMs), which have fully coupled general circulation models for the atmosphere and ocean and often include some atmospheric chemistry. ESMs additionally include terrestrial, sea ice, ocean biology, and sediment representations for carbon and nutrient cycling. These models are the current state-of-the-art models for representing how the Earth system changes under emissions, but because of their computational complexity their results generally involve one model setup rather than multiple setups that span uncertainty ranges. Experiments using these high-complexity models can still be run out for multiple centuries, and so they play a major role in providing projections for the IPCC.

While climate models have generally become more complex over the decades, more simplified models have not lost their merits. Complex models involve a large amount of tuning and complex parametrisations, making identification of controls more difficult. Simplified models provide valuable independent means for testing hypotheses, as well as for expanding the field of testable hypotheses. In this thesis we focus not on exploiting one model to its limits but rather attempt to span the available model space by choosing the models most appropriate for each chapter's driving questions.

1.5.2 Uncertainties and errors within climate models

Climate models provide incomplete and imperfect representations of the Earth climate system. As such, decisions and concessions made while constructing different climate models impose a level of uncertainty. Substantial areas of climate modelling that induce errors and uncertainty are:

- **Model resolution:** The horizontal and depth resolutions of models impact the lower limit of processes which may be resolved; for example, a 1° ocean model will require eddy parameterisations, while a 1/12° ocean model will resolve eddies for the low and mid-latitudes. An increase in resolution normally results in shorter integrations, as the time step for models must decrease as the resolution increases. Both the atmosphere and ocean have important mixing processes that act on molecular scales, so any model within current computational abilities will be unable to resolve all length scales. The resolution of a model is normally chosen to balance the computational requirements with the timescales of highest interest.

- **Model spinups and equilibration procedures:** Models are ideally started from equilibrium, so that the response to an external forcing can be properly attributed and is not a result of the model reaching its original thermodynamic and biogeochemical equilibria. For slower climate processes, such as sediment carbon uptake and deep ocean heat uptake, reaching equilibrium conditions may take thousands of model years. If a model is not able to reach its equilibrium before external forcing begins, it may experience a drift that is unrelated to either variability or external forcing.

- **Initial or boundary conditions:** The chaotic nature of atmospheric and ocean dynamics means that the state of the climate system is heavily dependent on accurate knowledge of initial conditions and boundary conditions. The resolution of climate models and inevitable rounding errors needed to calculate and store the forward integration creates imperfect boundary conditions, which are then propagated through the dynamics and can change large-scale climate patterns. As a result, coupled climate models are generally unable to represent the history of climate modes of variability, although they may contain the patterns associated with these climate modes of variability.

- **Parameterised chemical or biological processes:** In addition to small-scale physical processes, chemical and biological processes within the climate system also require

parameterisations. Part of this requirement comes from the scale of chemistry or biology, but the complexity and heterogeneity inherent in chemistry or biology also complicates how these processes are represented. Biology is often represented through a few classes of terrestrial plant life or phytoplankton types, which use averaged empirical relationship for their impacts on air and soil moisture, nutrient levels, and carbon fluxes.

Each choice made in the creation and setup of a climate model run has implications for how the model will respond to thermal and carbon forcing. Model resolution has implications for air-sea coupling (Hewitt et al., 2017), and ocean model diffusivity parametrisations play a role in setting the inter-model spread in ocean heat uptake (Huber and Zanna, 2017). Model drift can lead to difficulties in attributing trends for ocean biochemistry (Séférian et al., 2016) or steric sea level rise (Sen Gupta et al., 2013). These uncertainties ultimately combine to create the uncertainty envelopes seen for future climate projections (see Fig. 1.1, grey shading and SSP lines) and carbon budget metrics like the TCRE.

The uncertainty seen in climate model results is a perpetual issue. While the science behind climate modelling has continually improved, not all areas of uncertainty have become better constrained. Simple models can be tuned so that they fit integral responses consistent with observations, although they fail to contain the more complex processes that set those integral responses. As a result, as models become more complex and represent more climate processes, the new responses may not be as constrained as previous, tuned responses. One example of an increase in uncertainty with newer models rests in the climate feedback parameter, which is an empirical value that connects the amount of surface warming to the level of radiative forcing and global heat uptake. In the CMIP6 model ensemble there is a distinct widening of the range of the climate feedback parameter, and some models show a shift towards a smaller climate feedback parameter compared to the previous generation of coupled climate models (i.e., a climate state that is more sensitive to radiative forcing) (Zelinka et al., 2020). The decrease in the climate feedback parameter is related to better representation of mid-latitude and mixed-phase cloud feedbacks (Andrews et al., 2019; Zelinka et al., 2020). In a world where Newtonian mechanics is the basis for climate system modelling, the Navier-Stokes equations are unsolved, and irrational numbers exist, it is near-impossible to eliminate these uncertainties.

There are various methods to account for the uncertainty within climate model results. Ensemble means have been used extensively to extract consistent responses from ranges of model responses. Ensembles may consist of perturbed parameter ensembles, initial con-

dition ensembles, or ensembles of opportunity. Due to the complex interactions between climate components, the uncertainties seen within model results may obscure the role of individual uncertainties. While the CMIP5 and CMIP6 model ensembles show similar uncertainty ranges for the TCRE, the drivers behind the uncertainty ranges differ: in CMIP5, the terrestrial system is the largest source of uncertainty in the modelled TCRE, while in CMIP6, further constrained terrestrial responses from the inclusion of nitrogen limitation (Jones and Friedlingstein, 2020) coupled with more uncertain climate feedbacks (from the aforementioned developments in cloud representation) result in the climate feedback parameter providing the largest source of uncertainty in the ensemble TCRE (Williams et al., 2020). Thus, in addition to ensemble means and ranges, it is useful to examine controlling processes, so that the uncertainty space of emergent responses may be understood as a combination of uncertainties, and the mechanisms behind these responses may be better quantified.

1.6 Thesis Scope and Overview

The objective of this thesis is to explore how the Earth system carbon responses and thermal responses impact the response to anthropogenic forcing. This thesis considers the Earth system response to both carbon emissions and carbon dioxide removal. We focus on the response on timescales ranging from decades to centuries and emphasise the ocean's role in setting this response. Here, we shortly introduce the themes and major questions and briefly outline the methods used in each chapter.

Chapter 2 - Representativeness of ocean carbon timeseries

Understanding the ocean's role in setting the historical response to emissions requires characterising how ocean heat and carbon have evolved. While the historical behaviour of upper ocean heat content is well-known, there is substantial uncertainty in historical upper ocean carbon behaviour. The difficulties in obtaining ocean interior measurements mean that time series of ocean carbon that cover decadal or longer timescales are spatially sparse. Time series sites have been an invaluable source of ocean biogeochemical observations, as water column properties are consistently sampled throughout the year. While these time series have been used to describe ocean carbon behaviour and regional variability, their representativeness of global and regional behaviours is poorly understood.

In this chapter, we take two long-running time series sites and explore how their observations may be extrapolated to account for regional or global ocean carbon variability. We

employ a mapping method that spreads information from data-rich to data-sparse regions through the use of covariance fields. We create these covariance fields using an ensemble of CMIP6 ESM ocean carbon outputs and diagnose the upper-ocean carbon inventory covariance fields for each time series site. We then explore the skill of our mapping method with synthetic model tests and examine the extent to which observations from time series sites might be expanded.

Chapter 3 - Reconstructing ocean carbon variability using ocean temperature and salinity

In this chapter we continue with the goal of reconstructing ocean carbon variability, but focus on the information that spatially-dense, non-carbon observations provide. Temperature and salinity changes have first-principle impacts on the ability of seawater to take up CO_2 and are drivers of the circulation processes that transport carbon into the ocean interior. We explore the potential in using temperature and salinity observations to reconstruct historical carbon changes using the same mapping method from Chapter 2. In this way, we attempt to add skill to our reconstruction method by drawing upon the relationships between ocean thermodynamics and ocean carbon and utilising these spatially-dense observations.

Using the same CMIP6 ensemble from Chapter 2, we create synthetic model experiments in which we use the ensemble pointwise covariances between temperature, salinity, and carbon to reconstruct the modelled historical carbon variability. We also test the impact of the irregular sampling from Argo floats to see if nearby profiles and the spatial correlation fields provide a more complete global reconstruction of carbon variability.

Chapter 4 - Hemispheric contributions towards the surface warming response to emissions

The approximate proportionality between surface warming and cumulative carbon emissions relies on the compensation between the global carbon and thermal responses. This compensation may not necessarily hold on regional scales, as both carbon uptake and heat uptake have strong regional variability. In this chapter, we explore how the compensation between thermal and carbon responses holds on a hemispheric scale, and how the hemispheric responses combine to create the global TCRE. To explore these questions we use an intermediate-complexity Earth System Model, the University of Victoria Earth System Climate Model (UVic ESCM), run under RCP 4.5 and 8.5 emissions scenarios.

We separate the proportion of warming to cumulative emissions into thermal, radiative forcing, and carbon terms that can be applied to both regional and global responses. With this mathematical separation we compare how the University of Victoria model global response compares to its hemispheric responses. We then diagnose the various aspects of the hemispheric thermal and carbon responses to see how differences in heat uptake, radiative forcing, climate feedbacks, and carbon uptake combine to set the hemispheres' responses to emissions and the level of symmetry.

Chapter 5 - Physical controls on the impacts of carbon dioxide removal

Carbon dioxide removal is likely necessary to offset continued carbon dioxide emissions and prevent further surface temperature increases. The amount of carbon dioxide removal necessary to achieve certain warming limits is a function of the Earth system's thermal and carbon responses to negative emissions, which can be different to the responses to positive emissions due to nonlinearities in the system. Uncertainties in the climate system may further enhance the asymmetry between positive and negative emissions by altering the strength of these nonlinearities. In Chapter 5 we explore how uncertainties in the physical climate system control the integrated responses associated with different carbon dioxide removal methods. Specifically, we explore how well-constrained the climate responses to carbon dioxide removal are to the integrated amount of carbon drawdown, in a direct comparison to the integrated response that forms the basis for the TCRE.

In this chapter, we use a large perturbed parameter ensemble of an idealised ocean-atmosphere model, which contains variations in ocean overturning strength, climate feedback parameter, carbon emissions, and carbon drawdown scenarios. We consider the impacts of enhanced weathering in an idealised way through surface ocean fluxes of carbon and alkalinity. We compare atmospheric CO_2 levels, surface warming, and surface ocean pH for converging CDR timelines at the year of equivalent carbon drawdown, and investigate how these responses can be understood through the global thermal and ocean carbon responses to atmospheric carbon dioxide removal.

Chapter 2

Representativeness of ocean carbon time series for regional variability

We begin by exploring the use of climate models to reconstruct historical ocean carbon variability. Understanding ocean carbon variability is necessary to connect the historical response to emissions to controlling ocean processes. However, ocean interior carbon measurements are sparse in space and time, which poses difficulties in determining the ocean's historical carbon inventory changes.

In this chapter we construct a method that reconstructs ocean carbon behaviour using CMIP6 model output to extrapolate sparse measurements into data-poor regions. The covariance fields derived from high-complexity models provide insights as to how observations may be extrapolated in space. With this method we explore how spatially-sparse time series sites may be used to represent larger-scale ocean carbon variability.

We begin in Section 2.1 with an overview of the current state of ocean interior carbon observations and how sparse observations hinder the current understanding of ocean carbon variability. We present our method to extend carbon interior observations using an Ensemble Optimal Interpolation scheme with CMIP6 historical output, with discussion on the assumptions and construction of the method, in Section 2.2. We then use pseudo-observations and synthetic model tests to explore the scales on which time series carbon data can be extrapolated from two test site locations in the subtropical North Pacific and North Atlantic gyres (Section 2.3). We explore the sensitivity of the method to depth horizons and time series locations (Section 2.4). Finally, we reflect upon the results of our extrapolation, discuss its potential and limitations, and propose steps for expanding our method to include real-world data (Section 2.5).

2.1 Recent trends and variability in ocean carbon

The ocean constitutes a major sink of anthropogenic carbon, currently taking up a quarter of emissions (Friedlingstein et al., 2020). Variability in the ocean carbon sink impacts the atmospheric fraction of carbon emissions, which is a major component in determining the radiative forcing and the resulting amount of warming associated with anthropogenic emissions. Understanding the behaviour of ocean carbon variability is therefore vital to understanding how the relationship between carbon emissions and surface warming has evolved in the historical period.

Trends and variability in the ocean carbon sink are generally measured through surface ocean partial pressure of CO_2 ($pCO_{2,oce}$). The value of $pCO_{2,oce}$ and the atmospheric CO_2 partial pressure pCO_2 determine the amount of carbon an ocean parcel can take up at the surface. Even under equilibrium climate conditions, surface heat and freshwater exchange, ocean circulation, and biology cause some ocean regions to act as carbon sinks or sources (Gruber et al., 2009). On top of this natural regionality, there exists substantial regionality in the ocean sink of anthropogenic carbon, with model studies revealing important contributions from the Southern Ocean and North Atlantic (Frölicher et al., 2015).

The historical ocean carbon sink response includes substantial amounts of temporal variability in addition to the long-term forced uptake of anthropogenic emissions. Interannual and decadal variability can be driven by climate modes of variability and changes in circulation (Bates, 2001; Lovenduski et al., 2008; Gruber et al., 2019a) as well as by external forcing (McKinley et al., 2020). The ocean carbon sink as determined by $pCO_{2,oce}$ shows a substantial amount of temporal variability that can obscure long-term trends (Le Quéré et al., 2007; McKinley et al., 2016; Li and Ilyina, 2018; Gruber et al., 2019a), and model studies show that regions of high carbon uptake can have heightened decadal variability (Li and Ilyina, 2018). The strength of decadal variability in the regional behaviour of $pCO_{2,oce}$ creates difficulties in quantifying the state and evolution of the ocean carbon sink. Detecting anthropogenically-induced trends in the ocean carbon sink may require decades of $pCO_{2,oce}$ measurements (McKinley et al., 2016); in the meantime, the uncertainty associated with the magnitude of ocean variability reduces the predictability of the global carbon system (Gruber et al., 2019a).

Ocean interior measurements can be used to complement $pCO_{2,oce}$ measurements in order to better understand and constrain historical ocean carbon uptake. By providing an integrated view of carbon fluxes with time, ocean interior carbon content measurements provide a useful constraint on cumulative carbon sink behaviour (McKinley et al., 2017).

Ocean interior carbon changes reflect the long-term uptake of anthropogenic carbon as well as the impacts of climate variability on anthropogenic carbon and natural carbon pools. Ocean interior carbon measurements exhibit regional and temporal variability related to a combination of changes in atmosphere-ocean CO_2 fluxes, biological activity, and pathways of subducted water parcels (DeVries, 2014). Attributing variability in ocean carbon inventories to natural or anthropogenic carbon changes is regionally specific: some regions, such as the Southern Ocean, have natural and anthropogenic carbon changes that to reinforce each other, whereas others have natural and anthropogenic carbon changes that counteract (Bronselaer and Zanna, 2020).

2.1.1 Availability of ocean interior carbon measurements

Being able to decompose changes in ocean carbon into anthropogenic carbon uptake and changes to natural carbon pools requires complete temporal sampling. The ocean interior is only sparsely sampled relative to the surface ocean, particularly when considering ocean carbon measurements. Additionally, unlike ocean temperature and salinity measurements, which are taken autonomously with sensors and have good spatial coverage in the 21st Century, ocean inorganic carbon measurements require laboratory tests on bottle samples, making measurements of ocean interior carbon significantly time- and resource-intensive.

The logistical requirements for these measurements result in relatively sparse coverage for ocean carbon. To gain a global view of ocean interior carbon variability, both repeat transects and time series site sampling methods have been used. Repeat transects provide good horizontal resolution for ocean carbon (Fig. 2.1a) but can be separated by a decade or more. This infrequent sampling prevents differences in sampled ocean carbon to be attributed to forcing or natural variability, and so these repeat transects are best suited for creating global climatologies over a decadal composite (e.g., GLODAP2, Lauvset et al. (2016)) or describing long-term, globally-coherent trends (Sabine et al., 2004).

Ocean time series, which are repeat visits to a specific location multiple times per year, are able to resolve seasonal and interannual variability in ocean carbon and other biogeochemical parameters. These time series are relatively new, with most having existed for less than 20 years, but have the ability to provide a comprehensive, multidecadal view of ocean biogeochemistry; for instance, the Bermuda Atlantic Time Series (BATS) has had monthly data collected since late 1988. However, only a few time series are currently in existence (Fig. 2.1b), and these time series are generally near the ocean boundaries to make repeat data collection more feasible. Using time series to determine trends and variability

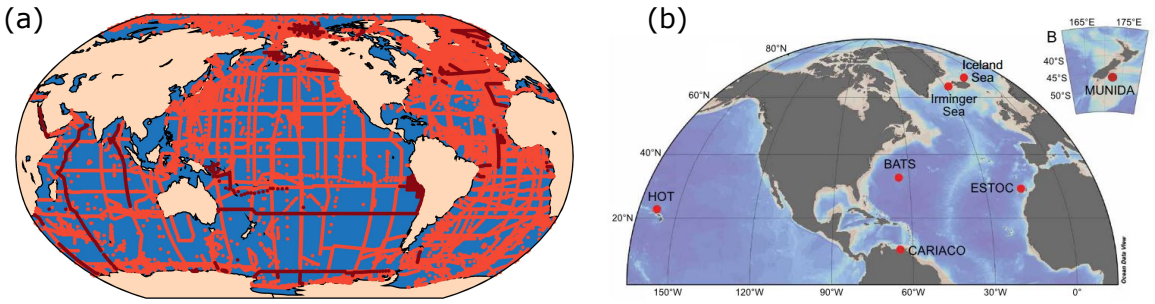

Figure 2.1: Location of current ocean interior carbon sampling. (a) Spatial coverage of repeat transects used in the GLODAP project up to 2021, with recent cruises in dark red. Figure from www.glodap.info, Lauvset et al. (2016). (b) Location of carbon timeseries. BATS and HOT are the locations used for the study in this chapter. Figure courtesy of Bates et al. (2014).

can be resource-intensive as the sampling of these sites should not be skewed towards specific times of year (Bates et al., 2014). In order to provide a diverse view of ocean carbon behaviour, these time series are generally separated by hundreds or thousands of kilometres, and thus little is known about how specific the results are to their location. Interannual variability signals are apparent at time series sites with varying levels of connection to modes of climate variability. The variability found at time series sites may be limited in how they represent regional or global carbon variability, although they have been used to describe variability in carbon changes in mode waters (Brix et al., 2004; Bates, 2012).

2.1.2 Expanding sparse observations using optimal interpolation

Resolving temporal and spatial variability typically requires high-resolution observations. As ocean carbon is highly variable, the sparse observations require some sort of extrapolation method. Basic linear interpolation of observations is likely to smooth out any variability, and so the use of these elementary methods to reconstruct global carbon is restricted to long-term trends or climatologies. Optimal interpolation methods aim to solve the problem of oversmoothing by using covariance fields to propagate information from data-rich to data-sparse regions. The use of optimal interpolation to extrapolate information requires accurate covariance fields that reflect the physical underpinnings between regions of observations and regions without (Daley, 1991; Oke et al., 2002; Smith and Murphy, 2007). These covariance fields are used to calculate the weights with which observations alter a background field,

which is generally taken to be a climatology.

In order to use such interpolation techniques, the covariance fields should be an accurate reflection of reality. Ideally, the covariance fields are calculated from observations; however, with sub-optimal observation coverage, observational covariances use averaged data, which decreases their accuracy by smoothing out small scale features as mentioned previously. Other methods for creating covariance fields use parametrizations or models. Atmospheric reconstructions and some oceanographic climatologies such as GLODAP use parameterized covariances based on characteristic correlation length scales (Jones et al., 2015; Lauvset et al., 2016). When considering variability in the ocean, however, autocorrelation length scales are likely to be anisotropic and vary by region, meaning that the typical constant and symmetric parametrizations are likely to induce significant errors.

Climate models have also been used to construct covariance fields for ocean data analyses. Realistic climate models have the advantage of providing high-resolution output across all ocean regions, so that covariance fields can be derived without any parametric assumptions and limited data interpolation. Model covariances have been able to replicate historical ocean heat content and salinity using ensembles of a single model (Oke et al., 2002; Smith and Murphy, 2007) or an ensemble of CMIP5 models (Cheng and Zhu, 2016). As models are an imperfect representation of the climate system, covariance fields derived from their output will include errors. Analyses can attempt to reduce these errors through localisation, iterative methods, or by using a sufficiently large ensemble so that uncertainties in farther-field covariances result in a dampening of information propagation.

Other studies focused on reconstructing ocean interior carbon use more indirect methods to avoid the issues associated with expanding sparse datasets. Inversion techniques use model ensembles (Gruber et al., 2009) or measurements of temperature, salinity, and other tracers (Khatiwala et al., 2009; DeVries, 2014) to create Green's functions for ocean ventilation rates and pathways. Assuming that anthropogenic carbon acts as a passive tracer and ventilation pathways are time-invariant, the ventilation function can then be used to solve for ocean carbon storage. One major caveat of this method is the assumption of constant circulation. This restriction can be circumvented by separate runs for different periods of time (DeVries et al., 2017), although this separation inherently removes any potential for different periods of variability to interact.

2.1.3 Driving questions and objectives

Historical reconstructions of global ocean carbon face the logistical difficulty of having very few observations with which the reconstructions can be validated. Ocean carbon time series sites attempt to solve one aspect of this scarcity by providing regularly-sampled water column measurements of ocean physical and biogeochemical variables. Studies using time series sites to characterise interannual and multidecadal carbon variability have attempted to expand their results to regional carbon variability through mode water arguments (Bates, 2012) or through similarities in surface ocean properties to temperature (Brix et al., 2004). However, these expansions may overshadow physical constraints on the extrapolation such as nonlinear ocean physics or uncertainty in small-scale processes.

In this chapter, we explore the limits to which time series for ocean carbon variability may be extrapolated. We construct covariance fields from CMIP6 model ensembles and use these fields to assess the limitations of using sparse observations to create larger-scale reconstructions. Within this chapter we explore the following questions:

- Do the CMIP6 Earth System Models provide useful covariance fields for extending ocean carbon measurements?
- What limits are there in extending carbon time series to regional or global scales?
- Do time series sites reflect any climate variability-based changes in the ocean carbon system?
- How does the location of carbon timeseries sites impact their potential for extrapolation?

We hypothesise that the CMIP6 models should be useful for extending ocean time series data on annual and longer time scales in regions where the physical drivers are well known and where biology plays less of a role in setting the ocean carbon content. We expect that the time series at BATS and HOT, being in the oligotrophic subtropical gyres, will have some skill in representing the gyre carbon variability. Additionally, we hypothesise that the covariances and correlations derived from the CMIP6 models will show a dampened version of large modes of climate variability like the El Niño-Southern Oscillation (ENSO) or the North Atlantic Oscillation (NAO), and that the locations within the centres of these modes of variability will have the furthest-reaching correlations due to the amplitude associated with those modes of variability.

In this study we explore these questions by creating an ensemble optimal interpolation method using CMIP6 models for integrated upper ocean carbon content (Section 2.2).

We conduct idealized model experiments in which we attempt to reconstruct the models' carbon variability with the data output at two time series locations in the North Atlantic and North Pacific subtropical gyres. With the CMIP6 model ensemble covariance fields, we characterise how the data from these time series locations may be extended and how this extension is sensitive to the depth integration and CMIP6 ensemble makeup (Section 2.3). We then explore the shortcomings of using super-sparse time series sites through the covariance fields' representation of climate variability patterns and a global correlation length scale analysis (Section 2.4), and discuss the merits of the optimal interpolation method in general and its application to carbon time series sites (Section 2.5).

2.2 The optimal interpolation method

The optimal interpolation method used in this chapter is similar to the methods used in Smith and Murphy (2007) and Cheng and Zhu (2016) to construct ocean subsurface temperature and salinity analyses. For this chapter we present an idealised form of optimal interpolation as described in Daley (1991).

Optimal interpolation propagates information from observations O_k at locations k to analysis locations i. The analysis values A_i are equal to a "first guess" or background value B_i and a weighted sum of increments to the background values $(O_k - B_k)$:

$$A_i = B_i + \sum_{k=1}^{K} w_k (O_k - B_k). \tag{2.1}$$

The weights are calculated to minimise the expected analysis error variance. Together, the matrix G of the weights w_k to translate observations to analysis space is known as the *gain matrix*. For schemes in which real-world observations are used, there are errors in both the observation and background values that must be accounted for within the algorithm. In this chapter, we conduct synthetic model experiments to test the general method, so there is an absence of additional errors from real-world observations. Thus this method is equivalent to solving an ordinary least squares regression for the weights w_k. Upon obtaining the weights w_k, observations may then be extrapolated to modify model output for a forward experiment (typical data assimilation schemes) or extrapolated onto a map that relaxes towards the first guess field as the impact of the observations decreases (what we use here as OI).

Ensemble Optimal Interpolation (EnOI) expands upon the OI method by calculating covariances and correlations from a set of concatenated model runs (Oke et al., 2002; Evensen,

2003; Smith and Murphy, 2007; Oke et al., 2010; Cheng and Zhu, 2016). The model ensemble can be created from multiple realisations of a single model or through multiple models that have been interpolated to a common grid. The use of models to calculate the covariance fields eliminates the need to make specific assumptions about the covariance fields that may be inconsistent with ocean dynamics. Early atmospheric data assimilation schemes often assumed that correlations are homogeneous and isotropic, which may not hold in ocean environments due to topography and current systems.

One advantage of using EnOI rather than other data assimilation techniques is that it is possible to use existing model output to create the covariance fields and solutions for the weighting coefficients w_k. The use of existing model output reduces the computational costs of this method significantly compared to methods which require updating model fields before integrating further in time and then recalculating the model covariances. As we cannot update the CMIP6 model output with observations and iterate, the method assumes that the model processes associated with ocean carbon variability are stationary in time. For time periods of up to a few decades, this stationarity should hold reasonably well; under longer timescales and large changes in CO_2, changes in the ocean buffer factor will impact the equilibrium storage capacity of the ocean.

2.2.1 The CMIP6 model ensemble

This work uses the output of the latest Coupled Model Intercomparison Project (CMIP6). CMIP6 provides standardised forcing protocols to explore the climate response to different aspects of forced and unforced variability, and to provide a platform with which model projections may be compared (IPCC, 2021b). We use CMIP6 output for the historical runs, in which the models are forced with prescribed atmospheric CO_2 concentrations and other greenhouse gases associated with the years 1850-2014 (Meinshausen et al., 2017). For this analysis, we extract the last 60 years of each realisation's output, corresponding with the years 1955-2014.

The ensemble used in this chapter is made up of ocean dissolved inorganic carbon (DIC, labeled `dissic`) output from 6 Earth system models (Table 2.1). The models have been chosen because they provide full-depth ocean carbon output for multiple realisations. Overviews of the models' components and level of complexity can be found in the latest IPCC report (2021b). Most models have a nominal horizontal resolution of 1° for their ocean component, with the exception of the MPI model which has a nominal horizontal resolution of 0.4°. To construct the ensemble, 5 realisations are taken from each model,

Model Name	Institute	Outputs used	Citation
ACCESS-ESM1.5	CSIRO, AUS	r(1,2,4,5,6)i1p1f1	Ziehn et al. (2020)
CanESM5	CCCma, CAN	r(10-14)i1p1f1	Swart et al. (2019)
CESM2	NCAR, USA	r(1-5)i1p1f1	Danabasoglu et al. (2020)
IPSL-CM6A-LR	IPSL, FRA	r(1-4,32)i1p1f1	Boucher et al. (2020)
MPI-ESM1-2-HR	MPI, GER	r(1-5)i1p1f1	Müller et al. (2018)
UKESM1.0-LL	MOHC, UK	r(1-4,8)i1p1f2	Sellar et al. (2019)
NorESM2-LM	NCC, NOR	r1i1p1f1	Seland et al. (2020)

Table 2.1: CMIP6 models and realisations used to create the analysis, with citations for models' documentation. Below the doubled line is the realisation of the Norwegian model used for testing the ensemble covariances.

resulting in a full ensemble size of 30 realisations and 1800 years. Realisations were chosen based off each model's available outputs.

Most CMIP6 model output is released as monthly averages. To focus on interannual and decadal variability and avoid potential complications of the seasonal cycle and mesoscale eddies, monthly average model output was reformatted into annual averages using `zsh` and `nco`. The output was then regridded to a 1° x 1° horizontal grid using a bilinear interpolation programme included in the python package `xesmf`. Ocean carbon, temperature, and salinity are then integrated from the surface to various depth levels. The use of integrated carbon measurements reduces the complexity of the covariance fields from 3-dimensional to 2-dimensional. The use of integrated carbon rather than surface fluxes has been used by Bronselaer and Zanna (2020) and Mazloff et al. (2018).

Testing the optimal interpolation

Once a model is included in the ensemble, the ensemble covariances should fit the model relatively well; therefore, it can be misleading to use models existing in the ensemble to test the covariance fields for errors. For this reason we employ two different methods for testing the covariance fields in this chapter and the following chapter:

- Sensitivity tests in which an entire model type is excluded from the covariance calculations - these tests can provide some idea of how each model within the ensemble is contributing towards the covariance fields.
- Tests in which we attempt to reconstruct DIC behaviour in the NorESM2-LM model

using the covariances calculated from the full ensemble. NorESM2 includes an isopycnic ocean model rather than a z-level ocean model, so this test can provide some insight into the general limitations of this method.

For these idealised tests we create pseudo-observations at time series sites using either the excluded model or the NorESM model. As we have the model truth, T_i, it is possible to quantify the improvement incurred by the interpolation method by calculating the root mean squared error (RMSE) at each location i over the historical period:

$$\mathrm{RMSE}_i = \sqrt{\frac{1}{60}\sum_{t=1955}^{2014}(T_i(t) - A_i(t))^2}. \tag{2.2}$$

The RMSE can then be compared to the variance of the true model minus the first guess to see where errors are reduced and how much of the existing variability is reconstructed through the mapping.

2.2.2 Choosing a background field

The background errors B_i should ideally be normally distributed, as non-normal background errors can lead to errors in the covariances (Daley, 1991). For temperature and salinity analyses by Smith and Murphy (2007) and Cheng and Zhu (2016), the first guess field is taken to be a climatology.

We use a realisation of the UKESM model to show the impact of different background field choices on the associated anomaly time series at BATS (33°N 64°W), along with correlations of the residual integrated ocean carbon against that at BATS (Fig. 2.2). We compare the correlation fields using first guesses of climatology, linear trend, and a trend proportional to atmospheric CO_2 concentrations:

$$B_i = f_i \times p_{CO_2}(t). \tag{2.3}$$

We provide this background field as an option as, on annual time scales, atmospheric CO_2 is globally well-mixed, and continuous measurements of atmospheric CO_2 from the mid-1950s means that such information can be included without further restricting suitable observations.

Is there a physical basis for removing a pCO_2 regression?

Ocean DIC may be decomposed as

$$DIC = C_{sat} + C_{bio} + C_{dis}, \tag{2.4}$$

where C_{sat} is the *saturated carbon pool*, which describes how much CO_2 a water parcel would have if brought to chemical equilibrium with the current-day atmosphere, C_{bio} is the *biological carbon pool*, which accounts for the regeneration of DIC from organic carbon through remineralisation, and C_{dis} is the *disequilibrium carbon pool*, which describes the amount of carbon a parcel would need to take up in order to reach chemical equilibrium (Williams and Follows, 2011; Katavouta and Williams, 2021).

The saturated carbon pool can be calculated as a function of atmospheric CO_2 partial pressure pCO_2, CO_2 gas exchange and chemical equilibrium coefficients, and ocean hydrogen ion concentrations (Williams and Follows, 2011):

$$C_{sat} = pCO_2\, K_0 \Big(1 + \frac{K_1}{[H^+]} + \frac{K_1 K_2}{[H^+]^2}\Big). \tag{2.5}$$

To first order, C_{sat} is determined by atmospheric pCO_2; changes in ocean temperature, salinity, and alkalinity induce secondary changes in C_{sat}, so we break down C_{sat} into pCO_2 and non-pCO_2 terms:

$$C_{sat} = f_{sat}(x, y) pCO_2(t) + \varepsilon_{sat}. \tag{2.6}$$

As an ocean parcel is subducted and isolated from the atmosphere, carbon is added by the regeneration of inorganic carbon from the breakdown of organic carbon, either through heterotrophic respiration or by the dissolution of calcium carbonate. Over longer periods of carbon emissions, the global change in regenerated carbon is expected to be a lower-order response in comparison to the global change of the saturated and disequilibrium carbon pools (Katavouta and Williams, 2021).

The disequilibrium carbon content of an interior ocean parcel is set by the instantaneous saturated carbon content of that parcel, the carbon content of the parcel when it was last in contact with the atmosphere, and the amount of carbon regenerated through biological processes. Under assumptions of constant ocean circulation and biology, the disequilibrium carbon content can be simplified to be proportional to the difference in C_{sat} calculated at the instantaneous time t from that calculated at the

time of subduction $t_{sub} = t - a$, where a is the age of the water parcel:

$$C_{dis} \approx C_{sat}(t_{sub}) - C_{sat}(t). \tag{2.7}$$

For simple emissions scenarios (linear and/or exponential behaviour), $C_{sat}(t_{sub})$ is proportional to $C_{sat}(t)$. Thus for historical emissions, for which atmospheric pCO_2 is near-exponential, a first-order component of C_{dis} can also be understood as proportional to instantaneous pCO_2 via its relationship to the saturated carbon pool:

$$C_{dis} = f_{dis}(x, y)pCO_2(t) + \varepsilon_{dis}. \tag{2.8}$$

Thus, under these assumptions, there are physical understandings to a decomposition of integrated DIC into a pCO_2 term and a residual:

$$DIC(x, y, t) = f(x, y)p_{CO_2}(t) + DIC'(x, y, t), \text{ where} \tag{2.9}$$

$$f(x, y) = f_{sat}(x, y) + f_{dis}(x, y) \text{ and} \tag{2.10}$$

$$DIC'(x, y, t) = C_{bio} + \varepsilon_{sat} + \varepsilon_{dis}. \tag{2.11}$$

The pCO_2 regression accounts for most of the changes in C_{sat} and the background field of C_{dis} for interior water parcels. The calculated residual should include the impacts of carbon-climate feedbacks on C_{sat}, variability in the biological carbon pump, and changes in the pathways of ocean parcel subduction either due to climate variability or carbon emissions. By removing $f(x, y)p_{CO_2}(t)$ as the background field B_i as in (2.3), we isolate the variability in historical carbon inventory observations and model outputs to focus on how interannual and decadal changes in carbon content can be extended on regional scales.

Note that, when atmospheric CO_2 experiences atmospheric growth, this separation of DIC into the first-guess pCO_2 term and a residual term is similar to the expected anthropogenic carbon flux as estimated for inverse modelling experiments (Mikaloff Fletcher et al., 2006; Lovenduski et al., 2008). Thus our separation is highly similar to common separations between anthropogenic and natural carbon, although the two are not identical as we include potential feedbacks within this decomposition.

All first guesses result in background error fields that have near-zero biases globally as all methods use a least-squares estimation: the largest biases in the model ensemble are $\mathcal{O}(10^{-4})$ mol C m^{-2} using the atmospheric CO_2 regression, orders of magnitude less than the anomaly range of $\mathcal{O}(10^{0})$ mol C m^{-2} (Fig. 2.2c). When a constant climatology value is removed, the carbon anomalies at BATS are highly positively correlated with most global

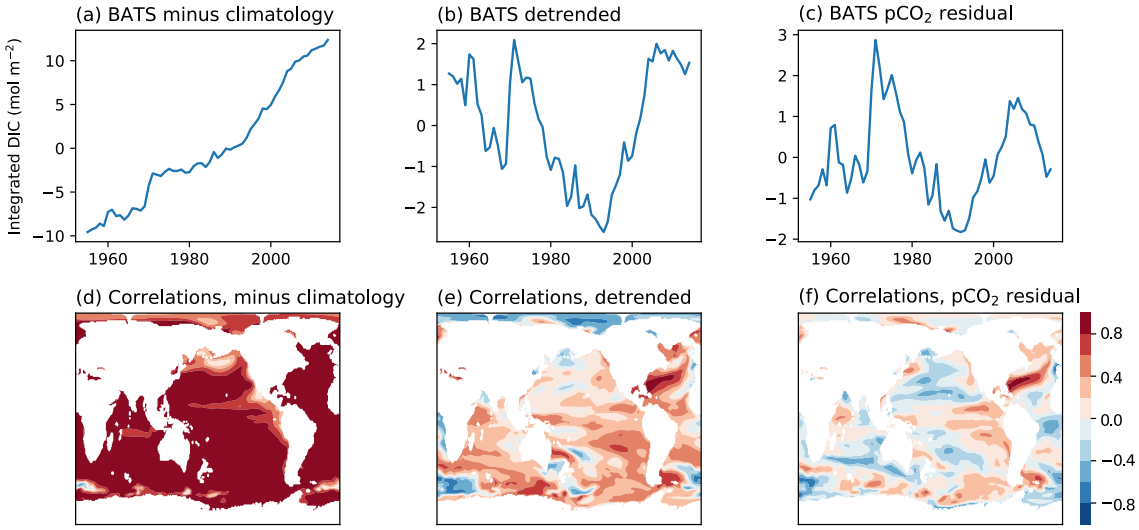

Figure 2.2: (a-c) Comparison of BATS upper ocean (0-500m) carbon behaviour for a UKESM model realisation when removing various first guesses: (a) the climatological mean from 1955-2014, (b) the linear trend from 1955-2014, and (c) a regression against atmospheric CO_2 concentrations. (d-f) Pointwise correlations with carbon at BATS for residuals from (d) the global climatology, (e) global detrended carbon data, and (f) global carbon residuals calculated from a regression against atmospheric CO_2 concentrations.

anomalies (Fig. 2.2a,d). These correlations reflect a consistent rise in ocean DIC associated with the long-term increase in atmospheric CO_2. Removing a linear trend reduces the far-field error correlations, but large areas maintain relatively strong correlations (magnitudes ≥ 0.4, Fig. 2.2e). Modelled upper-ocean DIC inventories generally show some slight positive second-order behaviour when a linear trend is removed, visible as concave residual behaviour in Fig. 2.2b. Lastly, removing a least-squares regression against atmospheric CO_2 concentrations eliminates most of the higher-order behaviour in integrated DIC (Fig. 2.2c). Correlations against DIC at BATS are now heavily localised around BATS (Fig. 2.2f); far-field correlations, while still present in this realisation, are substantially smaller in most regions than those when considering the residual from a linear fit. The correlation field for the example realisation shows a potential connection between the BATS time series DIC residual and the Pacific, with an ENSO/PDO-type correlation pattern in the equatorial Pacific.

The near-zero bias in the background errors, coupled with the weakening of correlations far away from observation locations, indicate that the background field as defined in (2.9)

for the entire model ensemble can provide useful spatial limits for the extension of DIC observations through the background error covariances. The regression coefficients used to calculate the first-guess response show enhanced uptake with raised CO_2 concentrations in the subtropics, particularly in the Southern Hemisphere (Fig. 2.3a). This regionality mirrors the regionality in anthropogenic carbon storage found in Frölicher et al. (2015), suggesting that the first-guess field in our decomposition may provide an approximation of ocean added carbon. The standard deviation of the residual DIC content (Fig. 2.3b) is highest in the Pacific Ocean; the subtropical gyres, in contrast, are local minima for DIC variability.

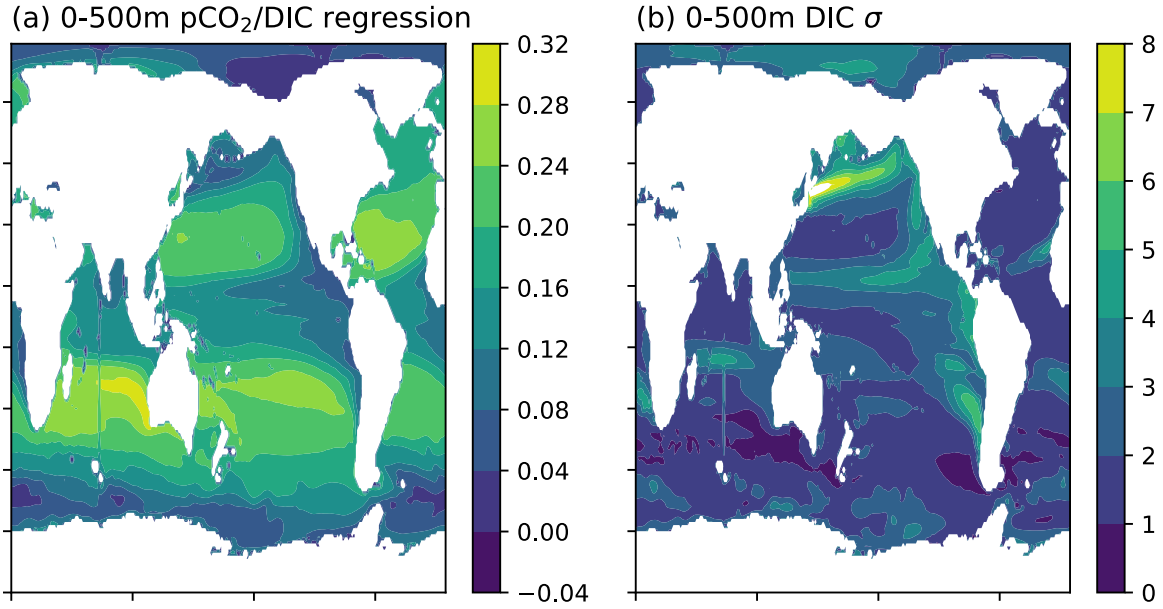

Figure 2.3: (a) Model ensemble regression coefficients f_i for 0-500m integrated DIC against atmospheric pCO_2 following (2.9), in units mol m^{-2} (μmol CO_2)$^{-1}$ and (b) standard deviation of the residual DIC, in units mol m^{-2}.

2.3 Representativeness of carbon timeseries

To better understand how data from observation sites is propagated, we consider the covariance/correlation fields for two time series sites corresponding with the locaton of BATS and the Hawaii Ocean Time series (HOT), two of the longest-running ocean biogeochemical time series. Both time series are located within Northern Hemisphere subtropical gyres and have the benefits of being close to existing oceanographic research centres, having deep

enough waters (over 4000m) so that coastal and shelf dynamics play little role in setting the physics and biogeochemistry, and being located within the oligotrophic subtropical gyres, which are regions of high DIC storage.

BATS: Located at 31°41'N, 64°10'W, BATS has been the site of water column biogeochemical observations since 1988 around 14-16 times each year. BATS is located within the North Atlantic Subtropical Gyre and is downstream of Subtropical Mode Water formation. Surface measurements of salinity-normalised DIC have found a substantial increase in accord with rising atmospheric CO_2 concentrations consistent with trends in other time series around the world (Bates et al., 2014). The North Atlantic Oscillation has been found to impact variability in interior temperature, salinity, and density (Phillips and Joyce, 2007) and surface DIC concentrations (Bates, 2001).

HOT: Located near the centre of the North Pacific subtropical gyre at 22°45'N, 158°W, HOT is a time series made up of near-monthly observations of the full water column depth at Station ALOHA (Fujieki et al., 2018). Observations of full water column hydrography, chemistry, and biology have been conducted since October 1988. Like BATS, surface DIC at HOT shows a long-term increase associated with rising atmospheric CO_2 concentrations (Keeling et al., 2004; Bates et al., 2014), as well as interannual variability (Brix et al., 2004) and longer-timescale variability (Keeling et al., 2004). The Pacific Decadal Oscillation is believed to play a role in setting mode water properties that helps to set the carbon variability at HOT.

As mentioned previously, in this chapter we conduct idealized experiments to see what the CMIP6 climate models reveal about the representativeness of carbon time series sites. We first aim to characterise the behaviour of DIC at BATS and HOT for the model realisations used in our ensemble. The 60-year time series at BATS is characterised by a strong positive trend (Fig. 2.4a). The models show varying levels of interannual and decadal variability. At HOT, the 60-year trend in integrated DIC is lower than that for BATS for some of the models (Fig. 2.4b). The modelled DIC at HOT shows more variability on interannual and decadal timescales, although the level of variability still appears to depend on the model.

2.3.1 Correlation fields for carbon at BATS and HOT

For the figures used, we provide correlation fields rather than covariance fields so that they may be more easily understood. The correlation fields are for the ensemble background er-

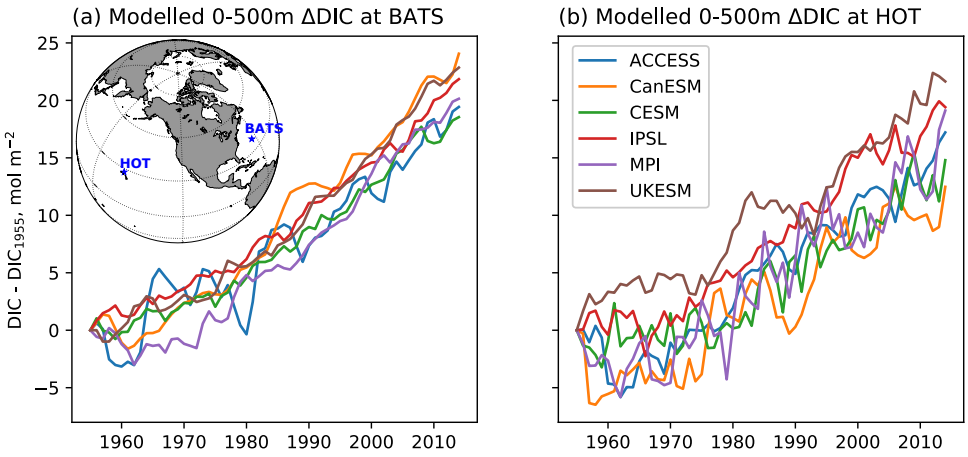

Figure 2.4: Integrated upper 500m DIC anomalies relative to 1955, in units mol C m^{-2}, at (a) BATS and (b) HOT. Time series shown are representive runs for each model within the ensemble. Inset is a map showing the locations of BATS and HOT.

rors DIC', calculated using the atmospheric pCO_2 regression in Section 2.2. The correlation fields provide the spatial limits for extrapolating observations; for regions with a small DIC' correlation against the observation locations, little to no information from the observations is included in the reconstruction. In this chapter, the synthetic reconstructions attempt to reconstruct modelled historical carbon variability using perfect information from the locations, so the correlations are the sole determining factor in setting the improvement in the analysis.

The use of the model ensemble to construct carbon content covariances, in comparison to one model realisation, results in near-zero far-field correlations against the DIC responses at both BATS and HOT (Fig. 2.5). There are strong positive correlations in the region surrounding each of the time series, indicating that the time series results may indeed be representative for the western subtropical North Atlantic for BATS and the central subtropical North Pacific for HOT. The correlation functions for both time series are anisotropic: the region of high correlation to both time series sites is oriented more zonally than meridionally, although the regions of equal correlations have major axes that are not purely zonal. The presence of this anisotropy indicates that the use of model ensembles can provide extra skill in reconstructing regional DIC changes that would not be found when using parametrised length scales seen in other interpolation schemes, similar to the temperature covariance results in Smith and Murphy (2007). Both BATS and HOT correlation

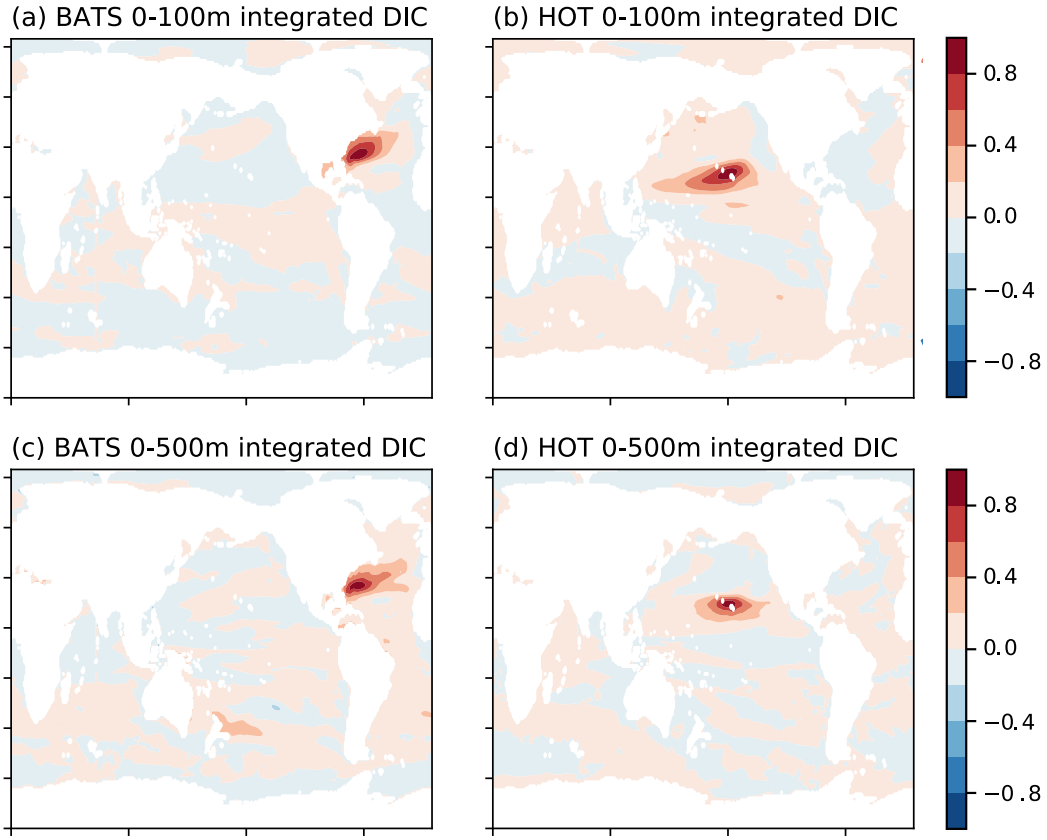

Figure 2.5: Ensemble correlations against integrated DIC at BATS and HOT for different depth horizons: 0-100m (a,b) and 0-500m (c,d).

fields show structure that is dependent on depth. The extent of high correlations decreases as the depth integration increases, and for the time series at HOT the correlations become more zonally-oriented when the integration is performed over the top 500m.

2.3.2 Error reduction from expanding BATS and HOT time series data

By construction, the areas with the highest background error correlations will show the largest improvement by adding observations to the first-guess field. We continue by using the NorESM model as a test to see how the inclusion of time series sites impacts the errors generated by the analysis. If the covariance field is perfect, the errors everywhere should decrease; however, as each model is different, the covariances from the ensemble likely contain some errors relative to the NorESM truth. Using the NorESM model can be an interesting test as its ocean model uses density rather than depth coordinates, making its physical setup substantially different from the other models included in the ensemble.

Integrated carbon in the NorESM model over the top 500m has a different regional pattern of variability compared to the model ensemble (Fig. 2.6a vs. Fig. 2.3b). The root mean square error of the interpolated product (Fig. 2.6b) is similar to the standard deviation of the model truth, which is expected as the regions of high correlation are closely confined to the time series' locations. As a result, the areas of highest improvement are small (Fig. 2.6c). The NorESM also shows a region where the error is increased by the analysis in the Southern Ocean. This area is of a similar size to the areas of improvement; however, the method used has no localisation procedure that would further dampen far-field correlations towards 0. The Southern Ocean is a region where the use of an isopycnal vs. a z-level ocean model can make substantial differences in the circulation, so the area of increased error could indicate the shortcomings of using a relatively small and similar selection of models to create the ensemble.

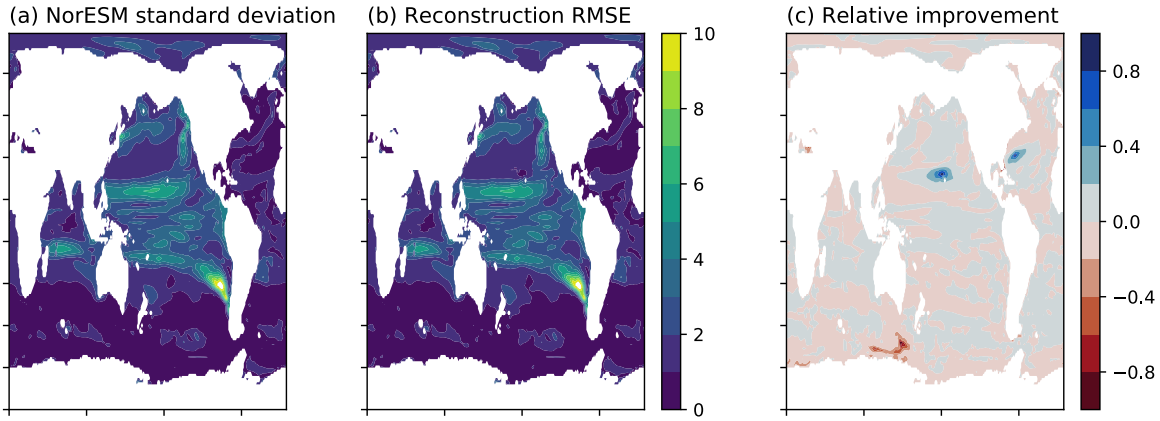

Figure 2.6: (a) Standard deviation of the NorESM residual, the "model truth", (b) the root mean squared error (RMSE) for the analysis using NorESM 0-500m DIC data at BATS and HOT, and (c) the relative improvement of the analysis over the first-guess field, $(RMSE(T_i - B_i) - RMSE(A_i))/RMSE(T_i - B_i)$.

2.4 Extrapolation limitations of super-sparse datasets

Our model experiments show a limited capacity for ocean time series to represent basin-wide or global carbon variability. When including pseudo-observations at BATS and HOT, small improvements to the carbon behaviour are apparent in the immediate vicinity of the time series, but more than 10° away from the time series any improvement is negligible. In order to better understand how this limitation occurs, we attempt to alter our method to

focus on near-surface ocean carbon variability to see if patterns associated with modes of variability become visible, and conduct a length scale analysis to see if this limited extension is characteristic of the global ocean or a product of the time series' locations.

2.4.1 Lack of patterns of climate variability

Hydrographic and biogeochemical variables at BATS have been found to be correlated with some of the main climate modes of variability (Bates, 2001; Brix et al., 2004). Using neural networks, internal variability in detrended ocean CO_2 partial pressure $pCO_{2,ocean}$ has also been found to correlate with major modes of climate variability (Landschützer et al., 2019). It is therefore surprising that the regions with strong correlations against carbon variability at BATS and HOT (Fig. 2.5) are mainly local and do not reflect any patterns associated with modes of climate variability. Unlike the bipolar patterns of variability for ENSO and NAO, correlations of carbon content against the time series sites show no regions that are significantly anti-correlated.

Modes of climate variability such as the NAO, ENSO, and the PDO have far-reaching teleconnections through their alteration of surface pressure and the associated geostrophic wind velocities. Through altering the wind fields, climate variability can impact time series carbon through air-sea fluxes or rates of entrainment. Time series data and climate models indicate that surface DIC at BATS decreases (increases) during a positive (negative) winter NAO phase, although the models do not agree on the magnitude of this pattern (Keller et al., 2012b). A stronger NAO is associated with decreased wind stress and mixed layer depth in the region of BATS, which hinders the upwelling of carbon-dense interior waters to the surface. On multidecadal timescales, the Atlantic Multidecadal Oscillation (AMO) may have similar impacts on North Atlantic subtropical gyre DIC through its impact on vertical mixing (Breeden and McKinley, 2016). For HOT, a positive Pacific Decadal Oscillation (PDO) is associated with warmer SST and lower wind speeds. The decrease in wind speed weakens vertical entrainment, lateral mixing, and net community production, creating lower near-surface DIC Brix et al. (2004). Significance testing of the upper ocean carbon cycle at HOT indicates that, on interannual time scales, the influence of ENSO and the PDO on climate variability at HOT is relatively small.

The lack of correlation patterns consistent with modes of variability in the time series correlation fields may be a result of the integration depth used when considering changes in ocean carbon content. By integrating from the surface to 100m or 500m, the carbon content anomalies contain carbon in the thermocline at BATS and HOT. These thermocline

waters are older than surface waters and are unable to respond to the atmospheric forcing associated with these modes of variability; rather, they will contain the imprint of these modes of variability from the year they were last at the surface. Thus the integration of carbon anomalies with depth involves the integration of multiple years of atmospheric forcing: as these modes of variability are noisy, the integration has the potential to dampen any patterns associated with the climate variability forcing. Reducing the depth level for the carbon anomaly integrations may produce stronger signals associated with these atmospheric-ocean coupled modes of variability, as it restricts the covariances to waters that have been in contact with the atmosphere at each year.

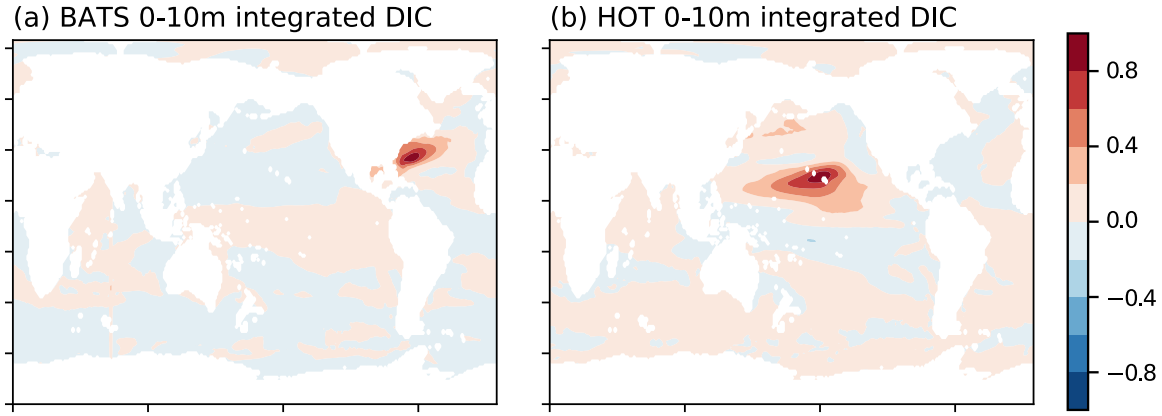

Figure 2.7: Ensemble correlations for 0-10m integrated DIC against integrated DIC at (a) BATS and (b) HOT.

Restricting the depth integration to the top 10m (Fig. 2.7) continues to produce correlation fields similar to those for the upper 100m (Fig 2.5). The correlation field for HOT has a larger zonal extent than the correlations including deeper carbon variability (Fig. 2.5b,d). Outside of the time series region, correlations remain close to zero. Even when considering the near-surface DIC content, which should be set mainly by atmosphere-ocean fluxes, the correlation fields still show no further-field structure that could be associated with of any modes of climate variability on the correlation fields. The regional behaviour of the correlation fields is thus robust in terms of the depth level used for the DIC integration.

Calculating the surface DIC correlations for each model within the ensemble can provide insight into whether this strong regional behaviour is robust across CMIP6 models or is a result of the ensemble's dampening of more uncertain relationships. At BATS, each model within the ensemble shows the strongest correlations in the immediate vicinity of the time series (Fig. 2.8). The extent of high correlations is highly model-dependent, with ACCESS

and MPI showing a sharp decay in the correlation field, whereas CanESM, CESM, IPSL, and UKESM show high correlations for most of the western North Atlantic. Some models, like CanESM and IPSL, show moderate correlations between BATS surface carbon and Pacific or Southern Ocean carbon; however, the sign and magnitude varies among the different models.

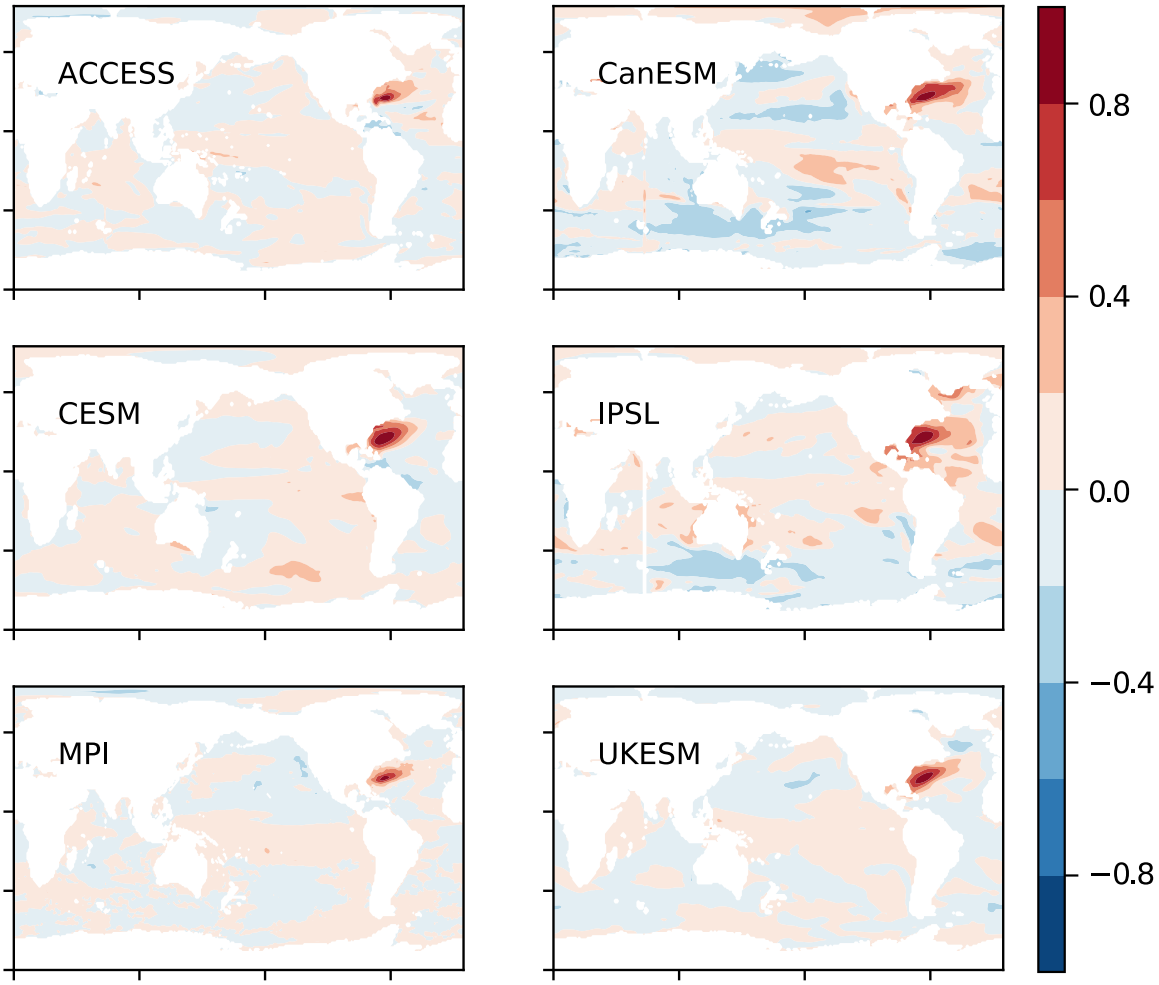

Figure 2.8: Correlations against BATS 0-10m integrated DIC for each model within our model ensemble.

Single model far-field correlations against surface DIC at HOT are generally stronger and show more structure than those at BATS (Fig. 2.9). However, there is still a substantial amount of inter-model uncertainty within the ensemble. The extent of high positive correlations around HOT ranges from almost entirely zonal (ACCESS) to covering most of the low-latitude North Pacific (CESM and UKESM). Most models, with the exception of

ACCESS, show moderate negative correlations between surface DIC at HOT and surface DIC in the equatorial Pacific. Individual models show moderate correlations between DIC at HOT and DIC in the Southern Ocean (positive in UKESM) and between DIC at HOT and DIC in the North Atlantic (negative, in CanESM and to a lesser extent IPSL). These weak correlations, combined with the ensemble method used in our optimal interpolation scheme, effectively cancel out any far-field correlations so that the impact of time series information farther away from the observations is insignificant.

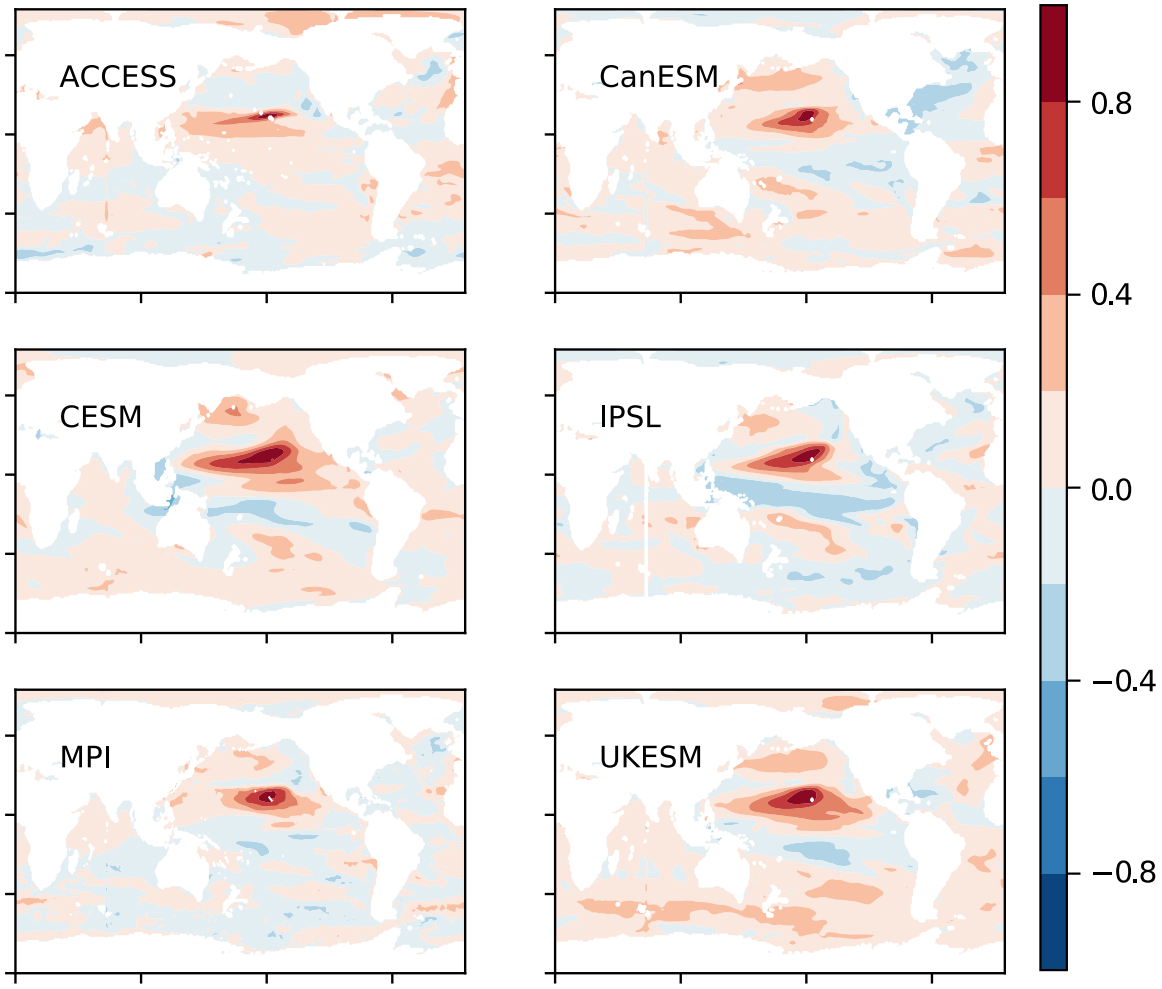

Figure 2.9: Correlations against HOT 0-10m integrated DIC for each model within our model ensemble.

Our results indicate an important distinction between drivers of carbon variability and the extension of that variability to other locations. Real-world observations find statistically significant correlations between carbon behaviour at the BATS and HOT time series sites,

but these correlations are typically of moderate strength ($\pm 0.3 - 0.5$). Weaker correlations can be used to compare historical time series variability with known modes of variability such as the NAO or ENSO. However, in a predictive and inverse sense, when using the EnOI method to extend observations, moderate correlations in the $\pm 0.3 - 0.5$ range result in small improvements in the reconstructed error (compare Fig. 2.5 with Fig. 2.6). Our method thus requires information to be strongly correlated in order for it to be extrapolated.

There are multiple physical reasons for the EnOI method to be poorly-suited for extending carbon observations to patterns of climate variability. The covariances calculated within our EnOI method use 60-year periods and will therefore resolve both decadal and some multidecadal variability alongside interannual variability. There has been no low-pass filtering applied to the method to isolate the extension of certain types of variability, so the covariance fields contain information about multiple timescales for carbon variability. Additionally, the use of climate models within our EnOI method may limit the connection between time series data extension and climate variability patterns, as climate models generally exhibit too low of a signal-to-noise ratio for some modes of variability such as the NAO (Smith et al., 2020). More predictability in climate modes of variability may lead to far-field extension of DIC variability as the ocean DIC content should act to integrate the impact of interannual modes on longer timescales.

The lack of correlation patterns associated with near-surface DIC measurements at BATS and HOT indicates that the regional extrapolation limits are robust across integration depth. As BATS and HOT are both located within subtropical gyres, this robustness may be a feature of their dynamic regime. Thus we continue by extending our scope globally to see if the regional representativeness seen at BATS and HOT is characteristic of the entire ocean, or if it is a product of the timeseries' locations.

2.4.2 Length scale analysis

Although the carbon variability at the BATS and HOT time series sites have limited extensions to regional carbon behaviour in these synthetic tests, this result may be more reflective of the locations chosen for the time series rather than the general capability of ocean time series to represent basin-wide behaviours. We can quantify the extent to which point observations may be expanded using a *correlation length scale*. If the locations of BATS and HOT are poorly-suited for extrapolating farther-field carbon variability, we expect the correlation length scales at these locations to be on the lower end of the range of values globally. The correlation length scale can be calculated in a variety of ways depend-

ing on the system being studied; for our work we calculate the length scale from a Gaussian fit as this aligns with previous studies considering hydrographic and carbon length scales arising from other ensemble-based optimal interpolation (Smith and Murphy, 2007) and ocean data assimilation experiments (Mazloff et al., 2018).

To calculate the correlation length scale, we follow a method similar to those of Smith and Murphy (2007) and Mazloff et al. (2018) and separate the correlation length scales into zonal and meridional length scales. We assume that the zonal and meridional correlation functions of integrated carbon against carbon at a location (x_0, y_0) follow a Gaussian, a negative exponential with regards to the distance between the two points:

$$\rho_{zon}(x, x_0) = e^{-\frac{(x-x_0)^2}{b_x^2}}, \tag{2.12}$$

$$\rho_{merid}(y, y_0) = e^{-\frac{(y-y_0)^2}{b_y^2}}. \tag{2.13}$$

The zonal and meridional length scales b_x, b_y, respectively, can then be approximated. At a zonal distance of b_x or a meridional distance of b_y from the original sample location, the correlation between the two points is $\rho = e^{-1} \approx 0.4$. This method assumes that the correlations are generally zonally and meridionally symmetric. To calculate the length scales b_x, b_y, we restrict the correlation domain to points near the original sample location, within about 5°, so that the correlations used for the length scale calculation remain positive and are not close to zero. Thus it should be noted that the correlation length scales over a few hundred kilometers should not be taken literally, as points at that distance have not been included in the calculation. However, because of the squared exponential, the correlations scale with distance, so the length scales are still applicable for points closer to the observation location. For instance, for a length scale b_x, points at a distance $0.5 \times b_x$ from an observation location will have a correlation of $\rho \approx 0.8$ to the observation location.

Using the CMIP6 ensemble, the zonal and meridional length scales show regional structure (Fig. 2.10). Zonal length scales are on the order of 1000 km and generally increase equatorward, with the maximum length scale of around 5000 km found in the equatorial Pacific for the integrated top 100m DIC. The meridional length scales are smaller, reaching up to about 1200 km, but have more complicated regional structures. The model ensemble shows larger meridional length scales within equatorial upwelling regions and in the subpolar South Pacific. The range of the zonal and meridional length scales and complicated structure compares well with the low-pass heat and carbon content length scales of Mazloff et al. (2018) for the Southern Ocean, even though their input of 90-day running means accounts for more intra-annual variability than the annual average output used here. The

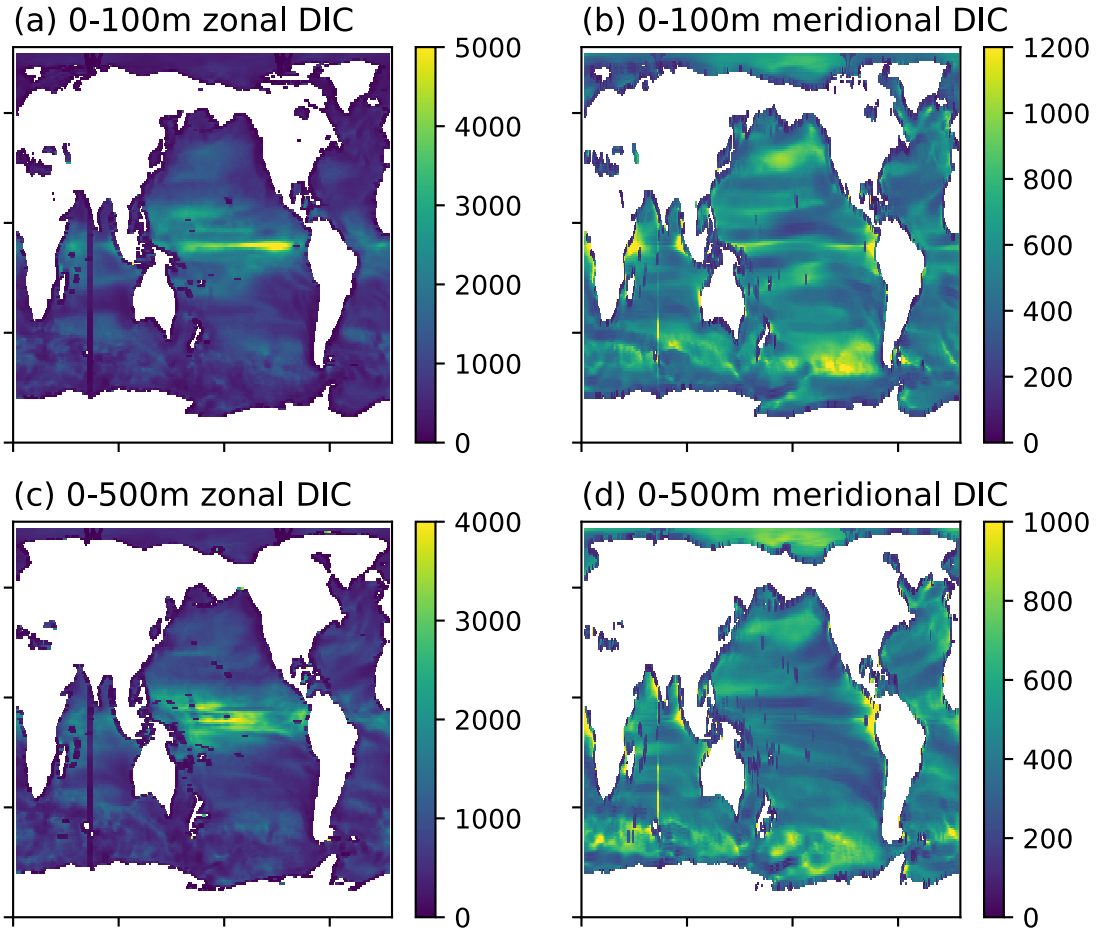

Figure 2.10: Characteristic zonal and meridional length scales, b_x ad b_y respectively, in units km, for 0-100m integrated DIC residuals (a,b) and 0-500m integrated DIC residuals (c,d).

correlation length scales for both BATS and HOT are relatively small compared to the rest of the ocean. The zonal and meridional length scales for the top 100m integrated carbon at BATS are $b_x = 1142$ km and $b_y = 527$ km, respectively; for the top 500m integrated carbon these length scales reduce slightly to $b_x = 1004$ km and $b_y = 474$ km. The zonal and meridional length scales at HOT are similar ($b_x = 1291$ km, $b_y = 582$ km for the top 100m; $b_x = 1111$ km, $b_y = 416$ km for the top 500m).

Overall, the global length scales diagnosed from the model ensemble indicate that the time series sites are too sparse to properly construct any sort of larger-scale variability. There is a potential for equatorial time series sites to be representative of larger areas of the ocean, but even this limiting case would still only cover part of the equatorial carbon

variability. High enough correlations necessary to propagate information from the time series will generally only be found within a few hundred kilometers of any time series, not just those time series locations we have considered up to now.

2.5 Discussion and Conclusions

In this chapter, the potential to expand time series carbon observations at BATS and HOT to regional and global scales has been explored using an ensemble optimal interpolation (EnOI) scheme with CMIP6 models. Although ocean carbon has substantial regional and temporal variability, the sparseness of carbon observations complicates the characterisation of such variability. The EnOI scheme uses model covariances to extrapolate information from observations into data-sparse regions and uses uncertainties within the model ensemble to prevent the spurious extension of information to far-removed locations.

Synthetic model tests have been conducted to quantify how the errors in carbon content are reduced when perfect data from the time series sites is extrapolated. Compared to a first-guess ocean carbon inventory set by atmospheric pCO_2 increases, the reconstructions show error reductions in the regions surrounding BATS and HOT when pseudo-observations from those locations are included. The error improvements are mainly restricted to the immediate areas surrounding BATS and HOT; in most other regions the error reduction is close to zero, and sensitivity tests show the potential for errors to increase in some regions. Thus, although time series provide point-wise measurements of the ocean interior that contain good coverage in depth and time, our results indicate that their observed variability is limited in how it represents variability in nearby ocean regions.

In this discussion we consider the two separate themes of work in this chapter. Firstly, we consider the EnOI scheme itself, and how we understand the assumptions and construction necessary to create appropriate covariance fields to extend ocean carbon data. Secondly, we consider the potential for time series sites to represent historic carbon variability for their regions and beyond. We end the discussion with some speculation on how the method can be expanded to include real data in a general sense, ignoring the limitations of the time series data.

2.5.1 Suitability of the Ensemble Optimal Interpolation method for ocean carbon

Ocean interior temperature and salinity fields have been successfully reconstructed using EnOI with a climatology first guess (Cheng and Zhu, 2016; Smith and Murphy, 2007). The-

oretically, a climatology first-guess field is ideal because it makes the fewest assumptions about the reconstructed variable's behaviour in time. However, we find that reconstructing ocean interior carbon content with a climatology first guess results in strong positive correlations with the time series sites almost everywhere Fig. 2.2 a,d). As this chapter considers the improvement from data at only two locations, a reconstruction using a climatology first guess would reconstruct the historical trend in integrated ocean carbon. This result arises from our choice to reconstruct integrated ocean carbon content rather than DIC concentrations at specific depth levels. We have made this choice as the integrated DIC content avoids some of the difficulties in separating historical changes into responses to forcing and climate variability seen on longer timescales. In addition, the decision to reconstruct integrated upper-ocean carbon avoids difficulties that may arise from model differences in mode water formation or depth structure.

As the correlation fields associated with a climatological first guess represent the long-term trend in ocean carbon rather than the variability, in this chapter we have employed a novel first-guess field. In lieu of a climatology first-guess, we separate ocean carbon into a term proportional to atmospheric CO_2 mixing ratio and a residual (2.9). The residual component includes changes of carbon concentration feedbacks on the saturated carbon pool, changes in the biological carbon pump, and changes in ocean circulation and ventilation. Testing the EnOI method with this separation shows a reduction in the influence of time series data on far-removed locations, consistent with the idea that an observation's influence on the analysis should decrease as the distance between the two locations increases (Fig. 2.2 c,f). While the removal of an atmospheric pCO_2 term from integrated DIC increases the complexity of the reconstruction, the well-mixed nature of atmospheric CO_2 globally on annual and longer timescales allows us to treat atmospheric CO_2 as a separate, independent variable for the entire domain.

2.5.2 Extrapolating observations using CMIP6 carbon variability

In this chapter we have used an ensemble of CMIP6 historical runs to construct the covariance fields that are used to extrapolate the time series data. The CMIP6 models were chosen because they include the most up-to-date and comprehensive representations of global climate and carbon dynamics, and the specific models were chosen because their multiple realisations allowed the ensemble to include uncertainties related to climate variability. However, these models sample a subset of the full uncertainty within the climate system. Thus some features found to be robust in the background error covariance fields may be

artefacts of the ensemble choice. We speculate that the risk of such artefacts is small, as the ensemble correlation fields tend towards zero correlation as the distance from the time series sites increases, and there are few far-field maxima or minima in the correlation fields (Fig. 2.5, 2.7). Before applying the background error covariance fields to observational carbon data products, sufficient ensemble sensitivity testing will be necessary; even without sufficient observational coverage, the covariance fields can still be tested for convergence.

The extrapolation of carbon observations using model ensemble covariances assumes that the variability within the models is a reasonably accurate representation of real-world variability in both space and time. In order to create the most accurate covariance fields possible, the models used should reflect real-world relationships as accurately as possible. In the previous generation of CMIP5 models, the sensitivity of carbon uptake to climate variability was found to be smaller than that in observationally-constrained products (DeVries et al., 2019). While the method in this chapter focuses on regional covariances of ocean carbon and uses the new generation of climate models to construct its ensemble, there is still the potential that the CMIP6 models are not sufficiently sensitive and thus overly dampening of real variability. As such validation studies have not yet been published for this generation of models, this aspect is still entirely speculative. However, if the CMIP6 models are also found to be overly dampening, it may be useful to conduct similar covariance analyses using observationally-constrained products for the background error covariances to see how the strength and patterns of variability impact carbon mapping products. The model covariances can also be improved using iterative procedures that combine observational data with model data so that such errors in the model ensemble are mitigated (Allison et al., 2019).

2.5.3 Potential and limits of extrapolating time series data

Reconstructing model carbon variability using synthetic profiles and CMIP6 covariance fields shows that the data from pointwise observations can be extrapolated to account for variability in a wider area. The extent of this wider area, however, is relatively small for existing time series sites. The model ensemble used in this chapter shows that time series sites BATS and HOT have a local region in which their information can improve carbon reconstructions (Fig. 2.6c). In addition, the CMIP6 model ensemble indicates that, when using the decomposition with atmospheric pCO_2, the remaining variability at BATS and HOT is relatively small (Fig. 2.3b). Thus the carbon time series data at BATS and HOT is poorly suited for understanding basin-scale or global carbon variability. However, the

dominance of an atmospheric pCO_2 trend at these locations, coupled with the far-field strongly positive correlations found when taking a climatology first-guess field, substantiate claims that the time series locations are well-suited for characterising the long-term changes in ocean carbon (Bates et al., 2014; Heinze et al., 2015).

The extension of carbon data from BATS and HOT, as made visible by the ensemble correlation, do not reflect larger-scale modes of climate variability (Fig. 2.5, 2.7). The lack of correlations that reflect these modes may partially be a result of the intermodel uncertainty in the pattern and strength of DIC anomalies, even for surface DIC (Fig. 2.8, 2.9). However, the absence of larger climate-mode-type correlation structures may be a result of the decisions made in the analysis. In regridding and concatenating the model outputs, we made the choice to calculate annual average DIC fields to minimise the impact of the seasonal cycle and associated ocean biology. Similarly, we could have isolated specific seasons, such as the winter months (December - March). Isolating the winter months could potentially lead to farther-field correlations for BATS, as the impact of the NAO on the North Atlantic is strongest in the winter (Keller et al., 2012b). While such an analysis would be less useful in quantifying global carbon variability, the potential far-field correlations would provide valuable insight as to how winter variability at BATS could be representative of larger-scale winter carbon uptake.

Globally, the small (generally < 1000 km) zonal and meridional length scales found in the CMIP6 ensemble indicate that, in order to create a global reconstruction of carbon variability, substantially more observation locations are required (Fig. 2.10). However, the regional skill that time series carbon data provides may still be useful for understanding smaller-scale processes and impacts of variability. The optimal interpolation method may also be used with regional biogeochemical ocean models to better characterise how carbon measurements may be extrapolated. As regional models resolve smaller-scale processes, some of the assumptions and simplifications used in this chapter may no longer be appropriate for such a study; for instance, the use of annual averaging will smooth out finer-scale features such as eddies that exist on seasonal to intraseasonal timescales. With a regional model, the carbonate chemistry variables and hydrographic variables measured at these time series sites may provide some insight to the state of the surrounding upper ocean, particularly with regards to the existence and persistence of extreme events such as short-lived periods of high upper ocean acidity (Burger et al., 2020).

Implications for ocean carbon networks

The length scale results from the CMIP6 ensemble indicate that a stationary ocean carbon monitoring network would need to include a large number of stations. If it were possible to eliminate all noise from carbon observations (i.e., there were no errors inherent in the observation network), and station frequency of 1 per lengthscale were adequate to resolve interannual and longer-timescale variability, over 1000 stations would be necessary to cover the global ocean. As observation include both errors and noise from other processes, the number of stations needed would likely be much larger, making such a network as determined by the model covariances logistically unfeasible. While concessions on coverage could be made to reduce the number of profiles so that only main regions of interannual variability were covered, there would still be budget and logistical difficulties in organising regular and frequent sampling missions. In this study we have necessarily ignored important benefits of current time series programmes such as proximity to strong oceanographic centres, rigorous sampling protocols, and the momentum that longer programmes have in renewing funding. Ultimately, any real-world sampling scheme must consider the economic feasibility as well as the scope of information obtainable from the observed locations.

Expansion to real world carbon data

Regardless of the scale at which time series data can be extrapolated, there is potential to gain insight from the combination of observations and model covariances. Time series measurements can be expanded to regional scales using regional models, or larger-scale reconstructions can be created when including other ocean carbon measurements from repeat transects and other observations within the GLODAP dataset.

The further development of this method to include real observations requires considering observational errors. In full EnOI schemes, the gain matrix considers both errors within the covariances and those within the observations themselves. The inclusion of observational errors will result in an algorithm that is no longer a least-squares solution; rather, our gain matrix will be calculated by the model ensemble covariance matrix P^b (exactly the covariance field for DIC') and observation error covariance matrix R alongside the mapping operator between model space and observation space H (Oke et al., 2010):

$$G = P^b H^T (H P^b H^T + R)^{-1}. \tag{2.14}$$

The observational error covariances will include errors from the measurements themselves but will also need to include errors induce by any choices made to bin, integrate, or average

the observations. Optimising the observational error covariances will require sensitivity tests and expert judgement.

In summary, the use of CMIP6 model ensembles covariance fields can help expand sparse interior carbon measurements to regions that have little to no carbon observations. The impact of carbon time series is dependent on whether the carbon trend or variability is being studied. The trend in upper ocean carbon at time series sites is highly positively correlated with the trend in upper ocean carbon globally; this trend can be approximated as a function of atmospheric CO_2 mixing ratios, emphasising the role that historical emissions have played in setting the historical ocean carbon response. The background error covariance fields for ocean carbon against measurements at any specific point are highly localised and show little extension to basin-wide or global carbon variability. Thus the carbon variability noticeable in multidecadal time series such as BATS and HOT may be expanded to nearby regions. The extension of time series carbon data to wider regions is limited because of uncertainties in far-field correlations. Modes of climate variability, while they impact carbon variability on location, are of limited use when expanding this induced variability. The uncertainties and small-scale features of ocean dynamics couple together to reveal that the difficulties of extremely sparse ocean carbon measurements outweigh the advantages of the highly-resolved temporal variability that time series sites provide.

Our results using the ensemble optimal interpolation method with time series sites indicate that the CMIP6 models may be well-suited for expanding observations, but the time series data is simply too sparse to provide a broad view of historical carbon variability. Thus we continue our work with this method in the next chapter using more widely-spread ARGO data. By using ARGO data as inputs, we can explore how more frequently-sampled oceanographic data may be used to characterise carbon variability indirectly through their statistical relationships.

Chapter 3

Using temperature and salinity fields to reconstruct historical ocean carbon changes

This next chapter considers another method for reconstructing ocean carbon variability using an ensemble of CMIP6 models. As sparse, high-quality time series show a limited ability to represent larger-scale carbon variability, we explore how dense non-carbon measurements may be used to reconstruct carbon content. In this chapter we explore how well ocean carbon variability can be represented by Argo data, which provides temperature and salinity measurements up to 2000m depth from all over the world.

We begin this chapter by providing an overview as to how non-carbon ocean variables have been used to reconstruct ocean carbon, and the physical relationships between ocean temperature, salinity, and carbon solubility (Section 3.1). In Section 3.2, we expand upon the method using the CMIP6 ensemble introduced in Chapter 2. We investigate the spatial patterns in correlations between integrated carbon, temperature, and salinity, and find skill in reconstructing carbon by using same-location temperature and salinity data (Section 3.3). We then explore how the distribution of Argo measurements impacts the accuracy of the reconstructions using different search radii for creating our reconstruction scheme (Section 3.4) and discuss the implications of these results in Section 3.5.

3.1 Introduction

As ocean carbon measurements are sparse in both space and time, using them to characterise historical variability requires some level of extrapolation. In the previous chapter, we explored whether time series sites with dissolved inorganic carbon (DIC) observations can be used to characterise basin-scale or global carbon variability. While the data at these time series sites are regularly-sampled and high-quality, the paucity of such time series prevents their variability from being extended to larger scales, as the length scales associated with carbon variability are much smaller than the distances between time series sites (Fig. 2.10). Thus reconstructing ocean carbon variability requires measurements that are relatively dense. As autonomous DIC sensors are still in development (Bushinsky et al., 2019), carbon measurements themselves remain resource-intensive, and current reconstructions of ocean carbon variability rely on more frequently sampled oceanographic variables.

When considering historical ocean carbon changes, the ideal observations to extrapolate will be those of ocean carbon. However, if there are observations of other ocean properties, the relationships between these variables can be used to reconstruct ocean carbon. This type of indirect reconstruction requires some assumptions for how ocean carbon is related to other variables and introduces an extra source of error. The use of multiple linear regression is well-suited for including other variables as their solutions are set by the covariances among all the variables considered. The covariance field can then be used to explain the potential usefulness or limits of these types of reconstructions.

3.1.1 Oceanographic inputs for reconstructing carbon

To avoid the issue of sparse data points, other reconstructions of the ocean carbonate system have used a variety of methods and inputs. Global $pCO_{2,oce}$ reconstructions have been made using multiple linear regression against observed sea surface temperature and salinity, mixed layer depth, and chlorophyll concentrations (Rödenbeck et al., 2015). The GLODAPv2 database has been used to reconstruct alkalinity and pH via multiple linear regression trained to specific regions (Carter et al., 2018). The recent inclusion of pH sensors as part of the biogeochemical Argo programme has allowed for further reconstructions of ocean pCO_2 with substantial impacts on winter fluxes in the Southern Ocean (Williams et al., 2017a).

Machine learning is another technique used for reconstructing the ocean carbonate system from sparse measurements. Machine learning algorithms are more general than algorithms for linear regression and are able to represent higher-order relationships. Regional

and decadal variability in surface $pCO_{2,oce}$ has been reconstructed by neural networks, which self-organise the ocean into regions and solve for non-linear regressions between measured variables (Landschützer et al., 2016; Landschützer et al., 2019; Gloege et al., subm). Neural networks have been used to reconstruct ocean alkalinity, carbon, and pH using temperature, salinity, and oxygen from the GLODAPv2 dataset (Bittig et al., 2018). While these techniques have proven powerful for expanding sparse ocean datasets, it is more difficult to gain mechanistic insights or to check the relationships against well-known physical or biogeochemical principles. Additionally, the nonlinearity allowed within machine learning techniques can overestimate internal variability when data is extrapolated (Rödenbeck et al., 2015; Gloege et al., 2021).

The ensemble optimal interpolation method described previously (EnOI, Chapter 2) can also accept multiple inputs, as the method essentially operates the same way as multiple linear regression. EnOI bridges a gap between multiple linear regression using observations and machine learning methods, as the climate models provide the spatial covariances and covariances between variables. While the EnOI method is less flexible than more complex machine learning methods, its solutions can be directly related to covariance and correlation fields between variables. The simplicity of the method allows for more straightforward connections to physical processes and cross-checking against observations. As EnOI tends towards a first-guess field, overestimating internal variability does not occur as it can in nonlinear machine learning methods.

3.1.2 First principles relationships between ocean carbon, temperature, and salinity

In this chapter we consider how ocean temperature and salinity measurements may be used to reconstruct ocean carbon variability. While the use of temperature and salinity as proxy data is relatively simple compared to other published reconstruction methods, temperature and salinity measurements have been taken across the globe. The introduction of autonomous CTD profilers in the Argo programme has resulted in millions of profiles being taken over the 21st Century (Wong et al., 2020).

Ocean temperature and salinity have well-established relationships with ocean carbon. Carbon emissions induce thermal and carbon disequilibria between the atmosphere and ocean, leading to an oceanic uptake of both carbon and heat. Natural variability of temperature and salinity within the ocean system also impacts the climatological uptake of carbon. Under constant atmospheric CO_2 concentrations, temperature and salinity impact

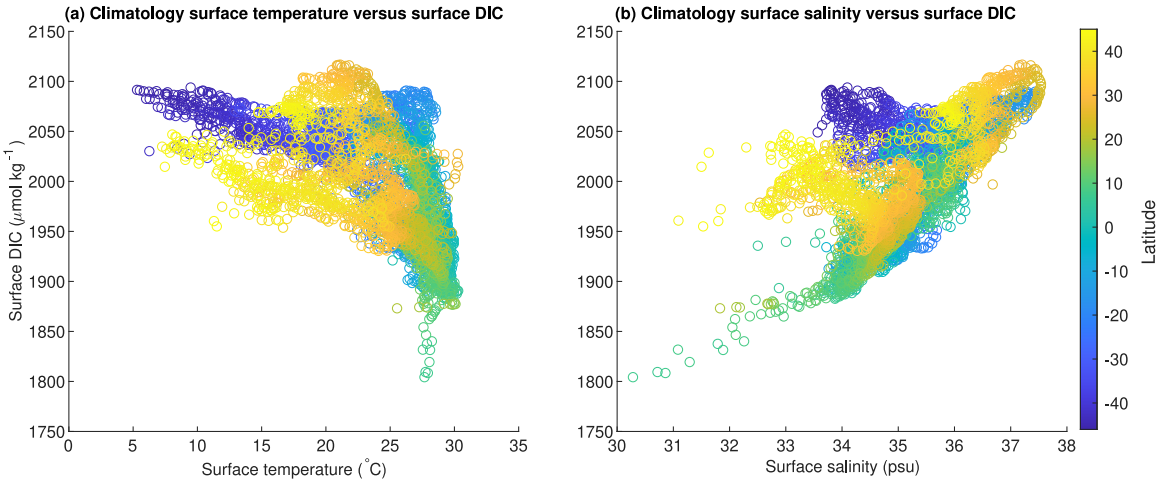

Figure 3.1: GLODAP gridded 1986-2013 climatology surface DIC, in units μmol kg^{-1}, against a) surface temperatures, in units °C, and b) surface salinity, in units psu. Data are plotted between the latitudes of 45° S and 45° N, with colour of markers indicating the latitude of the surface variables. Data from (Lauvset et al., 2016).

the amount of CO_2 that can be taken up until the water parcel is saturated. Higher water temperatures lower the solubility of CO_2 as the greater kinetic energy allows the gas to escape more easily (Fig. 3.1a). For ocean salinity the relationship is opposite: carbon dioxide is more soluble in more saline waters as the extra salinity raises the alkalinity (Fig. 3.1b). A higher alkalinity allows for more CO_2 to be taken up by the water parcel as the negative charges balance the acidification of the water parcel from CO_2 uptake. Constraining anthropogenic carbon uptake with sea surface salinity has been founded to lower the uncertainty in the Southern Ocean carbon sink (Terhaar et al., 2021).

Additionally, when considering integrated carbon storage, the carbon pump is a product of both ocean surface fluxes of CO_2 or biological uptake and regeneration of CO_2 and carbon divergence caused by ocean circulation. As ocean circulation is largely controlled by the impacts of temperature and salinity on ocean density and water masses spread along isopycnals to first order, temperature and salinity variability can alter the pattern of carbon content variability. Much of the anthropogenic carbon uptake and variability in ocean carbon is controlled by ocean circulation (Iudicone et al., 2016; DeVries et al., 2017; Gruber et al., 2019b); as such, temperature and salinity are promising proxies to help reconstruct ocean carbon variability.

Thus a linear reconstruction using temperature and salinity fields may provide a good first-order guess for reconstructing historical ocean carbon content. If carbon can be suffi-

ciently reconstructed using temperature and salinity measurements, Argo profiles can provide a large-scale view of ocean carbon variability not possible with ocean carbon time series sites. Conversely, a lack of skill in reconstructing carbon from these variables may indicate that carbon inventories have a larger biological control, or that the dynamics are more non-linear or poorly represented by the models used.

The Argo dataset

If ocean carbon can be sufficiently reconstructed using temperature and salinity measurements, then the extensive Argo dataset may be used to provide a large-scale view of ocean carbon variability. The Argo programme, in which autonomous floats record temperature and salinity in the upper 2000m of the ocean, includes over two million profiles from the year 1999 onwards (Wong et al., 2020). Coverage from Argo floats has significantly increased over the 21st Century and now spans most of the ocean (Fig. 3.2) The Argo sensors measure ocean temperatures accurate to ± 0.002°C and salinity accurate to ± 0.01 psu. Thus the Argo dataset presents the opportunity to use high-quality and horizontally-dense ocean measurements within an optimal interpolation scheme. While the correlations will necessarily be weaker than the carbon spatial correlations explored in Chapter 2, as they will take into account both spatial correlation length scales and the relationship between variables, the quantity of observations may still provide better estimates of carbon variability.

The ability to reconstruct ocean carbon variability from Argo temperature and salinity profiles relies on the relationships between these variables and ocean carbon and on the coverage of Argo floats in a given year. Argo floats are designed to descend to a depth of 2000m and follow the ocean interior currents for 2 weeks at a time. On their ascent to the surface, they measure temperature, salinity, and pressure; these measurements are then communicated via satellite transmission. As such, observations from the Argo programme are spatially irregular but cover areas of the ocean that are only rarely sampled by ships because of their logistical difficulties. The number of floats in the Argo programme has increased gradually to around 3000 floats in the ocean at any given time. The number of profiles taken by the programme has accordingly increased, and now covers most of the global oceans (Fig. 3.2).

3.1.3 Driving questions and objectives

As carbon observations are too sparse to be representative of global carbon variability, measurements of other variables related to the ocean carbonate system need to be used to

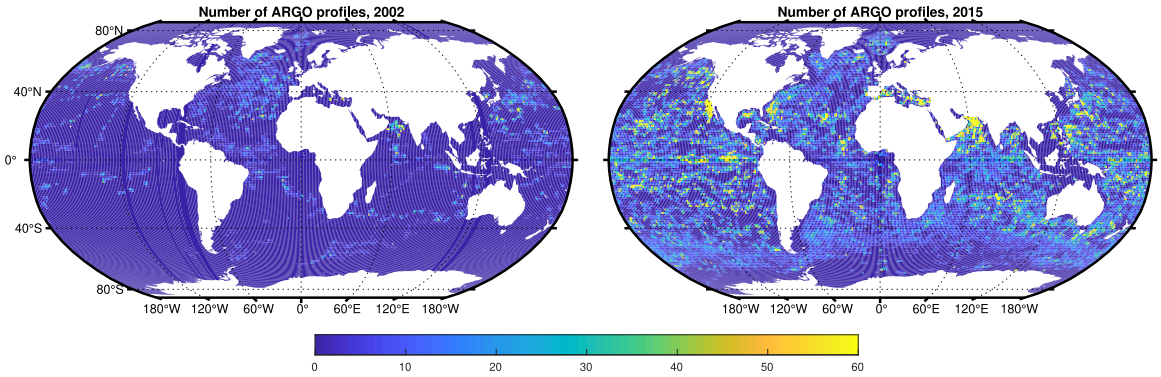

Figure 3.2: Binned Argo profiles for the years (left) 2002 and (right) 2015. Profiles were collected into $1° \times 1°$ bins for 01 January - 31 December of each year.

infer ocean carbon inventories. In this chapter, we continue with the EnOI method using the CMIP6 models from Chapter 2 to explore the potential for using Argo-style temperature and salinity profiles to reconstruct ocean carbon inventories. We continue with the use of synthetic model tests, in which we take the covariance fields between modelled carbon, temperature, and salinity to solve for the regression coefficients, and then calculate the reconstruction improvement over the first-guess field using modelled ocean temperature and salinity. In this chapter we aim to answer the following questions:

- Is there skill in reconstructing ocean carbon using temperature and salinity variability within the CMIP6 ensembles?
- Which regions are best suited for reconstructing ocean carbon using temperature and salinity variability?
- How does the horizontal distribution of Argo measurements impact the ability to reconstruct ocean carbon variability?

We hypothesise that the physical controls on CO_2 solubility will dictate the correlations and regression coefficients for both temperature and salinity (i.e., negative correlations and regression coefficients for temperature, positive correlations and regression coefficients for salinity). We hypothesise that there will be skill in reconstructing ocean carbon using ocean temperature and salinity content due to the role of the solubility pump in setting ocean carbon content. We expect this skill to be highest in regions with little biological activity such as the subtropical gyres, whereas in areas with high biological activity (such as upwelling regions) or significant model uncertainty (such as the Southern Ocean) the reconstruction will show less skill. Lastly, we hypothesise that using observations at Argo

locations will increase the skill of the EnOI method relative to pointwise reconstructions, as the inclusion of more data can allow for reconstruction of more complicated relationships between ocean thermal and carbon variability.

3.2 Ensemble Optimal Interpolation with temperature and salinity

The EnOI method in this chapter uses the same building blocks as that in Section 2.2, but with inputs (O_k in (2.1)) of temperature and salinity in lieu of carbon. The models and realisations used in this chapter are the same as those in Table 2.1. Monthly potential temperature and practical salinity output files (`otemp` and `so` in the CMIP6 database) were downloaded, converted to annual averages, and regridded to the $1^\circ \times 1^\circ$ horizontal grid.

As with DIC in Section 2.2, we separate integrated temperatures into a term directly proportional to changes in atmospheric CO_2 concentrations and a residual, in which the term proportional to atmospheric CO_2 is taken to be the background field:

$$T(x, y, t) = g(x, y)p_{CO_2}(t) + T'(x, y, t). \tag{3.1}$$

For salinity, we take the background field to be climatology. This choice for the salinity background field is made for two reasons: trends in salinity, while present on regional scales, are much smaller than trends in integrated temperature or carbon, and carbon emissions do not impact salinity directly as seen with DIC and temperature. Thus in lieu of a good physical reason for calculating a atmospheric pCO_2-related term we revert to the simplest first guess.

3.2.1 Spatial coverage of Argo floats

The Argo programme uses thousands of floats that measure profiles of temperature and salinity of the upper 2000m of the water column every 10 days. As the floats are passive drifters, their location is determined by ocean dynamics, and so the distribution of Argo measurements is irregular and time-dependent. In this chapter, we conduct synthetic tests in which the sampling distribution of Argo floats is replicated to explore how the sampling distribution impacts the EnOI error reductions.

At each month, the location of Argo profiles (Argo, 2021) were collected and binned in the $1^\circ \times 1^\circ$ horizontal grid. A bin is considered "observed" in a specific year if it contains Argo profiles for at least 6 months, and that bin is kept in the mask for the synthetic tests.

This estimate of "usable" Argo locations attempts to account for representation error in bins where the observations may be biased towards certain months or seasons, although this crude method will not account for all the types of representation errors associated with the use of passive drifters.

Reconstructions of ocean carbon variability were conducted using different search radii for Argo locations. As a first cut, we perform the test using pointwise temperature and salinity to provide a local view of the carbon reconstruction skill. Additionally, we conduct reconstructions in which the carbon variability at each point is solved as a function of temperature and salinity at all "usable" Argo locations within a certain radius. In this chapter we use search radii examples of $\pm 1°$ and $\pm 5°$ as these can show some of the abilities of localisation while still generally staying within the modelled correlation length scales for DIC (Fig. 2.10). For this chapter we use the binned Argo coverage from 2015 as a representative year. Global Argo coverage in 2015 using the binning and selection method described above is generally good (Fig. 3.3, dark blue), and when considering a radius of $\pm 5°$ most of the global ocean is covered, with only the high-latitude Southern Ocean and Arctic Ocean showing gaps in coverage (Fig. 3.3, light blue).

3.3 Local reconstructions of ocean carbon variability

We begin by exploring how temperature and salinity fields may be used to reconstruct ocean carbon in a pointwise manner. This reconstruction relies on the local covariances between temperature, salinity, and carbon, which set the weights in the EnOI scheme. Considering the pointwise relationships between variables temporarily avoids the complexities of irregular coverage, as multiple profiles within a search radius will likely be highly correlated. In this section, similar to Chapter 2, we illustrate these relationships using correlations rather than covariances, as the correlation fields take into account the temporal variance of carbon, temperature, and salinity at these points.

3.3.1 Correlations between upper-ocean temperature, salinity, and carbon

The ensemble correlations between the temperature residuals T' and carbon residuals DIC' (Fig. 3.4a) are negative in most regions, consistent with the physical control of temperature on CO_2 solubility. All ensemble members show the strongest negative correlations in the equatorial Pacific Ocean and equatorial Indian Ocean (not shown). The weak correlations in the Southern Ocean are a result of weak correlations within each model combined with a

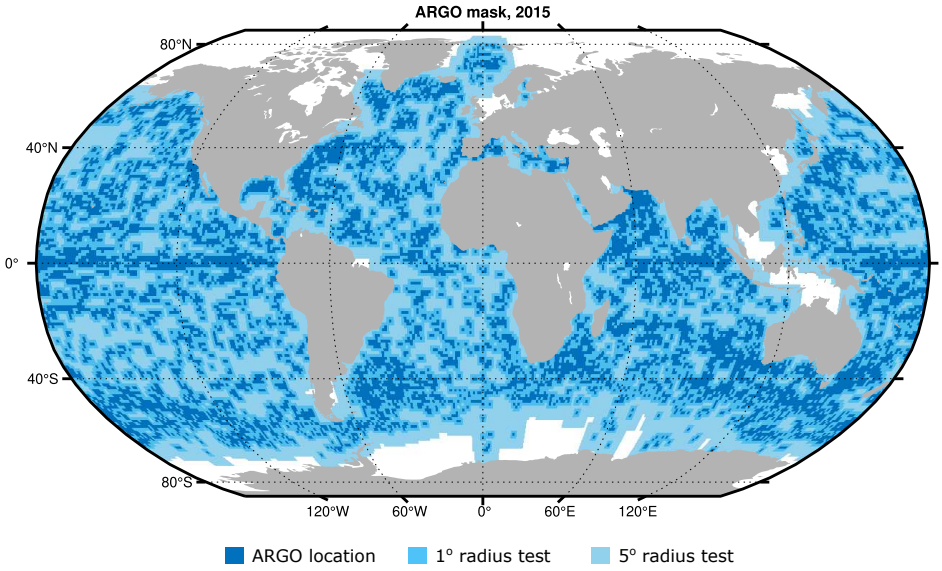

Figure 3.3: Coverage of binning procedure for Argo reconstructions using observation locations at the year 2015 (dark blue), with overage of experiments with search radii of 1° (bright medium blue) or 5° (light blue). Binning procedure and bin selection conducted from an original count of $n = 334881$ profile locations.

wide inter-model uncertainty: in each model, there is a region where correlations between temperature and carbon are positive, but the models differ on whether this region is confined to the Antarctic coast or extends further into the Antarctic Circumpolar Current region. There are significant regions where the correlation is positive in the ensemble and almost all the models, in particular the North Atlantic, the Labrador Sea, and the southeast Pacific and Southwest Atlantic.

The positive correlations between temperature and DIC anomalies are strongest in the Irminger Sea and Labrador Sea. These correlations contrast with the results of Fröb et al. (2016) and Rhein et al. (2017), which find that anthropogenic carbon storage in the Irminger and Labrador Seas increases with decreases in temperature. In a vertically-integrated view, the stronger winds and stronger heat fluxes from the ocean to the atmosphere lead to a strengthening of wintertime air-sea carbon fluxes and thus carbon storage. However, for this analysis we have chosen to focus on a static depth horizon and chosen a depth horizon that is particularly shallow when considering the convection depths in these regions of the North Atlantic. Thus, it is possible that, when considering near-surface carbon storage,

the increases in total carbon storage induced by the region's thermodynamics are offset by the dilution effect from the redistribution carbon from the upper 100m during periods of stronger convection.

Correlations between salinity residuals and carbon residuals (Fig. 3.4b) are mainly positive, consistent with the physical control of salinity on CO_2 solubility via ocean alkalinity. These positive correlations are strongest in the western tropical and subtropical Pacific, the North Atlantic, and the high-latitude Southern Ocean, although the correlations for the Southern Ocean are less certain among the ensemble members (not shown). The models all show a region of negative correlations between temperature and salinity in the equatorial Pacific upwelling region, which transports fresher, carbon-rich waters upwards from the deep ocean. Additionally, the correlation between salinity and carbon shows a local minimum in the Southern Ocean subduction region.

Lastly, residual temperature and salinity correlations (Fig. 3.4c) show more structure in their sign and magnitude than their respective correlations with DIC (Fig. 3.4a,b). The models show consistent positive correlations in the Atlantic, equatorial East and Central Pacific, and the Southern Ocean, while the correlations are consistently negative in the equatorial West Pacific and Indian Oceans. Notably, in most regions the correlations between temperature and salinity are weaker than the correlations that temperature and salinity have with carbon.

The pointwise correlations between carbon and temperature and salinity indicate that there is promise in reconstructing carbon variability through the EnOI scheme. Most ocean regions have one relationship with a strong correlation to carbon variability, and the cross-correlations between temperature and salinity are moderate. Thus there is less of a risk of the multiple linear regression solution being overly sensitive as it would be if temperature and salinity were strongly correlated or anticorrelated, as that would reduce the effective rank of the matrix being solved. The differences in structure for the correlation fields with carbon and temperature or salinity indicate that the use of both of these variables in our methodology provides added value, in contrast to if the method were to use only one value to reconstruct carbon.

3.3.2 Reconstructing carbon with pointwise temperature and salinity

The CMIP6 ensemble is used to solve for pointwise temperature and salinity weights w_T, w_S according to the equation

$$DIC'(x,y,t) = w_T(x,y)T'(x,y,t) + w_S(x,y)S'(x,y,t). \tag{3.2}$$

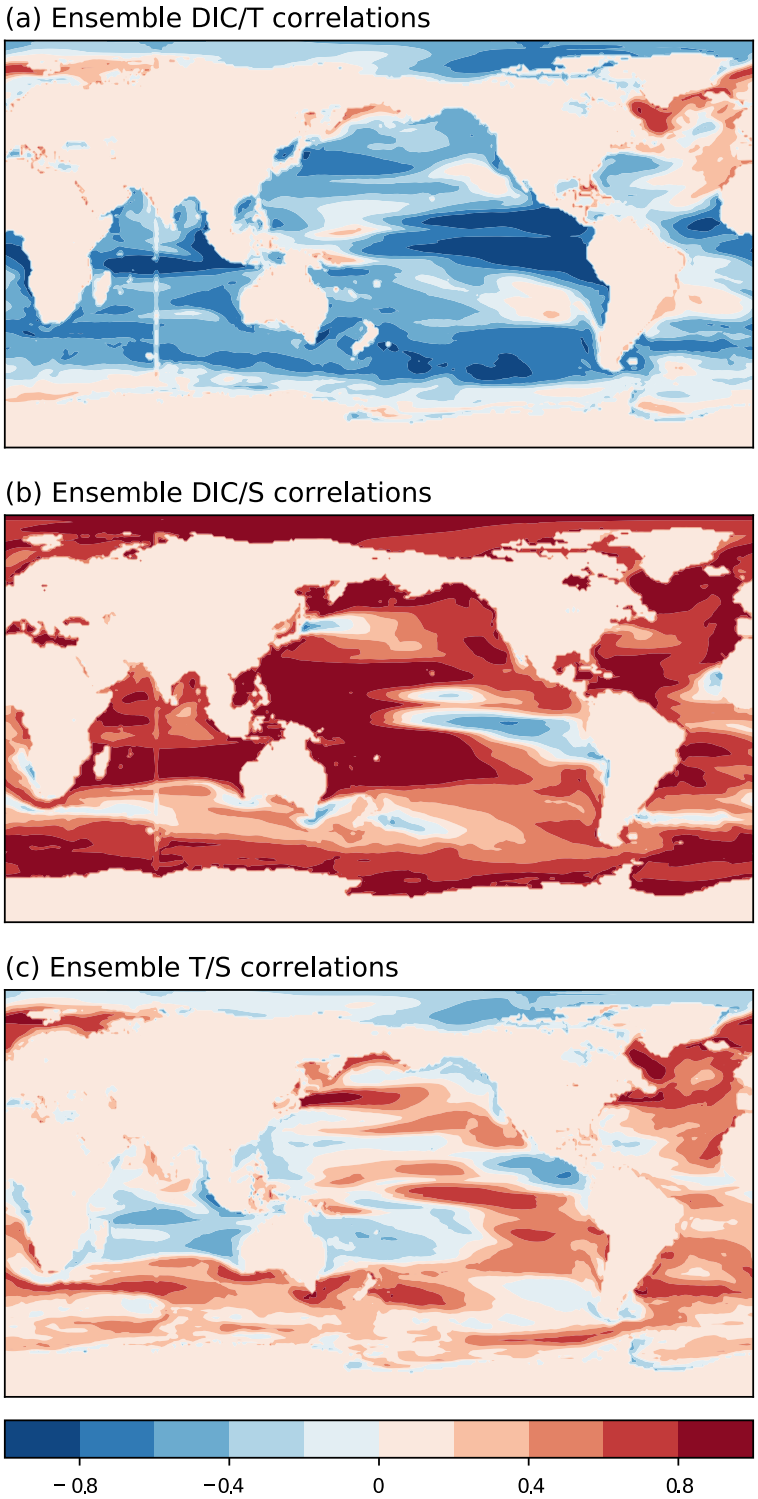

Figure 3.4: Pointwise correlations for upper 100m residuals of a) DIC and temperature, b) DIC and salinity, and c) temperature and salinity taken from the CMIP6 ensemble. Ensemble consists of 6 models and contains 5 historical runs for each model. Contours occur at correlation increments of 0.2.

It is these coefficients that transfer information from temperature and salinity observations into the mapped carbon output. To show the spread in the ensemble members and how the ensemble solution differs from that of the specific models within the ensemble, we calculate the weights for each model within the ensemble and normalise their standard deviation by the ensemble weights. This metric is akin to the coefficient of variation, although the ensemble weights are not the same as the ensemble average weights. Similar to the coefficient of variation, this metric can highlight regions where the EnOI method includes a high level of uncertainty.

The weights for pointwise T and S have structure but have a consistent global sign (Fig. 3.5a,c). Coefficients for temperature are negative in most of the ocean, whereas coefficients for salinity are positive. The lack of spatial noise in the weights emphasises that the rank of the solution is sufficient, as in a rank-deficient system we could expect more noise in the best-fit weights. It is interesting to note that the coefficients retain their sign throughout most of the ocean, even in regions where the pointwise correlation fields are the opposite sign. This result occurs because the cross-correlations between temperature and salinity in these regions are positive (Fig. 3.4c), so in our EnOI solution the impact of temperature in these regions is to decrease the carbon residual, and the impact of salinity in these regions is still to increase the carbon residual. Small regions where the coefficients change sign are visible, mainly in the weights for temperature in the Southern Ocean; however, the high-latitude Southern Ocean has a wide spread within the models used (Fig. 3.5b). The ensemble shows the most variability in the weights where the analogous cross-correlations are moderate, for instance in the Southern Ocean and Eastern subtropical North Atlantic for temperature (Fig. 3.5b) and the Eastern Pacific tropical upwelling zone for salinity (Fig. 3.5d).

3.3.3 Model tests of pointwise carbon reconstructions

We continue by testing these pointwise weights ((3.2), Fig. 3.5a,c) for the carbon reconstruction using NorESM temperature and salinity profiles. The NorESM model provides a useful test case for the constructed CMIP6 ensemble, as the NorESM model uses density depth coordinates rather than z-level coordinates. Using the NorESM model can show if the background error covariances in the ensemble are overly constrained by the similar physical setups of the included models. As in Chapter 2, we compare the root mean squared error (RMSE, (2.2)) for the NorESM first guess, in which carbon content is strictly a function of atmospheric CO_2 levels, to the RMSE generated that from the EnOI carbon reconstruction.

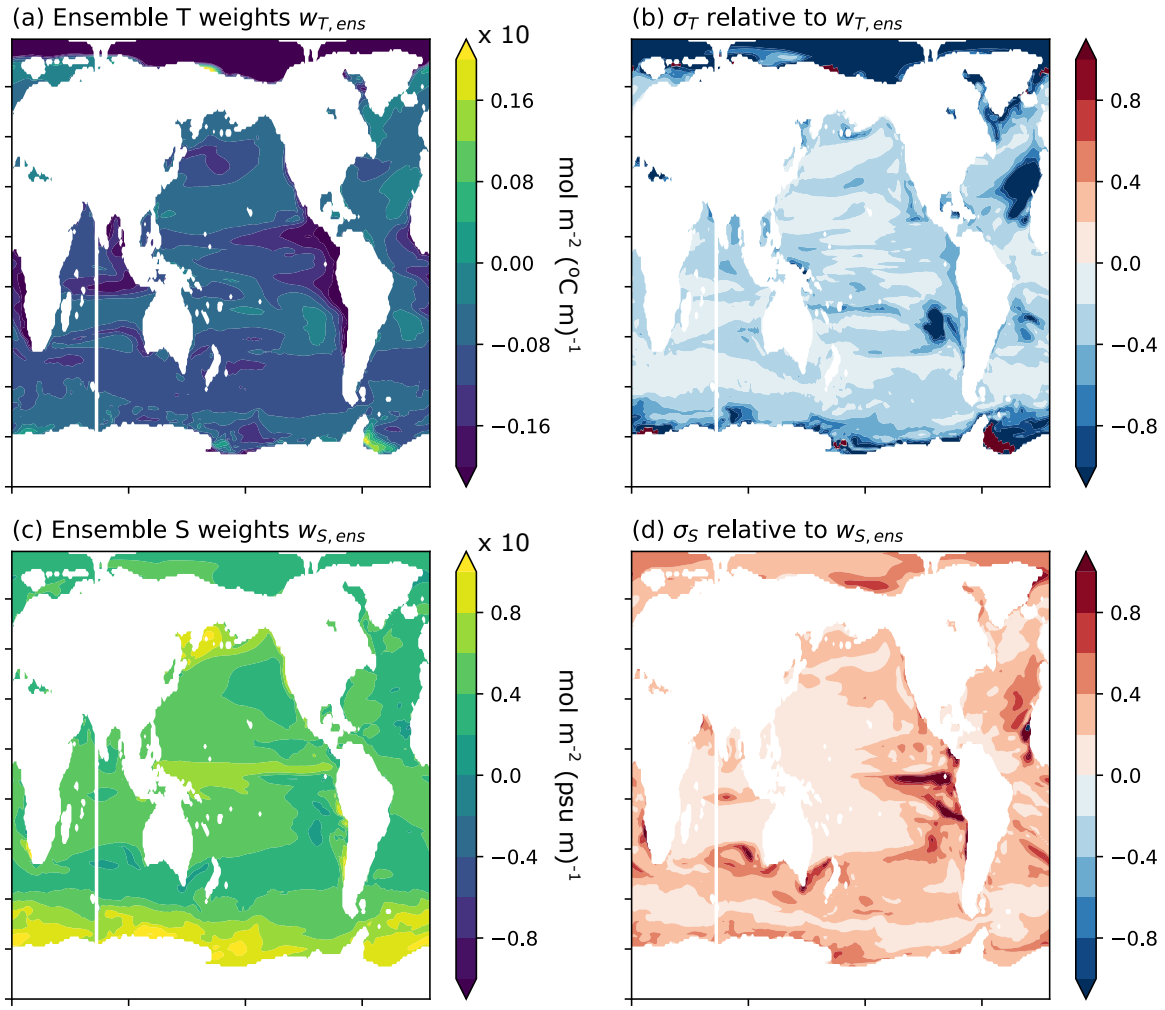

Figure 3.5: CMIP6 ensemble weights for reconstructing 0-100m integrated carbon variability: a) ensemble temperature weights, in units 10 mol m^{-2} (°C m)$^{-1}$ and b) the ensemble spread divided by the ensemble temperature weights; c) ensemble salinity weights, in units 10 mol m^{-2} (psu m)$^{-1}$, and d) the ensemble spread divided by the ensemble salinity weights.

The improvement in the RMSE relative to the first-guess field is equivalent to the amount of carbon content variability that is captured by the reconstruction.

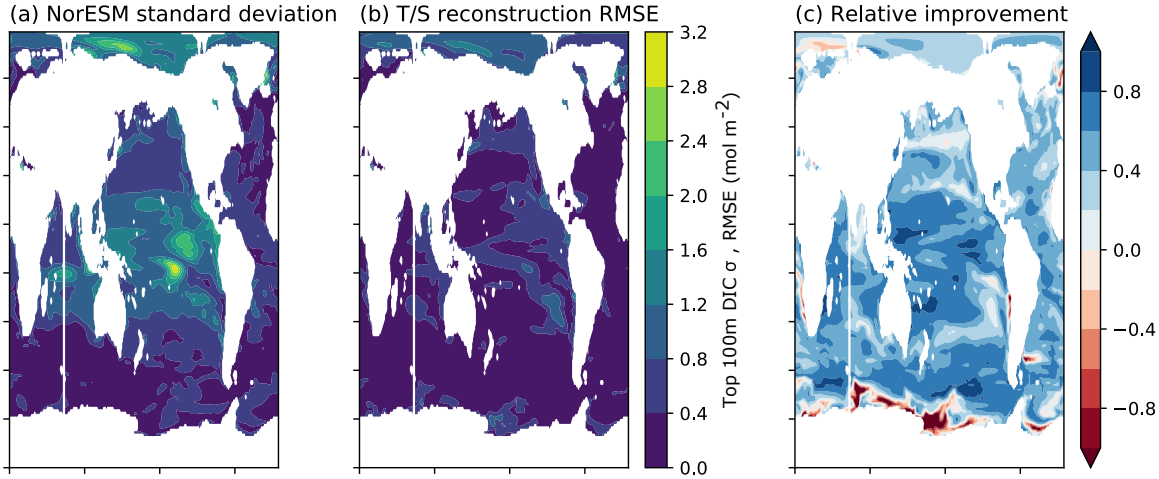

Figure 3.6: Diagnostics for testing the carbon reconstruction using pointwise temperature and salinity with the NorESM model: (a) standard deviation of the NorESM top 100m integrated carbon residual, in units mol m^{-2}, (b) RMSE of the top 100m integrated carbon reconstruction, in units mol m^{-2}, and (c) relative improvement of the reconstruction to the prior guess $(RMSE_{prior} - RMSE_{reconstruct})/RMSE_{prior}$, in which blue areas indicate an improvement, where an improvement of 1 indicates a perfect reconstruction, and red areas indicate that errors are increased under the reconstruction method.

The ensemble weights show substantial skill in reconstructing NorESM carbon content variability in the top 100m (Fig. 3.6). Most regions in the reconstruction experience an RMSE reduction compared to the first-guess field (Fig. 3.6b versus Fig. 3.6a), with a notable exception of the high-latitude Southern Ocean. The RMSE error increase in the high-latitude Southern Ocean is substantial, m ore than doubling the error of the first-guess field in some regions (dark red regions in Fig. 3.6c). However, for most of the ocean the pointwise reconstructions result in an RMSE reduction between 40% and 80% relative to the first-guess field. The skill in this pointwise reconstruction in highest in the tropical and subtropical Pacific Ocean, which is the region with the highest variability in ocean carbon content, and the Antarctic Circumpolar Current region.

In addition, we conduct sensitivity tests to see the impact of each individual model within the ensemble. For these sensitivity tests, we exclude an entire model from the ensemble, re-calculate the ensemble covariance fields, and then reconstruct the excluded model's carbon using its temperature and salinity fields. The resulting RMSE improvements provide insight

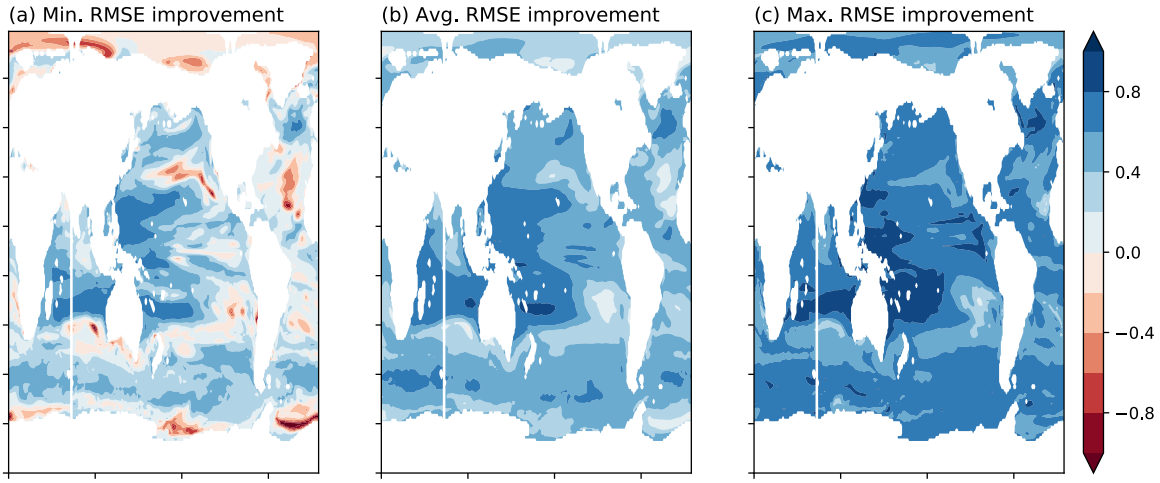

Figure 3.7: Sensitivity tests using the model ensemble reconstructions of 0-100m carbon content using pointwise temperature and salinity. RMSE improvement is relative to a first-guess of carbon content being proportional to atmospheric CO_2 concentrations. a) Pointwise minimum improvement seen in the sensitivity tests, b) Average RMSE improvement seen across the model ensemble, and c) Pointwise maximum RMSE improvement seen in the sensitivity tests.

as to the errors in the ensemble covariance fields as well as to which regions are generally well-suited for our DIC reconstruction method. As our ensemble is an ensemble of opportunity rather than one chosen to span certain parameter spaces, we present the minimum, average, and maximum RMSE improvements at each point to illustrate the potential limits of a pointwise reconstruction.

On average, the pointwise reconstruction of DIC using temperature and salinity residuals reduces the RMSE relative to a first-guess field dictated by atmospheric pCO_2 (Fig. 3.7b). For most regions the analysis captures between 40-80% of the variability not captured by the first-guess field. The largest average improvements are found in the western tropical and subtropical Pacific and the Indian Ocean, which are regions with the strongest correlations between DIC and salinity (Fig. 3.4b). The map of minimum RMSE impacts still show large areas where in the worst-case sensitivity test, RMSE improvements are still possible (Fig. 3.7a). It should be noted that the regions in which RMSE increases (red areas in Fig. 3.7a) show this behaviour in one model out of the entire ensemble, indicating that the models do have regions where the relationships between temperature, salinity, and carbon are relatively uncertain or nonlinear and further emphasising the need to include a variety of models in the ensemble. Lastly, the maximum RMSE improvement is high; the maximum improvement

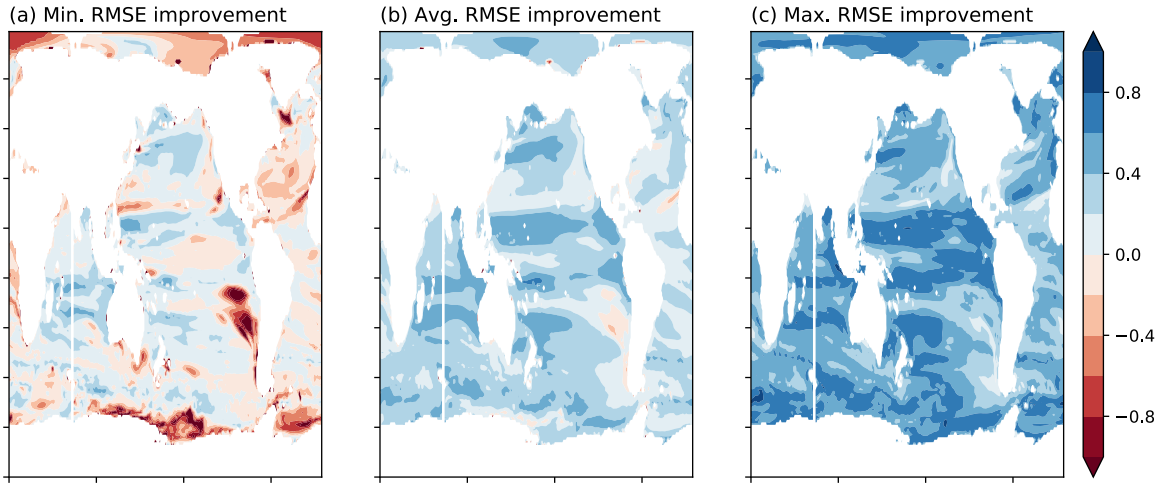

Figure 3.8: Sensitivity tests using the model ensemble reconstructions of 0-500m carbon content using pointwise temperature and salinity. RMSE improvement is relative to a first-guess of carbon content being proportional to atmospheric CO_2 concentrations. a) Pointwise minimum improvement seen in the sensitivity tests, b) Average RMSE improvement seen across the model ensemble, and c) Pointwise maximum RMSE improvement seen in the sensitivity tests.

of the carbon reconstruction method eliminates over 40% of the RMSE in most areas (Fig. 3.7c). This may mean that our ensemble insufficiently covers model uncertainty and that there is some artificial overlap between models, but may also simply be an indication that there is a substantial amount of skill in using only temperature and salinity to reconstruct carbon variability.

The ability of the EnOI scheme to reconstruct ocean carbon decreases as the depth level for integration increases. When expanding the sensitivity tests to reconstruct the upper 500m carbon content variability, the average RMSE error reduction is more modest everywhere, and reconstructs between 20%-60% of the residual variability in most regions (Fig. 3.8b). The minimum RMSE improvements show that almost every region experiences an increase in errors in at least one of the sensitivity tests (Fig. 3.8a). The regions where the upper 500m reconstruction works best is broadly similar to the results for the top 100m reconstructions: the RMSE improvements are greatest in the tropical Pacific, although this maximum appears much patchier (Fig. 3.8c). As the pointwise reconstructions imply that skill is greatest when considering the near-surface carbon content, we continue by considering only the top 100m for our carbon reconstructions when expanding our experiments to consider the Argo profile distribution.

3.4 Impacts of Argo distribution

Carbon reconstructions using same-location temperature and salinity show that the relationships between these variables are suitable for reconstructing a substantial amount of carbon content variability. Constructing a global product from these pointwise relationships requires global coverage of temperature and salinity; while ocean temperature and salinity measurements are much more common than carbon measurements, global coverage of these variables has yet to be achieved. Thus, any reconstruction of carbon variability must consider how irregular and incomplete sampling impacts the accuracy of the reconstruction.

The main goal of optimal interpolation schemes is to extend information from data-rich regions into the rest of the domain. The interpolated values are often a function of multiple observations. If the background error covariance field are correct, the inclusion of multiple observations will improve the reconstruction as each observation is able to contribute more information to the data-sparse region. However, as covariance fields always include some errors, extending the method to include more observational inputs may increase the overall error of the reconstruction.

In this section we continue by considering a primitive localisation procedure to compare reconstructions using Argo-type temperature and salinity sampling to the pointwise reconstructions in the previous section. We consider the distribution of Argo measurements for the year 2015 (Fig.3.2) and reconstruct global carbon inventories by using observations within search radii of $\pm 1°, \pm 2°$, and $\pm 5°$ from the reconstructed locations. The search radius of $\pm 5°$ is able to cover most of the global ocean except for the Arctic Ocean and parts of the Southern Ocean near the Antarctic shelf (Fig. 3.3, lightest blue). As the pointwise reconstructions for the top 100m integrated carbon show more robust improvements in the sensitivity tests than reconstructions of top 500m integrated carbon, this section focuses on the shallower 100m depth horizon for the reconstructions.

Using the full CMIP6 ensemble and NorESM temperature and salinity observations at Argo locations, reconstructions of NorESM carbon using all of the search radii show near-global reductions in the RMSE relative to the first guess field (Fig. 3.9). Similar to the pointwise reconstructions, this improvement is greatest in the western tropical and subtropical Pacific, while the ensemble background covariances are not able to reconstruct the carbon content variability in some regions in the Southern Ocean and some coastal areas.

As the search radius for the reconstructions increases, the regions in which the reconstruction increases errors rather than decreases them become larger (Fig. 3.9e-g). However

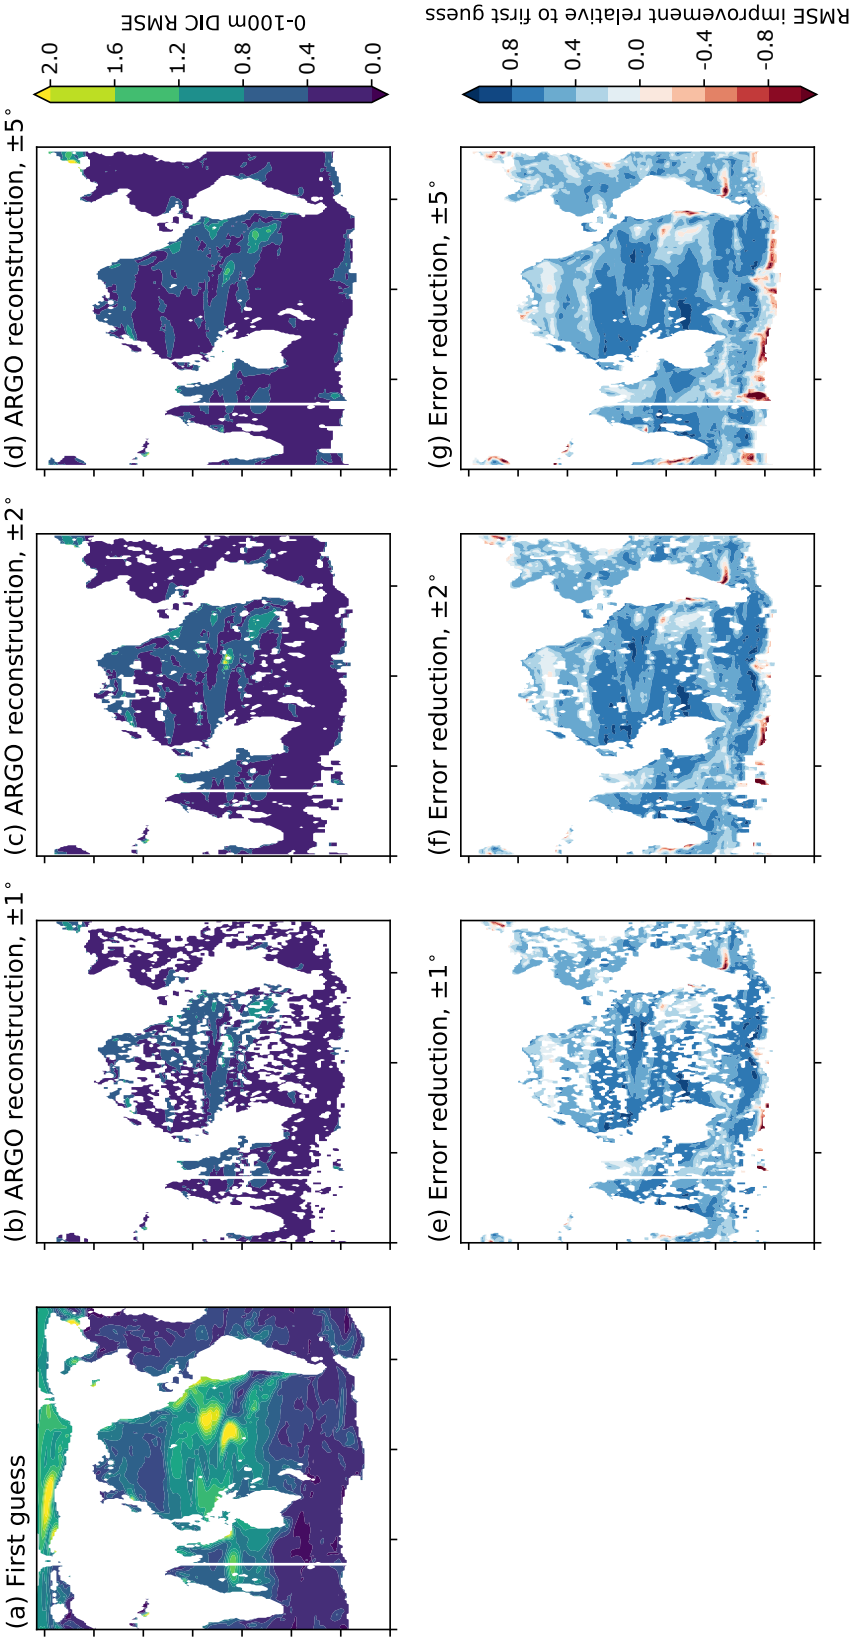

Figure 3.9: NorESM carbon reconstruction tests using synthetic Argo profiles of temperature and salinity. (a-d) Root mean squared errors for reconstructions a) assuming a linear response to atmospheric CO_2 (the first-guess field), and using Argo profile information within a radius of b)1°, c)2°, and d) 5°. (e-g) RMSE improvement relative to the first-guess RMSE for reconstructions using search radii of e)1°, f)2°, and g) 5°.

these regions are mostly confined to areas in which the pointwise reconstructions show increased errors (Fig. 3.6c), suggesting that the enhanced errors originate from regions in which the temperature and salinity relationships are not well-replicated within the ensemble. For most of the domain there is little change in the RMSE reduction when the reconstruction method uses observations within the different search radii. In regions where the reconstructions improve upon the first guess, the RMSE differences with the search radii are mainly noisy and show little consistent regional behaviour. Overall the NorESM experiment shows that the benefits of using multiple observations in the vicinity of a reconstruction point to reconstruct larger areas of the global ocean generally outweigh the disavantages of propagating errors in the covariance fields. Unlike the sparse carbon time series measurements, the density of Argo profiles allows for a higher-quality reconstruction, even when considering the irregular distribution of profile and the weaker correlations between carbon and temperature or salinity.

Testing each of the models within the ensemble shows a clear discrepancy between the MPI model and the other models within the ensemble (Fig. 3.10). For most models, the reconstruction using a 5° Argo search radius captures a substantial part of the carbon variability. Reconstruction error increases when using the 5° Argo search radius are restricted to areas in which the sensitivity tests using pointwise covariance fields also showed enhanced errors (not shown). However, for the MPI model, using an expanding search radius for the reconstruction leads to large regions of increased errors in the Southern Ocean, North Atlantic Ocean, and eastern North Pacific Ocean. These errors increase dramatically as the search radius increases in both magnitude and regional extent (not shown). While this result could indicate that the MPI model is an outlier that should be eliminated from the ensemble, the exclusion of the MPI model from the ensemble leads to larger errors in the NorESM reconstructions in the Southern Ocean. Thus including the MPI model in the ensemble provides additional skill for this region.

The MPI model is the highest-resolution model in the CMIP6 ensemble used, with a nominal ocean model horizontal resolution of 0.4°; in comparison, the other models within the ensemble have nominal horizontal resolutions of 1°. Snapshots of residual carbon DIC' show more small-scale structure in the MPI model than in the other models within the ensemble (Fig. 3.11). The increased errors in the North Atlantic, eastern North Pacific, and throughout most of the Southern Ocean may be a result of the different physical representations inherent in a higher-resolution model, as these are the regions that also show the most small-scale variability in the DIC' snapshot. Although the MPI model does not resolve the mid-latitude Rossby radius, its eddy parameterisations will differ from the

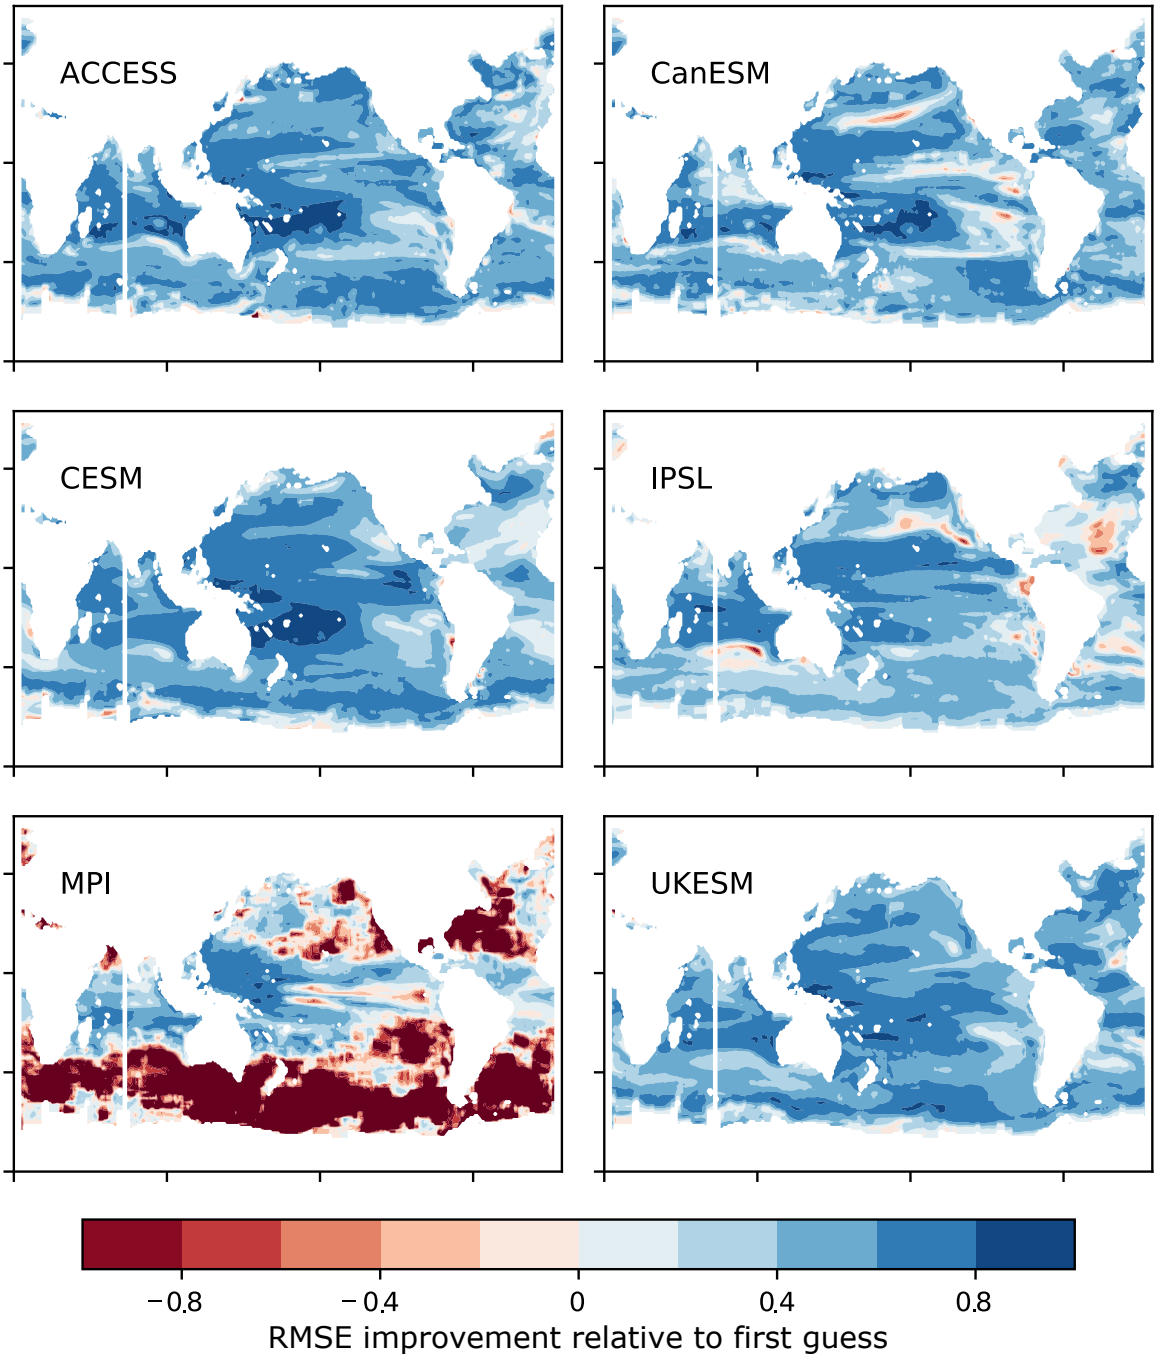

Figure 3.10: Sensitivity tests for reconstructing 0-100m carbon content variability using Argo profiles within a 5° search radius. Error improvement is relative to the RMSE when using the first-guess that carbon is proportional to atmospheric CO_2 concentrations. Blue areas indicate regions where the Argo-style reconstruction improves upon the first guess (up to 1 = 100% of variability is accounted for in the reconstruction), whereas red areas indicate regions where the Argo-style reconstruction is worse than the first guess.

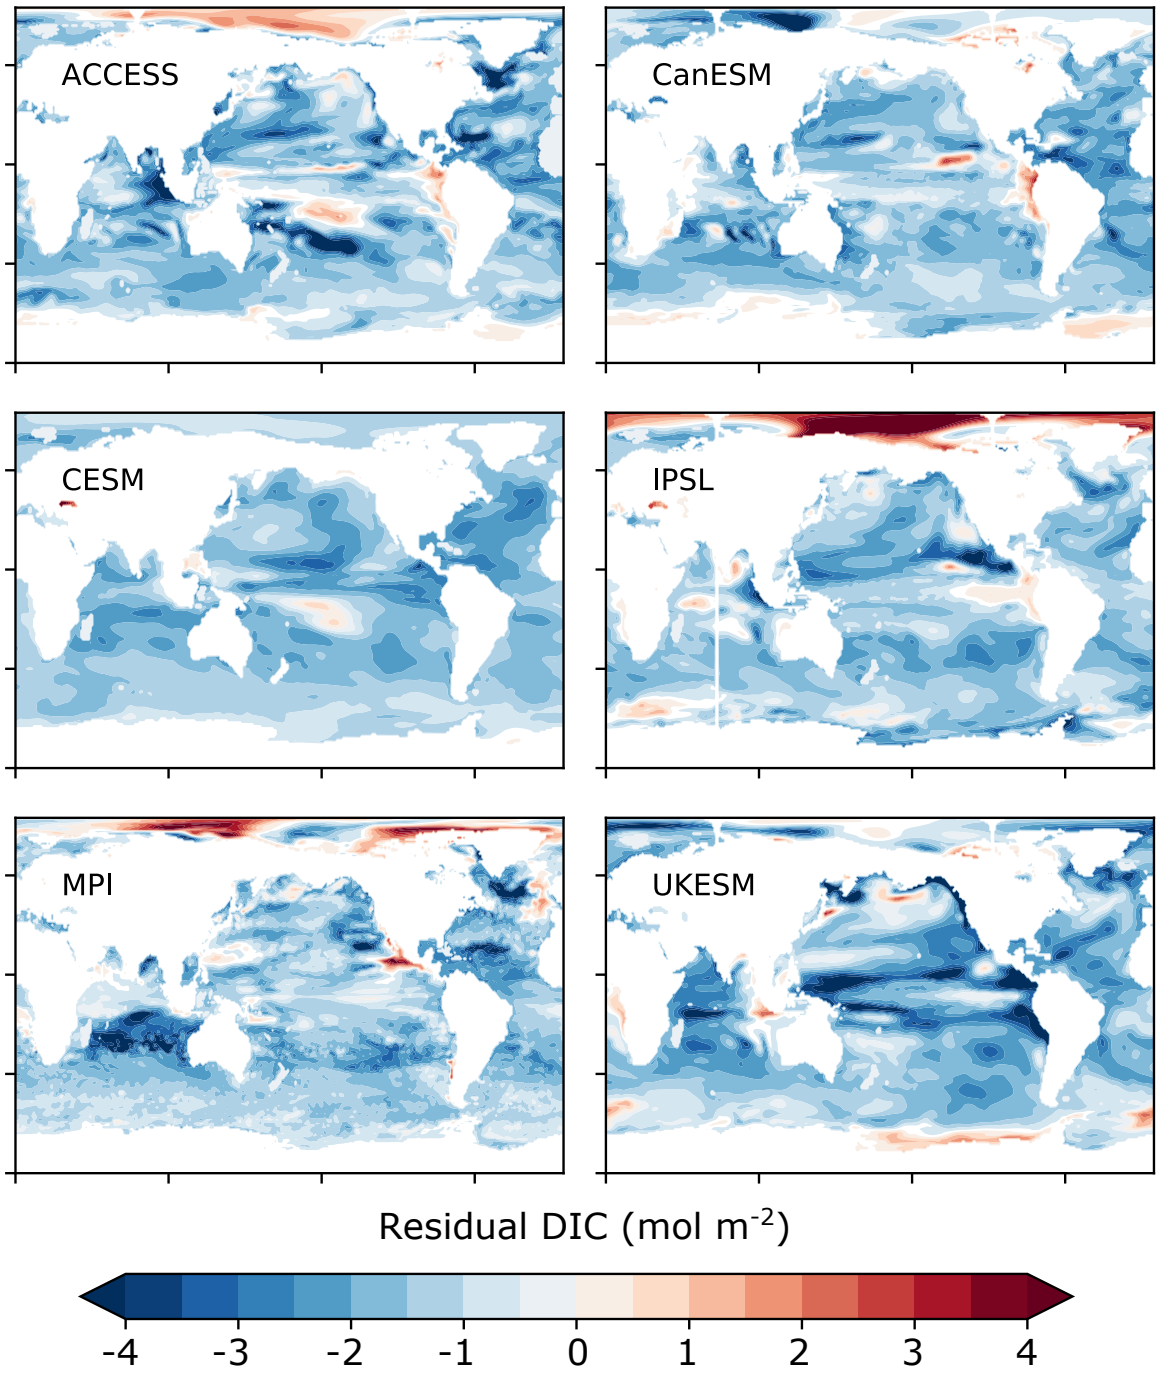

Figure 3.11: Snapshots of integrated 0-100m residual carbon DIC' for one realisation of each model, taken at year 1950. Note that the patterns of DIC' should not line up as the atmospheric components of each model are not set to match up with reanalysis.

parameterisations used in the coarser models. The regridding method used to put all the model outputs on a standard grid may also play a role, as there is the potential for smaller-scale structure in the MPI annual averages to be aliased through the relatively simple bilinear regridding. The 5° search radius sensitivity results suggest that the makeup of the CMIP6 ensemble should be evaluated carefully before the method is applied to Argo profile data, as the assumptions within the model ensemble can impact the amount of skill available when reconstructing ocean carbon using nearby temperature and salinity.

3.5 Discussion and Conclusions

The paucity of ocean carbon observations presents major complications when trying to understand historical ocean carbon variability. Other ocean variables can be used to reconstruct ocean carbon variability if the relationships between variables can be ascertained. Ocean temperature and salinity have well-understood relationships with carbon solubility and have imprints on ocean circulation through the equation of state. The density of temperature and salinity observations in both space and time have the potential to overcome the complications seen in expanding carbon observations from sparse datasets. In this chapter, we conducted synthetic model tests using EnOI and a set of high-complexity CMIP6 coupled models to explore how carbon variability may be understood through temperature and salinity variability. We have conducted tests using local relationships with temperature and salinity and using covariance fields that account for the irregular distribution of the current-day Argo profiles.

3.5.1 Reconstructing ocean carbon using temperature and salinity correlations

EnOI uses covariances to translate information from one variable to another. The ensemble shows large regions of strong correlations between ocean carbon and ocean temperature or salinity (Fig. 3.4), which leads to well-constrained weights for solving for ocean carbon as a linear system with temperature and salinity (Fig. 3.5). Sensitivity tests indicate that there is potential to reconstruct ocean carbon variability using ocean temperature and salinity measurements for large areas of the ocean (Fig. 3.6, 3.7). This skill is highest when considering near-surface carbon content, where the sensitivity tests show an average 50% reduction in RMSE relative to the first-guess field. Smith and Murphy (2007) estimate that EnOI coupled with the full Argo array has the potential to capture around 70% of the variability in temperature and salinity. Thus, the estimate of 50%, while smaller than that

for temperature and salinity alone, is remarkable as it relies on the relationships between variables as well as spatial extrapolation and does not explicitly include any contributions from the biological carbon pump. Regions which show the highest error reductions from the reconstructions, the western tropical Pacific and subtropical Indian Ocean, are those which have strong positive correlations with salinity and weaker correlations with temperature.

Areas which show less potential in reconstructing ocean carbon variability using temperature and salinity, such as the eastern subtropical North Atlantic, the Arctic Ocean, and Antarctic shelf area, are marked by larger spreads the carbon-temperature covariance fields, as seen in the temperature weight standard deviation (Fig. 3.5b). This uncertainty ultimately gets translated into weak ensemble correlations between temperature and carbon (Fig. 3.4a); as a result, the reconstructions in these areas are dampened to the first-guess field. The uncertainty seen in the model ensemble may be a result of uncertainties in the models' dynamics. Diagnostics on the strength of carbon concentration and carbon-climate feedbacks show large ranges in both the Arctic Ocean and Antarctic shelf areas (Katavouta and Williams, 2021). Additionally, uncertainty in the covariance fields and temperature weighting coefficients could arise because these areas have other controls on carbon variability not considered by our method. For instance, biological impacts on upper-ocean carbon variability may play a larger role in these regions, and so the modest RMSE improvement may indicate that our reconstruction setup with temperature and salinity could be improved by including other oceanographic variables.

Expanding the EnOI method to include other oceanographic variables

While reconstructions of carbon variability through temperature and salinity variability show promise for much of the global ocean, the RMSE reductions in our synthetic experiments are relatively modest in some regions. The flexibility of the EnOI method allows for the inclusion of other oceanographic variables. Similar to the expansion to use multiple observations to reconstruct carbon in Section 3.4, the inclusion of more variables should decrease the RMSE if the covariance fields do not contain errors.

While biology plays a small role in setting carbon variability globally, it may still play a role regionally. Biology impacts water column DIC through the *regenerated pool*: as organisms die and sink, their organic matter is respired by bacteria into inorganic carbon. The role of biology could potentially be added into the EnOI scheme by including oxygen measurements, which are related to the ventilation pathways and age of a water parcel and thus the amount of regenerated DIC it holds, or chlorophyll measurements, which are

related to the current amount of biological activity in the upper water column.

Biogeochemical sensors are currently being added to Argo floats within the Biogeochemical-Argo programme (BGC-Argo, Claustre et al. (2020)). The BGC-Argo programme aims to equip around a quarter of all Argo floats with sensors for oxygen, nitrate, pH, chlorophyll, and suspended particles and irradiance. Synthetic reconstructions of ocean $pCO_{2,oce}$ using BGC-Argo locations have found substantial added skill when the observed variables are added to a machine learning algorithm (Denvil-Sommer et al., 2021). There may be added value in including these variables to our carbon inventory reconstructions, particularly as our method works best for near-surface ocean carbon content. However, the number of BGC-Argo floats required to improve carbon reconstructions may be more than what is currently available, and the current skill provided by the BGC-Argo measurements may be limited by a similar sparseness problem found with the ocean time series sites (Section 2.4).

One major issue with including biological variables in the CMIP EnOI method is how model architecture impacts the covariance fields. Unlike ocean dynamics or the carbonate chemistry system, which have well-known analytical formulae, most of the biological activity within climate models is heavily parameterised with empirical formulae. The use of an ensemble to avoid overfitting may not be as effective with biological parameters, as the models may have quite similar parameterisations, which could lead to spuriously well-constrained model responses. In the future, if more biogeochemical variables are to be included in our method, sensitivity tests will need to be conducted to ensure that any additional variables add skill to the system rather than overfit the solutions.

3.5.2 Ability to reconstruct global carbon variability using Argo-like inputs

The Argo programme has resulted in near-global coverage of upper ocean temperature and salinity over the 21st Century. While the number of Argo profiles available can increase the reconstruction skill through the amount of information they provide, it is possible that the EnOI method used with these profiles might increase the error in some regions if the background covariance fields contain errors. Preliminary results using sampling radii of up to 5° from reconstruction locations indicate that there is still skill in reconstructing ocean carbon variability from Argo-style measurements. Relative to the first-guess field, all reconstructions reduce the errors in most ocean regions, although for regions in which the pointwise covariance fields contain errors, the errors tend to propagate as the sampling becomes farther removed from the reconstruction location. For most of the ocean and

most sensitivity setups, though, the error reductions are insensitive to the search radius for contributing Argo profiles. These results indicate that the sheer number of Argo profiles, coupled with the relationships between temperature, salinity, and carbon variability, provide a substantial basis on which carbon variability may be reconstructed from non-carbon ocean observations.

Sensitivity tests, in which one model is removed from the ensemble and then reconstructed using its temperature and salinity fields, show divergent behaviour between the MPI model and the other models in the ensemble. Unlike the other models, MPI carbon reconstructions using temperature and salinity profiles within a search radius result in large errors in the North Atlantic, North Pacific, and Southern Oceans. These errors indicate that the covariance fields formed by the ACCESS/CanESM/CESM/IPSL/UKESM models are insufficient for representing the variability in the MPI model. We speculate that there could be multiple reasons for these errors to appear in the sensitivity tests:

- The MPI model, having a higher resolution (0.4° nominal horizontal resolution versus 1° nominal horizontal resolution for the other ensemble members), is able to represent processes that are poorly parameterised or excluded from the other ocean models. If this is the case, the inclusion of the MPI model in the ensemble may be important for interpreting small-scale variability found in Argo profiles.

- The bilinear approach taken to regrid all of the models to a consistent $1^\circ \times 1^\circ$ rectilinear grid has allowed for aliasing of smaller-scale features from the MPI model, resulting in MPI covariance fields that will not be able to reflect the binning and averaging of Argo profiles within the area's bounds. If this is the case, the regridding procedure for the MPI model can be altered through smoothing, conservative mapping, or a weighted subsampling procedure to reduce the aliasing of smaller-scale features and produce model covariances more akin to those in the other models.

If neither of these methods work, the output from high-resolution MPI runs can be replaced with output from low-resolution MPI runs. However, as the inclusion of the MPI model within the ensemble does not increase the reconstruction errors for the other models in the sensitivity tests, and including the MPI model reduces the errors for the NorESM construction, we speculate that the high-resolution MPI model is able to add skill to the reconstruction method overall. Work will continue by checking the regridding method for the MPI model to ensure that the regridded product is tracer-conserving and comperable to coarser model outputs.

Nevertheless, these results emphasise the importance of careful selection of ensemble members. The models within this ensemble are only a small subset of the available models and were chosen as they had multiple historical realisations published at the time of data collection. Future work in this area will include more thorough testing of the ensemble to ensure that the background error covariance fields are well-suited for real-world carbon reconstructions.

Expanding tests to include real Argo observations

The synthetic carbon reconstructions in this chapter indicate that Argo profiles can be used to reconstruct a substantial amount of variability in historical carbon content. The obvious next step for this work is to expand the method to include real observations from the Argo programme. The ocean time series sites are ideal testing grounds for our method, as they provide concurrent measurements of ocean carbon, temperature, and salinity. The high sampling frequency of variables at the time series sites will allow us to test the assumptions made when calculating the carbon residual calculations, creating binning procedures, and choosing search radii with which to construct the covariance fields.

Expanding the method in this chapter to account for real profiles involves considering both model errors and observational errors. While the Argo profiles are highly accurate (Wong et al., 2020), there are still important questions about how the method translates from model world to the real world. Further work will require careful consideration of the following:

- **The pCO_2 regression first-guess:** The first guess field used in this chapter is a term proportional to atmospheric CO_2 increases. Calculating this proportion from data may be more difficult for real-world observations as the observations are irregularly spaced in time and there will be fewer years of data to work with than the 60 years of data taken from each model observation. The consistent measurements of atmospheric CO_2, coupled with its well-mixed behaviour on annual and longer timescales, should reduce the errors inherent in taking away this first guess field.

- **Annual average Argo profiles:** The method used in this chapter relies on CMIP6 annual average output to create the covariance matrices. Calculating binned annual averages with Argo floats ideally requires multiple observations for each month. In the inevitable case that bins do not have many observations within the year, or that these observations are skewed towards a specific time of year, an additional error term will need to be considered. While we have attempted to account for some of the

irregularities in Argo profiling through by requiring bins to have at least 6 months of observations, any scheme using real measurements should aim to maximise the number of observations it can allow. The bin sizes can also be adjusted, although there will be a trade-off between spatial resolution and number of Argo profiles per bin. We speculate that the representation problems inherent in creating annual average fields from point profiles will incur the largest new errors in this method.

- **Observation error covariances:** While in this chapter we have operated under the assumption of perfect model data, real observations will always have errors. These errors may be instrument-based or may reflect how the behaviour seen in the Argo profiles differs from the general environment (e.g., profiles taken within an eddy will have different temperatures and salinities to the surrounding environment). The Argo sensors are able to measure ocean temperatures accurate to ± 0.002°C and salinity accurate to ± 0.01 psu (Wong et al., 2020). EnOI schemes generally assume that the observations are independent, allowing for a diagonal observational error covariance matrix. The Lagrangian nature of Argo floats may not permit this assumption, although the binning and averaging procedures may reduce the impact of non-independent measurements on the system. The observation error covariance fields will need to be tested with various assumptions made about how binning procedures impact assumptions of independence.

In conclusion, the Ensemble Optimal Interpolation method shows promise in using CMIP6 covariance fields to reconstruct carbon variability from temperature and salinity variability. In synthetic model tests, our reconstruction method shows the potential to resolve around 50% of upper ocean carbon variability when long-term trends in ocean carbon related to atmospheric pCO_2 are removed. The success of the method arises from strong and physically-consistent correlations between temperature, salinity, and carbon in the model ensemble. The availability of temperature and salinity profiles through the Argo programme allows for reconstruction of most of the global upper-ocean carbon content. Expanding the EnOI method to include measurements close to the reconstruction location with errors similar to reconstructions shows similar potential to pointwise reconstructions, although sensitivity tests reveal that there needs to be further consideration on how different model complexities are accounted for in the methodology. The addition of other oceanographic variables, such as those observed in the BGC-Argo programme, and a more thorough examination of the makeup of the CMIP6 ensemble may further improve the potential to reconstruct carbon variability; nevertheless, our methodology shows promise in

using state-of-the-art climate model output and relationships between thermohaline and carbon variability to create a new, independent reconstruction of historical ocean carbon content.

Chapter 4

Hemispheric contrasts in the relationship between surface warming and cumulative carbon emissions

We continue this thesis by exploring how 21st Century emission projects are controlled by hemisphere thermal and carbon responses. This chapter analyses the hemispheric relationship between cumulative emissions and surface warming using the University of Victoria Earth System Climate Model (UVic ESCM). The UVic ESCM has been widely used to study the relationship between cumulative emissions and surface warming, but have generally focused on global average behaviours.

We begin in Section 4.1 by providing an overview to the transient climate response to emissions and its underlying theory. This theory is expanded to describe how the thermal and carbon controls operate for each hemisphere in Section 4.2, and we show how the global and hemispheric terms compare in the UVic ESCM in Sections 4.3 and 4.4. We then connect these diagnostics with transports and patterns associated with the heat budget (Section 4.5) and carbon budget (Section 4.6). We find that, while the Southern Hemisphere plays a larger role in heat and carbon uptake, the Northern Hemisphere plays a significant role in setting the thermal response, as the climate system tries to compensate for highly localised negative radiative forcing from sulfate aerosols. Finally, a discussion and reflection upon the results from the UVic ESCM is provided in Section 4.7.

4.1 The transient climate response to emissions

Model studies on the climate response to emissions have found that surface warming is approximately proportional to cumulative carbon dioxide emissions and is insensitive to the rate of emissions (Matthews et al., 2009; Allen et al., 2009). This proportionality has been found in Earth system models of varying complexities (MacDougall, 2017; Eby et al., 2013; Gillett et al., 2013) and forms the basis for the climate metric, the *Transient Climate Response to Emissions* (TCRE). The TCRE is defined as the ratio of changes in global-mean surface air temperature ΔT, in K, to cumulative carbon dioxide emissions ΔI_{em}, in 1000 PgC, where all changes are relative to the preindustrial (taken here to be year 1860):

$$\text{TCRE} = \frac{\Delta T}{\Delta I_{em}} = \frac{T(t) - T(t = 1860)}{I_{em}(t) - I_{em}(t = 1860)}. \tag{4.1}$$

This proportionality has led to a shift in climate policy from the stabilisation of atmospheric CO_2 to the creation of carbon budgets, and has been used to estimate emission ranges consistent with maintaining average surface warming below 1.5 °C or 2 °C (Meinshausen et al., 2009; Zickfeld et al., 2009; Millar et al., 2017; Goodwin et al., 2018; Mengis et al., 2018).

The near-constant proportionality between surface warming and cumulative emissions has been linked to a compensation between the planetary uptake of anthropogenic heat and carbon (Goodwin et al., 2014; Ehlert et al., 2017). On time scales up to many centuries, the impact of ocean ventilation on heat and carbon uptake nearly compensate in terms of the effects of the TCRE (Solomon et al., 2009; Goodwin et al., 2014; Williams et al., 2016). Both diffusive models (MacDougall, 2017) and advective models of the ocean (Katavouta et al., 2019) show this near-compensation in ocean heat and carbon uptake, although including carbonate chemistry can reduce this level of compensation (Katavouta et al., 2018).

4.1.1 Theory connecting carbon emissions to surface warming

For CO_2-only emissions, the TCRE was originally expanded by Matthews et al. (2009) through multiplicative identities with changes in the atmospheric carbon concentration ΔC_A, so that the proportionality between was set by the atmospheric fraction $\Delta C_A / \Delta I_{em}$ and the warming response to an increase in atmospheric CO_2 content:

$$\text{TCRE} = \frac{\Delta T}{\Delta I_{em}} = \frac{\Delta T}{\Delta C_A} \frac{\Delta C_A}{\Delta I_{em}}. \tag{4.2}$$

This separation considers how emissions are translated to warming through the original partition of carbon among reservoirs, and the term $\Delta T / \Delta C_A$ combines the thermal response

to radiative forcing with the relationship between radiative forcing and atmospheric carbon changes, which is nonlinear and dependent on the original carbon concentration. To more clearly show how the TCRE relates to the Earth system response to radiative forcing, Williams et al. (2016) expand the TCRE using global average radiative forcing ΔR, in W m^{-2},

$$\text{TCRE} = \frac{\Delta T}{\Delta I_{em}} = \frac{\Delta T}{\Delta R}\frac{\Delta R}{\Delta I_{em}}, \tag{4.3}$$

thus relating TCRE behaviour to a compensation between $\Delta T/\Delta R$, the thermal response, and $\Delta R/\Delta I_{em}$, the combined translation of emissions to atmospheric radiative forcing through ocean and terrestrial carbon uptake and the logarithmic radiative forcing response of atmospheric CO_2.

The TCRE can be generalised to include the impacts of non-CO_2 radiative forcing elements, which is defined as the *effective TCRE* (TCRE_{eff}). Further studies have subsequently expanded the effective TCRE to account for changes in globally-averaged CO_2-induced radiative forcing ΔR_{CO_2}, in Wm^{-2}, and the global atmospheric carbon inventory ΔI_{atm}, in 1000 PgC (Williams et al., 2016; Ehlert et al., 2017; Williams et al., 2017b; Katavouta et al., 2018; Williams et al., 2020):

$$\text{TCRE}_{eff} = \frac{\Delta T}{\Delta I_{em}} = \left(\frac{\Delta T}{\Delta R}\right)\left(\frac{\Delta R}{\Delta R_{CO_2}}\right)\left(\frac{\Delta R_{CO_2}}{\Delta I_{atm}}\right)\left(\frac{\Delta I_{atm}}{\Delta I_{em}}\right). \tag{4.4}$$

The effective TCRE is then the product of four dependencies:

1. the *thermal response*, $\Delta T/\Delta R$, which describes the temperature response ΔT to a radiative forcing ΔR and is related to the transient climate response (TCR),

2. a radiative response involving the relative importance of CO_2 in the total radiative forcing, $\Delta R/\Delta R_{CO_2}$,

3. a radiative response involving changes in radiative forcing caused by changes in atmospheric CO_2 concentrations, $\Delta R_{CO_2}/\Delta I_{atm}$, and

4. the *carbon response*, $\Delta I_{atm}/\Delta I_{em}$, which describes the evolution of the *airborne fraction*, the proportion of emitted carbon is retained in the atmosphere.

Empirical heat and carbon balances

The thermal response may be described by an empirical heat budget, in which radiative forcing ΔR drives a radiative response, written as a product of ΔT and a climate feedback parameter λ, in $\text{Wm}^{-2}\text{K}^{-1}$, and a net top-of-the-atmosphere heat flux $\mathcal{N}$, in Wm^{-2}, which

is equal to the planetary heat uptake and dominated by ocean heat uptake (Gregory et al., 2002, 2015):

$$\Delta R = \lambda \Delta T + \mathcal{N}. \tag{4.5}$$

The global climate feedback parameter λ is classically assumed to be a constant but has been found to vary with time (Armour et al., 2013; Ceppi and Gregory, 2017). An equivalent method to account for time variation in the climate feedback response to forcing is through including a time-dependent efficacy term $\varepsilon(t)$ to $\mathcal{N}$ to consider the impact of ocean heat uptake altering the radiative budget (Winton et al., 2010, 2013). For simplicity, we allow λ to vary with time, which eliminates a need for a Gregory-type regression to calculate a constant λ and time-varying $\varepsilon(t)$.

Rearranging (4.5) to fit the form of (4.4) yields the surface warming dependence on radiative forcing expressed in terms of λ, $\mathcal{N}$, and ΔR:

$$\frac{\Delta T}{\Delta R} = \frac{1}{\lambda}\left(1 - \frac{\mathcal{N}}{\Delta R}\right). \tag{4.6}$$

The carbon response is determined by the partition of emitted CO_2 among the terrestrial, oceanic, and atmospheric reservoirs:

$$\Delta I_{em} = \Delta I_{terr} + \Delta I_{ocean} + \Delta I_{atm}. \tag{4.7}$$

Rearranging (4.7) to fit the carbon response in (4.4) yields the expression for the airborne fraction

$$\frac{\Delta I_{atm}}{\Delta I_{em}} = 1 - \frac{\Delta I_{terr} + \Delta I_{ocean}}{\Delta I_{em}}. \tag{4.8}$$

The effective TCRE can thus be expanded with these empirical balances to reveal an explicit dependence on climate feedbacks, planetary heat uptake, and oceanic and terrestrial carbon uptake:

$$\mathrm{TCRE}_{eff} = \frac{\Delta T}{\Delta I_{em}} = \frac{1}{\lambda}\left(1 - \frac{\mathcal{N}}{\Delta R}\right)\left(\frac{\Delta R}{\Delta R_{CO_2}}\right)\left(\frac{\Delta R_{CO_2}}{\Delta I_{atm}}\right)\left(1 - \frac{\Delta I_{terr} + \Delta I_{ocean}}{\Delta I_{em}}\right). \tag{4.9}$$

Using (4.9), the compensation of ocean heat uptake and carbon uptake arises because the ocean's decreasing ability to take up excess heat causes a decline in the term $1 - \mathcal{N}/\Delta R$, that offsets a decrease in the term $\left(\Delta R_{CO_2}/\Delta I_{atm}\right)\left(1-(\Delta I_{terr}+\Delta I_{ocean})/\Delta I_{em}\right)$ caused by the decrease in ocean carbon uptake coupled with the logarithmic nature of CO_2-induced radiative forcing (Goodwin et al., 2014; Williams et al., 2016). A decline in the climate feedback parameter $\lambda(t)$ can also play a role by strengthening the surface warming response, which opposes the decrease in ocean carbon uptake (Williams et al., 2020).

4.1.2 Driving questions and objectives

The TCRE and effective TCRE are global-mean behaviours of the climate system that are connected to global fluxes and inventories (Matthews et al., 2009; MacDougall and Friedlingstein, 2015; Ehlert et al., 2017; Williams et al., 2017c). However, there are substantial regional differences in the climate system that contribute to the global mean. Oceanic uptake and storage of both heat and carbon have strong regional signals, with the Southern Ocean playing a dominant role in their uptake and redistribution (Marshall et al., 2014b; Frölicher et al., 2015). Variations in ocean circulation (Marshall et al., 2014b) and cloud and water vapor feedbacks (Armour et al., 2013; Rugenstein et al., 2016; Ceppi and Gregory, 2017; Andrews and Webb, 2018) may also respond to regional warming patterns and lead to spatial differences in climate sensitivity. Regionality of heat and carbon uptake has the potential to impact the current uncertainties in the TCRE, as climate and carbon feedbacks help set the inter-model spread in the TCRE (Spafford and MacDougall, 2020; Jones and Friedlingstein, 2020; Williams et al., 2020). Thus, the global compensation between the effects of heat and carbon uptake may not hold when considering regional responses to carbon emissions, and the roles that specific regions play in setting this global compensation may be important to understand the wider inter-model spread in the TCRE.

In this chapter we explore the hemispheric responses to cumulative carbon emissions using an intermediate-complexity Earth System model commonly used for TCRE studies. We aim to answer the following questions:

1. How does the response to cumulative carbon emissions behave when considered on a regional scale?

2. Do the thermal and carbon responses compensate on a regional scale? If not, what are their contributions to the global compensation between the thermal and carbon responses?

3. How does the transport of carbon or heat between regions set up the global compensation necessary for proportional TCRE behaviour?

We hypothesise that the level of compensation between heat and carbon uptake for both hemispheres is similar to the compensation globally. Contributions towards the global heat and carbon uptake will be set by the ocean surface area percentage of each hemisphere; thus, the Northern Hemisphere contribution towards the heat and carbon responses will be smaller than that for the Southern Hemisphere. The presence of regional radiative forcing elements such as aerosols or regional feedbacks that amplify warming in the Northern Hemisphere

should impose some hemisphere-specific nonlinearities in the local response to cumulative carbon emissions and enhance any nonlinearities present in the global response.

In this study, the global TCRE theory is expanded to represent the physical climate responses to emissions and localised to focus on hemispheric responses (Section 4.2). Our theory explicitly connects regional warming signals to local heat and carbon balances, which can be used to understand the processes and locations that drive both the global TCRE and regional TCRE (Sections 4.3 and 4.4). We then examine the hemispheric compensations driving the relationship between surface warming and cumulative carbon emissions with an Earth system model of intermediate complexity. The hemispheric thermal responses (Section 4.5) and carbon responses (Section 4.6) are examined together with a discussion on how the relationship between temperature changes and cumulative emissions arises from these regional contributions (Section 4.7).

4.2 Methods

4.2.1 Expansion of relationships to hemispheric behaviours

In order to understand the regional warming behaviour to cumulative carbon emissions, we first expand (4.9) to account for regional changes in both the thermal and the carbon responses. Previous work by Leduc et al. (2016) on the regional TCRE considers spatial variations in temperature changes relative to global carbon emissions; for the hemispheres, this local response is expressed as $\Delta T_{hem}/\Delta I_{em}$. This TCRE localisation emphasises rates and patterns of amplified warming under carbon emissions, although the carbon contribution is still expressed as a global inventory. This presentation of the local TCRE provides little insight as to how regional changes in carbon reservoirs contribute towards the atmospheric level of CO_2 after emissions.

In contrast, we choose to separate the terms for the local TCRE by localising the terms in (4.4). In this chapter we take our regions to be the Northern and Southern Hemisphere; however, this separation may be generalised to any regional partition of the Earth System. Our separation includes a fraction that relates the total hemispheric carbon changes $\Delta I_{tot,hem}$ to the global carbon emissions and accounts for differences in the hemispheric storage of emitted carbon through differences in the hemispheric land and ocean components:

$$\frac{\Delta T_{hem}}{\Delta I_{em}} = \left(\frac{\Delta T_{hem}}{\Delta R_{hem}}\right)\left(\frac{\Delta R_{hem}}{\Delta R_{CO_2,hem}}\right)\left(\frac{\Delta R_{CO_2,hem}}{\Delta I_{atm,hem}}\right)\left(\frac{\Delta I_{atm,hem}}{\Delta I_{tot,hem}}\right)\left(\frac{\Delta I_{tot,hem}}{\Delta I_{em}}\right). \quad (4.10)$$

The product of the first four terms, condensed as $\Delta T_{hem}/\Delta I_{tot,hem}$, represent the extent that the local heat and carbon responses compensate; similar trends in $\Delta T_{hem}/\Delta I_{tot,hem}$ imply that the hemispheres have similar compensations, whereas differences imply that the global compensation between the heat and carbon contributions occurs in spite of opposing hemispheric behaviors. This breakdown provides a regional effective TCRE analogue that can reveal physical insight to the regional thermal and carbon responses that determine the global effective TCRE.

4.2.2 The University of Victoria Earth System Climate Model

In this chapter we analyse output from the University of Victoria Earth System Climate Model, version 2.9 (UVic ESCM), an Earth system model of intermediate complexity (Weaver et al., 2001; Eby et al., 2013). The UVic ESCM consists of an ocean model with a horizontal resolution of 3.6° longitude ×1.8° latitude and 19 vertical levels coupled to a two-dimensional energy-moisture balance atmospheric model of the same horizontal resolution (Fanning and Weaver, 1996). The atmospheric model uses climatological wind fields that are allowed to adjust geostrophically to changes in temperature gradients and a constant climatological albedo pattern to account for the impact of clouds on the radiation budget. The model also contains representations for dynamic-thermodynamic sea ice (Bitz et al., 2001; Hunke and Dukowicz, 1997), ocean biogeochemistry (Keller et al., 2012a), dynamic vegetation (Meissner et al., 2003), and sediments (Archer, 1996). The UVic ESCM has been used in a multitude of studies on the proportion of surface warming to cumulative emissions (Matthews et al., 2009; Zickfeld et al., 2009; Eby et al., 2013; MacDougall and Friedlingstein, 2015; MacDougall et al., 2017), particularly those relating to the ocean response (Ehlert et al., 2017), as it includes complex carbon cycle dynamics while remaining computationally efficient due to its simplified atmospheric component.

The model is forced with historical greenhouse gas emissions, land-use changes, and volcanic aerosol forcing up to year 2005, then run under emissions scenarios consistent with Representative Concentration Pathway (RCP) 4.5 and 8.5 to year 2100 (Meinshausen et al., 2011a) (Fig. 4.1). RCP 4.5 represents a "medium mitigation scenario", whereas RCP 8.5 represents a "high baseline scenario" where mitigation of carbon emissions is kept at a minimum. The use of two emissions scenarios allows for interrogation into the dependency of hemispheric symmetry or asymmetry on the background emissions rates as emissions in RCP 4.5 decrease and stabilise below 5 PgC yr^{-1}, whereas RCP 8.5 experiences accelerating emissions up to near 30 PgC yr^{-1}. As some models experiences different ratios of warming to

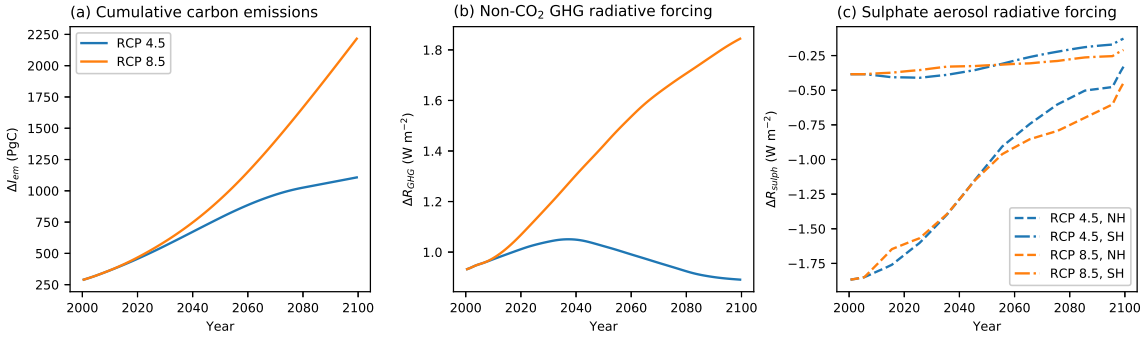

Figure 4.1: RCP 4.5 (blue lines) and RCP 8.5 (orange lines) forcing agents: (a) cumulative CO_2 emissions ΔI_{em}, in Pg C, (b) global average radiative forcing from non-CO_2 greenhouse gases, ΔR_{GHG}, in W m^{-2}, (c) hemispheric-averaged radiative forcings from sulphate aerosols, ΔR_{sulph}, in W m^{-2}, for the Northern Hemisphere (dashed lines) and the Southern Hemisphere (dot-dashed lines).

cumulative emissions depending on the background emissions rates (Krasting et al., 2014), there is the potential for the hemispheric response to also be dependent on the background emissions. In the case of the UVic ESCM, this dependence would occur in spite of the general path-independence found on a global scale (Herrington and Zickfeld, 2014).

An 11-year solar cycle is included in the insolation forcing for both the historical and RCP periods. The RCP forcing components include land-use changes, CO_2 emissions, non-CO_2 greenhouse gases, and sulfate aerosols. In the model, CO_2 and non-CO_2 greenhouse gases are well-mixed, while sulfate aerosol concentrations are spatially nonuniform. We also consider the output of a model run in which sulphate aerosol emissions are not included in the historical and RCP forcings, to explore the role of forcing on any potential hemispheric asymmetries in the radiative forcing. As output, 20-year averages of spatial fields are provided from 1860 to 2005, and annual averages of spatial fields and global quantities are provided from 2000 to 2100.

Since all the UVic ESCM runs provided include non-CO_2 radiative forcing elements, and the CO_2 is forced by emissions scenarios rather than a 1% atmospheric concentration increase, we will henceforth refrain from using the TCRE metric. Instead, we consider the effective TCRE given by the ratio of surface warming to cumulative carbon emissions, $\Delta T/\Delta I_{em}$, and the terms referenced in (4.10).

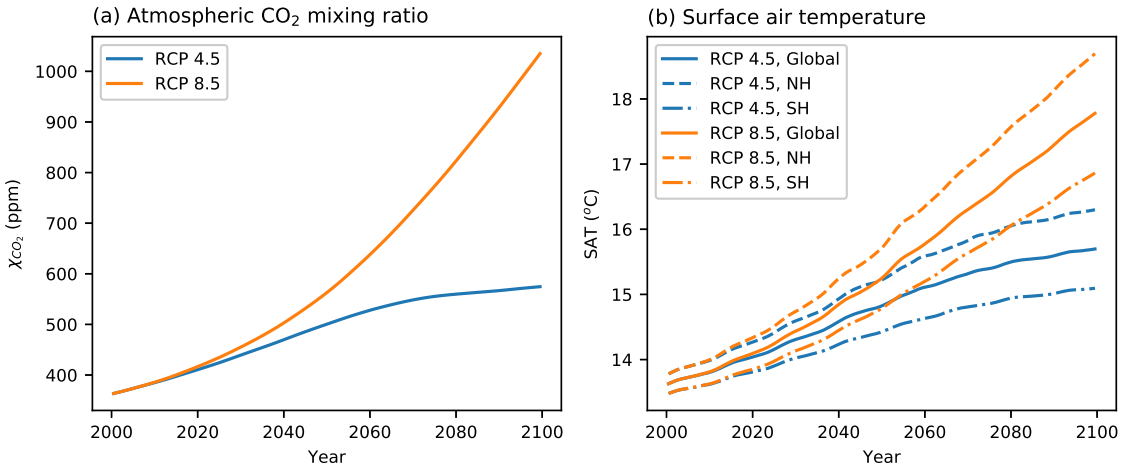

Figure 4.2: (a) Atmospheric CO_2 mixing ratios for RCP 4.5 (blue) and RCP 8.5 (orange). (b) Surface air temperature for RCP 4.5 (blue) and RCP 8.5 (orange), for the global average (solid lines), Northern Hemisphere averages (dashed lines), and Southern Hemisphere averages (dot-dashed lines).

4.3 Global and hemispheric UVic ESCM responses to emissions

4.3.1 Global UVic ESCM response to RCP forcings

We present the UVic ESCM response to the RCP scenarios for the control runs. The control RCP 4.5 and RCP 8.5 emissions scenarios result in an atmospheric carbon increase from 365 ppm at year 2000 to 574 ppm and 1035 ppm by year 2100, respectively. This increase in atmospheric carbon is associated with an increase in global average surface temperatures of 2.5 °C in RCP 4.5 and 4.5 °C in RCP 8.5 relative to the preindustrial. The model shows an amplification of surface warming in the Northern Hemisphere for both scenarios. At year 2000 the Northern Hemisphere is 0.3 °C warmer than the Southern Hemisphere; by 2100 this difference increases to 1.2 °C in RCP 4.5 and 1.8 °C in RCP 8.5 (Fig. 4.2b). When aerosols forcing is excluded, there is a slight enhancement in atmospheric carbon levels, from a year 2000 concentration of 375 ppm to 582 ppm in RCP 4.5 and 1162 ppm in RCP 8.5 (not shown). The lack of aerosol forcing leads to surface temperatures around $0.3-0.5$°C warmer than those in the control runs; however, the differences between the two hemispheres are highly similar to those found in the control runs (not shown).

Ocean average temperatures (Fig. 4.3a) increase from 3.3 °C at 2000 to 3.7 °C in RCP

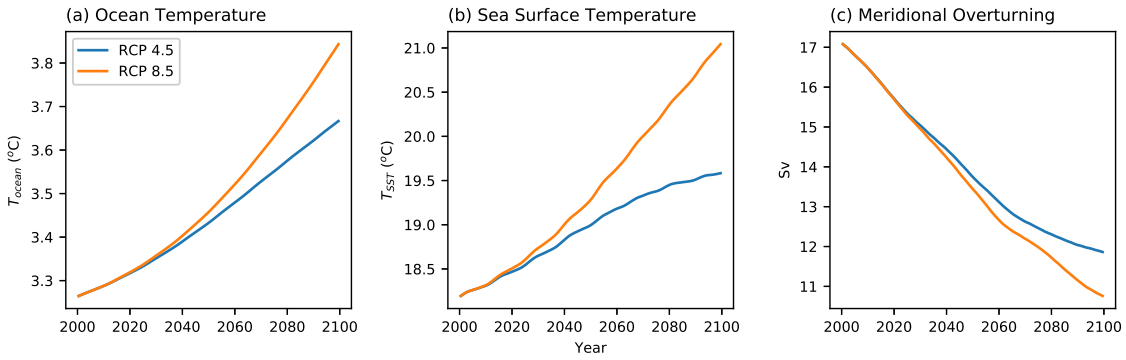

Figure 4.3: 21st Century global ocean response for RCP 4.5 (blue) and RCP 8.5 (orange): (a) Average ocean temperature, in °C, (b) Average sea surface temperature, in °C, and (c) Meridional overturning strength in Sv.

4.5 and 3.8 °C by the end of RCP 8.5. Sea surface temperatures (Fig. 4.3a) increase from 18.2 °C at 2000 to 19.6 °C in RCP 4.5 and 21.0 °C by the end of RCP 8.5. There is a weakening in the meridional overturning streamfunction throughout the RCP scenarios, from a year 2000 strength of 17 Sv (1 Sv = 10^6 m^3 s^{-1}) to 11.9 Sv at the end of RCP 4.5 and 10.8 Sv at the end of RCP 8.5. These gross responses to the RCP scenarios are consistent with the likely estimates from the CMIP5 RCP experiments using higher-complexity coupled climate models, although it should be noted that the confidence in the meridional overturning evolution was low in the IPCC 5th Assessment Report (Collins et al., 2013).

Climate feedback parameter diagnostics

We can diagnose the time-dependent global climate feedback parameter λ using (4.5). The global radiative forcing ΔR is split unevenly between the radiative response $\lambda \Delta T$ and the net top-of-the-atmosphere heat flux $\mathcal{N}$ (Fig. 4.4a,b). The net top-of-the-atmosphere heat flux uptake initially dominates the balance, but as emissions continue the radiative response grows more quickly. The diagnosed climate feedback parameters for RCP 4.5 and 8.5 reveal small variations from an average of about 1.1 W m^{-2} K^{-1} (Fig. 4.4c,d). Cyclic variations in λ arise mainly from the 11-year solar cycle included in the boundary conditions; an 11-year filter applied on λ eliminates most of this variability. Both control runs show a similar asymptotic increase in λ. The asymptotic behaviour and decreasing amplitude of the 11-year signal are both results of our diagnostic approach towards $\lambda(t)$, as the evolution of λ with time shows the integrated impact of all climate feedbacks involved; thus, on the centennial

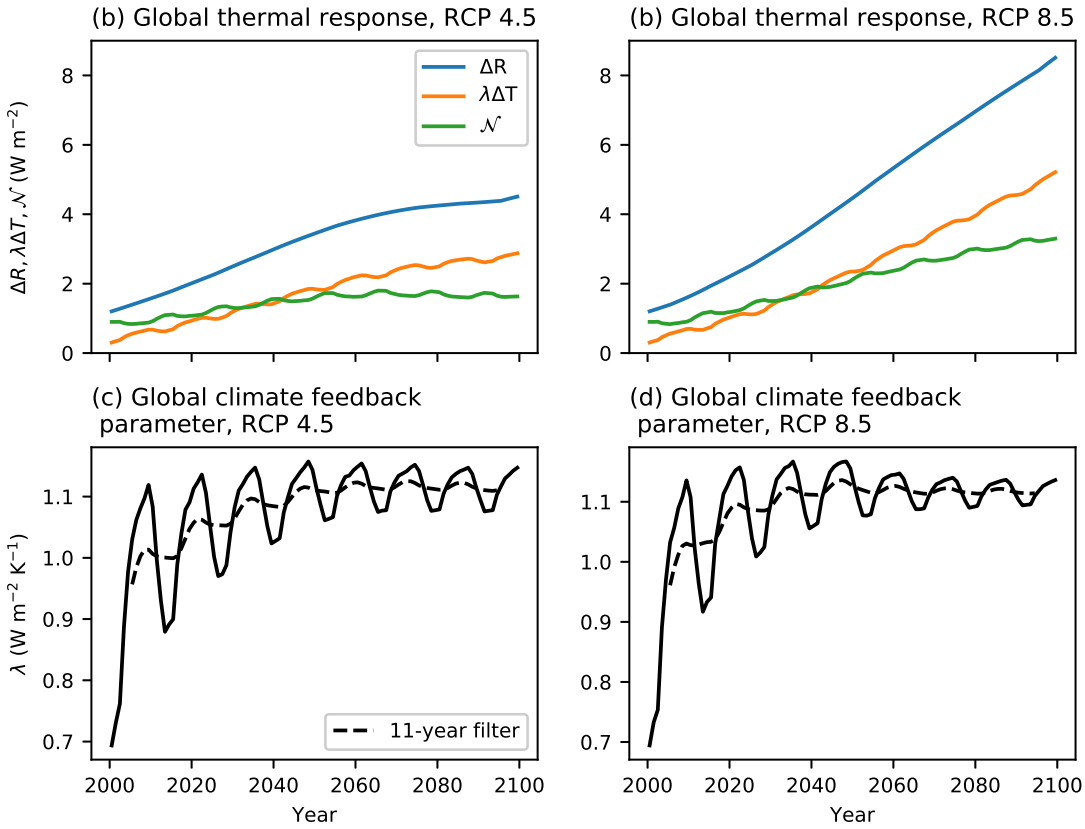

Figure 4.4: Globally-averaged radiative forcing ΔR, radiative response $\lambda \Delta T$, and planetary heat uptake $\mathcal{N}$, in units W m^{-2}, for (a) RCP 4.5 and (b) RCP 8.5. Diagnosed climate feedback parameter λ (solid lines), with 11-year average (dashed lines), for (c) RCP 4.5 and (d) RCP 8.5.

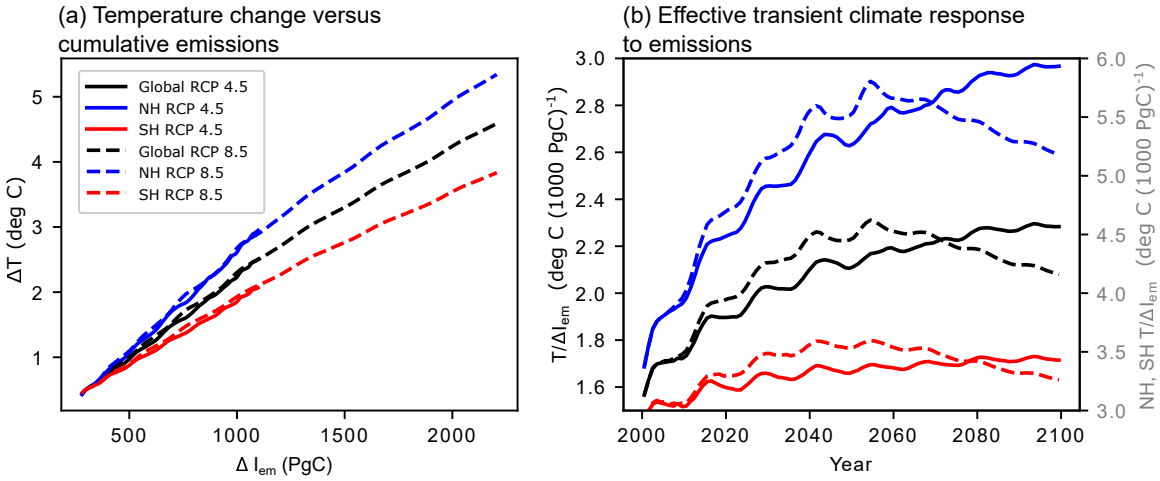

Figure 4.5: (a) Average global and hemispheric temperature changes versus cumulative CO_2 emissions for RCP 4.5 and RCP 8.5. (b) Global and hemispheric effective transient climate responses to emissions, with units °C (1000 Pg C)$^{-1}$, for RCP 4.5 and RCP 8.5.

timescales explored here, the dominance of CO_2 in setting the thermal response reduces the impact of other external forcings. This long-term global-mean behavior is consistent with previous diagnostics involving the UVic ESCM, which assume a constant climate feedback parameter (Eby et al., 2013; Ehlert et al., 2017). As there is still some variation with time, we continue to employ the time-variant version of λ for our analysis.

4.3.2 Proportion of surface warming to cumulative emissions

Previous studies using the UVic ESCM show a robust linear relationship between cumulative carbon dioxide emissions and globally-averaged surface warming (Matthews et al., 2009; Zickfeld et al., 2009; MacDougall et al., 2017). In the UVic ESCM runs with non-CO_2 forcings included, a near-constant proportionality exists for global and hemispheric temperature changes against global cumulative CO_2 emissions (Fig. 4.5a). Inspection of the ratio of warming to cumulative CO_2 emissions with time reveals consistent nonlinearities, particularly within the first 50 years during which the ratio increases in both scenarios and for each hemisphere (Fig. 4.5b). Globally, the increase continues over RCP 4.5 but reverses during the second half of RCP 8.5.

The nonlinearity in the global response to cumulative carbon emissions arises mainly from the Northern Hemisphere response to cumulative carbon emissions. From 2000-2050, the ratio of warming to cumulative CO_2 emissions in the Northern Hemisphere almost doubles; in comparison, the Southern Hemisphere ratio of warming increases only slightly. Over

the second half of RCP 8.5, although the reversal is seen in both hemispheres, the decrease in the Northern Hemisphere ratio of warming to cumulative CO_2 emissions is stronger than that in the Southern Hemisphere. The differences in the hemispheric relationships between warming and cumulative CO_2 emissions compare with those of Leduc et al. (2016), who found that while most regions had a relatively linear local TCRE, there are noticeable regions of nonlinear behaviour, particularly in the North Atlantic and Arctic. When sulfate emissions are removed from the forcings, the hemispheric responses to cumulative CO_2 emissions are more symmetric. The Northern Hemisphere response to cumulative emissions is higher in the beginning of the RCP scenarios, and its decline over the century has approximately the same magnitude as the Southern Hemisphere response for both RCP scenarios (not shown).

4.4 Contributions to the hemispheric proportion of surface warming to cumulative emissions

The contributions towards the global and hemispheric responses to cumulative CO_2 emissions (4.4 and (4.10) can be normalised as logarithmic changes (e.g. $\Delta T/\Delta R \rightarrow \Delta \ln(T/R)$), so that compensation from factors can be visible as an additive property (Fig. 4.6). Globally, the UVic ESCM reveals a near complete compensation between temporal changes in the thermal response $\Delta T/\Delta R$, CO_2 radiative forcing ratio $\Delta R_{CO_2}/\Delta I_{atm}$, and atmospheric fraction $\Delta I_{atm}/\Delta I_{em}$ (Fig. 4.6a,b). Changes in the response to cumulative CO_2 emissions are thus set by the strength of non-CO_2 radiative forcings through the term $\Delta R/\Delta R_{CO_2}$, in line with Mengis et al. (2018). There are slight differences between the response in RCP 4.5 and 8.5. For instance, in the last 20 years of RCP 8.5 the decrease in $\Delta R_{CO_2}/\Delta I_{atm}$ begins to dominate, leading to a decrease in $\Delta T/\Delta I_{em}$; in RCP 4.5 a near-constant $\Delta R_{CO_2}/\Delta I_{atm}$ in the last 20 years allows for the thermal response $\Delta T/\Delta R$ to maintain a slight increase in $\Delta T/\Delta I_{em}$.

The near-compensation among the thermal and carbon contributions found in the global response to CO_2 emissions breaks down when considering hemispheric responses. Besides the ratio of CO_2 radiative forcing to atmospheric carbon changes, which is identical in each hemisphere as atmospheric CO_2 is well mixed, there is little similarity in the behaviour of terms within the hemispheric response to cumulative CO_2 emissions. In the Northern Hemisphere (Fig. 4.6c-d), the increase in the response to cumulative CO_2 emissions is caused by the increase in the radiative forcing fraction $\Delta R/\Delta R_{CO_2}$ which occurs as aerosol emis-

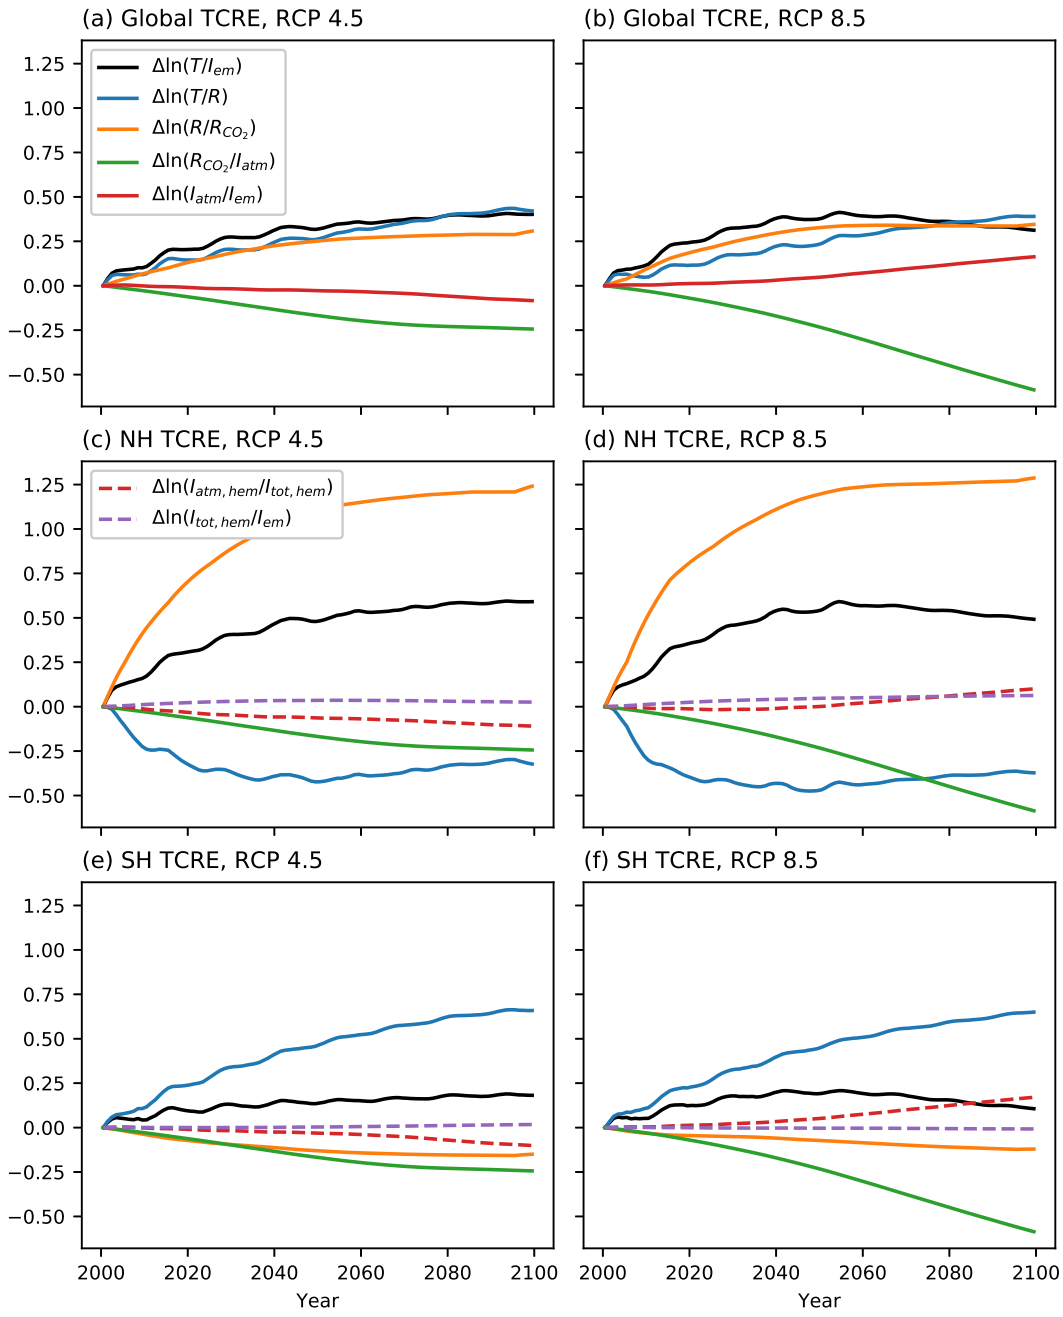

Figure 4.6: Logarithmic changes in the terms setting the responses to cumulative carbon emissions for the global response under (a) RCP 4.5 and (b) RCP 8.5, the Northern Hemisphere response under (c) RCP 4.5 and (d) RCP 8.5, and the Southern Hemisphere response for (e) RCP 4.5 and (f) RCP 8.5. Note that the hemispheric response to emissions has two terms different from the global carbon response term $\Delta I_{atm}/\Delta I_{em}$; these are set as dashed lines.

sions decline over the RCP scenarios. The thermal response $\Delta T/\Delta R$ decreases to partially offset the increase in $\Delta R/\Delta R_{CO_2}$. In the Southern Hemisphere (Fig. 4.6e-f), the thermal response $\Delta T/\Delta R$ is positive. The radiative forcing fraction $\Delta R/\Delta R_{CO_2}$ decreases here as it is mainly set by the evolution of non-CO_2 greenhouse gases in the RCP scenarios. The more modest increase in the response to cumulative CO_2 thus occurs because the radiative forcing fraction and carbon response together compensate for most of the hemispheric thermal response.

Unlike the thermal responses which are set by the hemisphere and non-CO_2 radiative forcing, the hemispheric carbon responses $\Delta I_{atm,hem}/\Delta I_{tot,hem}$ (red dashed lines) are mainly set by the emissions scenario, although the values are hemisphere-dependent. The partition of carbon emissions between the hemispheres $\Delta I_{tot,hem}/\Delta I_{em}$ is nearly constant, with the Northern Hemisphere (Southern Hemisphere) containing about 46% (54%) of total emissions (lilac dashed lines, Fig. 4.6 c-f). Thus $\Delta T_{hem}/\Delta I_{tot,hem} \approx \gamma \Delta T_{hem}/\Delta I_{em}$ for some hemisphere-dependent scaling term γ, implying that our process-based understanding of the hemispheric response to (hemispheric) carbon changes is analogous to the local TCRE defined by Leduc et al. (2016) using global cumulative carbon emissions. Thus, to investigate the controlling mechanisms behind these hemispheric differences, we can next diagnose the regional heat balances (Section 4.5) and carbon balances (Section 4.6).

4.5 Hemispheric contributions to the thermal response

The global thermal response to radiative forcing (4.6) can be generalised as the sum of the radiative response and an area-average heat content tendency term dQ/dt, in W m^{-2}, which consists of the vertically-integrated atmospheric and ocean heat tendencies. This empirical balance is extended to a regional view for a local area A:

$$\Delta R_A = \lambda_A \Delta T_A + \frac{dQ_A}{dt}. \tag{4.11}$$

By Gauss's theorem the heat content tendency dQ_A/dt is equal to the divergence of the heat flux through the boundary of A: for a region bounded in the vertical by the atmosphere-land boundary and top of the atmosphere, the equivalent to the top-of-the-atmosphere heat flux $\mathcal{N}_A$ and horizontal divergence of the vertically-integrated heat transport $\int \nabla \cdot F dz$ (Armour et al., 2013):

$$\frac{dQ_A}{dt} = \mathcal{N}_A + \int_z \nabla \cdot F \, dz. \tag{4.12}$$

When considering the hemispheric breakdown, the divergence of the vertically-integrated heat transport reduces to the cross-equatorial heat transport F_{eq} averaged over the hemi-

spheric surface area (to obtain units Wm^{-2}), so the hemispheric thermal response from (4.10) can be expanded as

$$\frac{\Delta T_{hem}}{\Delta R_{hem}} = \frac{1}{\lambda_{hem}}\left(1 - \frac{dQ_{hem}/dt}{\Delta R_{hem}}\right) = \frac{1}{\lambda_{hem}}\left(1 - \frac{\mathcal{N}_{hem} \pm F_{eq}}{\Delta R_{hem}}\right). \tag{4.13}$$

The hemispheric heat balance terms (4.11) are now presented (Fig. 4.7) to see how they compare to the global heat balance (Fig. 4.4 a, b). There is a more rapid increase in the Northern Hemisphere radiative forcing caused by decreasing sulfate aerosol concentrations (Fig. 4.1c). The amplified radiative forcing signal in the Northern Hemisphere is accompanied by larger changes in the radiative response over the RCP scenarios. In the Northern Hemisphere, the more rapid increase in radiative forcing, $\int \Delta R_{NH}\, dA$, and slower heat gain, $\int dQ/dt|_{NH}\, dA$, lead to the radiative response, $\lambda_{NH} \int \Delta T_{NH}\, dA$, becoming the dominant contribution to the thermal response (Fig. 4.7 a,b). Conversely, in the Southern Hemisphere, the radiative forcing is almost evenly split between the radiative response and hemispheric heat content tendency, particularly for RCP 8.5 (Fig. 4.7 c,d).

4.5.1 Hemispheric impacts on the climate feedback parameter

The hemispheric climate feedback parameters λ_{NH} and λ_{SH} are diagnosed using the hemispheric heat balances in (4.11). The model reveals a clear hemispheric asymmetry in the climate feedback parameter. While the global-mean λ is nearly constant, λ_{NH} increases and λ_{SH} decreases over the century; neither hemispheric value reaches the global-mean λ by year 2100 (Fig. 4.8a). This evolution of the climate feedback parameter implies that warming in the Southern Hemisphere becomes more sensitive to radiative forcing over time, while Northern Hemisphere warming becomes less sensitive to radiative forcing.

To understand the impact of non-uniform radiative forcing on the climate feedback parameter, global-mean and hemispheric λ are diagnosed for RCP scenarios in which sulfate aerosol forcing is excluded. In this case, all three λ decrease slightly over the century, and the model climate feedback parameters have a similar spread at year 2100 (Fig. 4.8b). This decline indicates an increased sensitivity to forcing, which is stronger in the Northern Hemisphere but also occurs over the Southern Hemisphere and is consistent with other studies on the climate feedback parameter (Andrews et al., 2015; Andrews and Webb, 2018) The Northern Hemisphere surface warming remains more sensitive to radiative forcing as $\lambda_{NH} < \lambda_{SH}$. Thus, the hemispheric differences in radiative forcing from sulfate aerosols provide the primary control on the evolution of the hemispheric climate feedback parameters λ_{NH} and λ_{SH}.

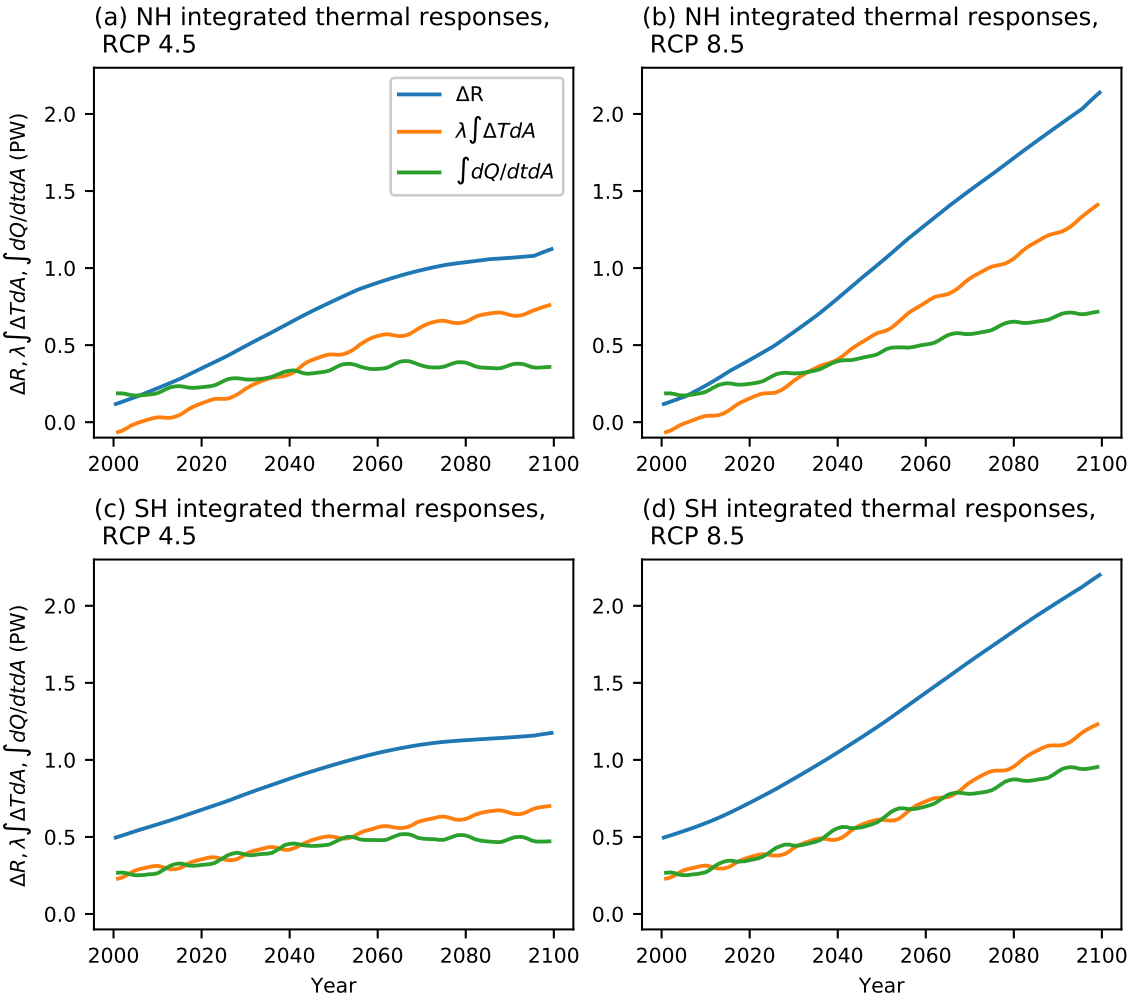

Figure 4.7: Integrated thermal response terms - Radiative forcing ΔR, radiative response $\lambda \Delta T$, and heat content tendency dQ/dt, for (a) RCP 4.5 Northern Hemisphere, (b) RCP 8.5 Northern Hemisphere, (c) RCP 4.5 Southern Hemisphere, and (d) RCP 8.5 Southern Hemisphere.

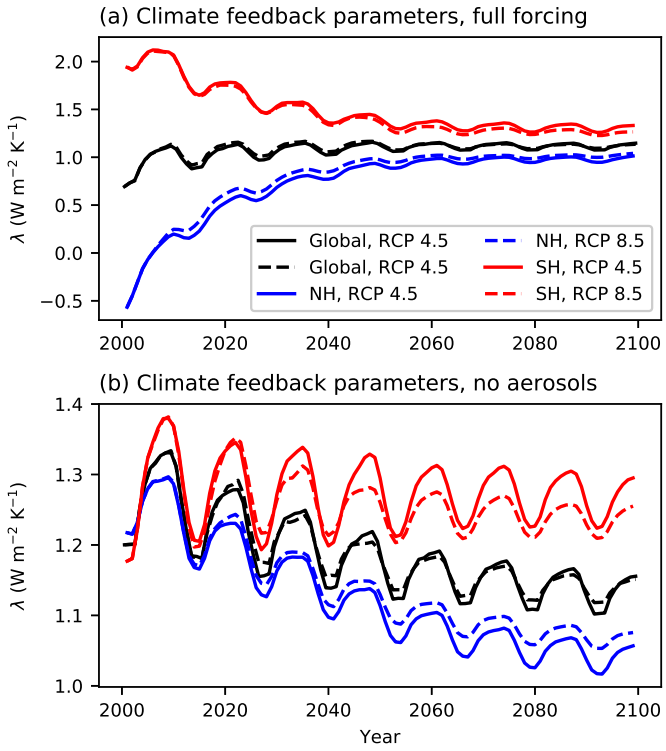

Figure 4.8: (a) Diagnosed global and hemispheric climate feedback parameters for RCP 4.5 and 8.5, with full forcing sets including CO_2 emissions, GHG radiative forcing, and sulfate aerosol radiative forcing (b) As for (a), but using runs with RCP 4.5 CO_2 emissions and GHG radiative forcing only.

4.5.2 Diagnosing the UVic ESCM cross-equatorial heat flux

The hemispheric heat tendency term dQ_{hem}/dt in (4.11) and (4.13) results from a combination of the local top-of-the-atmosphere heat flux and heat transported across the equator. The UVic ESCM output includes ocean-atmosphere heat fluxes (F_{sea}), top-of-the-atmosphere heat fluxes ($\mathcal{N}$), as well as surface air temperature and ocean potential temperature and ocean velocities. Thus there are two methods for calculating the cross-equatorial heat transport F_{eq}: through the difference between the atmosphere and ocean heat fluxes and heat tendencies, or through the direct calculation of the meridional heat transport $v \cdot T$ for both the atmosphere and ocean at the equator. These two methods should be identical; however, errors may arise as the UVic ESCM uses an Arakawa c-grid, meaning that model meridional velocities are found on the northern and southern boundaries of the grid cell. In order to avoid errors from averaging and issues from the z-levels used for the ocean model, we opt to calculate the cross-equatorial heat fluxes by integrating $\mathcal{N}$ across the hemispheres and integrating the atmosphere and ocean heat content changes in each hemisphere.

The UVic ESCM sets constants for atmosphere and the ocean specific heat and density. The product of specific heat and density of seawater in the model is set as $c_{p,oce}\rho_{oce} = 4.185 \times 10^6$ J m^{-3} K^{-1}. The slab atmospheric component approximates vertically-integrated atmospheric properties by using surface atmospheric properties and multiplying them with an atmospheric scale height (Gill, 1982). For the atmosphere, specific heat is set as $c_{p,atm} = 1 \times 10^3$ J (kg K)$^{-1}$, density is set as $\rho_{atm} = 1.25 \times 10^{-3}$ kg m^{-3}, with an atmospheric scale height of $H = 8.4$ km. Taking both the atmosphere-ocean flux F_{sea} and the top-of-the-atmosphere flux $\mathcal{N}$ to be positive when pointing down, the calculations for the cross-equatorial fluxes are then

$$F_{eq,atm} = \int_{A_{hem}} \Big(F_{sea} - \mathcal{N} + H c_{p,atm} \rho_{atm} \frac{dT}{dt} \Big) dA_{hem} \text{ and} \tag{4.14}$$

$$F_{eq,oce} = -\int_{A_{hem}} F_{sea} dA_{hem} + c_{p,oce}\rho_{oce} \int_{V_{hem}} \frac{dT_{oce}}{dt} dV_{hem}. \tag{4.15}$$

In the UVic ESCM, there is a net northward transport of heat, which occurs in the atmosphere. The ocean transports of heat southwards, but the amount of heat transported by the ocean is about 10% that transported by the atmosphere. The relative strengths and directions of cross-equatorial heat transport in the UVic ESCM are rather unrealistic: observations and reanalyses show a smaller northwards cross-equatorial heat transport set mainly by oceanic transport (Marshall et al., 2014a). The tendencies in heat transport in the atmosphere and ocean partially compensate, with the atmosphere cross-equatorial

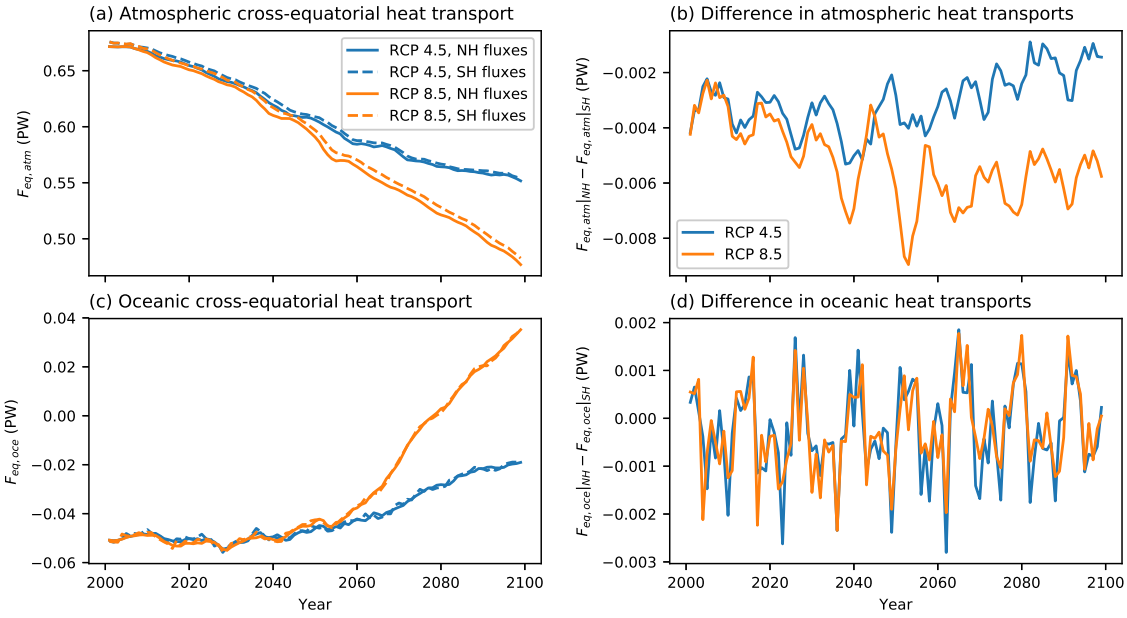

Figure 4.9: (a) Atmospheric cross-equatorial heat transport calculated by heat flux convergence for RCP 4.5 (blue) and RCP 8.5 (orange). Values derived from Northern Hemisphere output are solid lines, values derived from Southern Hemisphere output are dashed lines. (b) Difference in cross-equatorial heat transports between calculations using Northern Hemisphere and Southern Hemisphere values. (c) As in (a), but calculated for the ocean cross-equatorial heat transport. (d) As in (b), but for the ocean heat transport.

transport decreasing 0.08-0.18 PW over the 21st Century and the ocean transport increasing by 0.04-0.09 PW over the 21st Century.

Calculating the atmospheric and oceanic cross-equatorial heat transports using (4.14) and (4.15) should result in the same values regardless of which hemisphere has its heat fluxes and heat tendencies integrated. The calculations of the cross-equatorial heat fluxes show small errors relative to their magnitudes (errors $\mathcal{O}(10^{12}\text{W})$ versus magnitude $\mathcal{O}(10^{14}\text{W})$) (Fig. 4.9). However, the differences in the calculations show periodic behaviour, the atmospheric errors show a scenario-dependency that the ocean errors do not.

4.5.3 Hemispheric energy balances with asymmetric and symmetric radiative forcing

Hemispheric top-of-the-atmosphere heat imbalances in the UVic ESCM are asymmetric for forcings with and without sulfate aerosols at the start of the 21st Century (Fig. 4.10, left column). When aerosols are excluded, the Southern Hemisphere has a net influx of energy through the top of the atmosphere of 2.8×10^{22} J per year, and the Northern Hemisphere has a net outflux of energy of -0.7×10^{22} J per year; when aerosols are included, the Southern Hemisphere influx stays about the same and the Northern Hemisphere outflux almost doubles to -1.4×10^{22} J per year. In both forcing setups, heat is redistributed mainly by a northward atmospheric transport, resulting in the Northern Hemisphere ocean heat uptake being positive and similar in magnitude to that in the Southern Hemisphere. A small southward ocean heat transport results in ocean heat content changes that are close to volumetric responses. Thus, in the Northern Hemisphere, the heat convergence created by the cross-equatorial heat transport creates a hemispheric heat tendency that is larger than the local radiative forcing; this heat tendency results in surface warming, and through the closure of (4.11) results in a negative climate feedback parameter λ_{NH}. When no aerosols are included, the weaker heat convergence in the Northern Hemisphere results in a hemispheric heat tendency smaller than the local radiative forcing, thus creating a positive Northern Hemisphere climate feedback parameter (Fig. 4.8).

Over the 21st Century, the dominance of CO_2 in emissions over RCP 4.5 and 8.5 increases the top-of-the-atmosphere heat flux for both runs with and without sulfate aerosols. The hemispheric asymmetry decreases at the top of the atmosphere as the integrated Northern Hemisphere flux, while still negative in most cases, is much smaller in magnitude relative to the integrated Southern Hemisphere flux. The northward heat transport through the atmosphere declines, so that the hemispheric air-sea heat fluxes (and, to first order, the

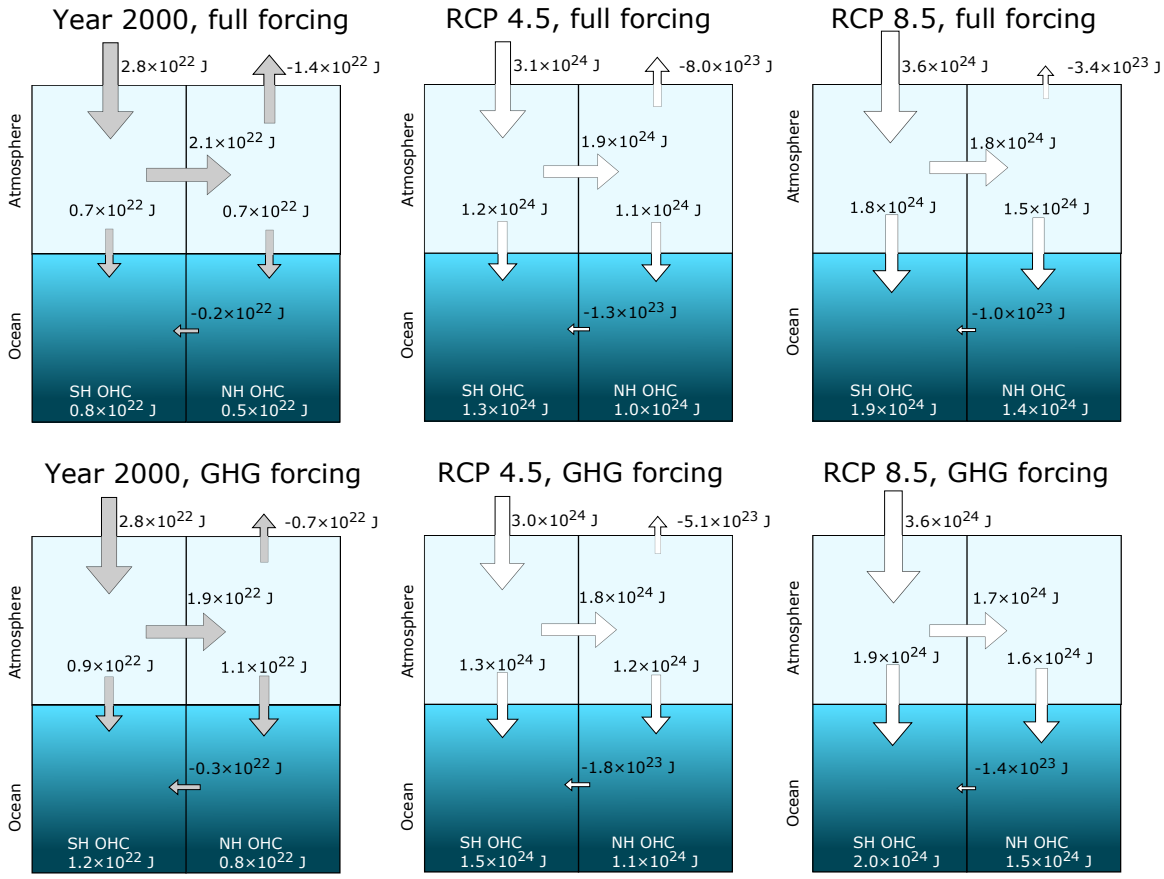

Figure 4.10: Integrated ocean heat content and top-of-the-atmosphere, air-sea, and cross-equatorial heat fluxes over the 21st Century for RCP 4.5 and 8.5 scenarios, with and without sulfate aerosol forcings (top and bottom rows, respectively). Arrows denote direction of flux (with positive values indicating downwards/northwards) and area of arrows scale with magnitude. Atmospheric heat content changes are not included as they are 1-2 orders of magnitude less than the energy transfers and ocean heat content changes. Year 2000 values are integrated for that year only, and the relative sizes of the grey arrows do not correspond with the sizes of the white arrows for the integrated RCP responses.

hemispheric ocean heat content) approach a ratio more similar to the ratio of their surface areas. The more localised response is mirrored in the tendency of the hemispheric climate feedback parameters, as they asymptote to the same values independent on the forcing. Thus, within the UVic ESCM the cross-equatorial transport of heat plays a substantial role in setting each hemisphere's thermal response to emissions throughout the 21st Century.

4.6 Hemispheric contributions to the global carbon response

As CO_2 is well-mixed in the UVic ESCM forcing files, the fraction of emitted carbon residing in each hemisphere, $\Delta I_{tot,hem}/\Delta I_{em}$ obscures the transport that spreads carbon emissions from their point sources. The Northern Hemisphere homes the vast majority of the human population and industrial centres and therefore is the source of most anthropogenic CO_2; in 2005, the United States, Europe, and China together produced almost 60% of all carbon dioxide emissions (International Energy Agency, 2021). It follows that there is a massive transport of CO_2 from the Northern to the Southern Hemisphere; however, in the model forcing set, emissions are considered to be spatially uniform (Meinshausen et al., 2011a), and the UVic ESCM setup does not include an atmospheric CO_2 tracer. Subsequently, for this section we consider the bulk amounts of carbon residing in each hemispheric reservoir, conceding that this view of the hemispheric carbon response lacks a component necessary to understand the real-world response.

4.6.1 Hemispheric terrestrial and oceanic carbon responses

The net supply of carbon from fossil fuel emissions are partitioned among the hemisphere's atmosphere, ocean, and terrestrial reservoirs:

$$\Delta I_{tot,hem} = \Delta I_{atm,hem} + \Delta I_{ocean,hem} + \Delta I_{terr,hem}, \tag{4.16}$$

where $I_{em} = I_{tot,NH} + I_{tot,SH}$ does not necessarily have to be an even partitioning of carbon between the hemispheres. The hemispheric reservoir changes include both the response to emissions and any forcings and responses due to land use change, which the UVic ESCM runs include within the RCP 4.5 and 8.5 experiments. However, as land use change constitutes a reallocation of carbon rather than an introduction of carbon, without additional runs we must consider the net changes to each hemispheric reservoir.

As atmospheric CO_2 concentrations are spatially uniform, hemispheric contributions to the carbon response are set by differences in terrestrial and oceanic carbon storage. The

hemispheric carbon response $\Delta I_{atm,hem}/\Delta I_{tot,hem}$ is then equivalent to

$$\frac{\Delta I_{atm,hem}}{\Delta I_{tot,hem}} = 1 - \frac{\Delta I_{terr,hem} + \Delta I_{ocean,hem}}{\Delta I_{tot,hem}}. \tag{4.17}$$

Between the non-atmospheric hemisphere carbon reservoirs, ocean carbon changes dominate the response for both RCP 4.5 and RCP 8.5 (Fig. 4.11a). By 2100 the ocean has taken up 36% of the added carbon in RCP 4.5 and 22% of the added carbon in RCP 8.5 (Fig. 4.11b). The Southern Hemisphere has a larger ocean carbon uptake than the Northern Hemisphere as it holds most of the ocean surface area. The Southern Hemisphere plays a disproportionate role in ocean carbon uptake; while it contains around 60% of the global ocean surface area, its ocean carbon uptake is over 65% the global total for both emissions scenarios.

Terrestrial carbon uptake plays a relatively small role in setting the hemispheric carbon response to emissions. Terrestrial carbon uptake is, however, more impacted by the representation of land-use changes, which alter the carbon uptake potential through the vegetation type (typically from forest vegetation types to grass vegetation types). The impact of land-use changes can be seen in the Southern Hemisphere terrestrial carbon response, which is initially slightly negative (Fig. 4.11a). Land carbon does increase due to raised atmospheric CO_2 levels, but the uptake of added carbon never reaches over 5% in either hemisphere (Fig. 4.11b). As the Northern Hemisphere contains a higher fraction of global land area (about 66%), its terrestrial carbon uptake is higher. The partition of terrestrial carbon uptake is more scenario-dependent than the partition of ocean carbon uptake: by the end of RCP 4.5 Northern Hemisphere carbon uptake makes up 61% of the total terrestrial carbon uptake, while in RCP 8.5 it makes up 74%.

4.6.2 Structure of ocean carbon pools

Ocean interior dissolved inorganic carbon changes can be split into carbon pools to better understand the physical and biological controls on carbon content (Williams and Follows, 2011):

$$C_{ocean} = C_{sat} + C_{soft} + C_{carb} + C_{res}. \tag{4.18}$$

The saturated component C_{sat} constitutes the amount of carbon that a water parcel would have if it were brought adiabiatically to the surface and allowed to equilibrate with atmospheric CO_2 concentrations at that time. The soft tissue component C_{soft} and carbonate dissolution component C_{carb} constitute the inorganic carbon that is added after the water

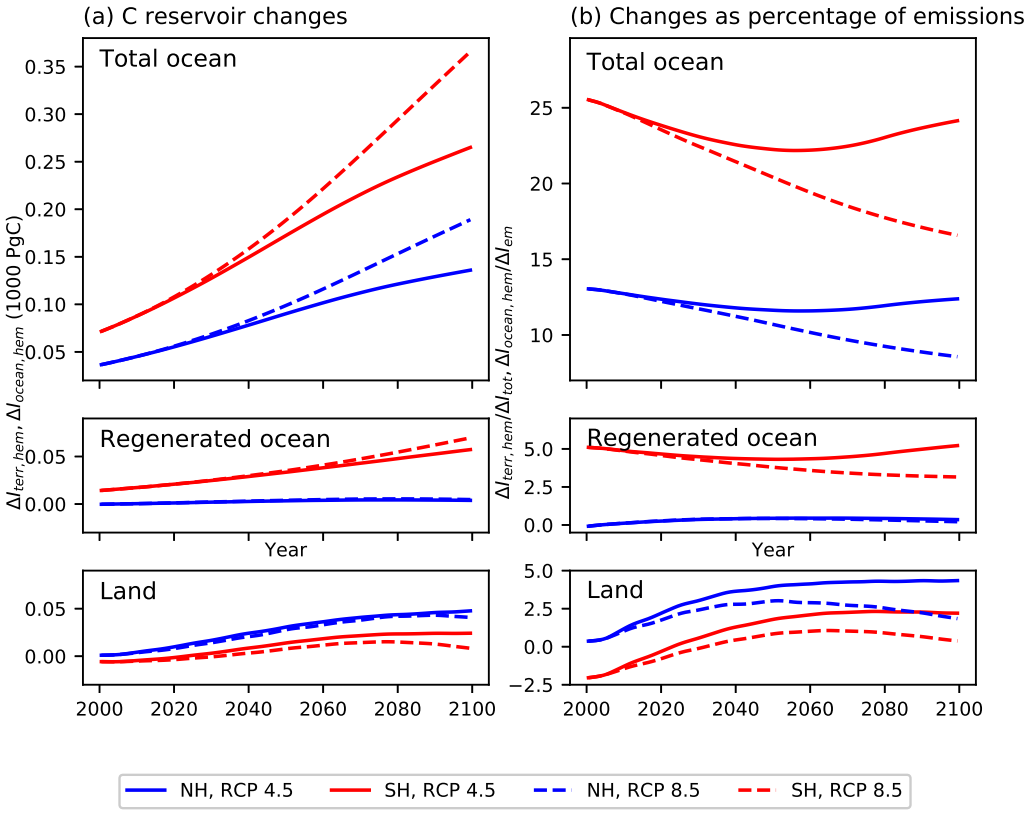

Figure 4.11: (Left hand) Total reservoir changes relative to 1860 for (top) ocean carbon, (middle) regenerated ocean carbon, and (bottom) terrestrial carbon. (Right hand) As for the left-hand column but reservoir changes as a percentage of cumulative emissions.

parcel has been subducted into the ocean interior by regenerating soft tissue or calcium carbonate, respectively. The residual component C_{res} then describes how far removed the parcel is from chemical equilibrium with the system. Together, $C_{sat} + C_{res}$ make up the *preformed carbon pool* and $C_{soft} + C_{carb}$ make up the *regenerated carbon pool*.

As the UVic ESCM ocean model contains a biogeochemical model with nutrients and oxygen as tracers, it is possible to diagnose each of these ocean carbon pools separately to see how they impact the hemispheric carbon changes (Ito and Follows, 2005; Williams and Follows, 2011; Lauderdale et al., 2013). For simplicity, we compare the total hemispheric carbon changes to regenerated carbon changes, as it gives us an indication about the strength of biology in setting the hemispheric ocean carbon response. Regenerated ocean carbon makes up a small part of ocean carbon changes over RCP 4.5 and 8.5, similar in magnitude to the terrestrial carbon response (Figure 4.11). The role that regenerated carbon plays is hemisphere-dependent; in the Northern Hemisphere, regenerated carbon plays almost no role in ocean carbon uptake, whereas in the Southern Hemisphere regenerated carbon can make up about 15% of the total carbon uptake. The small role of regenerated carbon uptake means that the carbon response is dominated by preformed carbon and is a product of air-sea carbon fluxes and ventilation rates.

Looking at the horizontal and depth structure of ocean carbon and regenerated carbon changes may provide further insights into how the hemispheric carbon responses are set (Fig. 4.12). Although the total ocean carbon changes are somewhat similar to a surface-area-weighted response, the horizontal distribution of carbon uptake is highly hemisphere-dependent. In the Northern Hemisphere, carbon uptake is concentrated in the subpolar and subtropical North Atlantic. In the Southern Hemisphere, carbon uptake is spread throughout the midlatitudes. In terms of structure with depth, ocean carbon changes are concentrated in the upper 1000m for both scenarios with a near-symmetric ventilated thermocline structure (Fig. 4.12c,d). The depth structure is similar between RCP 4.5 and 8.5, and mirrors the structure of ocean carbon seen in other models after being forced with increased atmospheric carbon (Ehlert et al., 2017; Frölicher et al., 2015). The Southern Hemisphere's dominant role in setting the carbon response thus occurs mainly from its horizontal extent rather than from differences in the penetration of carbon into the interior.

The largest changes in regenerated carbon, on the other hand, are generally located below the depth horizon of the largest total carbon changes (Fig. 4.12e,f). Regeneration carbon is related to the strength of the overturning, as a slowdown in overturning results in longer residence times for subducted water parcels, thus accumulating more carbon from the breakdown of marine organisms. In the Southern Hemisphere, a slight increase in

regenerated carbon in RCP 8.5 is associated with slightly more regenerated carbon along the intermediate water pathway. Regenerated carbon also increases at the depths and latitudes where North Atlantic Deep Water is found, but its impact across a small zonal extent is compensated by declining regenerated carbon concentrations in the subtropical interior. However, in both hemispheres, this increase is small, resulting in even more dominance of preformed ocean carbon in the hemispheric carbon response.

4.7 Discussion and Conclusions

The climate response to emissions can be understood as a ratio of surface warming to cumulative carbon dioxide emissions $\Delta T/\Delta I_{em}$, and the behaviour of this ratio has been used to understand the processes that determine current carbon budgets. In this chapter, the relationship between surface warming and emissions is reframed to include hemispheric contributions towards the thermal and carbon responses. The framework developed in this chapter is general and can be modified for any partitioning of the Earth system (e.g., the Southern Ocean or the tropics) to isolate the local thermal and carbon factors that contribute to the climate response to emissions. In this chapter, the contributions of the Northern and Southern Hemispheres are assessed using an Earth System model of intermediate complexity forced with emissions consistent with RCP 4.5 and 8.5.

The global evolution of the TCRE results from a partial compensation between the increasing thermal response and decreasing carbon response (Fig. 4.6a,b). The level of partial compensation seen in the global TCRE does not occur for either hemispheric response. There is a stronger compensation in the Southern Hemisphere response, resulting in a hemispheric TCRE that changes less over the 21st Century (Fig. 4.6e,f). However, the Northern Hemisphere TCRE increases more strongly than the global TCRE, and in the Northern Hemisphere the thermal and carbon responses both decrease to compensate the radiative forcing fraction increase (Fig. 4.6c,d). While the Southern Hemisphere plays the dominant role in both heat and carbon uptake due to its large ocean area and ventilation processes, the Northern Hemisphere plays an important role in setting the temporal evolution of the global TCRE through its radiative forcing and thermal responses.

The asymmetry found in the hemispheric responses is largely independent of the forcing scenario (Fig. 4.6), indicating that the UVic ESCM produces a largely path-independent response on hemispheric as well as global scales. The largest path dependence is found in the hemispheric atmospheric fractions, which are likely controlled by the different land use change forcings in RCP 4.5 and RCP 8.5 (Davies-Barnard et al., 2014)"

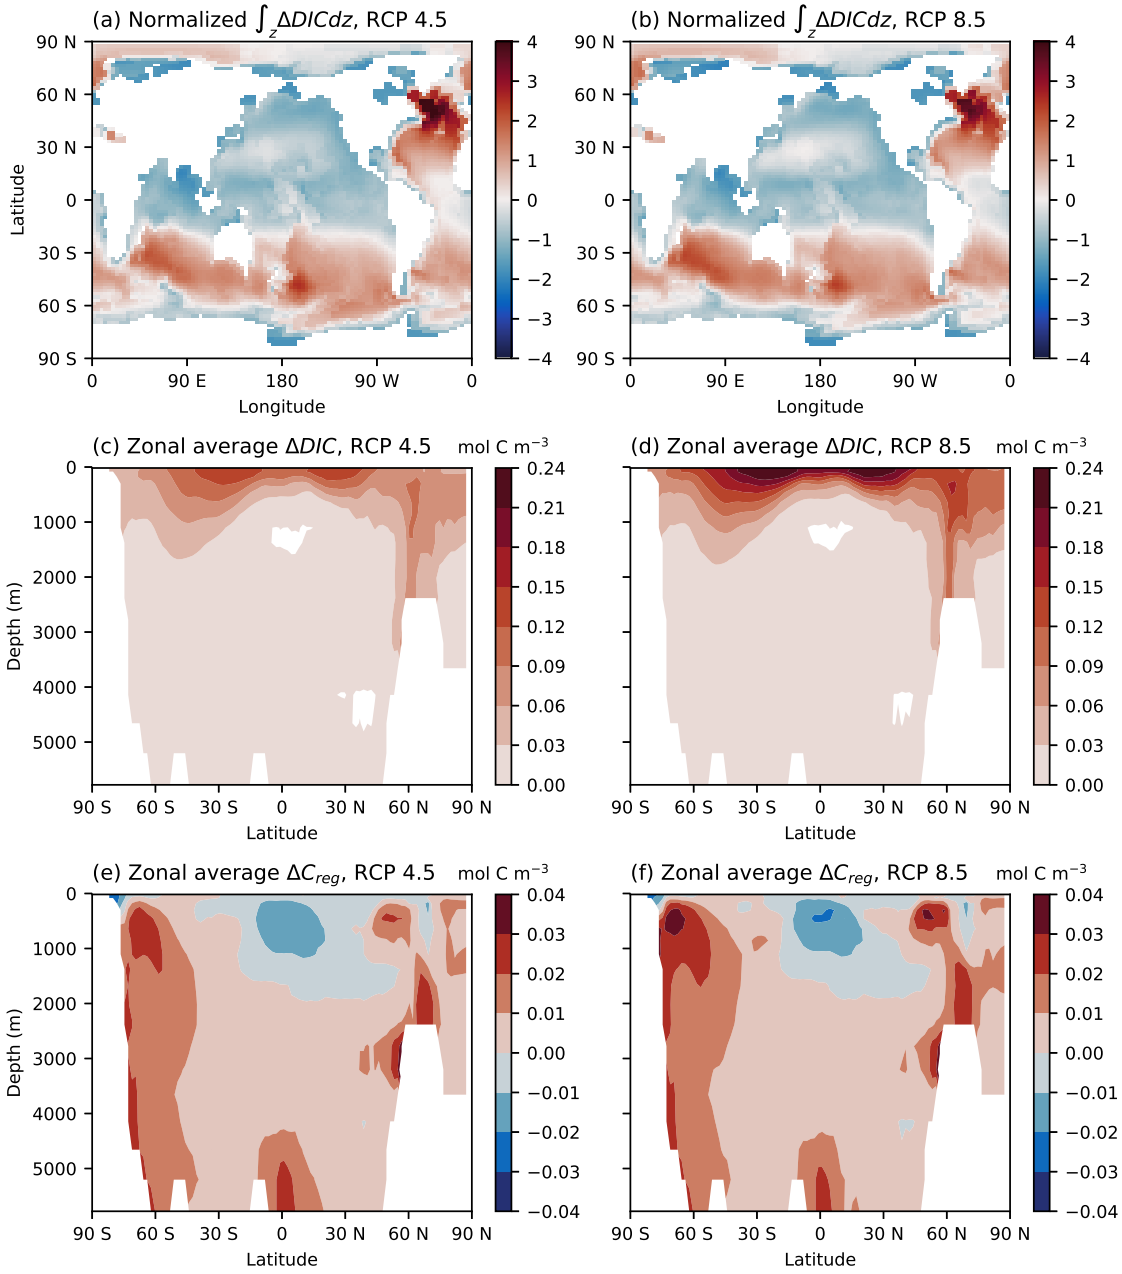

Figure 4.12: (a,b) Normalised depth-integrated DIC changes relative to the preindustrial for RCP 4.5 and 8.5. Values are normalised by the global mean and standard deviation of carbon uptake for each RCP scenario. (c,d) Zonal mean change in ocean DIC relative to 1860 for RCP 4.5 and 8.5 (e,f) Zonal mean change in regenerated carbon relative to 1860 for RCP 4.5 and 8.5. Note that the colourbars for (c,d) differ from those for (e,f).

4.7.1 Hemispheric thermal responses to local radiative forcing set by heat transport

The Northern and Southern Hemisphere thermal responses partially oppose each other: while the Southern Hemisphere thermal response increases in time similar to the global thermal response, the Northern Hemisphere response initially decreases. The hemispheric thermal responses are set by top-of-the-atmosphere imbalances, which are exacerbated by the presence of sulfate aerosols in the Northern Hemisphere, and by a northwards redistribution of heat. The heat redistribution results in the Northern Hemisphere experiencing warming and local heat storage that outweighs the local radiative forcing for the RCP scenarios that include sulfate aerosol forcing. This mismatch results in a negative hemispheric climate feedback parameter λ_{NH}, i.e., the hemispheric redistribution of heat acts as a positive feedback and enhances the amount of warming seen per unit of radiative forcing. As the century continues, the reduction of sulfate aerosols, slight decrease in cross-equatorial transport, and hemispheric symmetry in the additional top-of-the-atmosphere heat fluxes pushes the system towards a more symmetrical thermal response. The transition towards more symmetrical hemispheric thermal response can be seen through the asymptotic behaviour of $\lambda_{NH}, \lambda_{SH}$ to equilibrium values (Fig. 4.8) and the similar increases in the hemispheric thermal response in the latter part of the century (Fig. 4.6).

Understanding the impact of the hemispheric climate feedback parameter

Our results in diagnosing the hemispheric climate feedback parameters λ_{NH} and λ_{SH} indicate that local differences in radiative forcing can induce large differences in regional climate feedback parameters. These differences in the regional climate feedback parameters are significant as they occur with a highly simplified atmospheric physics and cloud representation. Clouds within the UVic ESCM are represented as a static albedo field, and the model does not permit any cloud feedbacks, which are instrumental in setting regional feedbacks (Ceppi and Gregory, 2017; Andrews and Webb, 2018) as well as responsive to the pattern of surface warming (Armour et al., 2013; Rugenstein et al., 2016) in more complex models.

As cloud representation in high-complexity climate models becomes more complicated in the CMIP6 models, the global climate feedback parameter has been found to play a larger role in setting the inter-model TCRE spread (Williams et al., 2020). Diagnosing the role of regional or hemispheric climate feedback parameters may potentially show even more inter-model differences within the CMIP6 models. There is also the potential to expand

the theory presented in this chapter to include linear radiative kernel decompositions of the climate feedback parameter (Armour et al., 2013), to see how the latest generation models compare in their regional longwave and shortwave cloud feedbacks and how the spreads in these regional behaviours ultimately imprint on the larger-scale response uncertainty.

While the atmospheric simplifications within the UVic ESCM restrict our ability to compare the climate feedback results to those found in other models or observations, our results do have implications for how the effective TCRE is understood in relationship to the CO_2-only TCRE diagnostics. When expanding the TCRE carbon budgets to account for non-CO_2 forcing elements, the TCRE is often scaled by the ratio of CO_2-induced and non-CO_2-induced radiative forcing (Matthews et al., 2021). Our results in this chapter concerning the feedback of equatorial heat transport to radiative forcing asymmetries lead us to caution against using such a linear approach towards carbon budgets, as the existence of such feedbacks are inherently nonlinear. Thus, a linear scaling of the TCRE to non-uniform radiative forcing elements runs the risk of misrepresenting how the climate dynamics set carbon budgets.

4.7.2 Hemispheric carbon responses dominated by ocean ventilation symmetry

In contrast to the asymmetric hemispheric thermal responses, the hemispheric carbon responses behave similarly to the global carbon response. Differences in the hemispheric carbon responses are linked to the magnitudes of the carbon inventory changes. The terrestrial system plays a relatively small role in setting the hemispheric carbon response, particularly in the Southern Hemisphere where the terrestrial area is small (Fig. 4.11). Accordingly, the dominance of the Southern Hemisphere in carbon response is a result of its larger ocean fraction. There are, however, regional differences in the ocean carbon uptake that are hidden by the choice of hemispheric averaging, as carbon uptake is highly localised in the North Atlantic in the Northern Hemisphere and more evenly spread in the Southern Hemisphere (Fig. 4.12a,b) (Frölicher et al., 2015). Ocean biology induces a slight hemispheric asymmetry in ocean dissolved inorganic carbon through relatively small changes in regenerated carbon. The asymmetry comes from the different pathways of deep water transport in the Northern and Southern Hemispheres as changes in regenerated carbon are largely caused by a slowdown in the overturning circulation rather than by changes in the marine biology.

4.7.3 Limitations of the UVic ESCM diagnostics

The hemispheric results within this chapter are heavily influenced by the model used. As an intermediate complexity Earth System model, the UVic ESCM has some necessary simplifications; however, these simplifications alter major contributors to the hemispheric global and climate response. We present some caveats associated with these simplifications here.

One major caveat of the UVic ESCM thermal responses involves the direction and pathways of cross-equatorial heat transport. While the UVic ESCM shows a global northward transport of heat, the transport is dominated by the atmosphere; the model diagnostics show an initial southward cross-equatorial ocean heat transport. In contrast, the ocean acts to transport heat northwards across the equator (Marshall et al., 2014a). Studies of heat transports during the historical period using more complex climate models show that the inclusion of aerosols drives an anomalous northwards oceanic transport (Paynter and Frölicher, 2015; Shi et al., 2018; Irving et al., 2019). As the heat transport within the UVic ESCM contrasts so strongly with these other lines of evidence, we take the thermal response in this chapter to be more illustrative of the UVic ESCM dynamics than prescriptive of future regional thermal responses. The dominance of the atmospheric cross-equatorial heat transport within the UVic ESCM is likely to make the air-sea heat fluxes unrealistic, which has further implications for how the UVic ESCM represents the ocean's response to radiative forcing and surface temperature changes.

The inclusion of land use changes in the RCP scenarios complicates the carbon response diagnostics as some terrestrial carbon changes relative to the preindustrial will be forced rather than a response to a combination of warming and elevated atmospheric CO_2. While changes in the terrestrial carbon inventories are relatively limited in the UVic ESCM, this is likely a result of the model setup, as there are large uncertainties in the response of the terrestrial system to carbon emissions (Friedlingstein et al., 2014; Arora et al., 2019). Higher-complexity CMIP models continue to show a wide spread in terrestrial carbon-concentration and carbon-climate feedbacks that contributes towards the inter-model TCRE spread (Jones and Friedlingstein, 2020). This version of the UVic ESCM does not include permafrost feedbacks, which can increase the proportion of surface warming to cumulative emissions and provide greater hemispheric carbon contrasts (MacDougall and Friedlingstein, 2015).

Interestingly, one of the main reasons for the UVic ESCM's widespread usage in the TCRE literature is its ability to re-create the first-order response to emissions. While

our results also find that the TCRE results from existing literature hold for these model experiments, the diagnostics of regional thermal and carbon components indicate that this first-order response overlooks some less desirable model behaviours. Our results suggest that intermediate-complexity models are better suited for the original TCRE understanding using global fluxes and inventories and likely have limited scope for understanding more complex and regional drivers of such behaviour.

4.7.4 Potential to extend theory to more complex regions and models

Despite the caveats inherent to the model used in this chapter, the framework presented can be used to understand regional thermal and carbon TCRE contributions within other Earth system models, perturbed parameter experiments, or different regional partitions. Our work argues that there are complex regional asymmetries that contribute to the well-defined global TCRE behaviour and that these contributions should be considered in discussions about the drivers of the TCRE, both in terms of its general behaviour and its uncertainties.

One potential continuation of this work involves using different separations of the Earth system. in particular, a separation of the Earth System to isolate the Southern Ocean region could be helpful in more clearly indicating how the Southern Ocean heat and carbon uptake relate to current and future warming. Additionally, the separation of the globe into areas of strong feedbacks (i.e., the high latitudes) versus the low latitude/subtropical areas could be useful in comparing the heat/carbon storage impacts of the low latitudes vs the high-latitude feedback impacts, both on the amount of warming projected in individual models but also in comparing responses across a suite of models. The mathematics of the theory (Section 4.2) for these experiments would remain largely the same, although there may be complications in closing budgets depending on the boundaries of each region and the model grid.

Additionally, as mentioned before, this work can also be extended to higher-complexity models. By using higher-complexity models, there is the potential to isolate aspects of the climate feedback parameter that are not represented within the UVic ESCM, and the method could be expanded to include linear radiative kernel results for each region's climate feedback parameter. The use of more complex models may also allow for further consideration into some of the more complex carbon-climate feedbacks, depending on the terrestrial model components. For this expansion, the main difficulties would lie in the diagnostics of some variables, such as the atmospheric heat content; as the UVic ESCM uses a slab atmosphere this diagnostic was relatively easy, but when working with a 3-dimensional

atmosphere these diagnostics will necessarily involve more variables and careful consideration. However, the mathematical building blocks of our theory would remain applicable regardless of region, model complexity, or even forcing setup.

In summary, the near compensation between the thermal response and carbon response in the warming response to emissions is a consequence of different hemispheric contributions. Changes in the hemispheric carbon responses are relatively similar, and the small differences are largely the result of the larger ocean surface area in the Southern Hemisphere. The RCP scenarios induce a strong hemispheric asymmetry in the non-CO_2 radiative forcing, which enhances the asymmetry in the hemispheric thermal responses. Under this non-uniform forcing, the Northern Hemisphere thermal response decreases as the radiative forcing ratio increases, while in the Southern Hemisphere the thermal response increases. The radiative forcing and thermal responses combine to create large differences in the climate feedback parameter. As the asymmetry in the hemispheric radiative forcing declines, the hemispheric thermal tendencies become more similar and the climate feedback parameters stabilise. However, there remain persistent offsets in the climate feedback parameter and carbon reservoir sizes that continue to affect the contributions of each hemisphere towards the global response to cumulative emissions. While the regional results of the UVic ESCM are limited by the dynamic shortcomings of an intermediate-complexity Earth system model, our results argue that there are important asymmetries that contribute to the global response to emissions and that can be studied in a variety of setups.

Chapter 5

Controls on climate change mitigation through Carbon Dioxide Removal

We now consider how carbon dioxide removal (CDR) impacts the relationship between emissions and surface warming. This chapter explores how the response to CDR technologies compares to the response to emissions, with an emphasis on the path-dependency of CDR signals. We use a large perturbed parameter ensemble of an idealised ocean-atmosphere model forced under various combinations of background emissions, CDR timelines, and CDR methods. With this large ensemble we investigate how different physical climate parameters interact with negative emissions technologies to produce mitigation signals and compare the evolution of thermal and chemical signals from different CDR timelines.

We begin in Section 5.1 by providing an overview of the current understanding of the response to negative emissions. We introduce an extension of the Gnanadesikan model ensemble in Section 5.2, with a full description in Appendix A, and provide a description of the combination of climate parameters and forcings used to make up the large (> 700 member) perturbed parameter ensemble. We consider how the surface temperature mitigation from CDR is realised under two timelines with the same amount of cumulative carbon drawdown in Sections 5.3 and 5.4. We then consider how the carbonate chemistry co-benefits of CDR are path dependent in Section 5.5, and finish with a discussion of the results and their connection to the wider context of CDR decision-making in Section 5.6.

5.1 Carbon dioxide removal as a potential climate change solution

Even as recent national and international agreements have been created to limit emissions in hopes of maintaining surface warming below 2 °C, and ideally below 1.5 °C, global carbon emissions continue to increase. At current emission rates, the estimated remaining carbon budgets for the Paris Agreement limits could be exceeded within the next two decades (Goodwin et al., 2018; Millar et al., 2017; Tokarska and Gillett, 2018). Within scientific and political circles, artificial CDR has been proposed to mitigate or even reverse some of the climate impacts associated with continued emissions (IPCC, 2018).

CDR occurs naturally in the Earth system, as carbon reservoirs slowly accumulate and sequester carbon through the burial of organic matter or rock weathering. Fossil fuels that create carbon emissions are themselves products of natural CDR over millions of years. As climate processes sequester carbon much more slowly than fossil fuel use emits it, CDR schemes for present-day climate change requires creating artificial carbon sinks or enhancing natural carbon sinks to take up carbon on shorter timescales.

A variety of CDR schemes have been proposed to maintain surface warming around 1.5-2 °C (IPCC, 2018). Carbon capture and storage (CCS) uses man-made materials to capture CO_2 from flue gases or the atmosphere, after which the CO_2 can be extracted and stored in geological reservoirs (Fuss et al., 2018). Bioenergy carbon capture and storage (BECCS) proposes that bioenergy crops be grown to capture atmospheric CO_2: their combustion for energy would then require similar mechanisms for CO_2 capture and storage as for CCS. Other CDR methods involve nature-based solutions such as reforestation, ocean alkalinity addition, ocean fertilisation, and enhanced rock weathering (EW). Current analyses suggest that large-scale CDR will require a portfolio of technologies, as each method has logistical limitations (Smith et al., 2016).

It has been suggested that large-scale CDR will need to draw down around 2 PgC yr^{-1} by 2050 in order to meet the Paris climate goals (Rockström et al., 2017). However, the drawdown of carbon through engineered methods will require large amounts of resources. An average drawdown of 1 PgC yr^{-1} using BECCS, for instance, would require cropland area about the size of India for the production of feedstock crops (IPCC, 2018). For EW, the drawdown of 1 PgC yr^{-1} would require around 24 Gt yr^{-1} of silicate materials and deployment of rock dust over the vast majority of global croplands (Beerling et al., 2020).

5.1.1 Thermal and carbon responses associated with carbon dioxide removal

As the addition of carbon impacts both the thermal and carbon responses (Fig. 1.2), so too will CDR schemes alter both the carbon balance between reservoirs and the radiative forcing in Earth's energy budget. Negative emissions studies have explored the symmetry in the response to carbon emissions and carbon removal in idealised peak-and-decline scenarios. These studies find a range of asymmetries in the Earth system response, including hysteresis in sea level rise (Tokarska and Zickfeld, 2015), carbon feedbacks (Schwinger and Tjiputra, 2018), and carbon reservoir magnitudes (Jones et al., 2016; Keller et al., 2018a). Slower components to the thermal response, such as thermosteric sea level rise from carbon emissions, continue under negative emissions which induces hysteresis in sharp emission drawdown scenarios. While the specifics of CDR technologies are not present in these studies, these asymmetries lead to negative emissions being less effective in reducing temperatures than the equivalent positive emissions are in raising temperatures (Jeltsch-Thömmes et al., 2020; Zickfeld et al., 2021). Thus, while the TCRE (Chapter 4) may be useful for determining remaining carbon emissions budgets, its potential to diagnose necessary carbon drawdown for overshoot scenarios is limited. While these idealised scenarios are useful for inter-model comparisons and hysteresis studies, their results include the short-term response to CDR alongside the longer timescale adjustments to previous emissions. Thus the response seen in these scenarios is not a pure comparison between two opposing forcings, but rather a path-dependent description of the response to emissions. As CDR must necessarily occur after the bulk of CO_2 emissions, the results of these model studies are still useful in understanding the potential efficacy of CDR.

Idealised studies on negative emissions consider only removal of atmospheric carbon and do not include any assumptions for how the carbon is removed. However, CDR methods provide additional impacts on the Earth system through the mechanisms by which they draw down carbon; these impacts can be particularly strong for methods that engineer natural carbon sinks to become more effective. For instance, EW involves the addition of finely-ground rocks such as basalt on agricultural or forested soils to draw down atmospheric CO_2; EW can be considered an terrestrial-based analogue to ocean alkalinisation. When in contact with water and CO_2, the basalt produces bicarbonate as well as cations and silica (Schuiling and Krijgsman, 2006). For forsterite (Mg_2SiO_4) this reaction can be simplified as:

$$4CO_2 + 2H_2O + Mg_2SiO_4 \rightarrow 4HCO_3^- + 2Mg^{2+} + SiO_2. \quad (5.1)$$

The bicarbonate end product can be transported by rivers into the ocean, adding carbon to the surface ocean but also increasing the pH of surface waters through the bicarbonate's negative charge. The addition of alkalinity through EW has the potential to reverse ocean acidification (Taylor et al., 2016), impact ocean biology and calcite precipitation (Renforth and Henderson, 2017), and decrease the ocean buffer factor. Additional terrestrial benefits of enhanced rock weathering include increased crop yields and enhanced soils (Beerling et al., 2018). Less desirable side effects such as enhanced ocean toxicity for certain organisms may also result from the flux of weathered particles into the ocean. For CDR methods that are limited in their carbon impact, the balance of co-benefits and side-effects of CDR may be the determining factor for how or where these methods are implemented.

5.1.2 Challenges for climate model studies of negative emissions

Large scale experimentation on CDR will require international agreement and deep ethical consideration; in contrast, climate models provide a low-risk means by which global CDR experiments may be analysed. With climate models, negative emissions may be represent a variety of methods, and studies can consider the efficacy, feedbacks, co-benefits, and side-effects of removal methods.

High-complexity climate models are created as a best "current guess" for how the climate system operates; as such, they provide realistic representations of the complex Earth system but also provide insight into current uncertainties in the Earth system. Although knowledge of the climate system has increased substantially, some areas remain highly uncertain. One such area is the climate feedback parameter λ, an empirical measure that describes how radiative forcing is partitioned between heat uptake and the Earth's radiative response (4.5). The current generation of CMIP6 models shows a shift towards a lower climate feedback parameter, indicating a higher sensitivity to emissions, related to new representations of cloud feedbacks (Zelinka et al., 2020). Uncertainty in the climate feedback parameter is the main source of uncertainty for the CMIP6 response to emissions; in contrast, uncertainty in the CMIP5 models' response to emissions was determined mainly by terrestrial carbon uptake (Williams et al., 2020). While the realism of a low climate feedback parameter λ is under debate (Zhu et al., 2020; Tokarska et al., 2020; Goodwin and Cael, 2021), it presents a low-probability, high-impact climate state (Forster et al., 2020) that warrants investigation and understanding.

Although high-complexity climate models represent the current best approximation for how different aspects of the climate system interact, their high computational require-

ments result in relatively few configurations being run. The range of responses spanned by high-complexity models may be artificially narrow. The use of lower-complexity models can complement high-complexity model results, as lower-complexity models can more easily cover wide parameter spaces. Climate emulators have recently been used to create estimates for climate responses to forcing, most significantly in the recent IPCC AR6 report (Meinshausen et al., 2011b; Smith et al., 2018; Goodwin et al., 2020). Climate emulators contain a few fundamental equations to represent climate processes; in lieu of high resolution and complex numerical schemes, they contain a number of parameters that can be varied. The results of climate emulators and their parameter values are generally retained or thrown out depending on how well they replicate well-understood aspects of the climate system such as historical temperatures or ocean heat content changes. While emulators have been helpful for providing an independent means of exploring the likely range of warming behaviours, the tuning inherent in their selection process means that for process-based understanding limits are not often reached. In contrast, large ensembles of lower-complexity models may contain results that do not agree with observations, but they can be used to test physical controls and limits of model behaviours (Katavouta et al., 2019).

One limitation of current climate model setups is the difficulty in separating the response to emissions from the response to CDR. The shared socioeconomic pathways used to force the CMIP6 models are generated by integrated assessment models, which determine the amount of emissions or carbon drawdown associated with social development pathways. The modelled responses therefore contain concurrent responses to positive emissions and carbon drawdown. The removal of carbon within these pathways is also conducted through idealised negative emissions, preventing any comparison of CDR methods. The Carbon Dioxide Removal Intercomparison Project provides protocols for comparing emissions pathways with specific types of carbon drawdown such as afforestation or ocean alkalinisation, and upcoming ensemble results can better provide spreads on the efficacy of carbon drawdown methods and their co-benefits and side-effects (Keller et al., 2018b).

5.1.3 Driving questions and objectives

The use of CDR technologies is likely necessary to maintain global average surface warming below 2 °C. The use of such technologies requires a thorough understanding of how the climate system responds to negative emissions, particularly in how the response differs from the response to CO_2 emissions. While there is a solid empirical relationship between the amount of carbon emitted and the amount of surface warming experienced, negative emis-

sions are generally less effective at lowering surface temperatures than positive emissions are at raising them. This asymmetry seen in negative emission studies may be exacerbated by the specific type of CDR used. The relationship between CDR scenarios and the resulting thermal and chemical impacts is important for understanding the level of asymmetry inherent in the Earth system response to external forcing, which will impact how CDR can be most effectively implemented. As the climate system still has areas of high uncertainty, these uncertainties must also be included in any diagnostics to determine sensitivities within the response to CDR.

In this chapter we explore the response to different CDR timelines under a wide physical climate parameter space. We focus on the response up to 200 years post-emissions and post-CDR. We aim to answer the following questions:

1. How constrained are mitigation signals by the cumulative amount of drawdown?
2. How are the thermal and chemical signals from CDR sensitive to the CDR method?
3. How are mitigation signals from CDR sensitive to uncertainties in the climate system?

We hypothesise that, similarly to the response to carbon emissions, the response to CDR will be proportional to the total amount of drawdown and similarly path-independent, although the two responses will have different magnitudes. We hypothesise that the type of CDR will impact the chemical signals rather than the thermal signals for moderate rates of carbon drawdown. Lastly, we hypothesise that uncertainties in the climate system (here, the overturning strength and climate feedback parameter) will mainly impact the thermal response and temperature reductions on our timescales of interest, in line with the results of Katavouta et al. (2019) under positive emissions.

In this study, we create a large ensemble using a computationally-efficient, idealised ocean-atmosphere model that includes a representation of global overturning circulation and tracers for heat, carbon, and alkalinity (Section 5.2). The ensemble is created through variations in overturning parameters and climate sensitivity, as well as the external forcing from CO_2 emissions and CDR methods. The end ensemble spans a large uncertainty space with over 700 perturbed parameter setus and 3000 model runs. We analyse the ensemble response to emissions and compare the response to emissions to the response to CDR. We analyse how the peak impacts of CDR on atmospheric CO_2 levels, surface temperatures, and surface ocean pH, vary according to the parameter space. We also test how constrained these responses are by the integral amount of carbon drawdown, and how these constraints can be understood through thermal and carbon responses. We finish the work by discussing

the implications of the ensemble response to CDR, both in terms of the importance of understanding certain aspects of climate uncertainty and in terms of supporting decisions about how CDR will be implemented in the real world.

5.2 The Gnanadesikan model ensemble

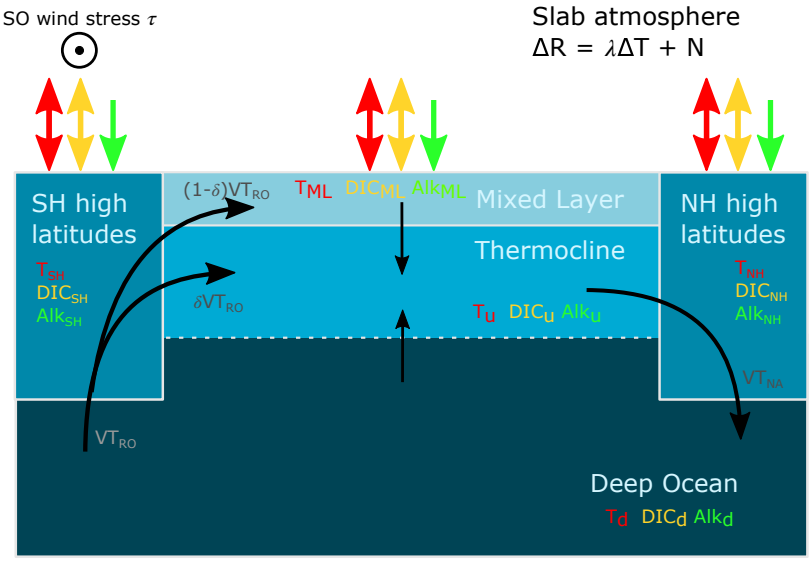

Figure 5.1: Schematic of the Gnanadesikan model, as expanded upon by Katavouta et al. (2018). Boxes indicate the ocean boxes within the model, with black arrows indicating the water transformations solved for. Fluxes (arrows) and tracer concentrations (coloured text) are shown for temperature (red), carbon (yellow), and alkalinity (green).

The idealised ocean-atmosphere model used in this chapter is the Gnanadesikan model of the ventilated thermocline (Gnanadesikan, 1999), which has been expanded to include time dependence, a slab atmosphere, carbon tracers, and an iterative ocean carbonate chemistry solver (Marshall and Zanna, 2014; Katavouta et al., 2019) (Fig. 5.1). The model consists of a slab atmosphere and five well-mixed ocean boxes, which represent the tropical and subtropical mixed layer, tropical and subtropical thermocline, Northern Hemisphere high latitudes, Southern Hemisphere high latitudes, and deep ocean. Dense water formation occurs in the Northern Hemisphere high latitudes, representing the creation of North Atlantic Deep Water. Upwelling of dense water to the upper ocean boxes occurs in the Southern Hemisphere high latitudes through a representation of the residual overturning, which includes the impacts of the Ekman and mesoscale eddy transports in the Southern Ocean.

Diapycnic transfer brings water from the deep ocean to the thermocline. Ocean transports are constrained by global ocean mass conservation, although the mass in the thermocline and deep ocean boxes are allowed to adjust as the overturning circulation evolves. The model includes tracers for potential temperature, dissolved inorganic carbon (DIC), and alkalinity; these tracers follow standard tracer equations, and pH and carbonate chemistry species fractions are calculated in each box using the iterative algorithm from Follows et al. (2006).

Coupling between the surface ocean boxes (the tropical mixed layer and high-latitude boxes) and the slab atmosphere allows for the exchange of heat and carbon. Global atmosphere-ocean heat and carbon fluxes are calculated using the area-weighted temperatures and DIC concentrations. Heat and carbon fluxes are partitioned unevenly among the three surface ocean boxes, with the high latitude boxes receiving 40% of the global flux each and the mixed layer box receiving 20% of the global flux. Full descriptions of the balances used to solve for the volume transports, heat transports, and carbon transports within the model can be found in Appendix A.

5.2.1 Model parameter space

One major advantage of this model setup is that the runs are computationally cheap, which allows for the creation of large ensembles that can span wide ranges in parameter uncertainty. The parameter uncertainty for the ocean model focuses on the constants that set the overturning strength and follows a similar protocol to Katavouta et al. (2019) (Table 5.1). The strength of the Northern Hemisphere sinking and the growth of the thermocline box are partially set by a constant Southern Ocean wind stress τ. The wind stress term ranges from 0.06 N m^{-2} to 0.14 N m^{-2}, which corresponds to preindustrial overturning strengths of 15 to 33 Sv. The model also includes a subduction fraction term, δ, which accounts for the proportion of upwelled Southern Ocean waters that are subducted into the thermocline rather than brought into the mixed layer through the residual overturning. In this work, δ ranges from 0.3 (i.e., 30% of the residual overturning is fluxed into the thermocline) to 0.8 (80% of the residual overturning is fluxed into the thermocline).

In addition to these ocean parameters, the ensemble includes variations in the climate feedback parameter λ. The climate feedback parameter λ controls the strength of the radiative response, which describes the re-emission of energy into space when the climate system undergoes a radiative forcing. In the Gnanadesikan model, the atmosphere follows a zero-dimensional heat balance, in which the radiative response is directly proportional to

Variable	Variable name	Units	Min.	Max.	Sampling interval
δ	Subduction fraction	N/A	0.3	0.8	0.1
τ	Southern Ocean wind stress	N m^{-2}	0.06	0.14	0.02
λ	Climate feedback parameter	W m^{-2} K^{-1}	0.6	2.0	0.2

Table 5.1: Perturbed variables and ranges used to create the Gnanadesikan ensemble.

the surface temperature increase ΔT_{air} (Gregory et al., 2002, 2015):

$$\Delta R = \underbrace{\lambda \Delta T_{air}}_{\text{radiative response}} + \underbrace{\mathcal{N}}_{\text{heat uptake}} . \tag{5.2}$$

The climate feedback parameter λ, with units W m^{-2} K^{-1}, is constant in each model run and ranges from 0.6 W m^{-2} K^{-1} to 2.0 W m^{-2} K^{-1}. This setup covers the range of climate feedback parameters diagnosed in Zelinka et al. (2020) from CMIP5 and CMIP6 models.

5.2.2 Model setups with CO_2 emissions and carbon dioxide removal

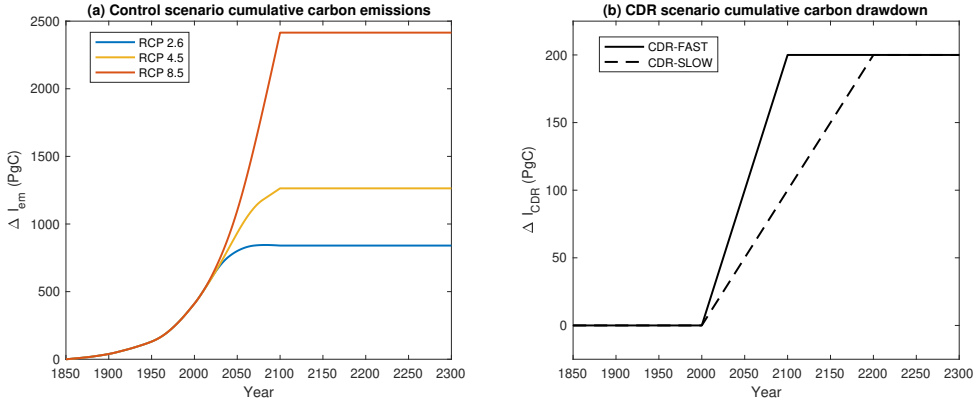

Figure 5.2: (a) Control Historical (1860-2005) and RCP (2005-2100) CO_2 cumulative emissions (ΔI_{em}, in PgC) used to force the Gnanadesikan model. Forcing starts at 1860 after a spin-up to equilibrium conditions under an atmospheric CO_2 concentration of 280 ppm. (b) Cumulative carbon drawdown (ΔI_{CDR}, in PgC) for the CDR-FAST and CDR-SLOW timelines.

Each model setup is run to equilibrium with a preindustrial atmospheric CO_2 mixing ratio of 280 ppm. Each setup is then run under a variety of emissions scenarios, which involve a combination of carbon emissions and CDR schemes (Table 5.2). For the control

runs, which involve carbon emissions but no CDR, the models are forced with historical emissions from years 1860 to2005, then carbon emissions consistent with either Representative Concentration Pathway (RCP) 2.6, 4.5, or 8.5 to year 2100 (Meinshausen et al., 2011a). Emissions are set to zero immediately thereafter and the system is run for another 300 years to allow for further climate adjustment. Under the control runs, 13 out of the 720 parameter/emissions setups show a complete reversal in the high-latitude deep water formation and upwelling of deep water into the Northern Hemisphere high-latitude box instead. We eliminate any runs in which this occurs so that our range of results is consistent with projected overturning changes, which can show a marked weakening but show low confidence for a complete shutdown (Weijer et al., 2020).

Name	Description	Type of drawdown	Years of drawdown	Rate of drawdown
Control	RCP 2.6, 4.5, or 8.5 CO_2 emissions	None	N/A	N/A
CCS-FAST	Control + 100 years CO_2 drawdown	Removal of atmospheric CO_2	2000-2100	2 PgC yr^{-1}
CCS-SLOW	Control + 200 years CO_2 drawdown	Removal of atmospheric CO_2	2000-2200	1 PgC yr^{-1}
EW-FAST	Control + 100 years CO_2 drawdown	Addition of ocean CO_2 and Alk	2000-2100	2 PgC yr^{-1}
EW-SLOW	Control + 200 years CO_2 drawdown	Addition of ocean CO_2 and Alk	2000-2200	1 PgC yr^{-1}

Table 5.2: Description of forcing scenarios used for the Gnanadesikan ensemble.

In addition to the control runs, we create forcing setups to consider the impacts of CDR. The first setup removes atmospheric CO_2 from the system entirely, representing CCS. The second, representing EW, removes CO_2 from the atmosphere and deposits the carbon uniformly into the surface ocean as equal parts dissolved inorganic carbon (DIC) and alkalinity. The EW scenario is consistent with the theory that enhanced weathering creates bicarbonate ions, which are then deposited by rivers into the ocean and stored on timescales of millennia (5.1). By including representations for EW, we explore whether the secondary impacts of CDR method have the potential to alter the Earth system response.

For both CDR methods, we test the climate system's sensitivity to the timing of CDR through two idealised scenarios. The "fast" timeline (CCS-FAST, EW-FAST) begins at year 2000 and draws down atmospheric CO_2 at a rate of 2 PgC yr^{-1} for 100 years. While this starting date is not applicable to real-world emissions, this setup provides a scenario that, for CCS, is also comparable to reducing carbon emissions. We include a scenario (CCS-SLOW, EW-SLOW) in which CDR is conducted at half the rate of that in the "fast" timeline, but for twice the amount of time (i.e., 1 PgC/yr from 2000 to 2200). As the two scenarios have the same cumulative carbon drawdown by 2200, the results can be compared to see how the climate system continues to adjust post-drawdown, and to explore how constrained the results are by the cumulative amount of drawdown. While these rates of

drawdown are low compared to the emissions in most of the RCP scenarios, they sit between the estimated drawdown potentials of CCS and EW (Beerling et al., 2018; Fuss et al., 2018; Beerling et al., 2020).

5.3 Model responses to forcings

The model parameter space creates a wide range of temperature responses to the control emissions scenarios (Fig. 5.3). The range of surface warming is similar to the range of surface warming seen in the CMIP5 models for the extended RCP scenarios, while the median surface warming is similar to the average RCP warming (Collins et al., 2013). Thus, while the Gnanadesikan model is highly simplified and does not include any representation of the terrestrial system, its basic response to emissions is still comparable to that of more complex models.

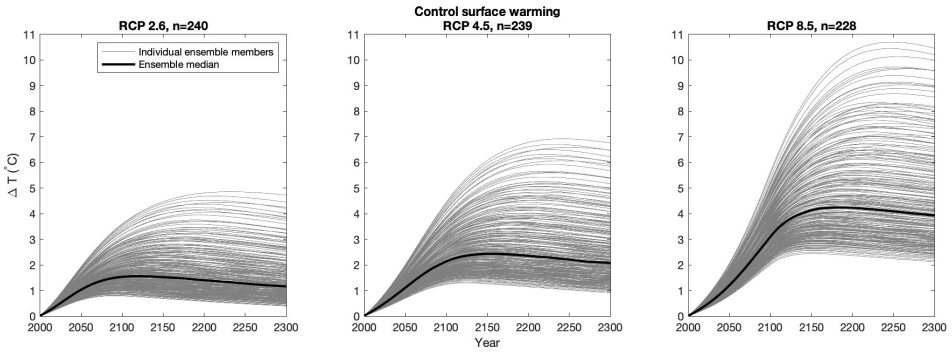

Figure 5.3: Control RCP surface temperature responses for all model setups (grey lines) and the ensemble median response (thick black lines). The number of runs conducted successfully under each forcing scenario is denoted by n.

One important aspect of the ensemble response is that almost all of the ensemble members experience continued surface warming after emissions cease at 2100. The continued warming is greatest for the ensemble members on the upper end of the warming ranges, both in terms of amount of continued warming and in terms of the time period in which warming continues. The continued warming seen in the Gnanadesikan model differs from most models in the zero emissions commitment intercomparison (MacDougall et al., 2020), although continued warming is an aspect of some high-complexity models (Williams et al., 2017b). However, MacDougall et al. (2020) find that the strength and behaviour of terrestrial carbon uptake plays a major role in determining the amount of warming post-emissions; as the Gnanadesikan model has no terrestrial representation, we retain the ensemble with the

caveat that the absolute behaviours may not be directly comparable to the behaviours of more complex systems. Instead, we focus on understanding the response to CDR relative to the response to emissions and the controlling processes within the model, as the controlling processes within the ensemble are likely to remain relevant for other climate models and the real world.

5.3.1 The climate response to carbon dioxide removal

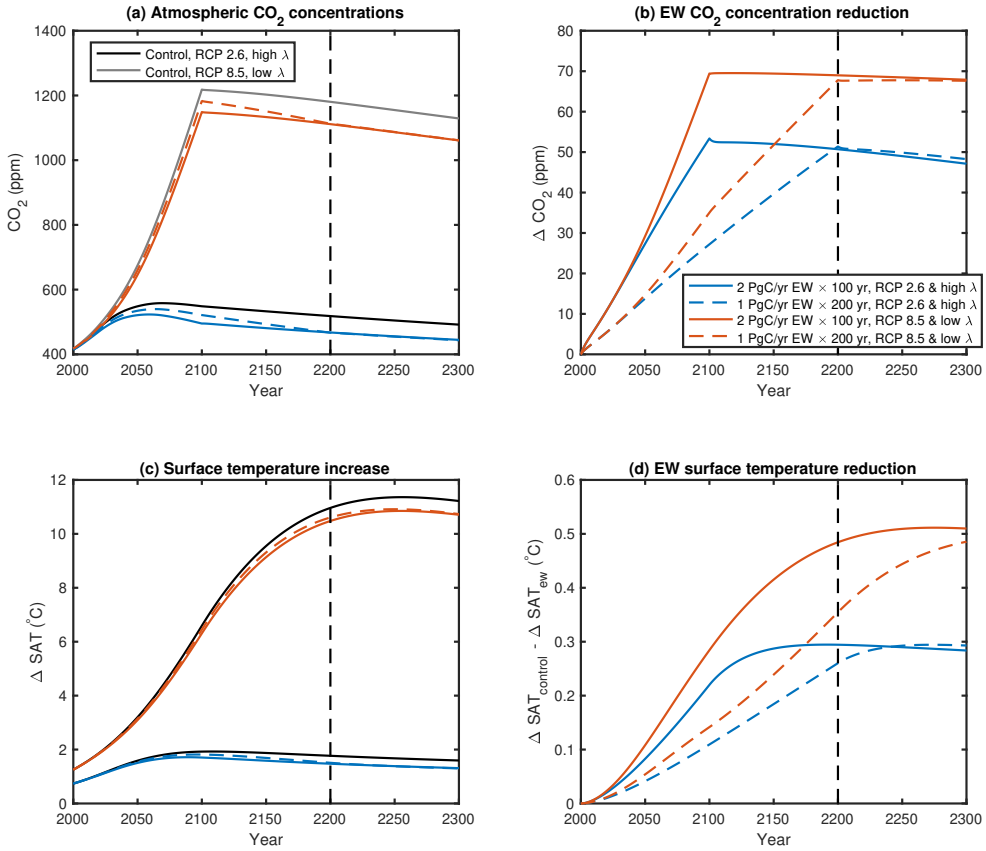

Figure 5.4: Illustration of atmospheric CO_2 concentrations, surface temperature increases, and impacts of CDR for two model ensemble members forced with enhanced weathering along the CDR-FAST and CDR-SLOW timelines.

Unlike previous studies on negative emissions, the scenarios here involve concurrent CO_2 emissions and drawdown. To explore the response to CDR, we consider the difference in variables between the runs with CDR and the control runs.

Although the ensemble has a wide range of temperature responses to emissions, the responses to CDR are more similar across ensemble members. For simplicity, we illustrate the response to CDR using two ensemble members that are highly dissimilar in terms of their warming behaviour: one with a high climate feedback parameter (low climate sensitivity) forced by RCP 2.6 emissions, and one with a low climate feedback parameter (high climate sensitivity) forced by RCP 8.5 emissions (Fig. 5.4). For simplicity of visualisation, the overturning parameters and CDR method are identical.

Relative to the control scenarios, enhanced weathering lowers atmospheric CO_2 levels (Fig. 5.4a,b) and average surface temperatures (Fig. 5.4c,d). Given that atmospheric CO_2 and temperature increases for the two control scenarios are an order of magnitude different, the CO_2 and temperature responses to CDR are relatively similar. A drawdown of 200 PgC through the EW scheme results in an atmospheric CO_2 reduction between 48 to 70 ppm; this reduction is highest at the final year of CDR and is near-constant during the following centuries. The reduction in atmospheric CO_2 is highly constrained by the integral amount of CDR: at year 2200, when the EW experiments converge in terms of their carbon drawdown, atmospheric CO_2 levels are approximately equal.

When considering the impacts of CDR on surface temperatures, the response no longer follows the integrated carbon drawdown. Similar to the Gnanadesikan model response to emissions, temperature mitigation is maximised after the EW scheme ends. At year 2200, when the CDR experiments have been forced with the same amount of weathering, the temperature mitigation is higher in EW-FAST than in EW-SLOW. The temperature adjustment period can be substantial before the system reaches an integrally-constrained equilibrium temperature, lasting for over a century as visible in the low climate feedback parameter, RCP 8.5 response.

5.3.2 Relationship between post-emissions warming and post-drawdown cooling

Continued warming post-emissions and continued cooling post-drawdown are robust features within the Gnanadesikan model ensemble. To see if and how these properties are related, we compare the year of peak warming from the control runs with the year of peak temperature reduction from the EW scheme (Fig. 5.5, left panel). The year of peak mitigation from CDR is positively correlated with the year of peak warming and occurs slightly later, indicating that the equivalent adjustment processes post-emissions are slower for CDR than they are for emissions. The offset between these peaks decreases as the back-

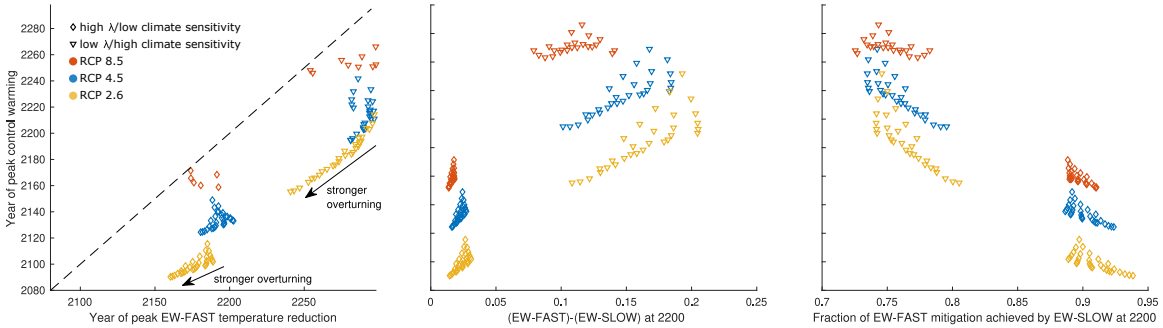

Figure 5.5: Scatterplots comparing the year of peak warming under the control runs with (left) the year of peak temperature mitigation from the EW-FAST scenario; (centre) the absolute temperature mitigation difference between EW-FAST and EW-SLOW at 2200 (5.3), in units °C; and (right) the relative temperature mitigation difference between EW-FAST and EW-SLOW at 2200 (5.4). Colours and markers show different emissions scenarios and climate feedback parameters, arrows indicate the direction that these relationships move in stronger overturning regimes.

ground emissions scenario increases, is approximately independent of the climate feedback parameter, and decreases as overturning strengthens.

The lags between CDR timelines at their year of convergence can be measured in two ways: the absolute temperature difference at year 2200,

$$\Delta T_{CDR-FAST} - \Delta T_{CDR-SLOW}, \tag{5.3}$$

or the temperature difference at year 2200 relative to the CDR-FAST scenario,

$$(\Delta T_{CDR-FAST} - \Delta T_{CDR-SLOW})/\Delta T_{CDR-FAST}. \tag{5.4}$$

Comparing both the absolute and relative temperature differences in EW runs against the year of peak control warming (Fig. 5.5, centre and right panels) results in similarly strong correlations. Models with longer periods of continued warming exhibit larger absolute and relative differences when comparing EW temperature reductions at the year the forcings have the same integrated drawdown. Thus the adjustment processes/time lag to peak temperature signals in the Gnanadesikan model work in a consistent way, delaying peak warming as well as peak mitigation, and maintaining this adjustment for slower mitigation scenarios. The post-forcing adjustment seen in the Gnanadesikan model is consistent across different diagnostics, indicating that the pace of adjustment towards the equilibrium is similar across different forcing setups.

5.4 Physical controls on the path-dependency of temperature mitigation signals

We continue by exploring how the model parameter space impacts the path-dependency of temperature mitigation. In order to better visualise runs with both low and high climate feedback parameters, we use the relative temperature difference defined in (5.4). In addition to being easier to visualise, the relative temperature difference is relevant to real-world CDR policy decisions, as regardless of the actual temperature mitigation possible from a specific technology, it is likely that any implementation will try to maximise benefits. As in the previous section, we illustrate the controls using the EW setup; models forced with the CCS setup show similar responses.

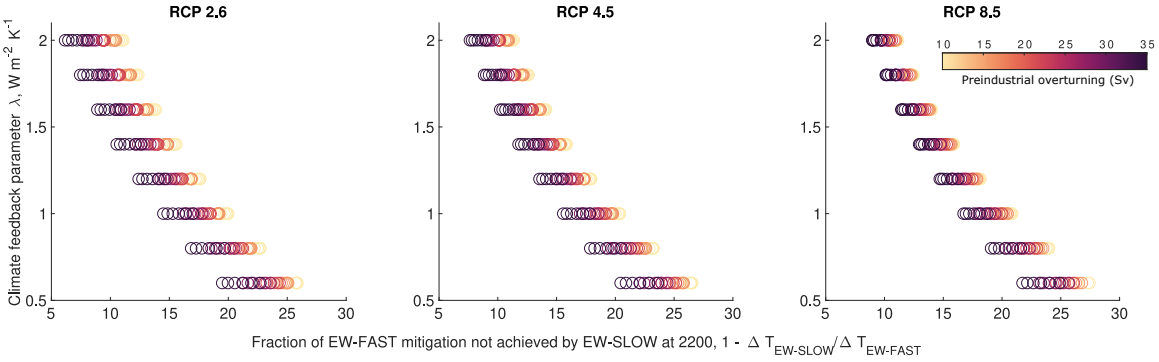

Figure 5.6: Percentage of EW-FAST temperature mitigation achieved by EW-SLOW at year 2200 for the RCP scenarios (separate boxes), climate feedback parameter λ (y-axis), and preindustrial overturning strength (marker colours).

The relative temperature difference between EW-FAST and EW-SLOW varies with the emissions scenario (Figure 5.6, separate panels) and model physics, particularly the climate feedback parameter (Figure 5.6, y-axis). The climate feedback parameter λ exerts the largest control on the relative temperature difference. Under a high climate feedback parameter (less sensitive climate), EW-SLOW attains about 90% of the temperature mitigation of EW-FAST at year 2200. As the climate feedback parameter decreases and the climate becomes more sensitive, EW-SLOW attains a smaller percentage of the mitigation of EW-FAST. As a more sensitive climate realises more warming, this larger relative difference between EW-FAST and EW-SLOW is also a larger absolute difference in temperatures.

The emissions scenario and overturning strength also play a part in setting this temperature lag, albeit smaller ones: for the emissions scenario, the temperature lag increases in scenarios with higher emissions (RCP 8.5 vs. RCP 2.6), and for overturning the setups with

a stronger preindustrial overturning show smaller temperature lags. The models also show how these sensitivities can compensate or interact; for instance, a higher climate feedback parameter or higher emissions reduces the sensitivity of the temperature lag to overturning strength.

5.4.1 Thermal and carbon controls on peak temperature mitigation

For mathematical simplicity, we return to the year of peak impact for the EW-FAST scenario. The peak temperature impact of CDR occurs when $(\Delta T_{control} - \Delta T_{CDR})$ reaches its local maximum. Under (5.2), as the climate feedback parameter is a constant λ, the local maximum is determined by the tendencies for both the heat uptake difference and the radiative forcing difference. At the time of maximum CDR impact:

$$0 = \frac{d}{dt}(T_{control} - T_{CDR}) = \frac{1}{\lambda}\frac{d}{dt}(\Delta R_{control} - \Delta R_{CDR} - \mathcal{N}_{control} - \mathcal{N}_{CDR}) \tag{5.5}$$

$$\Rightarrow 0 = \frac{d}{dt}(\Delta R_{control} - \Delta R_{CDR}) - \frac{d}{dt}(\mathcal{N}_{control} - \mathcal{N}_{CDR}) \tag{5.6}$$

$$\Rightarrow \frac{d}{dt}(\Delta R_{control} - \Delta R_{CDR}) = \frac{d}{dt}(\mathcal{N}_{control} - \mathcal{N}_{CDR}). \tag{5.7}$$

By comparing the radiative forcing tendencies and the heat uptake tendencies within our parameter space, we can ascribe changes in the adjustment to CDR to the carbon and thermal responses, respectively. To first order, $\Delta R_{control} - \Delta R_{CDR}$ is approximately constant post-emissions (Fig. 5.4); however, we retain the term when considering the sensitivities as under second-order controls even small tendencies in the radiative forcing can alter the year of maximum CDR impact. In this section, we will consider $\frac{d}{dt}(\Delta R_{control} - \Delta R_{CDR})$ to be the *carbon contribution* towards maximum CDR effectiveness and $\frac{d}{dt}(\mathcal{N}_{control} - \mathcal{N}_{CDR})$ to be the *thermal contribution* to maximum CDR effectiveness.

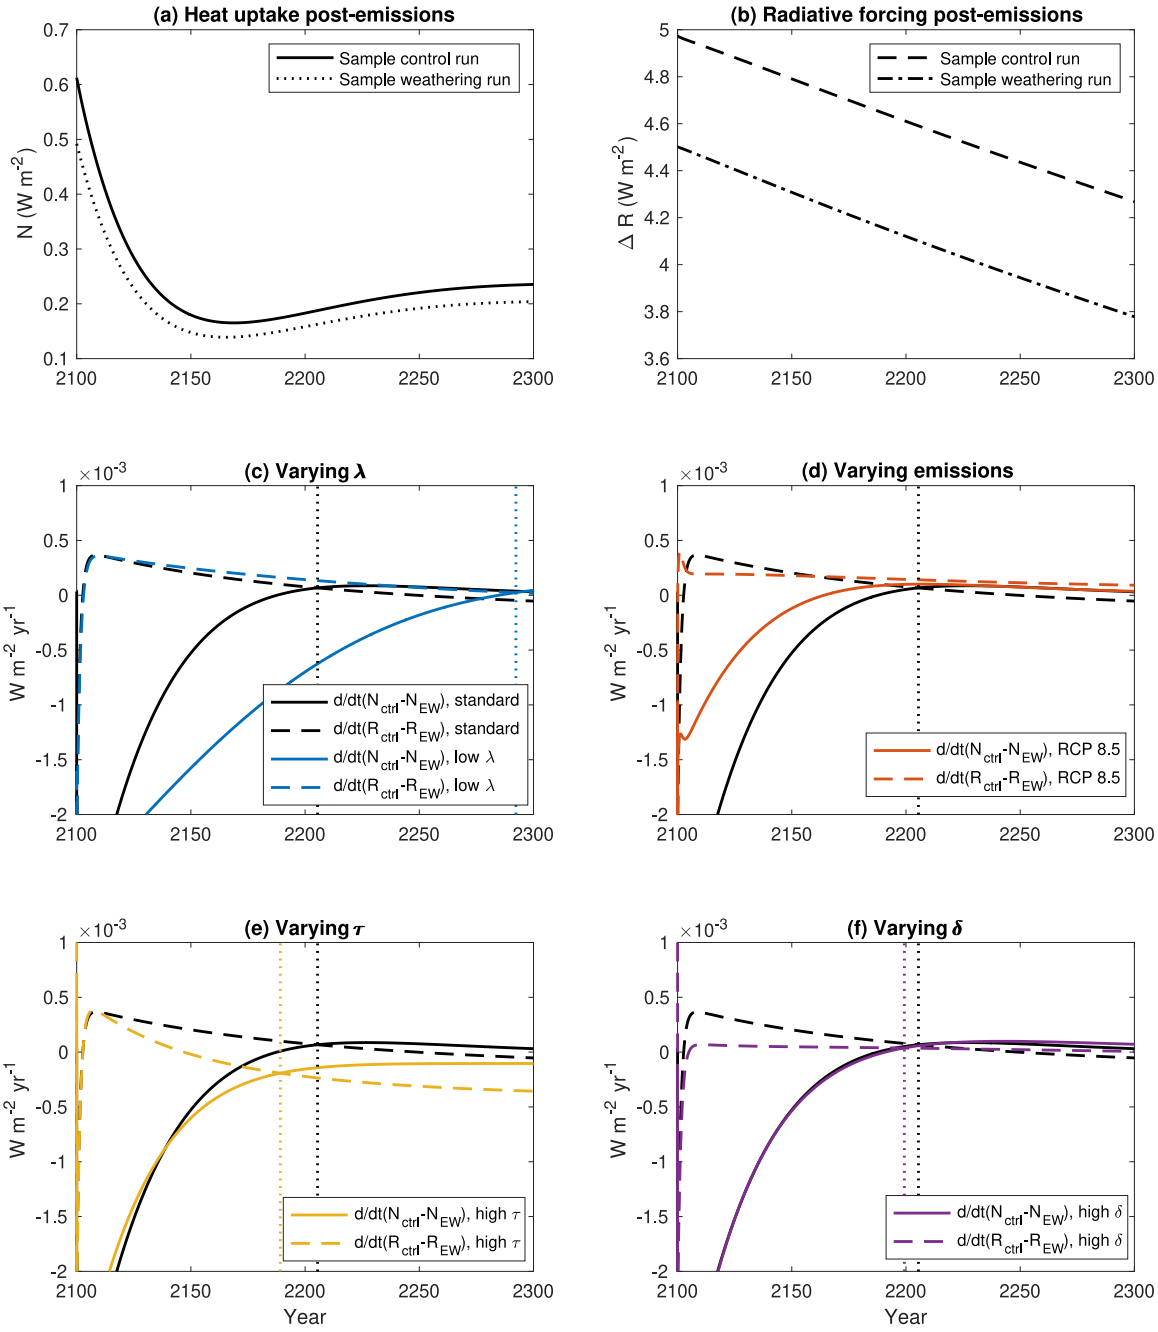

Figure 5.7: (a) Heat uptake post-emissions for a control run (solid line) and EW-FAST run (dotted line) for a reference setup of RCP 4.5, $\lambda = 1.8$ W m^{-2} K^{-1}, $\delta = 0.5$, $\tau = 0.08$ N m^{-2}.(b) Radiative forcing post-emissions for the reference control run (dashed line) and reference EW-FAST run (dashed-dotted line). (c-f) Thermal contributions (solid lines) and carbon contributions (dashed lines) towards the year of peak EW-FAST temperature impact for the reference weathering run (black lines) and runs with varying physical parameters (coloured lines).

To compare how the sensitivities arise through the thermal and carbon contributions, we take a base case of an RCP 4.5 run with a high climate feedback parameter and weak preindustrial overturning. After emissions and CDR cease at 2100, heat uptake and radiative forcing decline (Fig. 5.7 a,b). When weathering is included, both heat uptake and radiative forcing are lower, consistent with the removal of atmospheric carbon; the offset between the two setups is relatively small, but as the timing of peak EW effectiveness is a function of the derivatives, small changes can still alter how the timing of peak temperature mitigation. When considering the tendencies of these differences, the thermal contribution $d/dt(\mathcal{N}_{ctrl} - \mathcal{N}_{EW})$ is negative (i.e., the difference in heat uptake post-emissions decreases) and the carbon contribution $d/dt(R_{ctrl} - R_{EW})$ is positive (i.e., the difference in radiative forcing post-emissions increases), although both values approach values close to 0 as the system adjusts towards an equilibrium state (Figure 5.7c-f). These general behaviours hold when altering the physical parameters of the system; it is the slopes of the contributions that changes and ultimately shift the year at which the EW scheme maximises its temperature mitigation.

Climate feedback parameter λ: Compared to the standard run with a relatively high climate feedback parameter, the post-forcing adjustment under a low climate feedback parameter shows an decrease in the thermal contribution $d/dt(\mathcal{N}_{ctrl} - \mathcal{N}_{EW})$ (Fig. 5.7c, solid lines) and a increase in the carbon contribution $d/dt(R_{ctrl} - R_{EW})$ (dashed lines). As the system contains more heat under a lower climate feedback parameter, the system is further away from its equilibrium state after forcing ceases, and the extra heat in the system results in a slower thermal adjustment towards its equilibrium state. The carbon contribution is slightly higher under a lower λ as the enhanced temperatures result in more atmospheric carbon due to the model's carbon-climate feedbacks. Taken together, the more negative thermal contribution and more positive contribution combine to delay the year of peak EW effectiveness for the run with lower λ (vertical blue dashed line versus black dashed line for the standard run). This strong sensitivity to λ seen in our parameter space is set mainly by the modulation of the thermal contribution.

Emissions scenario: Compared to the standard run forced by RCP 4.5, the run forced under the higher-emissions RCP 8.5 scenario shows an increase in the thermal contribution $d/dt(\mathcal{N}_{ctrl} - \mathcal{N}_{EW})$ (Fig. 5.7d, solid red versus solid black lines) and a carbon contribution $d/dt(R_{ctrl} - R_{EW})$ that remains nearly constant in time, in contrast to the decreasing carbon contribution for the RCP 4.5 (dashed red versus black lines). The thermal contribution

for RCP 8.5 is higher for the first 100 years after the forcing ceases, but the two setups have convergent thermal contributions thereafter. The faster adjustment in the thermal contribution arises because the amount of weathering is relatively small compared to the amount of emissions, so the system forcings are more similar than the control and EW-FAST runs with RCP 4.5 forcings. In contrast, the flatter carbon contribution under RCP 8.5 results in divergence between the two RCP setups after about 70 years. Thus in this case the role of the emissions scenario in setting the year of peak temperature mitigation is set by the carbon response rather than the thermal response as seen in the climate feedback parameter controls. Under the physical model setups used here, the decline in the carbon contribution results in a peak EW mitigation occurring at around 2200, while under the higher RCP 8.5 scenario the temperature mitigation from EW continues to increase post-emissions.

Southern Ocean wind stress τ: Compared to the standard run, the run with a higher wind stress τ results shows a decrease in both the thermal contribution $d/dt(\mathcal{N}_{ctrl} - \mathcal{N}_{EW})$ (Fig. 5.7e, solid yellow versus solid black lines) and the carbon contribution $d/dt(R_{ctrl} - R_{EW})$ (dashed yellow versus dashed black lines). Both decreases are due to the impacts of a stronger preindustrial overturning circulation: the more rapid transport of heat and carbon away from the surface results in the control and EW scenarios' heat uptakes and atmospheric CO_2 mixing ratios converging more quickly, which allows for the thermal and carbon contributions towards the timing of peak EW temperature mitigation to both decrease. The decrease in the carbon contribution is stronger than the decrease in the thermal contribution, resulting in the year of peak temperature mitigation from EW occurring sooner under a higher τ (vertical dotted yellow versus black lines).

Subduction fraction δ: Compared to the standard run, increasing the subduction fraction has little impact on the thermal contribution $d/dt(\mathcal{N}_{ctrl} - \mathcal{N}_{EW})$ (Fig. 5.7f, solid purple versus solid black lines) but decreases the carbon contribution $d/dt(R_{ctrl} - R_{EW})$ slightly (Fig. 5.7f, dashed purple versus dashed black lines). The changes in the carbon contribution occur mainly in the first 100 years, after which both the thermal and carbon contributions are effectively the same between the two physical setups. As a result, there is little change in the year of peak EW temperature mitigation (vertical dotted purple versus black lines).

Thus, the sensitivity of temperature reductions from CDR, and the resulting path-dependency in the temperature signals, is most strongly controlled by the climate feedback parameter λ via its control on planetary heat uptake. In a low-λ, high-climate-sensitivity world, modest decreases in the radiative forcing can lead to larger temperature changes,

which in turn lengthens the time required to adjust the system and delays the year of peak EW effectiveness. The secondary roles of overturning circulation impact both the carbon and thermal contributions toward the timing of peak temperature mitigation. Stronger overturning circulation leads to faster subduction of heat and carbon, strengthening the adjustment to forcings. The faster subduction of carbon in particular leads to an earlier year for peak weathering impacts. Lastly, the emissions scenario alters the thermal and carbon contributions in ways that partially compensate. In the case of compensation under some physical regimes, this can result in a later year for peak impact (Fig. 5.6); for weaker overturning regimes as pictured in Figure 5.7 this compensation leads to divergence in the temperature trajectories of the control run and the EW run.

5.5 Integral constraints of carbon dioxide removal co-benefits

While temperature mitigation is the driving goal behind CDR proposals, any implementation must also consider other biogeochemical impacts of removing atmospheric CO_2 or altering existing biogeochemical cycles. We continue by exploring how a major co-benefit of CDR, reduced ocean acidification, behaves under the two CDR schemes. As the climate feedback parameter provides the main control on the path-dependence of temperature mitigation via its impact on the thermal response (Fig. 5.7), the pH benefits of CDR may be more path-independent. There is still potential for the climate feedback parameter to play a role in reducing ocean acidification, as ocean and air temperatures impact the ocean's ability to take up carbon. We also expect the type of CDR used to have an impact on pH co-benefits as the CCS forcing is a removal of carbon whereas the EW forcing is a redistribution of carbon.

Ocean carbon inventories vary greatly depending on the CDR technique used: under CCS there is a loss of ocean carbon, whereas for EW ocean carbon increases relative to the control (Fig. 5.8 a,b). The ocean carbon differences between all three forcing scenarios highlight the impact of removing carbon from the system in comparison to repartitioning emitted carbon among reservoirs. Like the atmospheric CO_2 mixing ratio, ocean carbon inventories are path-independent, and for both CCS and EW schemes the total amount of carbon is the same for both timelines at year 2200.

When considering surface average pH changes from EW, the behaviour is dissimilar to both carbon inventory changes and temperature changes (Fig. 5.8 c,d). Both removal of carbon through CCS and addition of carbon as alkalinity through EW result in higher surface pH compared to the control experiments. EW produces a higher surface average

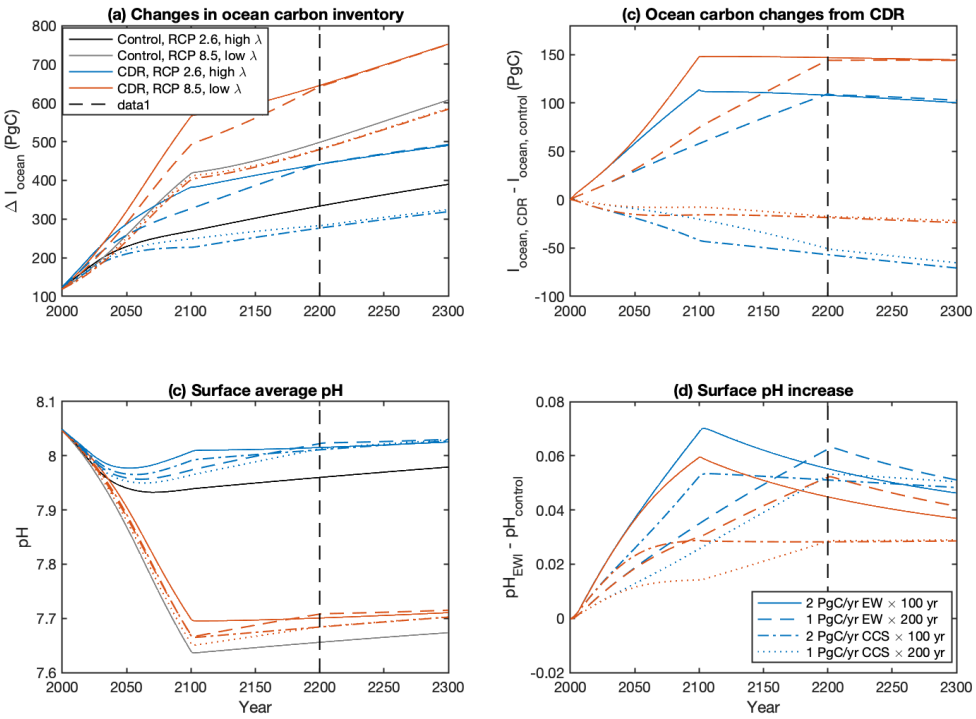

Figure 5.8: Impacts of CDR on ocean carbonate chemistry using two representative Gnanadesikan model runs: (a) Changes in ocean carbon inventory from control RCP scenarios (black/grey lines) and RCP scenarios with CDR representations (coloured lines), (b) Ocean carbon inventory differences between runs with CDR and control runs, (c) Surface average pH for control RCP runs and runs with CDR, and (d) Surface average pH difference between runs with CDR and control runs.

pH than CCS, as the alkalinity input dominates over the carbon input in setting surface carbonate chemistry. The surface pH increase peaks at the year that CDR ends and declines as the system adjusts. This peak-and-decline behaviour is more noticeable in the EW runs than in the CCS runs, which approach their long-term pH level at the year that CCS ends. The peak-and-decline behaviour in EW leads to an overshoot seen at 2200, where the surface ocean pH is slightly higher in EW-SLOW than in EW-FAST.

We can explore the sensitivity of the pH overshoot to the parameters used in our model setup (Fig. 5.9). As pH is nonlinear, we consider the absolute pH differences at 2200 rather than relative differences as seen in Fig. 5.6 for temperature. As the range of pH is more constrained than the range of temperatures the use of absolute pH differences does not

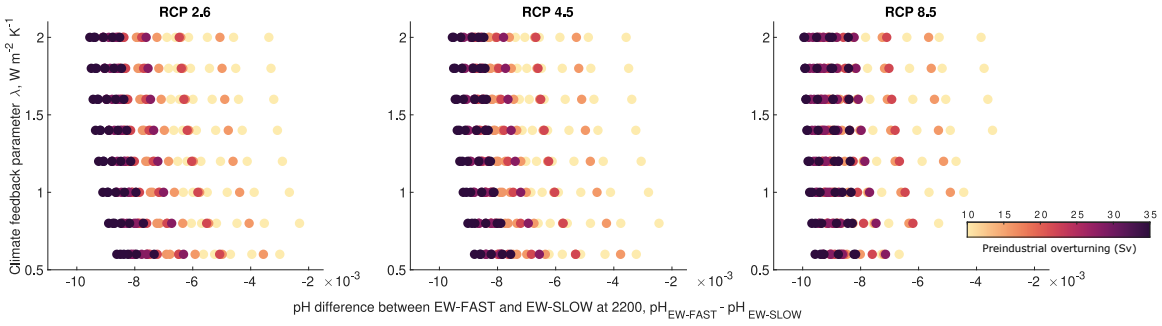

Figure 5.9: Surface pH difference between EW-FAST and EW-SLOW at year 2200 for the RCP scenarios (separate boxes), climate feedback parameter λ (y-axis), and preindustrial overturning strength (marker colours).

focus on the uncertainty space in the same way as absolute temperature differences. The magnitude of the pH overshoot $\text{pH}_{EW-FAST}$ - $\text{pH}_{EW-SLOW}$ at year 2200 is determined mainly by the emissions scenario and overturning strength. The pH overshoot is larger under higher-emissions scenarios; the differences in pH are approximately twice as large in RCP 8.5 as they are in RCP 2.6. The overturning strength induces a large spread in the magnitude of the overshoot as well. Under a stronger preindustrial overturning, the pH overshoot is higher. In contrast to the sensitivities of the temperature mitigation lag (Fig. 5.6), for the pH overshoot the climate feedback parameter plays a lower-order role. There is a slight control of the climate feedback parameter on the pH overshoot, in which a higher climate feedback parameter results in a slightly larger pH overshoot at 2200.

The nonlinearity of ocean carbonate chemistry prevents the use of analogous equations to (5.6) and (5.7) for the adjustment of surface pH induced by CDR. However, it is still possible to use tracer arguments for both temperature and carbon impacts. Thus in the final section we consider how the tracers for temperature, DIC, and alkalinity evolve under different overturning regimes and connect the results with the path-dependencies found in the surface temperature and surface pH results.

5.5.1 The role of overturning for thermal and carbon benefits

Notwithstanding the overwhelming role of the climate feedback parameter in setting the thermal benefits of CDR, the role of overturning has a similar impact on path dependency for both thermal and carbon benefits of CDR. Under stronger overturning regimes, slower CDR results in a smaller temperature lag (Fig. 5.6, 5.7e) and a larger pH overshoot (Fig. 5.9). This indicates that the similarities in the distribution of altered tracers for temperature,

heat, and carbon provide similar controls on CDR signals.

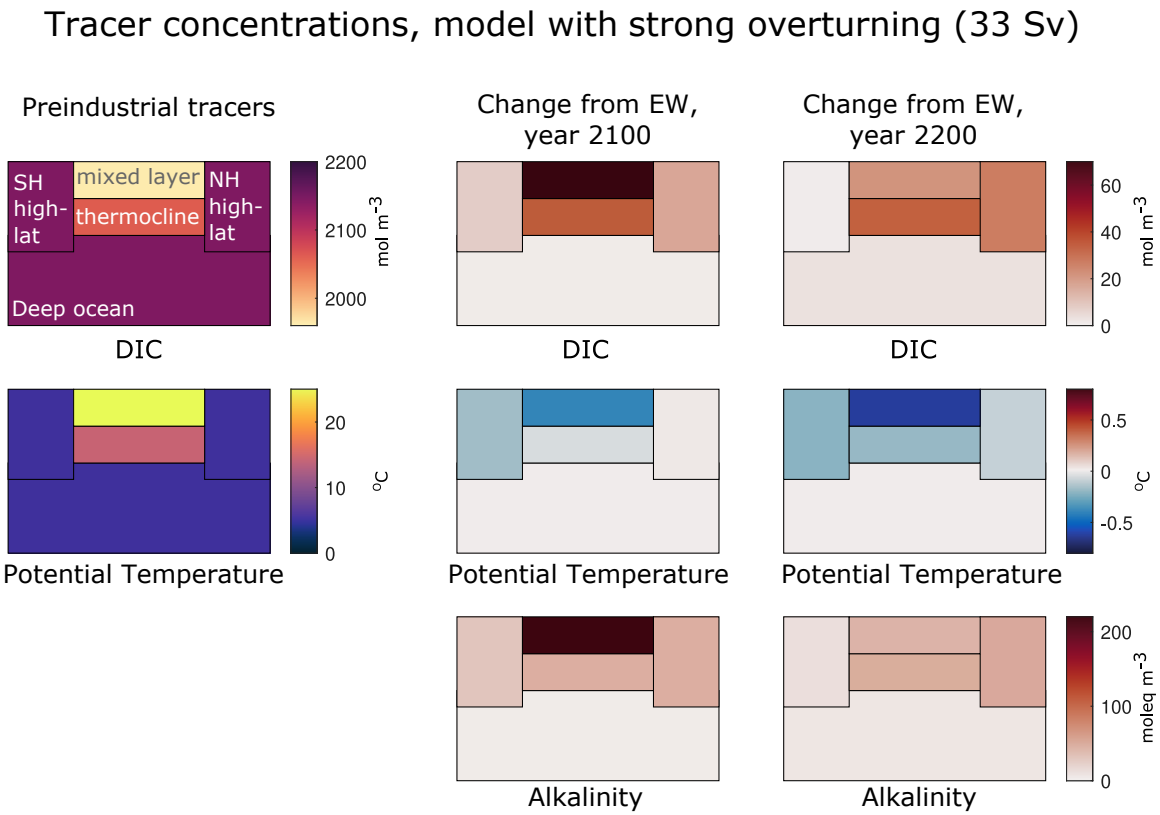

Figure 5.10: Schematic of preindustrial DIC and T, and tracer differences between EW and control runs under a strong overturning. Note that in these plots we have not visualised the change in volumes of the deep ocean box and the thermocline box.

Preindustrial distributions of DIC and temperature are relatively similar in model setups with strong and weak preindustrial overturning (Fig 5.10, 5.11 left-hand figures). We compare the differences in temperature, DIC, and alkalinity concentrations between the control runs and EW-FAST runs to see how the tracers that impact EW signals are distributed under different overturning regimes, and how that might reflect on the path dependence of the mitigation signals. The impact of EW on DIC, temperature, and alkalinity at year 2100 are concentrated in the upper ocean for both overturning states (Fig. 5.10, 5.11, middle figures). The largest concentration changes are located in the mixed layer box as that is the box with the largest surface area in contact with the atmosphere and the shallowest depth.

Post-emissions, the overturning transports added tracers from the surface ocean boxes into the ocean interior boxes. For alkalinity and DIC tracers, concentrations in the Northern Hemisphere high-latitudes, thermocline, and deep ocean increase after emissions and

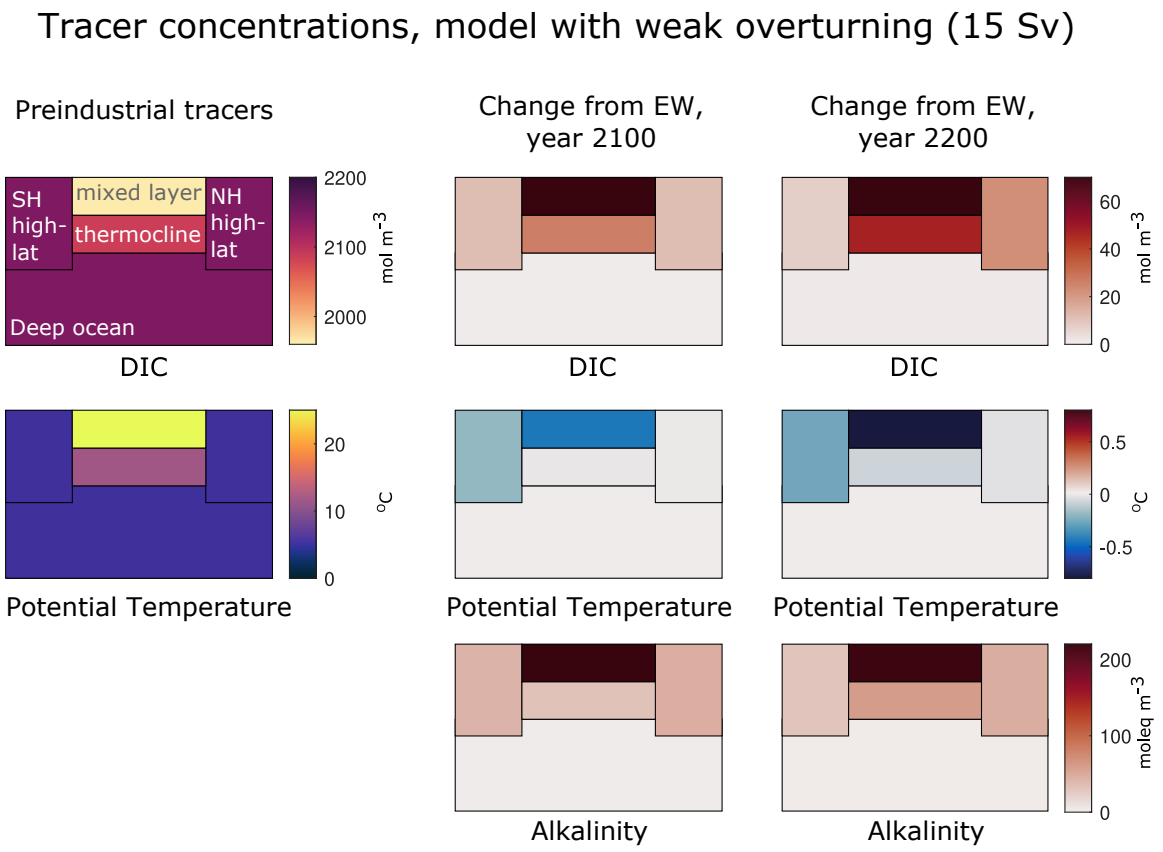

Figure 5.11: Schematic of preindustrial DIC and T, and tracer differences between EW and control runs under a weak overturning. Note that in these plots we have not visualised the change in volumes of the deep ocean box and the thermocline box.

weathering cease at 2100. For the added alkalinity tracer associated with EW, the enhanced surface concentrations are effectively eliminated through strong overturning. A slower implementation of EW maintains enhanced surface alkalinity concentrations for longer as the surface flux compensates for the removal of previously-added tracer, resulting in a larger year 2200 difference in surface pH induced by EW. Under a weaker overturning, tracers are more contained in the surface ocean boxes. Thus DIC and alkalinity remain concentrated in the mixed layer box even 100 years after EW-FAST finishes, and the surface pH impacts remain relatively similar to those under a EW-SLOW.

While EW increases the amount of carbon and alkalinity in the ocean boxes, the temperature impact on the ocean is negative as there is a reduction in the global radiative forcing. Similar to DIC and alkalinity increases, temperature decreases are strongest in the mixed layer box, but unlike the other tracers the continued lower radiative forcing results in

continued cooling post-emissions. In a weaker overturning regime, the mixed layer temperature decrease from EW is enhanced relative to the other ocean boxes, so that the uptake of heat by the ocean continues for longer, resulting in a longer time necessary until the EW scheme obtains its maximum temperature mitigation.

It is possible to consider the sensitivity to ocean overturning in terms of the total ocean adjustment to the various forcings within the CDR schemes. For any amount of carbon, temperature, or alkalinity added by CDR, there is a set equilibrium state for the ocean, which is managed by the balance of surface fluxes and mass transformation. As carbon, temperature, and alkalinity all enter the ocean through surface fluxes, the overturning acts to slowly propagate these qualities to the ocean interior. As the ocean overturning strengthens, it is able to redistribute these qualities more effectively; thus, differences in the oceanic input rates play a relatively smaller role in setting the ocean state.

5.6 Discussion and Conclusions

Successfully deploying CDR technologies depends on understanding the impacts of the technologies involved and how those impacts depend on the timing of CDR and the climate state. In this chapter we have created a large perturbed parameter ensemble using a low-complexity ocean-atmosphere model to explore how impacts of CDR are constrained by the cumulative amount of carbon drawdown. Our results indicate that there are contrasting path dependencies for the thermal and carbon products of CDR schemes.

5.6.1 Path dependency of temperature mitigation encourages early implementation

We find that the maximum temperature signal from CDR occurs after carbon drawdown finishes. The continued temperature adjustment post-drawdown leads to a clear path dependency: when two CDR timelines attain the same cumulative amount of carbon drawdown, the slower timelines lag in their temperature reductions relative to faster timelines (Fig. 5.6). Within our perturbed parameter ensemble the amount of lag can vary from around 5% to almost 30% of the temperature mitigation attained by the faster timeline. Our results emphasise that attaining the maximum benefit of CDR or carbon emission reductions policies requires early implementation.

The path dependency of temperature mitigation signals is inherently related to the amount of surface warming that continues post-emissions. Within the Gnanadesikan model

ensemble, each control model run also shows continued warming post-emissions, and models with longer periods of continued warming post-emissions also show longer periods of continued temperature mitigation post-CDR (Fig. 5.5). Throughout the ensemble, the lag to reach peak CDR effectiveness is longer than the lag to reach peak warming in the control runs. More complex models show little agreement in the behaviour of surface warming post emissions, with some showing near-zero warming post emissions while others show continued warming for up to a century (Williams et al., 2017b; MacDougall et al., 2020). The consistency of our model results suggest that, regardless of the degree to which continued warming occurs in control runs, the temperature response to CDR may be slower.

Role of the climate feedback parameter in determining temperature path dependency

The climate feedback parameter λ plays the dominant role in setting the path-dependency of temperature mitigation signals. Under a low climate feedback parameter, temperature mitigation signals show the highest level of path dependency, with temperature reductions from slower implementations lagging those from faster implementations by 19% to 27% (Fig. 5.6). The enhanced path dependency is related to the climate system's thermal response: under a lower climate feedback parameter, radiative forcing is more effectively maintained in the Earth system as added heat. Thus the same amount of CDR induces larger changes in the Earth's heat balance, requiring a longer adjustment period (Fig. 5.7c).

Thus, in a climate state with a low climate feedback parameter, the added benefit of early implementation is even further emphasised. Climate model studies show a potential for such a low climate feedback parameter, due in part to more complex representations of clouds (Zelinka et al., 2020); although our ensemble does not represent such processes, our sensitivity tests with λ indicate that this shift in climate models may also result in a shift of path-dependency for the response to emissions or CDR. The impact of the climate feedback parameter or climate sensitivity has only lightly been touched upon by model intercomparison studies, although there were some indications that more sensitive climate states have stronger levels of continued warming (MacDougall et al., 2020). Climate model studies on the climate feedback parameter find that λ becomes smaller with continued emissions (Andrews et al., 2015); in the case that the climate system becomes more sensitive under continued emissions, the importance of early CDR will increase further.

5.6.2 Temperature benefits and carbon co-benefits can have opposing path dependencies

In this perturbed parameter experiment, we have also been able to explore the path dependency of ocean pH increases arising from the addition of alkalinity through our enhanced weathering (EW) scheme. When considering the pH co-benefit of EW in the Gnanadesikan model, the path-dependency is opposite that of temperature reductions (Fig. 5.8). There is a peak-and-decline behaviour visible in surface pH, resulting in slower EW schemes having a higher surface pH than faster EW schemes when the two timelines have the same cumulative drawdown of carbon. Thus, in the case where EW is implemented less for its impact on temperatures and more for alkalinisation purposes, a continued, slower implementation of EW could be preferred.

Unlike temperature mitigation signals, surface ocean pH signals are only slightly sensitive to the climate feedback parameter λ (Fig. 5.9). Rather, the strength of the ocean overturning plays the dominant role in setting the peak-and-decline response of surface pH. A stronger ocean overturning acts to enhance the peak-and-decline behaviour for surface pH because the overturning more quickly sequesters added alkalinity away into the ocean interior, where they cannot impact the surface carbonate chemistry.

The role of overturning in setting responses to carbon dioxide removal

The strength of the ocean overturning impacts both the speed of temperature mitigation and pH mitigation. Because of the large range for the climate feedback parameter λ within our ensemble and the dominance of λ in setting the continued thermal response, the role of overturning in setting the temperature mitigation path-dependence is of secondary importance. However, the role of ocean overturning on temperature and pH signals path dependency works in the same direction: under stronger preindustrial overturning, the temperature mitigation for fast and slow timelines becomes less path-dependent, and the pH overshoot for the slow timeline increases. This is related to the rate at which the overturning removes tracers from the surface ocean, where they can no longer play a role in setting the ocean-atmosphere fluxes of heat and carbon.

The EW method examined in this chapter shows little "leakiness" of carbon dioxide from the ocean. The long-term storage of carbon from the EW scheme is due to its addition as bicarbonate, which is both dissolved inorganic carbon and alkalinity. The impact of the added alkalinity on the ocean charge balance allows the added carbon to be stored for the entire 300 years of the model integration. For other CDR techniques, particularly those

involving the enhancement of the ocean's biological carbon pump, there are limitations on how long carbon can be sequestered (Siegel et al., 2021). Enhanced weathering can also impact biology on shorter timescales through the iron fertilisation of diatoms (Renforth and Henderson, 2017). The strength of the ocean overturning is likely to play a major role in setting the timescales on which these responses occur. Although the construction of the model used in this chapter does not include ocean biology, and the inclusion of biology would likely have little impact on the temperature signals seen in our work, our results still suggest that understanding ocean overturning is important for understanding the myriad of ways that CDR will impact the Earth system.

5.6.3 Implications of results for carbon dioxide removal policy

In an ideal world, the implementation of a CDR method would maximise the method's benefits on the climate system, while minimising any potentially harmful side-effects. In this chapter we have considered path-dependency comparing results of different timelines at the year that cumulative drawdown is the same. Our results indicate that the "ideal" implementation is heavily dependent on the desired climate perturbation, such as reduced surface temperatures or reduced ocean acidification, and also the background climate state. While the method of comparing single-year values has proven useful for our analysis, it may not be the most appropriate metric for determining real-world CDR implementation. Indeed, some major aspects of climate change can be related to integrated impacts rather than impacts at a specific year. For instance, droughts or fisheries health may rely more on year-on-year conditions rather than the conditions in a specific year. For these aspects of the climate system, an overshoot scenario may be less ideal than one in which the benefits are moderate but more consistent.

Additionally, there may be harmful "rebound effects" associated with impacts that quickly decay after cessation. The surface ocean pH signals in this chapter show a slight rebound effect within our EW implementations. Thus we have attempted to avoid assigning any true value to the responses shown in this chapter, but rather suggested how these results may be interpreted if such responses were deemed desirable.

Caveats of the Gnanadesikan model ensemble

It is necessary to emphasize the main caveats of using an conceptual atmosphere-ocean model for this study. The Gnanadesikan model, being highly idealised and not representing many aspects of the climate system, is not suitable for creating policy recommendations.

Importantly, the Gnanadesikan model contains no terrestrial representations. MacDougall et al. (2020) find that terrestrial carbon dynamics can play a large role in determining the magnitude of warming post-emissions. Thus the inclusion of a terrestrial system may change the degree to which the temperature signals from CDR show path dependence. Still, our results indicate that the path-dependency of CDR impacts is partially related to the path-dependency of the response to emissions, and the sensitivity to the climate feedback parameter and ocean overturning strength found with the Gnanadesikan model provides context for further, in-depth study on how surface temperatures evolve after emissions or drawdown.

In summary, idealised model results show path dependency for the impacts of CDR schemes. Temperature mitigation produced by CDR continues after the forcing ceases, resulting in a temperature mitigation lag for slower CDR timelines. The mitigation lag is largely controlled by the climate feedback parameter, which alters how quickly the climate system can adjust towards its equilibrium climate state through continued heat uptake. Carbonate chemistry co-benefits such as ocean pH increases show an opposite path dependency to that for temperature mitigation. Slower CDR schemes show higher surface ocean pH at the year they have the same cumulative drawdown as faster CDR schemes, as the continued input of alkalinity from the schemes allows for higher concentrations of negative charges to be maintained in the surface layer. The radiative forcing and thermal responses combine to create large differences in the climate feedback parameter. While the results using an idealised, low-complexity model are of limited direct impact to CDR policy, they highlight how temperature benefits and carbonate chemistry benefits may need to be balanced through proper implementation of CDR schemes.

Chapter 6

Final Discussion and Concluding Remarks

The focus of this thesis has been to understand how the Earth system's carbon response and thermal response impact the response to anthropogenic forcing. The climate system is currently undergoing a level of change it has never seen before, thanks to the continued emission of CO_2 and other greenhouse gases from anthropogenic activities. To first order, the thermal and carbon responses have been assumed to compensate each other, leading to simple empirical relationships between carbon emissions and warming. Throughout this thesis we have examined the intricacies in both the carbon and thermal responses over the recent historical period and for future scenarios. We have found that the relationship between the thermal and carbon responses can be useful to understand the response to emissions, but care needs to be taken in assuming the relationship is purely compensatory, as climate feedbacks can strongly modulate the Earth system response.

6.1 A variety of methods to understand thermal and carbon processes

In this thesis we have used three different methods to examine the Earth system's responses to carbon emission forcings. Each method has involved using different climate models and diagnosing how their results hinge upon their representation of thermal and carbon processes. Each type of climate model has had its advantages and disadvantages: no one model would have been suitable for all three areas of work presented in this thesis. In this section we provide an overview of each method used, with a short discussion on how the

models and approaches have been able to help answer the research questions within the thesis.

6.1.1 High-complexity CMIP6 models to extrapolate observational datasets

In Chapters 2 and 3, we utilised CMIP6 model ensembles to better understand historical carbon variability from relatively sparse datasets. In Chapter 2 we focused on the ocean carbon system to determine how well carbon time series sites can represent historical ocean carbon variability, as they are some of the only sites with regular, depth-resolved observations of ocean interior carbon.

Our synthetic model tests revealed that CMIP6 ensembles can provide skillful extrapolation of observations. However, the potential to extrapolate time series information is heavily dependent on the type of information desired. If one is interested in the depth-integrated upper ocean carbon content, then the time series data can be extrapolated almost globally; however, this extrapolation only considers the long-term increase in ocean carbon under historical CO_2 emissions.

Removing a first-guess field proportional to the atmospheric CO_2 mixing ratio is able to isolate the historical ocean carbon content variability. With the variability isolated, the time series sites are representative of their immediate locales. The strong correlations with time series sites in the immediate area are a result of agreement among the ensemble members' spatial autocorrelations; further away from the time series sites, the correlations are less certain among the models, and the optimal interpolation scheme thus dampens the reconstruction towards the first-guess field. While the carbon time series provide highly accurate measurements and are able to resolve variability at seasonal, interannual, and decadal timescales, the scarcity of locations and the fact that current time series sites are located in similar dynamic regimes hinder their extension to basin-scale or global ocean carbon variability.

In Chapter 3 we considered another approach for reconstructing ocean carbon variability using the CMIP6 model ensembles with Argo temperature and salinity profiles. The covariance fields in this approach consider the CMIP6 models' relationship between ocean temperature, salinity, and carbon, as well as considering the models' spatial autocorrelations. Unlike the ocean time series sites, the available Argo profiles allow for good global coverage of variability for temperature and salinity. Such a reconstruction relies on how ocean circulation impacts the uptake and redistribution of heat and carbon, and so the correlations between temperature and salinity observations and carbon data elsewhere are

necessarily weaker than analogous carbon-to-carbon correlations. However, our synthetic model tests indicate that combining information from the model covariances with Argo-style temperature and salinity profiles is able to represent a significant portion of historical ocean carbon variability, even when considering the irregular sampling associated with the use of drifting profilers.

The results of these chapters emphasise how understanding the connection between thermal and carbon processes is useful for characterising the historical response to emissions. Whilst the reconstruction in Chapter 3 will be unable to fully represent ocean carbon variability, the strength of the error reductions using a relatively simple method with widespread observations of temperature and salinity is substantial. The ocean uptake of heat and carbon, being regulated by surface fluxes and ventilation pathways, are highly inter-related; this inter-relatedness has proven to be useful for reconstructing ocean carbon, which is notoriously poorly-sampled in the global ocean. There is promise that the constructed method with CMIP6 covariance fields and Argo data can provide an independent estimate for how upper ocean carbon has behaved over the past decades, and has the potential to complement previous work with the GLODAP dataset or machine learning results.

6.1.2 Intermediate-complexity models to understand regional thermal and carbon contributions to the Earth system response

In Chapter 4 we transitioned from an historical, observation-based viewpoint to explore how regionality in the Earth system thermal and carbon responses contribute to the future response to carbon emissions. Theory on the proportionality between surface warming and CO_2 emissions relies on the compensation between heat and carbon uptake on a global scale, although both heat and carbon uptake are highly spatially heterogeneous. In this chapter we employed an Earth System Model of Intermediate Complexity, the UVic ESCM, which has been widely used to describe the proportionality between surface warming and cumulative carbon emissions. We expanded the existing global theory to separate regional contributions within both the thermal and carbon responses.

Our results highlighted the intricacies of understanding global responses as a combination of local responses. The thermal and carbon responses were found to be substantially different in the Northern and Southern Hemispheres, and thus the compensation seen on global scales is not reflected on hemispheric scales. Regional variations in the climate feedback parameter and communication between the two hemispheres is vital for setting up the global response, particularly for the thermal response to radiative forcing.

Our work with the UVic ESCM highlights some of the shortcomings inherent in simplified models. The UVic ESCM has been used extensively to quantify the proportion of surface warming to cumulative carbon emissions and study how the oceanic uptake of heat and carbon control the proportionality. However, according to the results in this chapter, on hemispheric scales the UVic ESCM suffers from its highly simplified slab atmosphere componenent, which result in an unrealistic cross-equatorial heat transport and a lack of climate feedbacks. While the UVic ESCM results may not extend to other models or datasets due to these shortcomings, the results within the thesis do emphasise that the spatial heterogeneity in the thermal and carbon responses plays a role in setting the global response, particularly in the presence of non-uniform radiative forcing. The flexibility in our regional framework allows for further examination of the role of other regions, such as the Southern Ocean, to setting the relationship between cumulative carbon emissions and surface warming.

6.1.3 Idealised box model ensembles to understand the controlling processes behind the timing of CDR benefits

The final piece of work aimed to understand how the Earth system response to carbon dioxide removal (CDR) is related to the timing of these technologies. While the amount of surface warming to cumulative carbon emissions is approximately path-independent, nonlinearities in carbon storage and climate feedbacks have the potential to push the responses to carbon removal away from such integral contraints. In Chapter 5 we explored the limits of CDR responses through a large ensemble of a modified version of the Gnanadesikan model, which included tracers for heat, carbon, and alkalinity to account for the various thermal and carbon impacts of CDR.

The results across the Gnanadesikan ensemble showed that there are differing levels of path-dependence for the impacts of CDR technologies. The degree of path dependence was found to vary with the impact being examined and uncertainties within the climate system. While the reduction in atmospheric CO_2 from CDR was found to be path independent, both temperature reductions and surface ocean pH increases showed path dependence when two CDR timelines with the same cumulative drawdown were compared. Surface temperature reductions continue after CDR implementation ceases, resulting in slower CDR scenarios lagging behind faster scenarios when cumulative drawdown is equivalent. The ensemble results show a robust temperature benefit in early CDR implementation. Conversely, for surface ocean pH increases, there is a distinct peak at the year that CDR finishes, resulting in slower CDR scenarios having a higher surface pH than faster scenarios at the time

cumulative drawdown is equivalent. Thus decisions on how CDR is implemented will need to be very considerate of which impacts they would like to maximise from each method.

The strength of these path dependencies is determined by the system's physical climate parameters. The climate feedback parameter is the dominant control on how temperature reductions from CDR are realised, whilst the strength of the ocean overturning is the dominant control on the surface pH increases from CDR. Under a lower climate feedback parameter, the lag in temperature mitigation between convergent CDR timelines is exacerbated, further emphasising the benefits of early CDR implementation in a sensitive climate. The ocean overturning strength imposes similar changes for the temperature and pH impacts of CDR, as it modulates the rate at which surface tracers are transported into the ocean interior. In this way the thermal and carbon impacts of CDR behave in similar ways from the oceanic contribution, and the climate feedback parameter acts as a "decoupling" mechanism as it has a large impact on how atmospheric radiative forcing is partitioned into retained versus re-emitted heat.

Throughout this thesis, there has been a strong focus on climate model tests, regardless of the complexity of the model used. As models are all imperfect representations of the Earth system, these tests are vital to understand the scope of each chapter and provide context as to how the results within each chapter may be further expanded. Some model tests, like those using the CMIP6 model ensembles in Chapters 2 and 3, showed promise in using model results to understand observed historical phenomena. While the model diagnostics for the TCRE using the UVic ESCM highlighted multiple shortcomings of the intermediate-complexity model in Chapter 4, the results provide a good argument for limiting the use of such models to quantifying global-mean behaviour and controls. Lastly, while the specific results of the highly-idealised Gnanadesikan model used in Chapter 5 are not directly applicable to more complex systems, the limits on the controls of CDR path dependence will be useful to understand responses in more complex systems or provide suggestions for how to include uncertainty within policy decisions.

6.2 Future pathways for this work

The methods and work within this thesis have potential to be expanded upon post-submission, and some of the work within the thesis will be continued by the author. The reconstruction work presented in Chapters 2 and 3 will be expanded to reconstruct historical carbon variability using the CMIP6 covariances and real Argo observations, with validation using

interior carbon data from time series sites, GLODAP data, and other reconstructions using Biogeochemical Argo data. We aim to write up the results from this continuation for publication. The Gnanadesikan model results presented in Chapters 5 are also currently being written up for submission to Earth's Future. There are potential further avenues to extend the work using the Gnanadesikan model, including using more realistic CDR setups (e.g., fluxes of carbon and alkalinity more consistent with riverine fluxes) or full portfolios of CDR methods to test the linearity of the thermal and carbon responses. Additionally, the lag in the temperature response from CDR could be explored with higher-complexity models to see how it relates to the amount of continued warming post-emissions and to compare the controls on the two responses. While the results from Chapter 4 are not easily extendable to the wider climate community, as the model's representation of cross-equatorial heat transport and climate feedbacks are not extendable to other models or observations, the theory behind the work can potentially be used to understand regional contributions, such as the role of the Southern Ocean, to the surface warming response within the CMIP6 1% CO_2 experiments. In such an extension the shortcomings of the atmospheric component seen in the lower-complexity UVic ESCM will be avoided, and the diagnostics would instead highlight how model differences affect the regional drivers to the response to emissions.

6.3 Broader implications of the thesis

Ultimately, the work conducted within this thesis emphasises that understanding the level of asymmetry between the thermal and carbon responses is necessary to fully understand the response to emissions. The three main pieces of work in this thesis illustrate that these asymmetries operate within different time periods and forcing combinations, making them ubiquitous features in the Anthropocene climate system.

Over the historical period, the asymmetry between ocean carbon and heat anomalies for most of the global ocean allows for analysis methods to reconstruct historical ocean carbon from temperature and salinity profiles. The strength of the analysis belies the first-order linear relationships within the method and indicates that the sparseness of ocean carbon measurements can be overcome partially through the density of hydrographic measurements. This word has broader implications within the current field of historical carbon reconstructions as it can be used to cross-examine the results of other, potentially more complex analysis methods.

Under future emissions, decomposing the response to emissions into hemispheric carbon and thermal responses reveals that the two responses have different hemispheric behaviour.

The theory used in this work has broader implications for understanding uncertainties in the TCRE. While uncertainties in the TCRE can be understood as arising from uncertainties in regional temperature responses, they can also be understood as arising from uncertainties in regional heat and carbon balance. The theory presented in this thesis provides a method for translating process uncertainties into uncertainties in the emergent TCRE behaviour.

Lastly, under CDR scenarios, the thermal and biogeochemical benefits show asymmetric path dependencies. The asymmetric path dependencies occur as feedbacks and overturning work together to alter how added tracers from CDR impact surface properties. The asymmetry found in this work has major implications for how CDR policies should be determined, as the various benefits of CDR may be valued differently and thus affect which CDR timelines are determined to be ideal.

Appendix A

Construction of the Gnanadesikan ocean model

The following sections describe in more detail how the Gnanadesikan model represents mass transformation between the ocean boxes as well as how the model adjusts to added heat and carbon. The setup of the model mirrors that found in Marshall and Zanna (2014); Katavouta et al. (2019) and Williams et al. (2021) with conservative tracers for temperature and DIC. The model has been extended to include an alkalinity tracer that is allowed to evolve according to the external forcing from any present CDR methods and is redistributed as a passive tracer with the ocean circulation.

The Gnanadesikan model consists of a slab atmosphere model and an ocean model with boxes that represent the low- and mid-latitude mixed layer, the low- and mid-latitude thermocline, the Northern Hemisphere high-latitude upper ocean, the Southern Hemisphere high-latitude upper ocean, and the deep ocean. In its original incarnation, the Gnanadesikan model was used for steady-state solutions under different overturning forcings (Gnanadesikan, 1999). Marshall and Zanna (2014) extended the model to account for the transient response to heat forcing and expanded the model to include many layers. In Katavouta et al. (2019) the Gnanadesikan model was further extended so that changes in the overturning were linked to radiative forcing.

Constant	Constant name	Units	Value
A	Area of atmosphere and global ocean	m^2	3.4×10^{14}
A_{low}	Area of low latitude ocean	m^2	2×10^{14}
A_{NH}	Northern Hemisphere high latitude ocean area	m^2	7×10^{13}
A_{SH}	Southern Hemisphere high latitude ocean area	m^2	7×10^{13}
D_{atm}	Atmosphere depth	m	10^4
M_a	Atmospheric mass	kg	5.1352×10^{18}
D	Ocean depth	m	4000
D_{mix}	Mixed layer depth	m	100
D_{NH}	Depth of NH high latitude upper ocean box	m	1000
D_{SH}	Depth of SH high latitude upper ocean box	m	1000
L_x	Southern Ocean zonal extent	km	3×10^4
L_y	Southern Ocean meridional extent	km	1500

Table A.1: Dimensions within the Gnanadesikan atmosphere/ocean model. Note that the Southern Ocean zonal and meridional extent are used for calculating Southern Ocean intermediate water formation within the Antarctic Circumpolar Current and thus $A_{SH} > L_x \times L_y$.

Constant	Constant name	Units	Value
ρ_a	Air density	kg m^{-3}	1
ρ_o	Ocean density	kg m^{-3}	1025
S	Ocean salinity	psu	34.5
P	Ocean inorganic phosphate	mol kg^{-1}	2.157×10^{-6}
Sil	Ocean silicate	mol kg^{-1}	6.9024×10^{-5}
a	Radiative forcing coefficent	W m^{-2}	5.35
$C_{p,a}$	Atmosphere specific heat capacity	J kg^{-1} K^{-1}	1000
$C_{p,o}$	Ocean specific heat capacity	J kg^{-1} K^{-1}	4000
K_g	air-sea gas transfer coefficient	m s^{-1}	5×10^{-5}
k_v	diapycnal mixing coefficient	m^2 s^{-1}	10^{-5}
K_{ed}	eddy diffusion coefficient	m^2 s^{-1}	10^3
f	Midlatitude coriolis parameter	s^{-1}	10^{-4}
g'	Reduced gravity	m s^{-2}	0.02

Table A.2: Constants used in calculating volume, heat, and carbon transports within the Gnanadesikan model.

A.0.1 Mass transformation and overturning balances

Diapycnal upwelling in the model transforms thermocline waters into low-latitude mixed later waters and uses a constant diapycnal mixing coefficient k_v:

$$q_v(t) = \frac{A_{low} k_v}{D_{therm}(t)}. \tag{A.1}$$

The Southern Ocean residual overturning is set by an eddy volume flux q_{eddy} and an Ekman volume flux q_{Ek}, both with units $m^3 s^{-1}$. The integrated baroclinic eddy volume flux q_{eddy} induces a southward flux in the model that is a product of the horizontal eddy diffusivity K_{ed} and the ratio of the Southern Ocean zonal extent L_x to the Antarctic Circumpolar Current width L_y:

$$q_{eddy}(t) = -K_{ed} \times \frac{L_x}{L_y} \times D_{therm}(t). \tag{A.2}$$

The integrated northward Ekman flux q_{Ek} is a product of the Southern Ocean wind stress τ, Southern Hemisphere coriolis parameter $-f$, ocean density ρ_o, and Southern Ocean zonal extent L_x:

$$q_{Ek} = -\frac{\tau L_x}{-f \rho_o}. \tag{A.3}$$

The residual overturning is then the sum $q_{SO}(t) = q_{eddy}(t) + q_{Ek}$.

The North Atlantic deep water formation q_{NA} occurs in the Northern Hemisphere high-latitude box and responds to the global heat uptake $N(t)$ and vertical temperature gradient $T_{light} - T_{deep}$:

$$q_{NA}(t) = -\frac{g' D_{therm}(t_0)^2}{2f} + \frac{N(t) A_{low}}{\rho_o C_{p,o} [T_{light}(t) - T_{deep}(t)]}. \tag{A.4}$$

A stronger temperature gradient between the upper ocean and deep ocean boxes requires a stronger thermal forcing N to transform light waters into dense waters.

This version of the model is time-dependent and conserves volume. The low-latitude thermocline box has a variable depth $D_{therm}(t)$ that changes according to the volume convergence or divergence:

$$A_{low} \frac{dD_{therm}(t)}{dt} = q_{NA}(t) + q_v(t) + q_{SO}(t). \tag{A.5}$$

The model is spun up so that each box attains its equilibrium preindustrial volume for the given wind stress τ before external forcings from carbon emissions and carbon dioxide removal are imposed.

A.0.2 Forcings of carbon and heat into the ocean-atmosphere system

The models are spun up until they reach an equilibrium state in which, globally, there is no net heat or carbon flux between the ocean and atmosphere under a preindustrial atmospheric CO_2 mixing ratio of 280 ppm. After spinup the models show regional fluxes of heat and carbon (carbon uptake at high latitudes/release at low latitudes, heat uptake at low latitudes/release at high latitudes) that form vertical temperature and carbon gradients consistent with those found in the real world (Katavouta et al., 2019).

An addition of CO_2 to the atmosphere induces fluxes of both heat and carbon. Global fluxes of heat and carbon are calculated using surface average ocean heat and carbon. The ocean/atmosphere flux of carbon, $F_{DIC}(t)$, is a function of ocean density, the air-sea gas transfer coefficient K_g, and the difference between surface ocean DIC $[CO_2(t)]_{surf}$ and the product of the solubility coefficient K_0 and the atmospheric CO_2 mixing ratio, which approximates the air-sea gas transfer possible:

$$F_{DIC}(t) = \rho_0 K_g[K_0(t)\mathrm{CO}_2(t) - [\mathrm{CO}_2(t)]_{surf}]. \tag{A.6}$$

Changes to atmospheric carbon $C_{atm}(t)$ from emissions or carbon dioxide removal impact the atmospheric radiative forcing $R(t)$:

$$R(t) = a\ln(C_{atm}(t)/C_{atm}(t_0)), \tag{A.7}$$

where $C_{atm}(t_0) = 280$ppm is the preindustrial CO_2 mixing ratio used to spin up each model setup. The radiative forcing produces a radiative response, set to be a product of the climate feedback parameter λ and the atmospheric surface temperature anomaly $T'_{atm}(t)$, and a planetary heat uptake $N(t)$:

$$R(t) = \lambda T'_{atm}(t) + N(t). \tag{A.8}$$

The planetary heat uptake is partitioned between the atmosphere (N_{atm}) and ocean (N_o), and the atmospheric heat uptake is set by the air-sea heat transfer parameter c and the difference between the temperature anomalies in the surface ocean and the atmosphere, $T'_{surf}(t) - T'_{atm}(t)$:

$$N = N_{atm} + N_o, \tag{A.9}$$

$$N_{atm} = c(T'_{surf}(t) - T'_{atm}(t)). \tag{A.10}$$

The transfer of carbon and heat into the ocean is partitioned unequally between the low-latitude mixed layer box, Northern Hemisphere high latitude box, and Southern Hemisphere

high latitude box: each high latitude box receives 40% of the added heat and carbon, and the low-latitude mixed layer box receives the remaining 20%.

The evolution of heat, carbon, and alkalinity tracers within the ocean boxes follows the general tracer equation. For a scalar quantity C, the generalised tracer equation for an ocean box of (potentially-variable) volume $V(t)$ is

$$\frac{d}{dt}[C(t)V(t)] = S_c(t) - \int \nabla \cdot [C(t)\mathbf{v}(t)]dV, \tag{A.11}$$

where $S_c(t)$ denotes the source terms for the tracer and $\mathbf{v}(t)$ denotes the velocity within the ocean box. The change of volume is important for the deep ocean and thermocline boxes as the change in thermocline depth according to (A.5) impacts the partitioning of ocean volume between the deep ocean and thermocline boxes only. By Gauss' Theorem, a flux $\int \nabla \cdot C\mathbf{v}dV$ is equivalent to integrating the tracer transports over the box surface $\mathbf{S}$:

$$\int \nabla \cdot (C\mathbf{v})dV = \oint_S C\mathbf{v} \cdot d\mathbf{S}. \tag{A.12}$$

For the model, this is equivalent to summing up the product of the volume transports and tracer concentrations involved in each box, which are visible in the Gnanadesikan schematic (Fig. 5.1).

Bibliography

Allen, M. R., Frame, D. J., Huntingford, C., Jones, C. D., Lowe, J. A., Meinshausen, M., and Meinshausen, N. (2009). Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature*, 458(7242):1163–1166.

Allison, L. C., Roberts, C. D., Palmer, M. D., Hermanson, L., Killick, R. E., Rayner, N. A., Smith, D. M., and Andrews, M. B. (2019). Towards quantifying uncertainty in ocean heat content changes using synthetic profiles. 14(8):084037.

Andrews, T., Andrews, M. B., Bodas-Salcedo, A., Jones, G. S., Kuhlbrodt, T., Manners, J., Menary, M. B., Ridley, J., Ringer, M. A., Sellar, A. A., et al. (2019). Forcings, feedbacks, and climate sensitivity in HadGEM3-GC 3.1 and UKESM1. *Journal of Advances in Modeling Earth Systems*, 11(12):4377–4394.

Andrews, T., Gregory, J. M., and Webb, M. J. (2015). The dependence of radiative forcing and feedback on evolving patterns of surface temperature change in climate models. *Journal of Climate*, 28(4):1630–1648.

Andrews, T. and Webb, M. J. (2018). The dependence of global cloud and lapse rate feedbacks on the spatial structure of tropical Pacific warming. *Journal of Climate*, 31(2):641–654.

Archer, D. (1996). A data-driven model of the global calcite lysocline. *Global Biogeochemical Cycles*, 10(3):511–526.

Argo (2021). Argo float data and metadata from Global Data Assembly Centre (Argo GDAC).

Armour, K. C., Bitz, C. M., and Roe, G. H. (2013). Time-varying climate sensitivity from regional feedbacks. *Journal of Climate*, 26(13):4518–4534.

Arora, V. K., Katavouta, A., Williams, R. G., Jones, C. D., Brovkin, V., Friedlingstein, P., Schwinger, J., Bopp, L., Boucher, O., Cadule, P., Chamberlain, M. A., Christian, J. R., Delire, C., Fisher, R. A., Hajima, T., Ilyina, T., Joetzjer, E., Kawamiya, M., Koven, C., Krasting, J., Law, R. M., Lawrence, D. M., Lenton, A., Lindsay, K., Pongratz, J.,

Raddatz, T., Séférian, R., Tachiiri, K., Tjiputra, J. F., Wiltshire, A., Wu, T., and Ziehn, T. (2019). Carbon-concentration and carbon-climate feedbacks in CMIP6 models, and their comparison to CMIP5 models. *Biogeosciences Discussions*, 2019:1–124.

Arrhenius, S. (1896). On the influence of carbonic acid in the air upon the temperature of the ground. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 41(251):237–276.

Bach, L. T., Gill, S. J., Rickaby, R. E. M., Gore, S., and Renforth, P. (2019). CO_2 removal with enhanced weathering and ocean alkalinity enhancement: Potential risks and co-benefits for marine pelagic ecosystems. *Frontiers in Climate*, 1:7.

Bates, N. R. (2001). Interannual variability of oceanic CO_2 and biogeochemical properties in the Western North Atlantic subtropical gyre. *Deep Sea Research Part II: Topical Studies in Oceanography*, 48(8-9):1507–1528.

Bates, N. R. (2012). Multi-decadal uptake of carbon dioxide into subtropical mode water of the North Atlantic Ocean. *Biogeosciences*, 9(7):2649–2659.

Bates, N. R., Astor, Y. M., Church, M. J., Currie, K., Dore, J. E., González-Dávila, M., Lorenzoni, L., Muller-Karger, F., Olafsson, J., and Santana-Casiano, J. M. (2014). A time-series view of changing ocean chemistry due to ocean uptake of anthropogenic CO_2 and ocean acidification. *Oceanography*, 27.

Beerling, D. J., Kantzas, E. P., Lomas, M. R., Wade, P., Eufrasio, R. M., Renforth, P., Sarkar, B., Andrews, M. G., James, R. H., Pearce, C. R., et al. (2020). Potential for large-scale CO_2 removal via enhanced rock weathering with croplands. *Nature*, 583(7815):242–248.

Beerling, D. J., Leake, J. R., Long, S. P., Scholes, J. D., Ton, J., Nelson, P. N., Bird, M., Kantzas, E., Taylor, L. L., Sarkar, B., et al. (2018). Farming with crops and rocks to address global climate, food and soil security. *Nature Plants*, 4(3):138–147.

Bittig, H. C., Steinhoff, T., Claustre, H., Fiedler, B., Williams, N. L., Sauzède, R., Körtzinger, A., and Gattuso, J.-P. (2018). An alternative to static climatologies: Robust estimation of open ocean CO_2 variables and nutrient concentrations from T, S, and O_2 data using Bayesian neural networks. *Frontiers in Marine Science*, 5:328.

Bitz, C., Holland, M., Weaver, A., and Eby, M. (2001). Simulating the ice-thickness distribution in a coupled climate model. *Journal of Geophysical Research: Oceans*, 106(C2):2441–2463.

Boucher, O., Servonnat, J., Albright, A. L., Aumont, O., Balkanski, Y., Bastrikov, V., Bekki, S., Bonnet, R., Bony, S., Bopp, L., et al. (2020). Presentation and evaluation of the IPSL-CM6A-LR climate model. *Journal of Advances in Modeling Earth Systems*,

12(7):e2019MS002010.

Breeden, M. L. and McKinley, G. A. (2016). Climate impacts on multidecadal pco_2 variability in the
north atlantic: 1948–2009. *Biogeosciences*, 13(11):3387–3396.

Brix, H., Gruber, N., and Keeling, C. D. (2004). Interannual variability of the upper ocean carbon cycle at station ALOHA near Hawaii. *Global Biogeochemical Cycles*, 18(4).

Bronselaer, B. and Zanna, L. (2020). Heat and carbon coupling reveals ocean warming due to circulation changes. *Nature*, 584(7820):227–233.

Budyko, M. I. (1969). The effect of solar radiation variations on the climate of the earth. *Tellus*, 21(5):611–619.

Burger, F. A., John, J. G., and Frölicher, T. L. (2020). Increase in ocean acidity variability and extremes under increasing atmospheric CO_2. *Biogeosciences*, 17(18):4633–4662.

Bushinsky, S. M., Takeshita, Y., and Williams, N. L. (2019). Observing changes in ocean carbonate chemistry: our autonomous future. *Current climate change reports*, 5(3):207–220.

Canadell, J., Monteiro, P., Costa, M., Cotrim da Cunha, L., Cox, P., Eliseev, A., Henson, S., Ishii, M., Jaccard, S., Koven, C., Lohila, A., Patra, P., Piao, S., Rogelj, J., S., S., Zaehle, S., and Zickfeld, K. (2021). *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Global Carbon and other Biogeochemical Cycles and Feedbacks. Cambridge University Press.

Carter, B. R., Feely, R. A., Williams, N. L., Dickson, A. G., Fong, M. B., and Takeshita, Y. (2018). Updated methods for global locally interpolated estimation of alkalinity, pH, and nitrate. *Limnology and Oceanography: Methods*, 16(2):119–131.

Ceppi, P. and Gregory, J. M. (2017). Relationship of tropospheric stability to climate sensitivity and Earth's observed radiation budget. *Proceedings of the National Academy of Sciences*, 114(50):13126–13131.

Charney, J. G., Arakawa, A., Baker, D. J., Bolin, B., Dickinson, R. E., Goody, R. M., Leith, C. E., Stommel, H. M., and Wunsch, C. I. (1979). *Carbon dioxide and climate: a scientific assessment*. National Academy of Sciences, Washington, DC.

Cheng, L., Trenberth, K. E., Fasullo, J. T., Mayer, M., Balmaseda, M., and Zhu, J. (2019). Evolution of ocean heat content related to ENSO. *Journal of Climate*, 32(12):3529 – 3556.

Cheng, L. and Zhu, J. (2016). Benefits of CMIP5 multimodel ensemble in reconstructing historical ocean subsurface temperature variations. *Journal of Climate*, 29(15):5393–5416.

Claussen, M., Mysak, L., Weaver, A., Crucifix, M., Fichefet, T., Loutre, M.-F., Weber, S., Alcamo, J., Alexeev, V., Berger, A., et al. (2002). Earth system models of intermediate complexity: closing the gap in the spectrum of climate system models. *Climate dynamics*, 18(7):579–586.

Claustre, H., Johnson, K. S., and Takeshita, Y. (2020). Observing the global ocean with Biogeochemical-Argo. *Annual Review of Marine Science*, 12(1):23–48. PMID: 31433959.

Collins, M., Knutti, R., Arblaster, J., Dufresne, J.-L., Fichefet, T., Friedlingstein, P., Gao, X., Gutowski, W. J., Johns, T., Krinner, G., Shongwe, M., Tebaldi, C., Weaver, A. J., Wehner, M. F., Allen, M. R., Andrews, T., Beyerle, U., Bitz, C. M., Bony, S., and Booth, B. B. (2013). *Long-term Climate Change: Projections, Commitments and Irreversibility*, pages 1029–1136. Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, United Kingdom.

Cox, P. M., Betts, R. A., Jones, C. D., Spall, S. A., and Totterdell, I. J. (2000). Acceleration of global warming due to carbon-cycle feedbacks in a coupled climate model. *Nature*, 408(6809):184–187.

Daley, R. (1991). *Atmospheric data analysis*. Cambridge University Press, Cambridge.

Danabasoglu, G., Lamarque, J.-F., Bacmeister, J., Bailey, D., DuVivier, A., Edwards, J., Emmons, L., Fasullo, J., Garcia, R., Gettelman, A., et al. (2020). The Community Earth System Model version 2 (CESM2). *Journal of Advances in Modeling Earth Systems*, 12(2).

Davies-Barnard, T., Valdes, P. J., Singarayer, J. S., Pacifico, F. M., and Jones, C. D. (2014). Full effects of land use change in the representative concentration pathways. *Environmental Research Letters*, 9(11):114014.

Denvil-Sommer, A., Gehlen, M., and Vrac, M. (2021). Observation system simulation experiments in the Atlantic Ocean for enhanced surface ocean pCO_2 reconstructions. *Ocean Science*, 17(4):1011–1030.

DeVries, T. (2014). The oceanic anthropogenic CO_2 sink: Storage, air-sea fluxes, and transports over the industrial era. *Global Biogeochemical Cycles*, 28(7):631–647.

DeVries, T., Holzer, M., and Primeau, F. (2017). Recent increase in oceanic carbon uptake driven by weaker upper-ocean overturning. *Nature*, 542(7640):215–218.

DeVries, T., Le Quéré, C., Andrews, O., Berthet, S., Hauck, J., Ilyina, T., Landschützer, P., Lenton, A., Lima, I. D., Nowicki, M., Schwinger, J., and Séférian, R. (2019). Decadal trends in the ocean carbon sink. *Proceedings of the National Academy of Sciences*, 116(24):11646–11651.

Eby, M., Weaver, A. J., Alexander, K., Zickfeld, K., Abe-Ouchi, A., Cimatoribus, A. A.,

Crespin, E., Drijfhout, S. S., Edwards, N. R., Eliseev, A. V., and et al. (2013). Historical and idealized climate model experiments: an intercomparison of Earth system models of intermediate complexity. *Climate of the Past*, 9(3):1111–1140.

Ehlert, D., Zickfeld, K., Eby, M., and Gillett, N. (2017). The sensitivity of the proportionality between temperature change and cumulative CO_2 emissions to ocean mixing. *Journal of Climate*, 30(8):2921–2935.

Evensen, G. (2003). The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean dynamics*, 53(4):343–367.

Fanning, A. F. and Weaver, A. J. (1996). An atmospheric energy-moisture balance model: Climatology, interpentadal climate change, and coupling to an ocean general circulation model. *Journal of Geophysical Research: Atmospheres*, 101(D10):15111–15128.

Field, C. B. and Mach, K. J. (2017). Rightsizing carbon dioxide removal. *Science*, 356(6339):706–707.

Follows, M. J., Ito, T., and Dutkiewicz, S. (2006). On the solution of the carbonate chemistry system in ocean biogeochemistry models. *Ocean Modelling*, 12(3):290–301.

Forster, P. M., Maycock, A. C., McKenna, C. M., and Smith, C. J. (2020). Latest climate models confirm need for urgent mitigation. *Nature Climate Change*, 10(1):7–10.

Friedlingstein, P., Meinshausen, M., Arora, V. K., Jones, C. D., Anav, A., Liddicoat, S. K., and Knutti, R. (2014). Uncertainties in CMIP5 climate projections due to carbon cycle feedbacks. *Journal of Climate*, 27(2):511–526.

Friedlingstein, P., O'Sullivan, M., Jones, M. W., Andrew, R. M., Hauck, J., Olsen, A., Peters, G. P., Peters, W., Pongratz, J., Sitch, S., Le Quéré, C., Canadell, J. G., Ciais, P., Jackson, R. B., Alin, S., Aragão, L. E. O. C., Arneth, A., Arora, V., Bates, N. R., Becker, M., Benoit-Cattin, A., Bittig, H. C., Bopp, L., Bultan, S., Chandra, N., Chevallier, F., Chini, L. P., Evans, W., Florentie, L., Forster, P. M., Gasser, T., Gehlen, M., Gilfillan, D., Gkritzalis, T., Gregor, L., Gruber, N., Harris, I., Hartung, K., Haverd, V., Houghton, R. A., Ilyina, T., Jain, A. K., Joetzjer, E., Kadono, K., Kato, E., Kitidis, V., Korsbakken, J. I., Landschützer, P., Lefèvre, N., Lenton, A., Lienert, S., Liu, Z., Lombardozzi, D., Marland, G., Metzl, N., Munro, D. R., Nabel, J. E. M. S., Nakaoka, S.-I., Niwa, Y., O'Brien, K., Ono, T., Palmer, P. I., Pierrot, D., Poulter, B., Resplandy, L., Robertson, E., Rödenbeck, C., Schwinger, J., Séférian, R., Skjelvan, I., Smith, A. J. P., Sutton, A. J., Tanhua, T., Tans, P. P., Tian, H., Tilbrook, B., van der Werf, G., Vuichard, N., Walker, A. P., Wanninkhof, R., Watson, A. J., Willis, D., Wiltshire, A. J., Yuan, W., Yue, X., and Zaehle, S. (2020). Global carbon budget 2020. *Earth System Science Data*, 12(4):3269–3340.

Fröb, F., Olsen, A., Våge, K., Moore, G., Yashayaev, I., Jeansson, E., and Rajasakaren, B. (2016). Irminger sea deep convection injects oxygen and anthropogenic carbon to the ocean interior. *Nature communications*, 7(1):1–8.

Frölicher, T. L., Sarmiento, J. L., Paynter, D. J., Dunne, J. P., Krasting, J. P., and Winton, M. (2015). Dominance of the Southern Ocean in anthropogenic carbon and heat uptake in CMIP5 models. *Journal of Climate*, 28(2):862–886.

Fujieki, L. A., Santiago-Mandujano, F., Pacheco, F. C., Trifonova, K., Fitzgerald, D., Potemra, J., and White, A. (2018). Hawaii ocean time-series data report 30: 2018. Technical report, University of Hawaii, School of Ocean and Earth Science and Technology.

Fuss, S., Lamb, W. F., Callaghan, M. W., Hilaire, J., Creutzig, F., Amann, T., Beringer, T., de Oliveira Garcia, W., Hartmann, J., Khanna, T., Luderer, G., Nemet, G. F., Rogelj, J., Smith, P., Vicente, J. L. V., Wilcox, J., del Mar Zamora Dominguez, M., and Minx, J. C. (2018). Negative emissions—part 2: Costs, potentials and side effects. *Environmental Research Letters*, 13(6):063002.

Gidden, M. J., Riahi, K., Smith, S. J., Fujimori, S., Luderer, G., Kriegler, E., van Vuuren, D. P., van den Berg, M., Feng, L., Klein, D., Calvin, K., Doelman, J. C., Frank, S., Fricko, O., Harmsen, M., Hasegawa, T., Havlik, P., Hilaire, J., Hoesly, R., Horing, J., Popp, A., Stehfest, E., and Takahashi, K. (2019). Global emissions pathways under different socioeconomic scenarios for use in CMIP6: a dataset of harmonized emissions trajectories through the end of the century. *Geoscientific Model Development*, 12(4):1443–1475.

Gill, A. E. (1982). *Atmosphere-ocean dynamics*, volume 30. Academic press.

Gillett, N. P., Arora, V. K., Matthews, D., and Allen, M. R. (2013). Constraining the ratio of global warming to cumulative CO_2 emissions using CMIP5 simulations. *Journal of Climate*, 26(18):6844–6858.

Gloege, L., McKinley, G. A., Landschützer, P., Fay, A. R., Frölicher, T. L., Fyfe, J. C., Ilyina, T., Jones, S., Lovenduski, N. S., Rodgers, K. B., Schlunegger, S., and Takano, Y. (2021). Quantifying errors in observationally based estimates of ocean carbon sink variability. *Global Biogeochemical Cycles*, 35(4):e2020GB006788. e2020GB006788 2020GB006788.

Gloege, L., Yan, M., Zheng, T., and McKinley, G. A. (subm). Improved quantification of ocean carbon uptake by using machine learning to merge global models and pCO_2 data. *Earth and Space Science Open Archive*, page 28.

Gnanadesikan, A. (1999). A simple predictive model for the structure of the oceanic pycnocline. *Science*, 283(5410):2077–2079.

Goodwin, P. and Cael, B. B. (2021). Bayesian estimation of Earth's climate sensitivity and transient climate response from observational warming and heat content datasets. *Earth System Dynamics*, 12(2):709–723.

Goodwin, P., Katavouta, A., Roussenov, V. M., Foster, G. L., Rohling, E. J., and Williams, R. G. (2018). Pathways to 1.5 °C and 2 °C warming based on observational and geological constraints. *Nature Geoscience*, 11(2):102.

Goodwin, P., Leduc, M., Partanen, A.-I., Matthews, H. D., and Rogers, A. (2020). A computationally efficient method for probabilistic local warming projections constrained by history matching and pattern scaling, demonstrated by wasp–lgrtc-1.0. *Geoscientific Model Development*, 13(11):5389–5399.

Goodwin, P., Williams, R. G., and Ridgwell, A. (2014). Sensitivity of climate to cumulative carbon emissions due to compensation of ocean heat and carbon uptake. *Nature Geoscience*, 8(1):29–34.

Gregory, J., Andrews, T., and Good, P. (2015). The inconstancy of the transient climate response parameter under increasing CO_2. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373.

Gregory, J. M., Stouffer, R. J., Raper, S. C. B., Stott, P. A., and Rayner, N. A. (2002). An observationally based estimate of the climate sensitivity. *Journal of Climate*, 15(22):3117–3121.

Gruber, N., Clement, D., Carter, B. R., Feely, R. A., van Heuven, S., Hoppema, M., Ishii, M., Key, R. M., Kozyr, A., Lauvset, S. K., Lo Monaco, C., Mathis, J. T., Murata, A., Olsen, A., Perez, F. F., Sabine, C. L., Tanhua, T., and Wanninkhof, R. (2019a). The oceanic sink for anthropogenic CO_2 from 1994 to 2007. *Science*, 363(6432):1193–1199.

Gruber, N., Gloor, M., Mikaloff Fletcher, S. E., Doney, S. C., Dutkiewicz, S., Follows, M. J., Gerber, M., Jacobson, A. R., Joos, F., Lindsay, K., Menemenlis, D., Mouchet, A., Müller, S. A., Sarmiento, J. L., and Takahashi, T. (2009). Oceanic sources, sinks, and transport of atmospheric CO_2. *Global Biogeochemical Cycles*, 23(1).

Gruber, N., Landschützer, P., and Lovenduski, N. S. (2019b). The variable Southern Ocean carbon sink. *Annual Review of Marine Science*, 11(1):159–186. PMID: 30212259.

Heinze, C., Meyer, S., Goris, N., Anderson, L., Steinfeldt, R., Chang, N., Le Quéré, C., and Bakker, D. C. E. (2015). The ocean carbon sink – impacts, vulnerabilities and challenges. *Earth System Dynamics*, 6(1):327–358.

Herrington, T. and Zickfeld, K. (2014). Path independence of climate and carbon cycle response over a broad range of cumulative carbon emissions. *Earth System Dynamics*, 5(2):409–422.

Hewitt, H. T., Bell, M. J., Chassignet, E. P., Czaja, A., Ferreira, D., Griffies, S. M., Hyder, P., McClean, J. L., New, A. L., and Roberts, M. J. (2017). Will high-resolution global ocean models benefit coupled predictions on short-range to climate timescales? *Ocean Modelling*, 120:120–136.

Huber, M. B. and Zanna, L. (2017). Drivers of uncertainty in simulated ocean circulation and heat uptake. *Geophysical Research Letters*, 44(3):1402–1413.

Hunke, E. C. and Dukowicz, J. K. (1997). An elastic–viscous–plastic model for sea ice dynamics. *Journal of Physical Oceanography*, 27(9):1849 – 1867.

International Energy Agency (2021). Global energy review 2021.

IPCC (2018). *Global Warming of 1.5 °C. An IPCC Special Report on the impacts of global warming of 1.5 °C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty.* Cambridge University Press.

IPCC (2021a). *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Summary for Policymakers. Cambridge University Press.

IPCC (2021b). *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Annex II: Models. Cambridge University Press.

Irving, D., Wijffels, S., and Church, J. (2019). Anthropogenic aerosols, greenhouse gases, and the uptake, transport, and storage of excess heat in the climate system. *Geophysical Research Letters*, 46(9):4894–4903.

Ito, T. and Follows, M. J. (2005). Preformed phosphate, soft tissue pump and atmospheric CO_2. *Journal of Marine Research*, 63(4):813–839.

Iudicone, D., Rodgers, K. B., Plancherel, Y., Aumont, O., Ito, T., Key, R. M., Madec, G., and Ishii, M. (2016). The formation of the ocean's anthropogenic carbon reservoir. *Scientific Reports*, 6(1):1–16.

Jeltsch-Thömmes, A., Stocker, T. F., and Joos, F. (2020). Hysteresis of the Earth system under positive and negative CO2 emissions. *Environmental Research Letters*, 15(12):124026.

Jones, C. D., Ciais, P., Davis, S. J., Friedlingstein, P., Gasser, T., Peters, G. P., Rogelj, J., van Vuuren, D. P., Canadell, J. G., Cowie, A., Jackson, R. B., Jonas, M., Kriegler, E., Littleton, E., Lowe, J. A., Milne, J., Shrestha, G., Smith, P., Torvanger, A., and Wiltshire, A. (2016). Simulating the Earth system response to negative emissions. *Environmental Research Letters*, 11(9):095012.

Jones, C. D., Collins, M., Cox, P. M., and Spall, S. A. (2001). The carbon cycle response to ENSO: A coupled climate–carbon cycle model study. *Journal of Climate*, 14(21):4113 – 4129.

Jones, C. D. and Friedlingstein, P. (2020). Quantifying process-level uncertainty contributions to TCRE and carbon budgets for meeting paris agreement climate targets. 15(7):074019.

Jones, S. D., Le Quéré, C., Rödenbeck, C., Manning, A. C., and Olsen, A. (2015). A statistical gap-filling method to interpolate global monthly surface ocean carbon dioxide data. *Journal of Advances in Modeling Earth Systems*, 7(4):1554–1575.

Katavouta, A. and Williams, R. G. (2021). Ocean carbon cycle feedbacks in CMIP6 models: contributions from different basins. *Biogeosciences*, 18(10):3189–3218.

Katavouta, A., Williams, R. G., and Goodwin, P. (2019). The effect of ocean ventilation on the transient climate response to emissions. *Journal of Climate*, 32(16):5085 – 5105.

Katavouta, A., Williams, R. G., Goodwin, P., and Roussenov, V. M. (2018). Reconciling atmospheric and oceanic views of the transient climate response to emissions. *Geophysical Research Letters*, 45(12):6205–6214.

Keeling, C. D., Brix, H., and Gruber, N. (2004). Seasonal and long-term dynamics of the upper ocean carbon cycle at station ALOHA near Hawaii. *Global Biogeochemical Cycles*, 18(4).

Keller, D., Oschlies, A., and Eby, M. (2012a). A new marine ecosystem model for the University of Victoria Earth System Climate Model. *Geoscientific Model Development*, 5(5):1195–1220.

Keller, D. P., Lenton, A., Littleton, E. W., Oschlies, A., Scott, V., and Vaughan, N. E. (2018a). The effects of carbon dioxide removal on the carbon cycle. *Current climate change reports*, 4(3):250–265.

Keller, D. P., Lenton, A., Scott, V., Vaughan, N. E., Bauer, N., Ji, D., Jones, C. D., Kravitz, B., Muri, H., and Zickfeld, K. (2018b). The Carbon Dioxide Removal Model Intercomparison Project (CDRMIP): rationale and experimental protocol for CMIP6. *Geoscientific Model Development*, 11(3):1133–1160.

Keller, K. M., Joos, F., Raible, C. C., Cocco, V., Frölicher, T. L., Dunne, J. P., Gehlen, M., Bopp, L., Orr, J. C., Tjiputra, J., Heinze, C., Segschneider, J., Roy, T., and Metzl, N. (2012b). Variability of the ocean carbon cycle in response to the North Atlantic Oscillation. *Tellus B: Chemical and Physical Meteorology*, 64(1):18738.

Khatiwala, S., Primeau, F., and Hall, T. (2009). Reconstruction of the history of anthropogenic CO_2 concentrations in the ocean. *Nature*, 462(7271):346–349.

Krasting, J., Dunne, J., Shevliakova, E., and Stouffer, R. (2014). Trajectory sensitivity of the transient climate response to cumulative carbon emissions. *Geophysical Research Letters*, 41:2520–2527.

Landschützer, P., Gruber, N., and Bakker, D. C. E. (2016). Decadal variations and trends of the global ocean carbon sink. *Global Biogeochemical Cycles*, 30(10):1396–1417.

Landschützer, P., Ilyina, T., and Lovenduski, N. S. (2019). Detecting regional modes of variability in observation-based surface ocean pCO_2. *Geophysical Research Letters*, 46(5):2670–2679.

Lauderdale, J. M., Garabato, A. C. N., Oliver, K. I., Follows, M. J., and Williams, R. G. (2013). Wind-driven changes in southern ocean residual circulation, ocean carbon reservoirs and atmospheric CO_2. *Climate dynamics*, 41(7-8):2145–2164.

Lauvset, S. K., Key, R. M., Olsen, A., van Heuven, S., Velo, A., Lin, X., Schirnick, C., Kozyr, A., Tanhua, T., Hoppema, M., Jutterström, S., Steinfeldt, R., Jeansson, E., Ishii, M., Perez, F. F., Suzuki, T., and Watelet, S. (2016). A new global interior ocean mapped climatology: the $1^\circ \times 1^\circ$ GLODAP version 2. *Earth System Science Data*, 8(2):325–340.

Le Quéré, C., Rödenbeck, C., Buitenhuis, E. T., Conway, T. J., Langenfelds, R., Gomez, A., Labuschagne, C., Ramonet, M., Nakazawa, T., Metzl, N., Gillett, N., and Heimann, M. (2007). Saturation of the southern ocean CO_2 sink due to recent climate change. *Science*, 316(5832):1735–1738.

Leduc, M., Matthews, H. D., and de Elía, R. (2016). Regional estimates of the transient climate response to cumulative CO_2 emissions. *Nature Climate Change*, 6:474–478.

Leduc, M., Matthews, H. D., and de Elía, R. (2015). Quantifying the limits of a linear temperature response to cumulative co2 emissions. *Journal of Climate*, 28(24):9955 – 9968.

Li, H. and Ilyina, T. (2018). Current and future decadal trends in the oceanic carbon uptake are dominated by internal variability. *Geophysical Research Letters*, 45(2):916–925.

Liddicoat, S. K., Booth, B. B. B., and Joshi, M. M. (2016). An investigation into linearity with cumulative emissions of the climate and carbon cycle response in HadCM3lc. *Environmental Research Letters*, 11(6):065003.

Lovenduski, N. S., Gruber, N., and Doney, S. C. (2008). Toward a mechanistic understanding of the decadal trends in the southern ocean carbon sink. *Global Biogeochemical Cycles*, 22(3).

MacDougall, A. H. (2017). The oceanic origin of path-independent carbon budgets. *Scientific Reports*, 7(1):10373.

MacDougall, A. H. and Friedlingstein, P. (2015). The origin and limits of the near propor-

tionality between climate warming and cumulative CO_2 emissions. *Journal of Climate*, 28(10):4217–4230.

MacDougall, A. H., Frölicher, T. L., Jones, C. D., Rogelj, J., Matthews, H. D., Zickfeld, K., Arora, V. K., Barrett, N. J., Brovkin, V., Burger, F. A., Eby, M., Eliseev, A. V., Hajima, T., Holden, P. B., Jeltsch-Thömmes, A., Koven, C., Mengis, N., Menviel, L., Michou, M., Mokhov, I. I., Oka, A., Schwinger, J., Séférian, R., Shaffer, G., Sokolov, A., Tachiiri, K., Tjiputra, J., Wiltshire, A., and Ziehn, T. (2020). Is there warming in the pipeline? A multi-model analysis of the Zero Emissions Commitment from CO_2. *Biogeosciences*, 17(11):2987–3016.

MacDougall, A. H., Swart, N. C., and Knutti, R. (2017). The uncertainty in the transient climate response to cumulative CO_2 emissions arising from the uncertainty in physical climate parameters. *Journal of Climate*, 30(2):813–827.

Manabe, S., Bryan, K., and Spelman, M. J. (1975). A global ocean-atmosphere climate model. part i. the atmospheric circulation. *Journal of Physical Oceanography*, 5(1):3 – 29.

Manabe, S. and Wetherald, R. T. (1967). Thermal equilibrium of the atmosphere with a given distribution of relative humidity. *Journal of Atmospheric Sciences*, 24(3):241 – 259.

Marshall, D. P. and Zanna, L. (2014). A conceptual model of ocean heat uptake under climate change. *Journal of Climate*, 27(22):8444 – 8465.

Marshall, J., Donohoe, A., Ferreira, D., and McGee, D. (2014a). The ocean's role in setting the mean position of the inter-tropical convergence zone. *Climate Dynamics*, 42(7):1967–1979.

Marshall, J., Scott, J. R., Armour, K. C., Campin, J.-M., Kelley, M., and Romanou, A. (2014b). The ocean's role in the transient response of climate to abrupt greenhouse gas forcing. *Climate Dynamics*, 44(7-8):2287–2299.

Matthews, H. D., Gillett, N. P., Stott, P. A., and Zickfeld, K. (2009). The proportionality of global warming to cumulative carbon emissions. *Nature*, 459(7248):829–832.

Matthews, H. D., Landry, J.-S., Partanen, A.-I., Allen, M., Eby, M., Forster, P. M., Friedlingstein, P., and Zickfeld, K. (2017). Estimating carbon budgets for ambitious climate targets. *Current Climate Change Reports*, 3(1):69–77.

Matthews, H. D., Tokarska, K. B., Rogelj, J., Smith, C. J., MacDougall, A. H., Haustein, K., Mengis, N., Sippel, S., Forster, P. M., and Knutti, R. (2021). An integrated approach to quantifying uncertainties in the remaining carbon budget. *Communications Earth & Environment*, 2(1):1–11.

Mazloff, M. R., Cornuelle, B. D., Gille, S. T., and Verdy, A. (2018). Correlation lengths

for estimating the large-scale carbon and heat content of the southern ocean. *Journal of Geophysical Research: Oceans*, 123(2):883–901.

McKinley, G. A., Fay, A. R., Eddebbar, Y. A., Gloege, L., and Lovenduski, N. S. (2020). External forcing explains recent decadal variability of the ocean carbon sink. *AGU Advances*, 1(2):e2019AV000149. e2019AV000149 2019AV000149.

McKinley, G. A., Fay, A. R., Lovenduski, N. S., and Pilcher, D. J. (2017). Natural variability and anthropogenic trends in the ocean carbon sink. *Annual Review of Marine Science*, 9(1):125–150. PMID: 27620831.

McKinley, G. A., Pilcher, D. J., Fay, A. R., Lindsay, K., Long, M. C., and Lovenduski, N. S. (2016). Timescales for detection of trends in the ocean carbon sink. *Nature*, 530(7591):469–472.

Meinshausen, M., Meinshausen, N., Hare, W., Raper, S. C. B., Frieler, K., Knutti, R., Frame, D. J., and Allen, M. R. (2009). Greenhouse-gas emission targets for limiting global warming to 2 °C. *Nature*, 458(7242):1158–1162.

Meinshausen, M., Smith, S. J., Calvin, K., Daniel, J. S., Kainuma, M. L. T., Lamarque, J.-F., Matsumoto, K., Montzka, S. A., Raper, S. C. B., Riahi, K., and et al. (2011a). The RCP greenhouse gas concentrations and their extensions from 1765 to 2300. *Climatic Change*, 109(1-2):213–241.

Meinshausen, M., Vogel, E., Nauels, A., Lorbacher, K., Meinshausen, N., Etheridge, D. M., Fraser, P. J., Montzka, S. A., Rayner, P. J., Trudinger, C. M., Krummel, P. B., Beyerle, U., Canadell, J. G., Daniel, J. S., Enting, I. G., Law, R. M., Lunder, C. R., O'Doherty, S., Prinn, R. G., Reimann, S., Rubino, M., Velders, G. J. M., Vollmer, M. K., Wang, R. H. J., and Weiss, R. (2017). Historical greenhouse gas concentrations for climate modelling (CMIP6). *Geoscientific Model Development*, 10(5):2057–2116.

Meinshausen, M., Wigley, T. M. L., and Raper, S. C. B. (2011b). Emulating atmosphere-ocean and carbon cycle models with a simpler model, MAGICC6 – part 2: Applications. *Atmospheric Chemistry and Physics*, 11(4):1457–1471.

Meissner, K., Weaver, A., Matthews, H., and Cox, P. (2003). The role of land surface dynamics in glacial inception: a study with the UVic Earth System Model. *Climate Dynamics*, 21(7):515–537.

Mengis, N., Partanen, A.-I., Jalbert, J., and Matthews, H. D. (2018). 1.5 °c carbon budget dependent on carbon cycle uncertainty and future non-CO_2 forcing. *Scientific Reports*, 8(1).

Mikaloff Fletcher, S. E., Gruber, N., Jacobson, A. R., Doney, S. C., Dutkiewicz, S., Gerber, M., Follows, M., Joos, F., Lindsay, K., Menemenlis, D., Mouchet, A., Müller, S. A., and

Sarmiento, J. L. (2006). Inverse estimates of anthropogenic CO_2 uptake, transport, and storage by the ocean. *Global Biogeochemical Cycles*, 20(2).

Millar, R. J., Fuglestvedt, J. S., Friedlingstein, P., Rogelj, J., Grubb, M. J., Matthews, H. D., Skeie, R. B., Forster, P. M., Frame, D. J., and Allen, M. R. (2017). Emission budgets and pathways consistent with limiting warming to 1.5 °C. *Nature Geoscience*, 10(10):741–747.

Mitchell, J. F., Johns, T., Gregory, J. M., and Tett, S. (1995). Climate response to increasing levels of greenhouse gases and sulphate aerosols. *Nature*, 376(6540):501–504.

Müller, W. A., Jungclaus, J. H., Mauritsen, T., Baehr, J., Bittner, M., Budich, R., Bunzel, F., Esch, M., Ghosh, R., Haak, H., et al. (2018). A higher-resolution version of the Max Planck Institute Earth System Model (MPI-ESM1.2-HR). *Journal of Advances in Modeling Earth Systems*, 10(7):1383–1413.

Oke, P. R., Allen, J. S., Miller, R. N., Egbert, G. D., and Kosro, P. M. (2002). Assimilation of surface velocity data into a primitive equation coastal ocean model. *Journal of Geophysical Research: Oceans*, 107(C9):5–1–5–25.

Oke, P. R., Brassington, G. B., Griffin, D. A., and Schiller, A. (2010). Ocean data assimilation: a case for ensemble optimal interpolation. *Australian Meteorological and Oceanographic Journal*, 59(Sp. Iss.):67–76.

Paynter, D. J. and Frölicher, T. L. (2015). Sensitivity of radiative forcing, ocean heat uptake,and climate feedback to changes in anthropogenic greenhouse gases and aerosols. *Journal of Geophysical Research: Atmospheres*, 120:9837–9854.

Phillips, H. E. and Joyce, T. M. (2007). Bermuda's tale of two time series: Hydrostation S and BATS. *Journal of Physical Oceanography*, 37(3):554 – 571.

Phillips, N. A. (1956). The general circulation of the atmosphere: A numerical experiment. *Quarterly Journal of the Royal Meteorological Society*, 82(352):123–164.

Renforth, P. and Henderson, G. (2017). Assessing ocean alkalinity for carbon sequestration. *Reviews of Geophysics*, 55(3):636–674.

Rhein, M., Rintoul, S., Aoki, S., Campos, E., Chambers, D., Feely, R., Gulev, S., Johnson, G., Josey, S., Kostianoy, A., Mauritzen, C., Roemmich, D., Talley, L., and Wang, F. (2013). *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, book section Observations: Ocean, pages 255–316. Cambridge University Press, Cambridge, United Kingdom.

Rhein, M., Steinfeldt, R., Kieke, D., Stendardo, I., and Yashayaev, I. (2017). Ventilation variability of labrador sea water and its impact on oxygen and anthropogenic carbon: a

review. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 375(2102):20160321.

Richardson, L. F. (1922). *Weather prediction by numerical process*. Cambridge university press.

Rockström, J., Gaffney, O., Rogelj, J., Meinshausen, M., Nakicenovic, N., and Schellnhuber, H. J. (2017). A roadmap for rapid decarbonization. *Science*, 355(6331):1269–1271.

Rödenbeck, C., Bakker, D. C. E., Gruber, N., Iida, Y., Jacobson, A. R., Jones, S., Landschützer, P., Metzl, N., Nakaoka, S., Olsen, A., Park, G.-H., Peylin, P., Rodgers, K. B., Sasse, T. P., Schuster, U., Shutler, J. D., Valsala, V., Wanninkhof, R., and Zeng, J. (2015). Data-based estimates of the ocean carbon sink variability – first results of the Surface Ocean pCO_2 Mapping intercomparison (SOCOM). *Biogeosciences*, 12(23):7251–7278.

Rogelj, J., Forster, P. M., Kriegler, E., Smith, C. J., and Séférian, R. (2019). Estimating and tracking the remaining carbon budget for stringent climate targets. *Nature*, 571(7765):335–342.

Rogelj, J., Popp, A., Calvin, K. V., Luderer, G., Emmerling, J., Gernaat, D., Fujimori, S., Strefler, J., Hasegawa, T., Marangoni, G., et al. (2018). Scenarios towards limiting global mean temperature increase below 1.5 °C. *Nature Climate Change*, 8(4):325–332.

Rugenstein, M. A., Caldeira, K., and Knutti, R. (2016). Dependence of global radiative feedbacks on evolving patterns of surface heat fluxes. *Geophysical Research Letters*, 43(18):9877–9885.

Rugenstein, M. A. A., Winton, M., Stouffer, R. J., Griffies, S. M., and Hallberg, R. (2013). Northern high-latitude heat budget decomposition and transient warming. *Journal of Climate*, 26(2):609 – 621.

Sabine, C. L., Feely, R. A., Gruber, N., Key, R. M., Lee, K., Bullister, J. L., Wanninkhof, R., Wong, C. S., Wallace, D. W. R., Tilbrook, B., Millero, F. J., Peng, T.-H., Kozyr, A., Ono, T., and Rios, A. F. (2004). The oceanic sink for anthropogenic CO_2. *Science*, 305(5682):367–371.

Schuiling, R. and Krijgsman, P. (2006). Enhanced weathering: an effective and cheap tool to sequester CO_2. *Climatic Change*, 74(1):349–354.

Schwinger, J. and Tjiputra, J. (2018). Ocean carbon cycle feedbacks under negative emissions. *Geophysical Research Letters*, 45(10):5062–5070.

Séférian, R., Gehlen, M., Bopp, L., Resplandy, L., Orr, J. C., Marti, O., Dunne, J. P., Christian, J. R., Doney, S. C., Ilyina, T., Lindsay, K., Halloran, P. R., Heinze, C., Segschneider, J., Tjiputra, J., Aumont, O., and Romanou, A. (2016). Inconsistent strategies to spin up models in CMIP5: implications for ocean biogeochemical model performance assessment.

Geoscientific Model Development, 9(5):1827–1851.

Seland, Ø., Bentsen, M., Olivié, D., Toniazzo, T., Gjermundsen, A., Graff, L. S., Debernard, J. B., Gupta, A. K., He, Y.-C., Kirkevåg, A., et al. (2020). Overview of the Norwegian Earth System Model (NorESM2) and key climate response of CMIP6 DECK, historical, and scenario simulations. *Geoscientific Model Development*, 13(12):6165–6200.

Sellar, A. A., Jones, C. G., Mulcahy, J. P., Tang, Y., Yool, A., Wiltshire, A., O'connor, F. M., Stringer, M., Hill, R., Palmieri, J., et al. (2019). UKESM1: Description and evaluation of the UK Earth System Model. *Journal of Advances in Modeling Earth Systems*, 11(12):4513–4558.

Sen Gupta, A., Jourdain, N. C., Brown, J. N., and Monselesan, D. (2013). Climate drift in the CMIP5 models. *Journal of Climate*, 26(21):8597 – 8615.

Shi, J.-R., Xie, S.-P., and Talley, L. D. (2018). Evolving relative importance of the Southern Ocean and North Atlantic in anthropogenic ocean heat uptake. *Journal of Climate*, 31(18):7459–7479.

Siegel, D. A., DeVries, T., Doney, S. C., and Bell, T. (2021). Assessing the sequestration time scales of some ocean-based carbon dioxide reduction strategies. *Environmental Research Letters*, 16(10):104003.

Smith, C. J., Forster, P. M., Allen, M., Leach, N., Millar, R. J., Passerello, G. A., and Regayre, L. A. (2018). FAIR v1.3: a simple emissions-based impulse response and carbon cycle model. *Geoscientific Model Development*, 11(6):2273–2297.

Smith, D. M. and Murphy, J. M. (2007). An objective ocean temperature and salinity analysis using covariances from a global climate model. *Journal of Geophysical Research: Oceans*, 112(C2).

Smith, D. M., Scaife, A. A., Eade, R., Athanasiadis, P., Bellucci, A., Bethke, I., Bilbao, R., Borchert, L., Caron, L.-P., Counillon, F., et al. (2020). North Atlantic climate far more predictable than models imply. *Nature*, 583(7818):796–800.

Smith, P., Davis, S. J., Creutzig, F., Fuss, S., Minx, J., Gabrielle, B., Kato, E., Jackson, R. B., Cowie, A., Kriegler, E., et al. (2016). Biophysical and economic limits to negative CO_2 emissions. *Nature climate change*, 6(1):42–50.

Solomon, S., Plattner, G.-K., Knutti, R., and Friedlingstein, P. (2009). Irreversible climate change due to carbon dioxide emissions. *Proceedings of the National Academy of Sciences*, 106(6).

Spafford, L. and MacDougall, A. H. (2020). Quantifying the probability distribution function of the transient climate response to cumulative CO_2 emissions. *Environmental Research Letters*, 15.

Stouffer, R., Manabe, S., and Bryan, K. (1989). Interhemispheric asymmetry in climate response to a gradual increase of atmospheric co_2. *Nature*, 342(6250):660–662.

Swart, N. C., Cole, J. N., Kharin, V. V., Lazare, M., Scinocca, J. F., Gillett, N. P., Anstey, J., Arora, V., Christian, J. R., Hanna, S., et al. (2019). The Canadian Earth System Model version 5 (CanESM5.0.3). *Geoscientific Model Development*, 12(11):4823–4873.

Taylor, L. L., Quirk, J., Thorley, R. M., Kharecha, P. A., Hansen, J., Ridgwell, A., Lomas, M. R., Banwart, S. A., and Beerling, D. J. (2016). Enhanced weathering strategies for stabilizing climate and averting ocean acidification. *Nature Climate Change*, 6(4):402–406.

Terhaar, J., Frölicher, T. L., and Joos, F. (2021). Southern ocean anthropogenic carbon sink constrained by sea surface salinity. *Science Advances*, 7(18):eabd5964.

Thomas, J., Waugh, D., and Gnanadesikan, A. (2018). Relationship between ocean carbon and heat multidecadal variability. *Journal of Climate*, 31(4):1467 – 1482.

Tokarska, K. B. and Gillett, N. P. (2018). Cumulative carbon emissions budgets consistent with 1.5 °C global warming. *Nature Climate Change*, 8(4):296–299.

Tokarska, K. B., Gillett, N. P., Arora, V. K., Lee, W. G., and Zickfeld, K. (2018). The influence of non-CO_2 forcings on cumulative carbon emissions budgets. *Environmental Research Letters*, 13:034039.

Tokarska, K. B., Stolpe, M. B., Sippel, S., Fischer, E. M., Smith, C. J., Lehner, F., and Knutti, R. (2020). Past warming trend constrains future warming in CMIP6 models. *Science Advances*, 6(12):eaaz9549.

Tokarska, K. B. and Zickfeld, K. (2015). The effectiveness of net negative carbon dioxide emissions in reversing anthropogenic climate change. *Environmental Research Letters*, 10(9):094013.

Tokarska, K. B., Zickfeld, K., and Rogelj, J. (2019). Path independence of carbon budgets when meeting a stringent global mean temperature target after an overshoot. *Earth's Future*, 7(12):1283–1295.

Weaver, A. J., Eby, M., Wiebe, E. C., Bitz, C. M., Duffy, P. B., Ewen, T. L., Fanning, A. F., Holland, M. M., MacFayden, A., Matthews, H. D., Meissner, K. J., Saenko, O., Schmittner, A., Wang, H., and Yoshimori, M. (2001). The UVic Earth System Climate Model: Model description, climatology, and applications to past, present, and future climates. *Atmosphere-Ocean*, 39(4).

Weijer, W., Cheng, W., Garuba, O. A., Hu, A., and Nadiga, B. T. (2020). CMIP6 models predict significant 21st Century decline of the Atlantic Meridional Overturning Circulation. *Geophysical Research Letters*, 47(12):e2019GL086075. e2019GL086075

10.1029/2019GL086075.

Williams, N. L., Juranek, L. W., Feely, R. A., Johnson, K. S., Sarmiento, J. L., Talley, L. D., Dickson, A. G., Gray, A. R., Wanninkhof, R., Russell, J. L., Riser, S. C., and Takeshita, Y. (2017a). Calculating surface ocean pco_2 from biogeochemical Argo floats equipped with pH: An uncertainty analysis. *Global Biogeochemical Cycles*, 31(3):591–604.

Williams, R. G., Ceppi, P., and Katavouta, A. (2020). Controls of the transient climate response to emissions by physical feedbacks, heat uptake and carbon cycling. *Environmental Research Letters*, 15(9):0940c1.

Williams, R. G. and Follows, M. J. (2011). *Ocean Dynamics and the Carbon Cycle: Principles and Mechanisms*. Cambridge University Press.

Williams, R. G., Goodwin, P., Roussenov, V. M., and Bopp, L. (2016). A framework to understand the transient climate response to emissions. *Environmental Research Letters*, 11(1):015003.

Williams, R. G., Katavouta, A., and Roussenov, V. (2021). Regional asymmetries in ocean heat and carbon storage due to dynamic redistribution in climate model projections. *Journal of Climate*, 34(10):3907 – 3925.

Williams, R. G., Roussenov, V. M., Frölicher, T. L., and Goodwin, P. (2017b). Drivers of continued surface warming after cessation of carbon emissions. *Geophysical Research Letters*, 44:10633–10642.

Williams, R. G., Roussenov, V. M., Goodwin, P., Resplandy, L., and Bopp, L. (2017c). Sensitivity of global warming to carbon emissions: effects of heat and carbon uptake in a suite of Earth system models. *Journal of Climate*, 30:9343–9363.

Winton, M., Griffies, S. M., Samuels, B. L., Sarmiento, J. L., and Frölicher, T. L. (2013). Connecting changing ocean circulation with changing climate. *Journal of Climate*, 26(7):2268–2278.

Winton, M., Takahashi, K., and Held, I. M. (2010). Importance of ocean heat uptake efficacy to transient climate change. *Journal of Climate*, 23(9):2333–2344.

Wong, A. P. S., Wijffels, S. E., Riser, S. C., Pouliquen, S., Hosoda, S., Roemmich, D., Gilson, J., Johnson, G. C., Martini, K., Murphy, D. J., Scanderbeg, M., Bhaskar, T. V. S. U., Buck, J. J. H., Merceur, F., Carval, T., Maze, G., Cabanes, C., André, X., Poffa, N., Yashayaev, I., Barker, P. M., Guinehut, S., Belbéoch, M., Ignaszewski, M., Baringer, M. O., Schmid, C., Lyman, J. M., McTaggart, K. E., Purkey, S. G., Zilberman, N., Alkire, M. B., Swift, D., Owens, W. B., Jayne, S. R., Hersh, C., Robbins, P., West-Mack, D., Bahr, F., Yoshida, S., Sutton, P. J. H., Cancouët, R., Coatanoan, C., Dobbler, D., Juan, A. G., Gourrion, J., Kolodziejczyk, N., Bernard, V., Bourlès, B., Claustre, H.,

D'Ortenzio, F., Le Reste, S., Le Traon, P.-Y., Rannou, J.-P., Saout-Grit, C., Speich, S., Thierry, V., Verbrugge, N., Angel-Benavides, I. M., Klein, B., Notarstefano, G., Poulain, P.-M., Vélez-Belchí, P., Suga, T., Ando, K., Iwasaska, N., Kobayashi, T., Masuda, S., Oka, E., Sato, K., Nakamura, T., Sato, K., Takatsuki, Y., Yoshida, T., Cowley, R., Lovell, J. L., Oke, P. R., van Wijk, E. M., Carse, F., Donnelly, M., Gould, W. J., Gowers, K., King, B. A., Loch, S. G., Mowat, M., Turton, J., Rama Rao, E. P., Ravichandran, M., Freeland, H. J., Gaboury, I., Gilbert, D., Greenan, B. J. W., Ouellet, M., Ross, T., Tran, A., Dong, M., Liu, Z., Xu, J., Kang, K., Jo, H., Kim, S.-D., and Park, H.-M. (2020). Argo data 1999–2019: Two million temperature-salinity profiles and subsurface velocity observations from a global array of profiling floats. *Frontiers in Marine Science*, 7:700.

Zanna, L., Khatiwala, S., Gregory, J. M., Ison, J., and Heimbach, P. (2019). Global reconstruction of historical ocean heat storage and transport. *Proceedings of the National Academy of Sciences*, 116(4):1126–1131.

Zelinka, M. D., Myers, T. A., McCoy, D. T., Po-Chedley, S., Caldwell, P. M., Ceppi, P., Klein, S. A., and Taylor, K. E. (2020). Causes of higher climate sensitivity in CMIP6 models. *Geophysical Research Letters*, 47(1):e2019GL085782. e2019GL085782 10.1029/2019GL085782.

Zhu, J., Poulsen, C. J., and Otto-Bliesner, B. L. (2020). High climate sensitivity in CMIP6 model not supported by paleoclimate. *Nature Climate Change*, 10(5):378–379.

Zickfeld, K., Azevedo, D., Mathesius, S., and Matthews, H. D. (2021). Asymmetry in the climate–carbon cycle response to positive and negative CO_2 emissions. *Nature Climate Change*, pages 1–5.

Zickfeld, K., Eby, M., Matthews, H. D., and Weaver, A. J. (2009). Setting cumulative emissions targets to reduce the risk of dangerous climate change. *Proceedings of the National Academy of Sciences*, 106(38):16129–16134.

Ziehn, T., Chamberlain, M. A., Law, R. M., Lenton, A., Bodman, R. W., Dix, M., Stevens, L., Wang, Y.-P., and Srbinovsky, J. (2020). The Australian Earth System Model: ACCESS-ESM1.5. *Journal of Southern Hemisphere Earth Systems Science*, 70(1):193–214.