UNIVERSITY OF READING

Department of Meteorology

Simulating decadal variability in the North Atlantic Ocean

Matthew Burns Menary

A thesis submitted for the degree of Doctor of Philosophy

October 2015

Declaration

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

Matthew Menary

Abstract

Observations and climate models suggest significant decadal variability within the North Atlantic subpolar gyre (NA SPG), though observations are sparse and models disagree on the details of this variability. Therefore, it is important to understand 1) the mechanisms of simulated decadal variability, 2) which parts of simulated variability are more faithful representations of reality, and 3) the implications for climate predictions.

Here, we investigate the decadal variability in the NA SPG in the state-of-the-art, high resolution (0.25° ocean resolution), climate model 'HadGEM3'. We find a decadal mode with a period of 17 years that explains 30% of the annual variance in related indices.

The mode arises due to the advection of heat content anomalies, and shows asymmetries in the timescale of phase reversal between positive and negative phases. A negative feedback from temperature-driven density anomalies in the Labrador Sea (LS) allows for the phase reversal. The North Atlantic Oscillation (NAO), which exhibits the same periodicity, amplifies the mode. The atmosphere-ocean coupling is stronger during positive rather than negative NAO states, explaining the asymmetry. Within the NA SPG, there is potential predictability arising partly from this mode for up to 5 years.

There are important similarities between observed and simulated variability, such as the apparent role for the propagation of heat content anomalies. However, observations suggest interannual LS density anomalies are salinity-driven. Salinity control of density would change the temperature feedback to the south, possibly limiting real-world predictive skill in the southern NA SPG with this model.

Finally, to understand the diversity of behaviours, we analyse 42 present-generation climate models. Temperature and salinity biases are found to systematically influence the driver of density variability in the LS. Resolution is a good predictor of the biases. The dependence of variability on the background state has important implications for decadal predictions.

Acknowledgements

I have been incredibly fortunate to study for a PhD supervised by four very intelligent, distinct, and enthusiastic supervisors. You know who you are! But for the benefit of the reader, and in no particular order: Dan Hodson, Jon Robson, Rowan Sutton, and Richard Wood. Your intellect meant I had to prepare thoroughly for our meetings; your varied approaches/focus meant I could be quizzed about any aspect of my work; your passion meant I nonetheless looked forward to this opportunity. Thanks a lot. Thank you also to my monitoring committee, Keith and Bill for your guidance and advice. And special thanks in particular to Matt Palmer for getting the proverbial ball that was this PhD metaphorically rolling, and to Richard Wood and the Met Office as a whole for allowing — and paying for — the ensuing adventure. It's been great fun!

I would also like to thank my many friends at the Met Office who kindly discussed my work with me when they really should have been working on core deliverables or solving the southern ocean warm bias. A big thank you to Chris R, Laura, Leon, Tim, Matt Miz, Chris H, James, Jon, Daley amongst many, many others.

There have also been many people who have helped keep me sane during this time. 'Northern' folk (Phil, Mike, Si, and Rob) who can be relied upon for a distracting and inane conversation almost certainly involving Taylor Swift-based puns; the 'XCS Exiles' Exeter University postgraduate football team for indulging my (self-proclaimed) creative flair and not requiring me to track back; Roy and the crew for further reminding me not to take life too seriously. Indeed, these lists of names are not exhaustive and there have been many special people who have come and gone but nonetheless had a lasting effect; if you are reading this you are likely one of them.

Finally, thanks to all my family. My mum and dad for providing two excellent role models, my sister for proving how very far perseverance can take you, and my significantly younger and brighter brother for being gracious enough not to finish his PhD before me.

“Accordingly, seeing that our senses sometimes deceive us, I was willing to suppose that there existed nothing really such as they presented to us; and because some men err in reasoning, and fall into paralogisms, even on the simplest matters of geometry, I, convinced that I was as open to error as any other, rejected as false all the reasonings I had hitherto taken for demonstrations; and finally, when I considered that the very same thoughts (presentations) which we experience when awake may also be experienced when we are asleep, while there is at that time not one of them true, I supposed that all the objects (presentations) that had ever entered into my mind when awake, had in them no more truth than the illusions of my dreams. But immediately upon this I observed that, whilst I thus wished to think that all was false, it was absolutely necessary that I, who thus thought, should be somewhat; and I observed this truth, I think, therefore I am...

René Descartes (1596–1650)

Contents

Chapter 1

Introduction

1.1 Introduction to the introduction

This thesis presents analyses that are aimed at understanding the mechanisms of decadal climate variability in the North Atlantic ocean in order to help interpret decadal climate predictions, making use of a new high resolution coupled climate model, 'HadGEM3' (*Walters et al.*, 2011). Therefore, this first chapter begins with the motivation for making decadal climate predictions in Section 1.2. We then discuss the evidence for periodic decadal variability in the North Atlantic in Section 1.3 and the mechanisms of simulated decadal variability in Section 1.4. In Section 1.5 we describe the potential benefits of the high ocean and atmosphere resolution. Finally, we outline the key research questions we aim to address in Section 1.6 and the thesis structure in Section 1.7.

1.2 Near term/'decadal' climate prediction

In this first section we begin with the motivation for conducting near term (multi-annual, often specifically decadal) climate predictions (Section 1.2.1) before discussing the origins of predictability (Section 1.2.2) and the specific methods by which predictions are made (Section 1.2.3).

1.2.1 The motivation for near term climate prediction

The climate is changing (*Bindoff et al.*, 2013); carbon dioxide levels are now higher than at any point in almost one million years and global mean temperature is likely higher than at any point in the last 1400 years (*Hartmann et al.*, 2013). Despite this, it is still not clear precisely how much warming smaller regions, on the scale of countries, will encounter, and whether these regions will become wetter or drier (*Kirtman et al.*, 2013). In order to adapt to — or mitigate the effects of — climate change, policy planners and local people need reliable forecasts of how regional climate may evolve over the next years to decades. For example, if water is likely to be in short supply this might suggest investment in strategies for the efficient use of this resource, whereas if it is more likely there will be future large temperature extremes it may be more appropriate to invest in infrastructure to combat these. Key to these adaptation and mitigation strategies are the reliability of the climate models and prediction systems used to make these near term predictions.

Another motivation for near term climate prediction is related to long term climate projections (of order one century) with the same or similar models. On long timescales there exist large uncertainties in both the overall scenario (*i.e.* the future emissions of greenhouse gases and other climate pollutants, as well as natural forcing from volcanoes and changes in the solar cycles, *Moss et al.*, 2010) and the reliability of climate models in these 'out of sample' experiments. It may not be unreasonable to assume that models that more accurately predict present-day variability and warming may also be better at estimating future rates of warming (*Tebaldi et al.*, 2005) as well as simulating past climate states.

In summary, near term climate prediction may be useful for both informing regional adaptation/mitigation strategies, as well as improving the fidelity of climate models (that can then be used to investigate other climate states). In the next section we discuss the origins of this predictability.

1.2.2 Climate predictability

Given the potential utility of climate predictions we begin this section with a discussion of the fundamental origins of predictability within chaotic systems (such as the climate system) in Section 1.2.2.1. We then explore the physical drivers of this predictability (Section 1.2.2.2) before discussing the particular variables in which prediction systems currently show useful skill (Section 1.2.2.3).

1.2.2.1 Fundamental origins of predictability in chaotic systems

Chaos theory states that arbitrarily small perturbations in the initial conditions of a chaotic system can lead to very different outcomes, despite the system being innately deterministic (*Lorenz*, 1963). Such a feature may be expected to render these systems impossible to predict in the absence of perfect initial conditions (and an error free prediction system). Indeed, weather anomalies are generally not predictable more than a few weeks ahead and yet we would like to know whether regions of the planet will become warmer or cooler and wetter or dryer over the coming decades (Section 1.2.1). In order to make useful predictions we can make use of parts of the system that vary on longer (slower) timescales than others, *i.e.* that diverge from the initial state more slowly and are less susceptible to small perturbations. Assuming these slow parts of the system (*e.g.* the ocean) do interact with the fast parts of the system (*e.g.* the atmosphere/weather) they should then modify the distribution of these events (*e.g.* the climate) in some predictable way. The extent to which the 'slow' impacts the 'fast' can be estimated by conducting initial condition ensembles, in which small perturbations are made to the initial conditions. Taking the average over many ensemble members removes the chaotic part of the variability and leaves only the predictable component (if this exists). After *Lorenz* (1975), this represents predictability of the first kind (initial condition) and is separate from predictability of the second kind (boundary conditions). Here, we focus on the initial condition problem but see Section 1.2.3 for some discussion of the boundary condition problem.

The accuracy, or diagnostic 'skill', of a climate prediction can be assessed against either past real world observations, or against the climate model itself, the latter often described

as a perfect-model study that gives the *potential* skill of the system (*Kirtman et al.*, 2013). In either case, multiple initial conditions (start dates) are chosen and a perturbed ensemble of the prediction system is initialised and allowed to freely run forwards in time. The skill is the correlation between the ensemble mean forecast and the truth for a given variable over a given spatial domain and averaged over an appropriate time period, where a value of 1 represents perfect skill and a value of zero represents no skill. We now discuss the physical origins of predictability within the climate system, as well as the resulting skill in predicting certain climate variables.

1.2.2.2 Physical origins of climate predictability

The potential predictability of climate indices (in either the ocean or atmosphere) arises from a variety of sources, the relative importance of which is lead time dependant. For example, for one season ahead, a combination of sea surface temperatures (*Rodwell et al.*, 1999, SSTs,), El Niño/La Niña (*Bell et al.*, 2009), Arctic sea ice (*Yang and Christensen*, 2012), and the quasi-biennial oscillation (*Pascoe et al.*, 2006) have been shown to provide predictability in atmospheric pressure over the North Atlantic (*Scaife et al.*, 2014). On seasonal to annual timescales, El Niño/La Niña can provide regional predictability of some variables (see next) over most of the globe through its effect on the large scale atmospheric circulation (*Smith et al.*, 2012). El Niño/La Niña shows variability on multi-annual timescales, though the mere presence of large multi-annual variability is not a sufficient condition for long term predictability (*Wittenberg et al.*, 2014).

On multi-annual timescales, the potential predictability increasingly arises from large scale ocean processes, for which the thermal inertia is larger and for which there are effective mechanisms to 'store' anomalies from year to year (*Alexander and Deser*, 1995). Within the North Atlantic, the Atlantic Meridional Overturning Circulation (AMOC) provides a potential source of predictability on annual or longer timescales (*Collins*, 2002; *Robson*, 2010). However, the potential skill in predicting AMOC changes appears to be model dependant (*Collins et al.*, 2006) and so the translation to real-world skill (discussed next) is unclear and remains controversial (*Matei et al.*, 2012; *Vecchi et al.*, 2012).

1.2.2.3 Resulting skill in climate variables

Some variables and some regions of the globe yield more skilful predictions than others. In the atmosphere, global mean annual mean surface air temperature (SAT, a common metric of 'climate change') shows skill for several years along with regional SAT in many locations (*Doblas-Reyes et al.*, 2013). To begin with this skill arises from good initialisation of the prediction models (discussed in Section 1.2.3) with the skill from the boundary conditions (*e.g.* greenhouse gases) increasing through time (*Branstator and Teng*, 2012). In some regions, such as the North Atlantic, SAT skill remains dominated by the initial conditions for around a decade (*Branstator and Teng*, 2012). There is currently generally much less skill in precipitation than in SAT, though there is evidence of possible skill in precipitation over West Africa (*Gaetani and Mohino*, 2013).

In the ocean a paucity of long term observations makes assessing the actual skill of models difficult but there is potential skill in predicting sea surface temperatures (SSTs) in the North Atlantic for up to a decade (*Hawkins et al.*, 2011), possibly related to the large scale AMOC. In addition, skilful predictions of SSTs in the western Pacific (*Meehl et al.*, 2014) and Indian Oceans (*Corti et al.*, 2012) have been shown, although the latter of these is primarily due to external forcings (boundary conditions). Skilful predictions of future North Atlantic SST variability would be very valuable: North Atlantic SST variability has been linked to drought in the Sahel region (*Folland et al.*, 1986; *Zhang and Delworth*, 2006), Atlantic hurricane formation (*Goldenberg et al.*, 2001; *Smith et al.*, 2010; *Xie et al.*, 2005), precipitation over northern Europe (*Sutton and Hodson*, 2005), and the growth and persistence of Arctic sea ice, which could also affect the climate of northern Europe (*Screen*, 2013). In general, ocean heat content (OHC) is more predictable than atmospheric SAT and contributes to the skill in SAT at lead times longer than one year (*Hermanson and Sutton*, 2010). Once again, this near term predictability arises mostly from the initialisation of the climate model, which is now discussed.

1.2.3 Methods for initialising decadal predictions

Although much of the near term predictive skill of climate models is due to the initial conditions, provided by the initialisation procedure, the inherent biases in the underlying models mean that they cannot be relied upon to simply ingest real world observations and provide a skilful forecast. To reduce the impact of these biases, there are two main methods of initialising and post-processing the model output, which essentially differ on whether the climate model biases are removed before or after the forecast simulation (*Smith et al.*, 2013). These are 'full-field' and 'anomaly' assimilation methods, which are now described.

In the first method, the ocean model is initialised with observations of temperature and salinity at a given time (with the dynamics already spun up). Due to climate model biases, these 'full field' temperature and salinity observations are likely different to the climate model's (transient) equilibrium and, as such, the subsequent forecast exhibits large drifts away from the observed state back towards the model climatology (Figure 1.1, a). However, by using comparison of an assimilation simulation with hindcast simulations it is assumed that the lead time dependant drift term can be removed from the forecasts to give the bias corrected forecasts (Figure 1.1, c). In this situation, we assume that the signal we are attempting to (skilfully) detect evolves independently to the drift, which, given the resulting increase in skill after bias correction (*Smith et al.*, 2013), is at least partly true.

In the second method, the ocean model is initialised with observed anomalies of temperature and salinity at a given time, relative to some appropriate base state (*Robson*, 2010). As the climate model's mean state is already likely different from that observed, and it is only the observed variability that is being added, the forecasts are not intended to return absolute values similar to those observed (Figure 1.1, b). After the forecasts have been run, the difference between the observed and climate model mean states can be subtracted to give the bias corrected forecasts (Figure 1.1, d). In this situation, we assume that the signal and the mean state are independent (*i.e.* that the variability doesn't depend on the background state), which is again at least partly true, as evidenced by the increase in skill after the bias correction (*Smith et al.*, 2013).

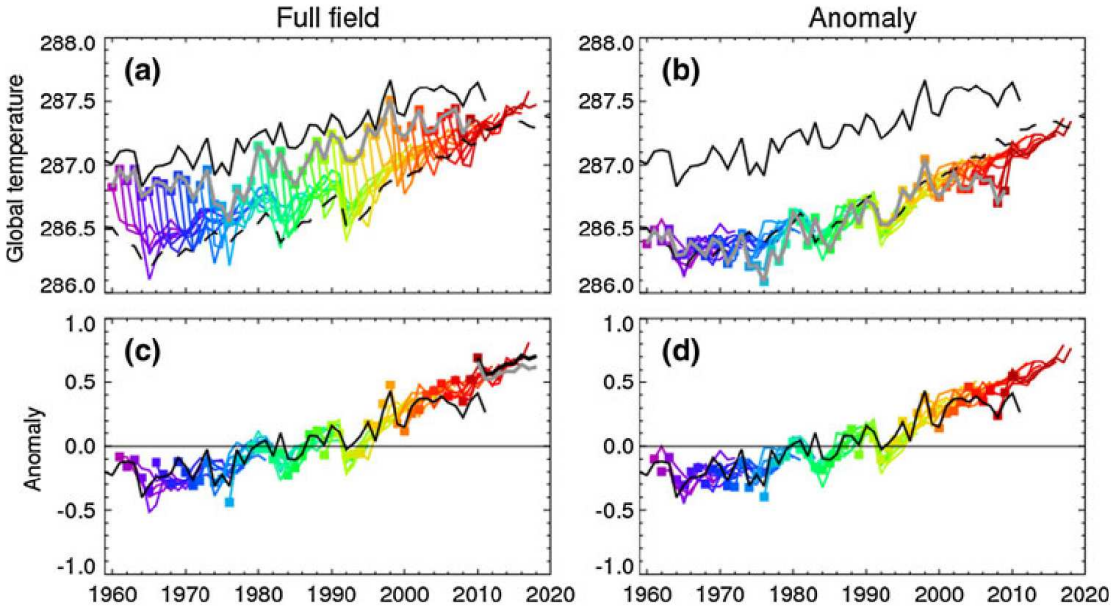

Figure 1.1: Illustration of bias correction for full field (left) and anomaly (right) initialisation. Thin black curves show the observed time series of annual mean global temperature (from HadCRUT3, *Brohan et al.*, 2006). Coloured curves show the ensemble mean hindcasts, with different colours showing different start dates. The upper panels (a, b) show absolute values (K), and the lower panels (c, d) show anomalies after adjustment for model biases. Figure and associated caption taken from *Smith et al.* (2013). ©Springer Science and Business Media. Used with permission.

Initialisation provides skill related to the fidelity of the model and due to the boundary conditions (*e.g.* external forcings such as greenhouse gases and aerosols, *cf.* predictability of the second kind), as well as relating to the internal variability of the climate system by attempting to initialise the model in the correct phase of the variability and trusting it to evolve correctly. For near term climate *projections* (projections here implies an uninitialised prediction that only attempts to capture the externally forced component of the variability), the combination of model uncertainty and internal variability dominates the variance in the projections (Figure 1.2a). On regional scales, such as for the British Isles, this is even more apparent (Figure 1.2b). It is clear that on annual to decadal timescales, improving the climate model's representation of the real world and the initialisation of internal variability could provide valuable increases in predictive skill, particularly in important regions such as the North Atlantic (*Hakkinen and Rhines*, 2004; *Msadek et al.*, 2010; *Dunstone et al.*, 2011). However, in order to achieve these aims we must first understand the multi-annual/decadal variability of the underlying system, which is the focus of the next section.

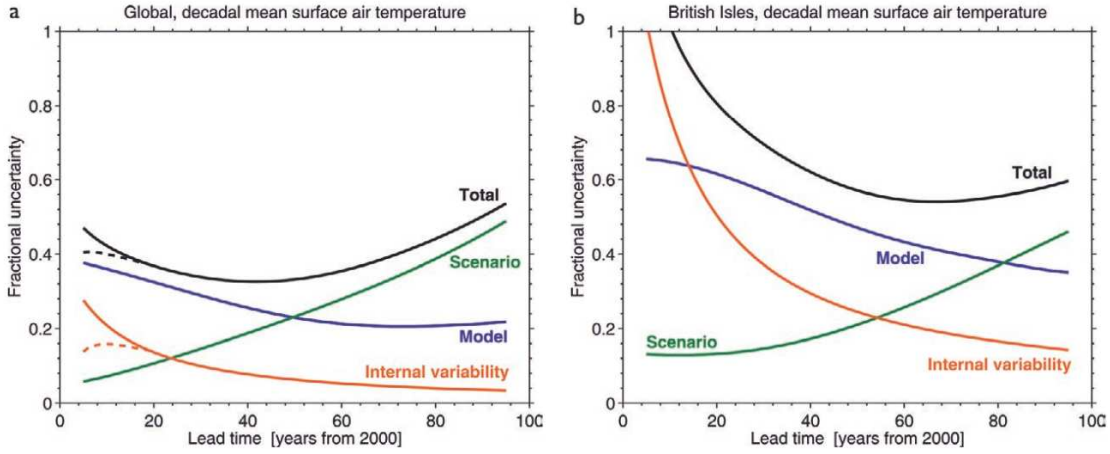

Figure 1.2: The relative importance of each source of uncertainty in decadal mean surface temperature projections is shown by the fractional uncertainty (the 90% confidence level divided by the mean prediction) for (a) the global mean, relative to the warming from the 1971–2000 mean, and (b) the British Isles mean, relative to the warming from the 1971–2000 mean. The importance of model uncertainty is clearly visible for all policy-relevant timescales. Internal variability grows in importance for the smaller region. Scenario uncertainty only becomes important at multi-decadal lead times. The dashed lines in (a) indicate reductions in internal variability, and hence total uncertainty, that may be possible through proper initialization of the predictions through assimilation of ocean observations (*Smith et al.*, 2007). Figure and associated caption taken from *Hawkins and Sutton* (2009). ©American Meteorological Society. Used with permission.

1.3 The evidence for periodic decadal variability in the North Atlantic

In this section we review the observational (Section 1.3.2) and model (Section 1.3.3) evidence for periodic decadal variability in and around the North Atlantic subpolar gyre. However, in order to put these reviews into context, we begin with a brief description of the region.

1.3.1 The North Atlantic subpolar gyre

Given the potential predictability of the North Atlantic subpolar gyre (NA SPG, *Hakkinen and Rhines*, 2004; *Wouters et al.*, 2013), and it's importance in particular for European climate (*e.g. Sutton and Hodson*, 2005), we focus the majority of our analysis on this

region. The NA SPG is a region between approximately 45°N and 80°N in the North Atlantic, roughly split in two by the Greenland-Iceland-Scotland (GIS) ridge at 65°N that separates the Arctic and Nordic Seas north of this from the classical Atlantic to the south (Figure 1.3). In terms of the large scale circulation, the region is characterised by the mean cyclonic circulation that brings warm and salty water from the subtropical Atlantic northwards, cooling as it traverses the basin, before sinking and returning south as North Atlantic Deep Water (NADW). This general meridional overturning circulation (MOC, or AMOC for Atlantic MOC) is another key feature of the region and is often the primary, explicit focus of studies of the NA SPG. Indeed, the region is often decomposed into a streamfunction in the meridional-depth plane (the AMOC) and a streamfunction in the horizontal plane (the gyre, *e.g. Dong and Sutton,* 2001), although to what extent this geometric decomposition is representative of the large scale thermohaline circulation (THC) at these latitudes is not clear (*Biastoch et al.*, 2008a; *Zhang*, 2010; *Kwon and Frankignoul*, 2014). In addition, the region feels strongly the influence of the atmosphere, lying as it does beneath the North Atlantic Oscillation (NAO, the index of which is defined as the pressure difference between the Azores and Iceland), which is the dominant mode of wintertime atmospheric variability in the region (*Hurrell et al.*, 2003). This complex interplay of drivers are what allow the NA SPG to exhibit prominent decadal variability, evidence for which is now discussed.

1.3.2 Instrumental and proxy evidence for (bi)decadal variability

Instrumental climate records within the NA SPG show decadal variability in the formation of NADW (*Mauritzen et al.*, 2012) and in the strength of the NA SPG itself (*Curry and McCartney*, 2001; *Rhein et al.*, 2011; *Roessler et al.*, 2015). In addition, there has been well documented multi-annual/decadal variability in the heat content of the NA SPG (*Robson et al.*, 2012) and potentially even in circulation indices such as the AMOC (*Smeed et al.*, 2014). This 'decadal' variability, here defined as variability on timescales of 10–30 years, is distinct from longer timescale multi-decadal (>30 years) variability that may also exist in the NA SPG (*Schlesinger and Ramankutty*, 1994). However, this multi-decadal variability — increasingly sensitive to forcing/scenario uncertainty in cli-

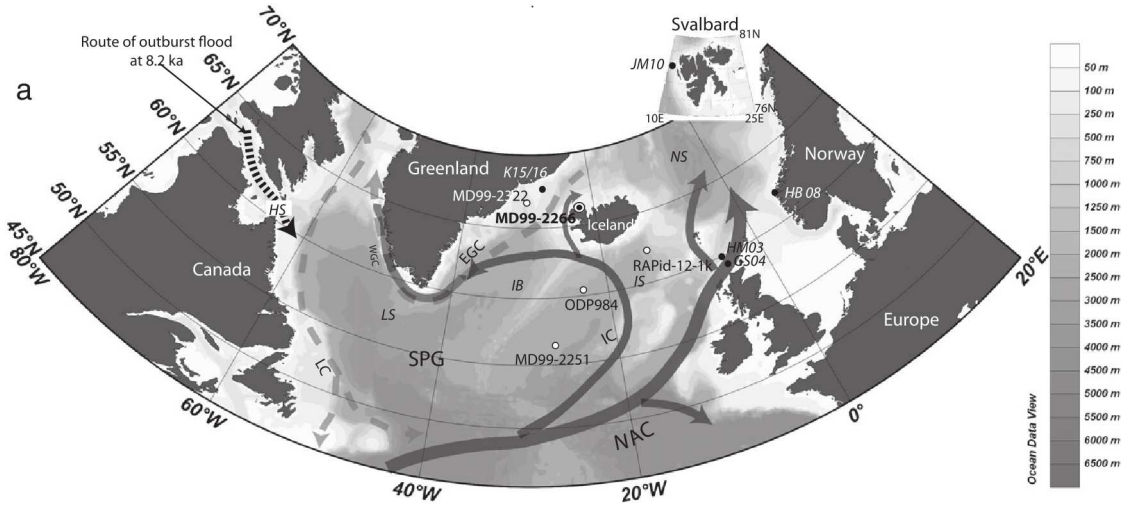

Figure 1.3: . Location map. Panel shows the high latitude North Atlantic with schematic surface current regime. Figure and associated caption taken from *Quillmann et al.* (2012). Shading represents the depth of the ocean. Selected acronyms: SPG: Subpolar gyre, NS: Nordic Seas, LS: Labrador Sea, NAC: North Atlantic Current, EGC: East Greenland Current, LC: Labrador Current. ©Elsevier. Used with permission.

mate projections (*cf.* Section 1.2.3 and Figure 1.2) rather than uncertainty related to the model or internal variability — is not the focus of this thesis. On decadal timescales, the relative contributions of internal variability and externally forced variability (*e.g.* the drivers of the 1960s cooling in the North Atlantic, *Baines and Folland*, 2007; *Hodson et al.*, 2014) has implications for the spectral characteristics of the variability — with self-sustaining internal variability likely to lead to spectral peaks (*Snoussi,* 1998) and externally forced variability likely to lead to episodic or transient changes (unless the external forcing has its own periodicity, *e.g.* the 11 year solar cycle or possibly anthropogenic or volcanic aerosols, *Otterå et al.*, 2010; *Booth et al.*, 2012).

Direct instrumental records in and around the NA SPG are too short to resolve reliably such spectral behaviour (*Kennedy et al.*, 2011). However, palaeo-proxy records have been shown to exhibit significant decadal variability, both globally (*Mann et al.*, 1995) and more recently using high temporal resolution palaeo proxies around the NA SPG (*Sicre et al.*, 2008; *Chylek et al.*, 2012). *Sicre et al.* (2008) used alkenones (37 Carbon atom organic molecules with either two or three Carbon-Carbon double bonds, the ratio of which is proportional to the temperature) from around Iceland to estimate sea surface temperatures (Figure 1.4) over the past 4500 years. Separately, *Chylek et al.* (2012) used Oxygen isotope ratios (which are affected by temperature as well as other environmental

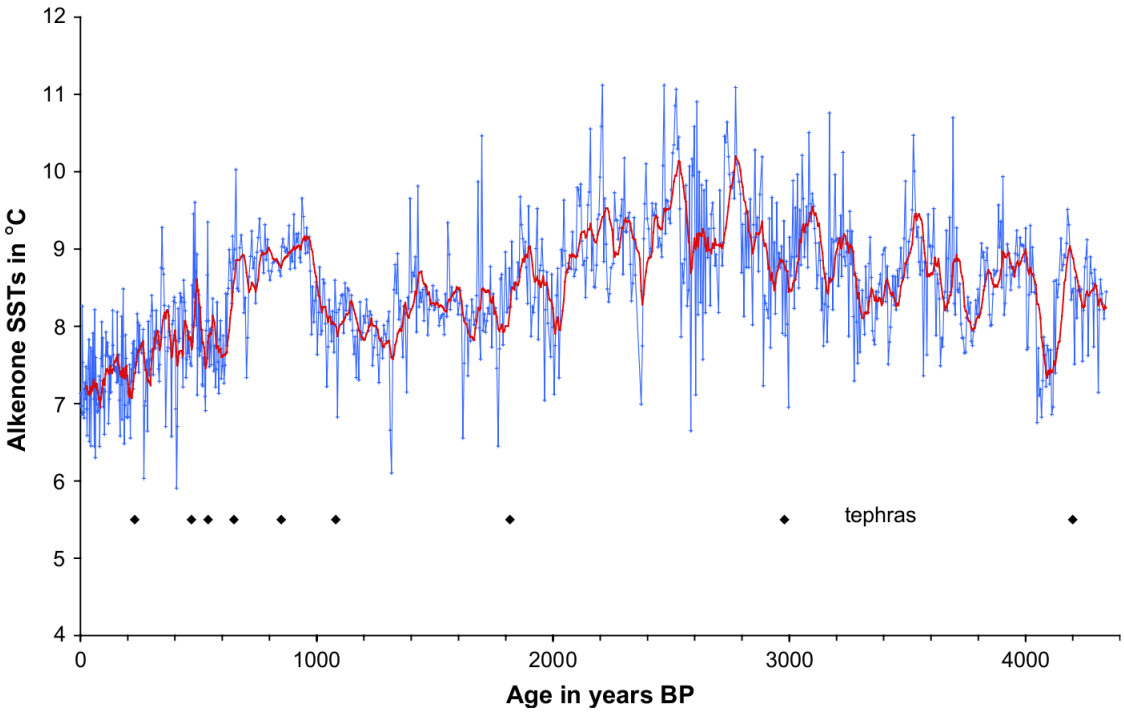

Figure 1.4: . Alkenone derived sea surface temperature (SST) estimates over the past 4500 years in the MD99-2275 core. The calibration established by *Prahl et al.* (1988) was used to convert $U_{37}^{K'}$ into SSTs. Black diamonds indicate tephra layers identified and used to build the age model. The red curve represents the 10-point running mean of the data. Figure and associated caption taken from *Sicre et al.* (2008).

factors) from Greenland ice cores to estimate the temperature of precipitation over the past 660 years. Both of these NA SPG records specify periodicity with a timescale of around 20 years, hereafter referred to as 'bidecadal'. In addition to these palaeoclimate proxies, another source of indirect evidence for decadal variability in the NA SPG comes from numerical models, discussed next.

1.3.3 Model evidence for decadal variability

As with paleoclimate proxies, and even direct observations (*Levitus et al.*, 2009), (uninitialised) numerical models of the climate must only be considered in light of their caveats, such as their inability to resolve potentially important spatial scales (*Penduff et al.*, 2010). Nonetheless, the sheer number of climate models means it is possible to investigate whether these independent simulations (although to what extent they are truly independent is an open question, *Knutti et al.*, 2013) give consistent results. The periodicity in

simulated mechanisms of variability (in climate model studies for which the NA SPG was a major focus) is documented in Figure 1.5 (along with other details that we describe in forthcoming sections). As can be seen, periodicities of around 20 years are very common in climate model simulations of this region. Indeed, as noted by *Frankcombe et al.* (2010), simulations of the North Atlantic tend to describe variability on either short (20 year/bidecadal) timescales, possibly related to the NAO, or longer (>70 year) timescales, related to advection from faraway sources (*Jungclaus et al.*, 2005; *Hawkins and Sutton*, 2007; *Park and Latif*, 2008; *Menary et al.*, 2012). The specific mechanisms of this bidecadal variability are discussed in the next section.

1.4 Origins and mechanisms of simulated decadal variability

In this section we discuss the mechanisms/processes by which models can simulate decadal variability in the NA SPG. Given the paucity of direct observations within the NA SPG, studies such as these are a useful tool to probe potential real world mechanisms of variability. This section is schematically summarised in Figure 1.5. We begin with the simplest intuitive model of a red noise ocean responding to a white noise atmosphere (Section 1.4.1) before specifically reviewing the role of the NAO in driving ocean variability via momentum and/or buoyancy fluxes (Section 1.4.2 and subsections therein). We then move on to the potential role of ocean advection in determining the timescale of the variability (Section 1.4.4) before discussing the possible role of wave processes (Section 1.4.5) and summarising alternative drivers of NA SPG variability (Section 1.4.6). We conclude this section with a review of the drivers of Labrador Sea density variability across the climate models (Section 1.4.7).

1.4.1 North Atlantic Ocean response to atmospheric noise

Since *Hasselmann* (1976), one paradigm for explaining decadal variability in the NA SPG has been that of an ocean integrating up white-noise (stochastic) forcing from the

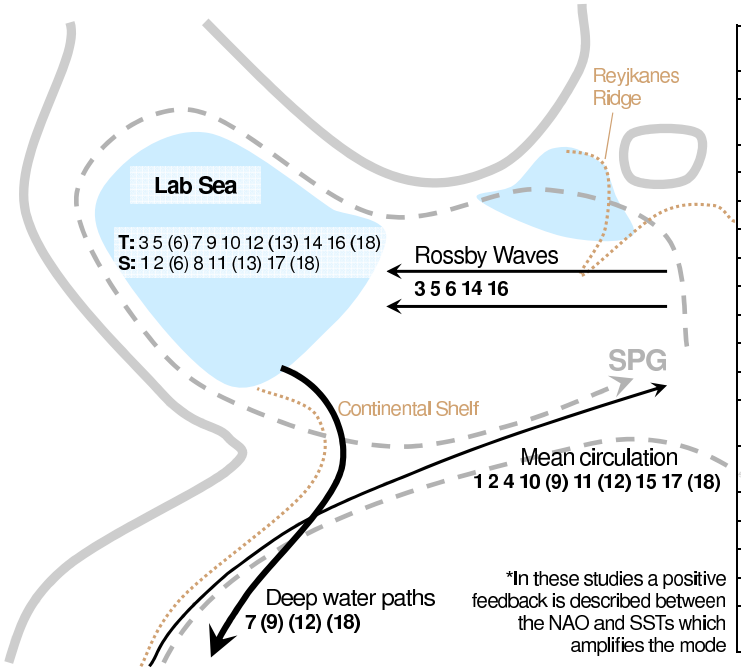

	Study	Period (years)	Mech-anism
1	Delworth '93	50	Ocean
2	Timmermann '98	35	Coupled
3	Grotzner '98	17	Coupled
4	Visbeck '98	12-16	Ocean*
5	Watanabe '99	~10	Ocean*
6	Dong '01	-	-
7	Eden '01	12-18	Coupled
8	Holland '01	20	Ocean
9	Eden '03	16	Ocean*
10	Dai '05	24	Ocean
11	Dong '05	25	Ocean
12	Àlvarez-Garcia '08	24	Ocean*
13	Danabasoglu '08	21	-
14	Frankcombe '10	20-30	Ocean
15	Born '12	15-20	Ocean
16	Sevellec '13	24	Ocean
17	Escudier '13	20	Ocean
18	Kwon '14	20 "regime 1"	Ocean*

Figure 1.5: A summary of some of the literature on simulated decadal variability in the North Atlantic subpolar gyre (NA SPG), with a particular emphasis on studies which found self-sustaining cyclical behaviour. Key regions of the NA SPG are marked. The figure legend (right) denotes the studies which we have attempted to synthesise and an associated numerical identifier. Where these studies report a significant peak in the power spectrum on decadal timescales this is noted as well as whether the mechanism is primarily ocean-only or inherently coupled. Studies where the atmosphere is postulated to amplify — but not explicitly propagate — the signal are marked with an asterisk. For each study the feedback or process which is reported as crucial in setting the timescale is marked on the map using a simple numbering system. These comprise: 1) Feedbacks relating to the deep water pathways and their interaction with the northward flowing western boundary current, 2) Rossby wave (or sometimes 'geostrophic self advection') transit times across the NA SPG, 3) the mean advection timescale for anomalies to propagate into the NA SPG from the tropics, or for small anomalies to integrate up over time. Lastly, using the same numerical key, the studies are split into which of temperature or salinity is reported to control decadal timescale density changes in the Labrador Sea. In all case, studies in brackets appear in more than one category. This represents a drastic simplification of each of these studies and the reader is referred to the original works for further details. In particular, the reported feedback/process that sets the overall timescale to some degree also subjectively reflects the precise focus of the particular study.

atmosphere[1]. *Frankignoul et al.* (1997) forced an idealised ocean basin representative of the Atlantic with stochastic wind stress forcing and found that the dominant timescale was related to the basin transit time of long baroclinic Rossby waves, although this idealised work compared more favourably with coupled coupled climate models when considering the subtropics rather than the subpolar regions. Later work with coupled models also highlighted the potential for the Atlantic Ocean, and in particular the AMOC, to respond to stochastic forcing with large decadal variability (*Holland et al.*, 2001; *Kwon and Frankignoul*, 2012), for example by preferentially transmitting atmospheric variability to the ocean during times of deep convection (*Dong and Sutton*, 2005; *Born and Mignot*, 2012).

In an ocean-only model, *Delworth and Greatbatch* (2000) found that the AMOC responded preferentially to the longer timescale ($>$20 year) component of the forcing, rather than necessarily integrating up the high frequency component. Similarly, *Holland et al.* (2001) showed decadal periodicity in their simulated AMOC (related to changes in Arctic ice export) but found that the ocean was generally insensitive to the origin of the stochasticity and that low frequency forcing had a larger effect. Indeed, whether the ocean integrates up atmospheric forcing or preferentially amplifies particular periods remains unclear. In the latter study, the presence of significant periodicity was due to a negative feedback on density provided by increasing heat transport associated with a stronger AMOC, and indeed a negative feedback is usually required in order for an otherwise white or red power spectrum to exhibit a sharp spectral peak at a preferred timescale (*Snoussi*, 1998). One possible exception to this is if the temporally stochastic forcing has sub-basin scale variability and the (advective) ocean has a deep enough mixed layer for advection to dominate over thermal damping. In such a situation, 'spatial resonance' could result in significant periodicity (*Saravanan and McWilliams*, 1997, 1998) although the preference for long timescale forcing in some models (*Delworth and Greatbatch*, 2000; *Holland et al.*, 2001) and the inability of forced ocean-only models to reproduce the periodicity (*Pierce*, 2001; *Mecking et al.*, 2014) suggests this mechanism is not of first order importance. Nonetheless, consistent with the requirements of the

[1]Indeed, this paradigm has recently been used to explain longer timescale variability over the whole North Atlantic (*Clement et al.*, 2015).

model of *Saravanan and McWilliams* (1998) often simulated white-noise forcing of the ocean is stochastic in time but spatially has the pattern of the NAO (*Delworth and Greatbatch*, 2000; *Mecking et al.*, 2014). As such, we next discuss the effect of specifically NAO-related forcing on the ocean.

1.4.2 North Atlantic Ocean response to the NAO

When considering the direct effect of the NAO on the North Atlantic Ocean we separate our discussion into momentum (wind) forcing (Section 1.4.2.1) and buoyancy (heat and freshwater) forcing (Section 1.4.2.2). Outside of the deep convection regions of the North Atlantic such as the Labrador Sea (see Section 1.4.7) the effect of momentum and buoyancy forcings are often found to add linearly (*Cabanes et al.*, 2008; *Biastoch et al.*, 2008a). Climate model studies in which the NAO is postulated to play a role in the simulated ocean decadal variability are noted on Figure 1.5 (final column).

1.4.2.1 Wind forcing

The response of a low resolution ocean model to wind forcing reminiscent of that associated with NAO variability was investigated by *Visbeck et al.* (1998) who found that the sign of the NA SPG temperature response depended on the timescale (annual or decadal) of the forcing. Subsequent work found that wind stress forcing could explain fast (annual) response times, with buoyancy fluxes implicated in longer timescale responses (see next section, *Eden and Willebrand*, 2001). The timescale of the response to wind forcing may be related to the mechanism by which the wind interacts with the ocean, with short timescale (less than 3 years) responses due to an Ekman response and long timescale (more than 3 years) responses due to the heaving of isopycnals and subsequent adjustment (*Cabanes et al.*, 2008). Similarly, using companion forward and adjoint ocean models to find the mechanisms by which optimal observations maximally affect the AMOC, *Kohl* (2005) found wind forcing affected the AMOC in three distinct ways: The first two (Ekman transport and Ekman pumping, relating to zonally integrated wind stress and wind stress curl respectively) were not as important as the third; wind driven coastal up-

welling/downwelling heaving isopycnals and resulting in alongshore currents known as "coastal upwelling jets". Buoyancy forcing was found to become increasingly important on multi-annual timescales. Nonetheless, wind forcing related to the NAO and the subsequent excitation of Rossby waves can contribute to multi-annual timescale variability (*Polo et al.*, 2014, see also Section 1.4.5).

1.4.2.2 Buoyancy forcing

The role of buoyancy fluxes in driving NA SPG variability has been investigated in targeted ocean-only experiments (*Eden and Willebrand*, 2001), diagnosed in coupled model studies (*Timmermann et al.*, 1998) and inferred from observational analyses (*Curry and McCartney*, 2001). Buoyancy fluxes generally explain variability on longer timescales to wind forcing (*Eden and Willebrand*, 2001; *Polo et al.*, 2014). Whilst in coupled model studies, the relative roles of heat and freshwater/salinity in buoyancy forcing remain model dependent (*e.g. Dai et al.*, 2005; *Dong and Sutton*, 2005, see Section 1.4.7), in ocean-only experiments it is generally the case that heat fluxes are the most important buoyancy flux (*Eden and Willebrand*, 2001; *Alvarez-Garcia et al.*, 2008; *Mecking et al.*, 2014; *Polo et al.*, 2014). For example, *Eden and Willebrand* (2001) suggested a simple NAO-AMOC feedback mechanism in which anomalously negative heat fluxes over the NA SPG associated with a positive NAO reduces temperatures in the SPG thereby increasing density and strengthening the AMOC. The subsequently strong AMOC brings more warm water northwards, reversing the cycle.

An important role for heat fluxes is also found in observational analyses, such as that of *Curry and McCartney* (2001) who investigated Potential Energy Anomalies (PEAs, an oceanic analogy to the NAO) at two sites approximately at the centre of the subpolar and subtropical gyres. On multi-annual/decadal timescales they found that the PEA and overturning were thermally driven insofar as temperatures changed twice as much as salinities in the sinking regions (after scaling by the thermal and haline expansion coefficients). A 10-year running sum of the NAO index correlated with the transport index from the PEAs indicating that the ocean integrates up the atmospheric signal. A similar running sum of Labrador Sea heat fluxes can also explain AMOC variability in a coupled

model (*Ortega et al.*, 2011). Although events such as 'Great Salinity Anomalies' (*Dickson et al.*, 1988) suggest a potential role for freshwater/salinity fluxes (not necessarily related to the NAO), they may only play a damping role in heat flux driven variability (*Curry and McCartney*, 2001), although it should be noted that coupled climate models can display decadal variability broadly consistent with both GSA events (in terms of the size of freshwater anomalies and their propagation pathways) and an important role for temperature variability (*Wohlleben and Weaver*, 1995; *Timmermann et al.*, 1998).

1.4.3 NAO response to the North Atlantic Ocean

We have highlighted the way in which the atmosphere/NAO can drive variability in the NA SPG, but it is also possible for the ocean to drive variability in the atmosphere (*Timmermann et al.*, 1998; *Eden and Willebrand*, 2001; *Gastineau et al.*, 2013), and indeed this must be the case for a coupled mode of variability to exist (in the absence of innate long timescale variability in the atmosphere). Investigations using uncoupled ocean or atmosphere simulations forced with appropriate boundary conditions, have shown both that 1) historical ocean SST variability can mostly be recreated from atmospheric heat fluxes alone (*Battisti et al.*, 1995), and 2) that historical NAO variability can be recreated from SSTs (*Rodwell et al.*, 1999) *i.e.* both that the atmosphere drives ocean variability, and the ocean apparently drives atmosphere variability, in models. Over the NA SPG, the atmospheric response may be a combination of forcing from the local SSTs as well as those in the tropical Atlantic (*Sutton et al.*, 2000), which provide some of the skill in seasonal forecasts of the NAO up to a season ahead (*Scaife et al.*, 2014). The ocean can drive the atmosphere via small-scale SST fronts that induce local wind anomalies (*Lindzen and Nigam*, 1987; *Minobe et al.*, 2008), or via latent/sensible heating over larger spatial scales (*Rodwell et al.*, 1999) possibly even involving communication with the stratosphere (*Omrani et al.*, 2014), but the precise mechanisms and their relative importance are still unclear. Nonetheless, much of the interannual variability in the NAO appears to be intrinsic to the atmosphere rather than driven by the ocean (*Bretherton and Battisti*, 2000).

1.4.4 Ocean advection

Often related to buoyancy forcing of the NA SPG, ocean advection has been shown to be an important regulator of the timescale of NA SPG variability in climate models (*Delworth et al.*, 1993; *Timmermann et al.*, 1998; *Dai et al.*, 2005; *Dong and Sutton*, 2005; *Born and Mignot*, 2012; *Escudier et al.*, 2013), in reanalyses of the real world (*Robson et al.*, 2012), and potentially in observations (*Sutton and Allen*, 1997). Heat and freshwater transport is often split into that due to mean and anomalous circulations, *e.g.*

$$OHT = \rho c_p \left(\int \overline{v}\overline{T} + \int v'\overline{T} + \int \overline{v}T' + \int v'T' \right) \mathrm{d}A \tag{1.1}$$

which describes the ocean heat transport (OHT) through a two dimensional section (A), where ρ is a constant reference density, c_p is the specific heat capacity of seawater, v and T are the circulation speed normal to the section and potential temperature on the section, and overbars and primes represent time means and anomalies respectively. When analysing variability in OHT, variability in $\overline{v}\overline{T}$ is zero and $v'T'$, which represents correlated changes in circulation and temperature, is often small for large enough areas (*Dong and Sutton*, 2005).

An alternative way to partition the OHT is to separate it into that due to horizontal (gyre) and vertical (AMOC) components (*Johns et al.*, 2011), *e.g.*

$$OHT_{AMOC} = \rho c_p \int_{-H}^{0} V \langle T \rangle \, \mathrm{d}z \tag{1.2}$$

$$OHT_{gyre} = \rho c_p \int_{-H}^{0} \int_{W}^{E} v^* T^* \, \mathrm{d}x \, \mathrm{d}z \tag{1.3}$$

where V is the zonally integrated circulation (*i.e.* the transport profile), angle brackets and asterisks represent zonal means and anomalies, and dz and dx represent vertical and zonal integrals over the full depth (depth, H, to the surface) and basin (West to East) respectively.

Although it is difficult to estimate these components from observations (except at some specific latitudes with trans-basin arrays, *Johns et al.*, 2011) they can be calculated in cou-

pled climate models. Both *Delworth et al.* (1993) and *Dong and Sutton* (2005) showed that, within the NA SPG, simulated advection of heat or freshwater anomalies into the sinking regions was due to gyre changes but the AMOC was implicated in related variability in the subtropical gyre (*Dong and Sutton*, 2005). Note that in the NA SPG, the gyre circulation may be more representative of the THC than the meridional circulation (AMOC, *Biastoch et al.*, 2008a; *Zhang*, 2010). The relative contributions of mean and anomalous circulation may depend on the specific region within the NA SPG, with mean circulation important throughout and anomalous circulation hypothesised to be important primarily for zonal flows (*Dong and Sutton*, 2001), although this is likely resolution and model dependent. *Watanabe et al.* (1999) found that different background mean states meant that simulated NA SPG heat content variability was more prone to the effects of anomalous circulation than may be the case in reality. Models for which advection by the mean or anomalous circulation and the subsequent accumulation of heat/freshwater anomalies are postulated to set the timescale of the variability are noted on Figure 1.5 (numbers — see caption).

1.4.5 Wave processes

Despite the benefits of high complexity in coupled climate models, which makes them a 'best guess' for understanding real world multi-annual/decadal variability, the importance of particular processes can be hard to elucidate. For example, wave processes, such as first mode baroclinic Rossby waves within the NA SPG, can have similar propagation timescales as ocean advection (*Watanabe et al.*, 1999). In these circumstances, simplified/idealised ocean models that allow analytical solutions can be useful to discern the precise drivers of change (*Johnson and Marshall*, 2002; *Sévellec and Fedorov*, 2013).

The contributions of Rossby and Kelvin waves to the propagation of AMOC signals was investigated in an idealised ocean by *Johnson and Marshall* (2002). They used a simple ocean model with smooth topography and no background circulation to investigate the ocean's response to volume flux forcing *i.e.* ignoring the specifics of deep water formation and how this is driven. They found that the response of this simple ocean to positive/negative volume flux anomalies was symmetric (though the response to temperature

or salinity anomalies might not be expected to be) and that the mechanisms and spatial extent of the ocean response varied with the period of the forcing. The mechanism involved southward propagation of Kelvin waves along the western boundary towards the equator followed by eastward propagation along the equator. Subsequently the Kelvin waves propagated poleward along the eastern boundaries in both hemispheres, all the while radiating Rossby waves into the ocean interior. Thus the response in the interior ocean in the opposite hemisphere was controlled by Rossby waves and, as the speed of these depends on the latitude, lower frequency forcing could penetrate signals further into the opposite hemisphere (in essence the equator acted as a low pass filter). Signals of Rossby wave propagation can often also be detected in more complex coupled climate models, although it is not clear that their existence implies an important role in driving decadal variability (*Eden and Greatbatch*, 2003; *MacMartin et al.*, 2013). Nonetheless, models for which Rossby or other trans-basin wave propagation (such as 'geostrophic self-advection' *Sévellec and Fedorov*, 2013) is invoked as an important mechanism setting the timescale of variability are noted on Figure 1.5 (numbers — see caption).

In a slightly more complex model, *Fevrier et al.* (2007) showed that the presence of a simulated Gulf Stream did not inhibit the propagation of Kelvin waves, as hypothesised by *Johnson and Marshall* (2002). Support for boundary wave propagation in observations could also be found in coherent sea surface height (SSH) signals around the North Atlantic (*Hughes and Meredith*, 2006). In this context, *Roussenov et al.* (2008) investigated the connection between boundary wave communication and the AMOC in an isopycnal ocean model with a realistic domain and on the multi-annual/decadal timescales important for climate. They found that fine scale (*i.e.* not smoothly varying) topography acted to inhibit wave communication, probably through scattering. As such, coarse resolution in a model may be expected to dampen the variability of the AMOC (see Sections 1.5 and 2.3). Finally, in contrast to *Johnson and Marshall* (2002), western boundary current anomalies were found in the opposite hemisphere within one year of the northern hemisphere forcing, *i.e.* much sooner. These may have been a response to barotropic (rather than baroclinic) Rossby waves rapidly propagating from the northern to southern hemisphere.

1.4.6 Alternative/non-local sources of variability

As well as surface forcing, the deep ocean may play a role in regulating near surface ocean variability. Complementary to the investigations of wave propagation, *Zhang and Vallis* (2007) investigated the reasons behind model temperature biases in the North Atlantic, and the potential role of vorticity dynamics, and found biases were vastly decreased if the model correctly simulated the Northern Recirculation Gyre (NRG), which helps to separate the Gulf Stream from the coast of North America. Correct representation of the NRG relies on the interaction between the southward flowing Deep Western Boundary Current (DWBC) and the northward flowing, surface Gulf Stream. The flow of the DWBC down the continental slope generates positive vorticity which helps form the NRG. As such, deep flows could exert some control on surface flows. Further evidence of this can be seen in an observational framework using a statistical model (a linear inverse model, LIM), and highlights how perturbations to the deep ocean can affect the predictability of the upper ocean (*Zanna*, 2012). Studies in which climate models have highlighted a role for interaction with the deep circulation as a possible regulator of the variability are noted on Figure 1.5 (numbers – see caption).

The processes we have outlined above are all local to the NA SPG but it is possible that processes outside of the NA SPG may drive variability within it. Within the ocean, simulated variability driven via advection from faraway sources such as the tropics (*Menary et al.*, 2012; *Park and Latif*, 2008) or the Arctic (*Jungclaus et al.*, 2005; *Hawkins and Sutton*, 2007) often has a much longer timescale. However, in the atmosphere, there is potential for subtropical Atlantic forcing of NAO variability (*Sutton et al.*, 2000), as well as forcing via the tropical Pacific (*Hoerling et al.*, 2001), both of which could then impact the NA SPG (as described in Section 1.4.2). In addition, the tropical Pacific is also a region of significant multi-annual variability and it has been shown that ocean-atmosphere interaction via El Niño/La Niña (*Trenberth*, 1997) may also drive temperature changes in the NA SPG region (*Ineson and Scaife*, 2009).

1.4.7 The role of the Labrador Sea

In most studies of multi-annual/decadal variability in the NA SPG the sinking regions, in which high density surface waters sink/convect leading to deep water formation, are found to be important for the subsequent evolution of the AMOC and by extension the variability within the NA SPG (Figure 1.5). The most important sinking region in the main NA SPG is the Labrador Sea, as well as the Irminger basin at the entrance to the Labrador Sea, where upper North Atlantic deep water (UNADW) is formed (*Toole et al.*, 2011). There is also deep water formation north of Iceland that flows out of the Nordic Seas and into the Labrador Sea, becoming lower North Atlantic deep water (LNADW, *Toole et al.*, 2011). In models, the interaction of deep water formed in the Nordic Seas and the main NA SPG, via the overflow sills, is poorly simulated except at ultra-high ($^1/_{12}^\circ$) resolution (*Chang et al.*, 2009).

Whether the density variability in these key sinking regions is driven by either temperature or salinity effects can be estimated by decomposing density changes into those due to temperature and those due to salinity (*Delworth et al.*, 1993). Indeed, *Delworth et al.* (1993) defined an index of the large scale variability in the North Atlantic and separately regressed this against density changes that were due to temperature and density changes that were due to salinity and found that salinity was the dominant driver (Figure 1.6). However, this is not the case in all climate models, which disagree on whether multi-annual/decadal temperature or salinity variability is most important for multi-annual/decadal density variability in the Labrador Sea (Figure 1.5, numbers — see caption). Given that warm/saline anomalies in the NA SPG often co-vary, and that they have opposing effects on density, this suggests that the manifestation of NA SPG variability within climate models may be somewhat sensitive to the particular density driver.

In these sections we have aimed to describe the key processes involved in simulated mechanisms of decadal variability in the NA SPG. There are clearly many potential manifestations of this variability and to what extent any of them are a more realistic representation of the real world is still unclear (*cf.* Figure 1.5). The numerical models are a discretised version of reality and so the potential benefits of improving the resolution of this discretisation is discussed next.

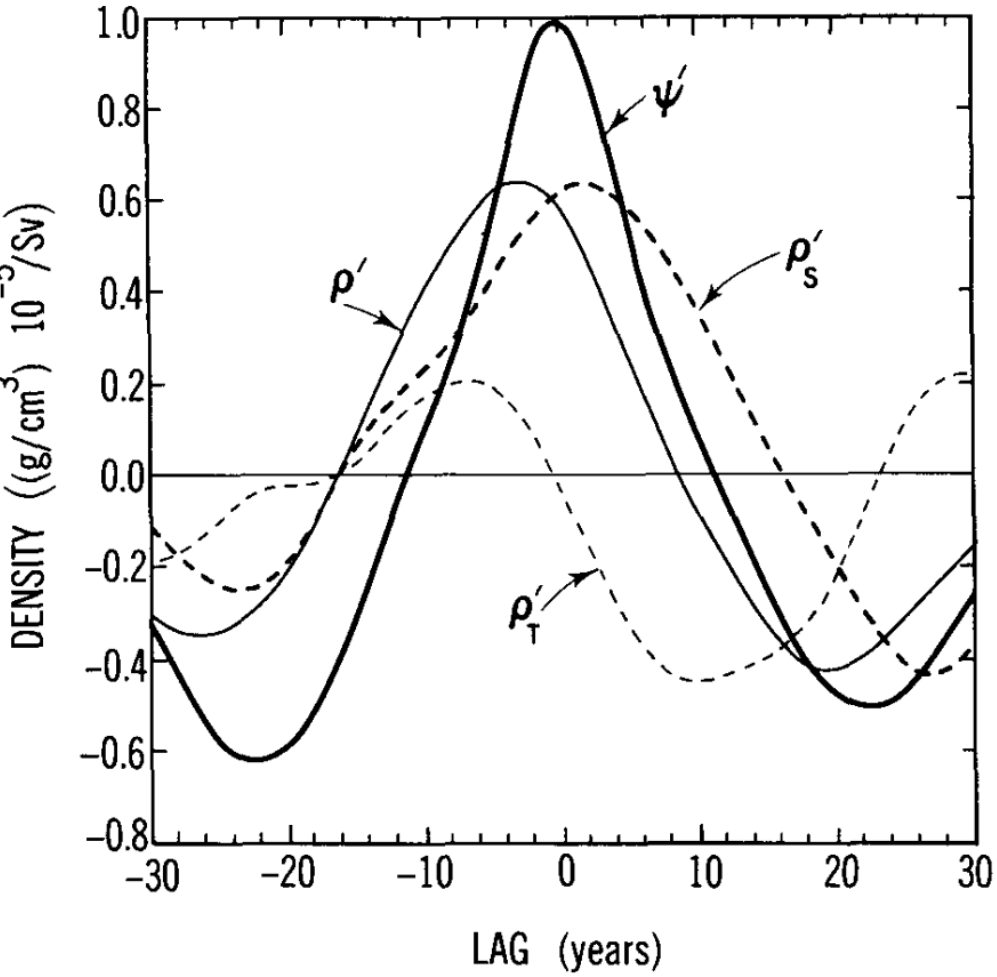

Figure 1.6: Regression coefficients between various quantities and the time series of the THC index. The heavy, solid line (Ψ') denotes the regression coefficients of the THC index with itself (thus representing a "typical" fluctuation). The thin, solid line (ρ') represents the regression coefficients between density and the THC index. The thick, dashed line (ρ'_S) denotes the regression coefficients for the density changes attributable solely to changes in salinity versus the THC index, while the thin, dashed line (ρ'_T) represents the regression coefficients for the density changes attributable solely to changes in temperature versus the THC index. The regression coefficients for ρ', (ρ'_S and (ρ'_T were averaged vertically and horizontally over the sinking region. The 'THC index' represents the maximum of the overturning streamfunction in the North Atlantic. This method of decomposing density changes has often been repeated and is used in this thesis (Figure 2.10b). Figure and associated caption taken from *Delworth et al.* (1993).

1.5 Potential benefits of increased model resolution

As analysing decadal variability requires many decades/centuries of integration these previous studies generally use low resolution coupled models ($>1^\circ$ ocean resolution, $>2^\circ$ atmosphere resolution) or higher resolution ocean-only models. However, there are reasons to suppose that improved atmospheric resolution could affect the amplitude of decadal variability (*Danabasoglu*, 2008), whilst improved ocean resolution — and associated representation of the variability within the Gulf Stream and other boundary currents (*Penduff et al.*, 2010) — may also affect the precise timescales of multi-annual/decadal variability (*Grotzner et al.*, 1998; *Gelderloos et al.*, 2011; *Hodson and Sutton*, 2012). Higher horizontal resolution may improve stratification and thus alter wave speeds (*Kirtman et al.*, 2012), which could also be affected by vertical resolution. Higher resolution topography may also be expected to affect the efficacy of wave processes as compared to idealised ocean models with smoothed/no topography (*Roussenov et al.*, 2008; *Zhang and Vallis*, 2007) and improve deep water pathways (*Spence et al.*, 2011).

At high ocean resolution eddies become permitted and then resolved (Figure 1.7) and eddy induced mixing can be left explicit, rather than parameterised (or the parameterisation significantly turned down), which may impact on the magnitude and variability of ocean heat and freshwater transports (*Volkov et al.*, 2008; *Tréguier et al.*, 2014). As a case in point, ultra-high resolution within the Agulhas region has been shown to affect the variability of the simulated low-latitude Atlantic overturning due to the shedding of large mesoscale eddies known as 'Agulhas rings' (*Biastoch et al.*, 2008b). In addition, reducing the need for parameterisation of eddies may affect the ability of the model to respond appropriately to 'out-of-sample' experiments, such as large scale atmospheric circulation changes under future climate change (*Hallberg and Gnanadesikan*, 2006). Stronger sea surface temperature gradients, associated with higher ocean resolution, may also improve the strength of atmosphere-ocean coupling (*Brayshaw et al.*, 2008; *Minobe et al.*, 2008; *Kirtman et al.*, 2012). In short, there are many reasons to suppose improved ocean and atmosphere resolution may change the details of simulated variability in the NA SPG. The model we use ('HadGEM3', *Walters et al.*, 2011) has an ocean resolution of 0.25°, an atmosphere resolution of N216 (92km at the equator), and is described in

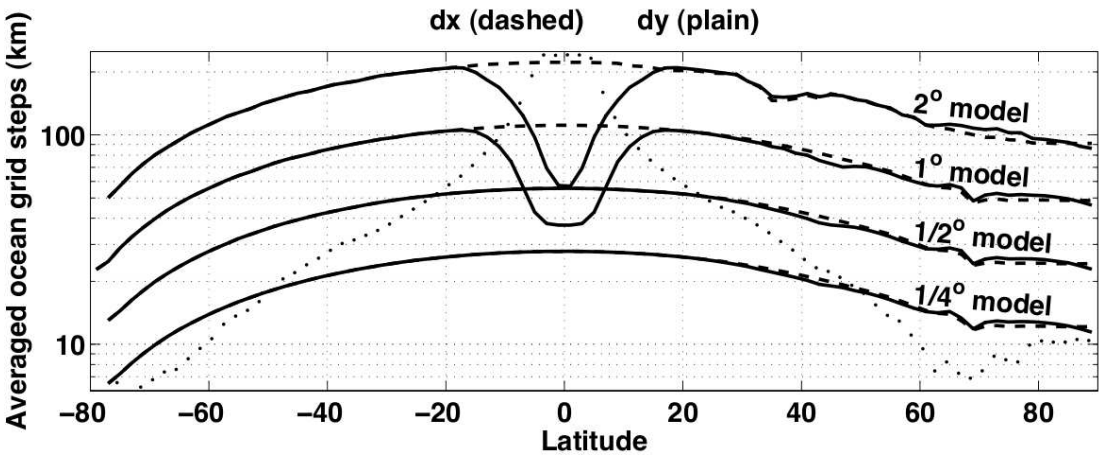

Figure 1.7: Zonally-averaged zonal (dashed) and meridional (plain) resolution of the four model grids (km, log scale). The meridional resolution is enhanced near the equator in the 2° and 1° models. Dots indicate the zonally-averaged first Rossby radius (*Chelton et al.*, 1998). Figure and associated caption taken from (*Penduff et al.*, 2010). Used under Creative Commons Attribution Licence 3.0.

Section 2.3.

1.6 Key research questions

In light of the material presented here, we now define the key research questions that we aim to address in this thesis. These are:

1. **What decadal variability exists in the NA SPG in HadGEM3, and how does this evolve in both space and time?** Given the plethora of previous lower resolution modelling studies that find many different manifestations of decadal variability in the NA SPG, and the enhanced resolution of HadGEM3, we first analyse whether HadGEM3 represents either a step-change in the simulated mechanisms of decadal variability, a refinement of these, or merely the same mechanisms more precisely resolved.

2. **To what extent is the decadal variability consistent with available observations?** Although there is a paucity of high quality long term observations within the North Atlantic there is nevertheless much analytical power in the observations that are available, which can be fully realised when combined with a detailed un-

derstanding of simulated variability. We conduct a critical comparison of simulated and observed variability to evaluate the plausibility of particular decadal variability.

3. **Does this variability provide potential skill for decadal predictions?** Having established the plausibility of various elements of the mechanism of decadal variability, we investigate to what extent they lead to increases in potential skill in decadal predictions with this model.

4. **Is there any systematic explanation for the diversity of simulated decadal variability within the NA SPG, as shown in Figure 1.5?** Recent multi-model archiving initiatives (*e.g.* the fifth coupled model intercomparison project, CMIP5) have provided a powerful resource to investigate inter-model relationships in a more consistent and controlled way than comparing reported results from studies with different aims. As such, although analysis of the decadal variability in HadGEM3 and its utility in predicting the real world is itself a worthy goal, we also aim to elucidate some of the potential causes of this multi-model spread and the implications for decadal predictions.

1.7 Thesis structure

In this thesis, we diagnose the drivers of NA SPG variability in a state-of-the-art coupled climate model that represents a rare combination of high resolution (in both ocean and atmosphere) and the multi-century length integration required to analyse decadal timescale modes (Chapter 2). We ask: Does high resolution, and the associated processes it allows, affect the nature of simulated decadal variability? We then critically compare this simulated mode of variability with instrumental observations in the North Atlantic in Chapter 3.

Most analyses of decadal variability in the North Atlantic have used the approach of lagged regression analysis to attempt to disentangle cause and effect — in which positive and negative phases are assumed symmetric. In Chapter 2 we note that there are signs of some asymmetry in the timescales of our proposed mechanism, depending on

the phase of the oscillation. These asymmetries are investigated in Chapter 4. Further testing of the robustness of our mechanism is conducted in Chapter 5, in which we test the predictability of various phases of our proposed mechanism using initial-condition ensembles.

Having thoroughly tested the mechanism of decadal variability in HadGEM3 we extend our analysis to other models using the CMIP5 archive in Chapter 6. Here, we address the question of whether there is a link between mean state biases and the representation of decadal variability, which would go some way to explain the wide variety of diagnosed mechanisms (*cf.* Figure 1.5).

Finally, the main conclusions of the thesis are summarised in Chapter 7, with recommendations for the development of decadal prediction systems and observational networks and a brief discussion of possible future work.

Chapter 2

The mechanism of North Atlantic internal decadal variability simulated in HadGEM3

2.1 Introduction

To begin to understand whether increased atmosphere and ocean resolution may affect the mechanisms of decadal variability we must first diagnose those mechanisms in a high resolution model. This chapter aims to elucidate the origin of decadal variability in the climate model 'HadGEM3'. The model's mean state and variability are characterised in Section 2.2 before diagnosing the mechanism of bidecadal variability in Section 2.4. A discussion of the proposed mechanism is directed in Section 2.5 before chapter conclusions are presented in Section 2.6. This chapter is reproduced in a similar form in *Menary et al.* (2015a).

2.2 Characterising the model

2.3 The model: HadGEM3

We examine a prototype of the Met Office Hadley Centre's state-of-the-art coupled ocean-atmosphere-land-ice global environment model, HadGEM3. 460 years of near present-day control simulation have been run at high resolution. The atmosphere component is the Met Office Unified Model version 7.7 (*Walters et al.*, 2011). It has a horizontal resolution of N216 (92km at the equator) and 85 levels in the vertical with a model top at 85km with at least 30 levels in or above the stratosphere. The ocean is resolved on the NEMO tripolar grid (0.25°, 75 depth levels, version 3.2, *Madec* (2008)), with a pole under Antarctica and poles either side of the Arctic Ocean in Asia and North America to resolve the Arctic Ocean. The ocean in HadGEM3 was initialised from rest at December 1st using the 2004–2008 time mean EN3 (*Ingleby and Huddleston*, 2007) December-time climatology and subsequently allowed to freely evolve with repeating 1978 external forcings in the atmosphere. For further details of the model configuration and other simulations see *Walters et al.* (2011).

HadGEM3 is a precursor to the model used in the Met Office global seasonal forecast, GloSea5 (*MacLachlan et al.*, 2015), which will also be similar to the new decadal prediction model. However, there are some differences between the HadGEM3 and GloSea5 models, as GloSea5 underwent additional development whilst the HadGEM3 control was running. Most importantly for the present study of the NA SPG is the more diffuse thermocline in the HadGEM3 ocean (NEMO version 3.2) as compared to GloSea5 (NEMO version 3.4, see discussion in Section 2.5) (*Megann et al.*, 2014). Despite this, the NA SPG biases in upper ocean temperature and salinity (compared to EN4), are small compared to many other coupled climate models used to study NA SPG variability (*Escudier et al.*, 2013; *Wang et al.*, 2014, see Section 2.3.1). Further details of global mean-state biases within the atmosphere and ocean in HadGEM3 can be found in *Walters et al.* (2011).

We use observed data from the EN4 objective analysis (*Good et al.*, 2013) which provides

infilled, optimally interpolated fields of temperature and salinity on a 1x1° grid from 1900 to present-day. EN4 is an updated version of EN3, with improved quality control and error estimates, but was not available when the climate model was initialised. We use the period 1900–2013 to construct a simple climatology for comparison with HadGEM3 and note that the biases in HadGEM3 are large enough (see Section 2.3.1) that the method used to construct the EN4 climatology is unlikely to be of first order importance. Unlike the HadGEM3 model, which is run with interannually constant forcings appropriate for the year 1978, this observational data also includes the effects of all other natural and anthropogenic forcings.

We now examine the NA mean state biases and signal of decadal variability in HadGEM3 in some more detail as a precursor to investigating the mechanisms of variability which exist on top of these biases.

2.3.1 NA SPG Mean state

Mean state biases in top 500m depth averaged temperatures (T500), salinities (S500), and densities (ρ500) in the NA SPG are less than $\pm$3°C, $\pm$0.4PSU, and $\pm$0.1kg/m^3 in the interior NA SPG, with larger +4°C, +0.6PSU, and $\pm$0.2kg/m^3 biases in the boundary current regions (Figure 2.1). The temperature and salinity biases are close to being density compensating in the NA SPG but in the subtropical gyre (not the focus of this study) temperature biases dominate resulting in lighter waters. The anomalously cold region in the western SPG, often attributed to the simulated Gulf Stream being too zonal (*Kwon et al.*, 2010), is not as large as in many coupled climate models (*Scaife et al.*, 2011). Warm anomalies exist all along the NA SPG northern boundary currents. These anomalies are associated with reduced ice distribution around southern Greenland and in the Labrador Sea (not shown). Within the NA SPG, deep convection, as estimated from the annual standard deviation in March mixed layer depths using the mixed layer estimate of *Kara et al.* (2000), is located in the Labrador Sea and Irminger Current.

The simulated Atlantic meridional overturning circulation (AMOC) streamfunction in the model is shallow compared to recent observations (Figure 2.2a). The zero streamfunction

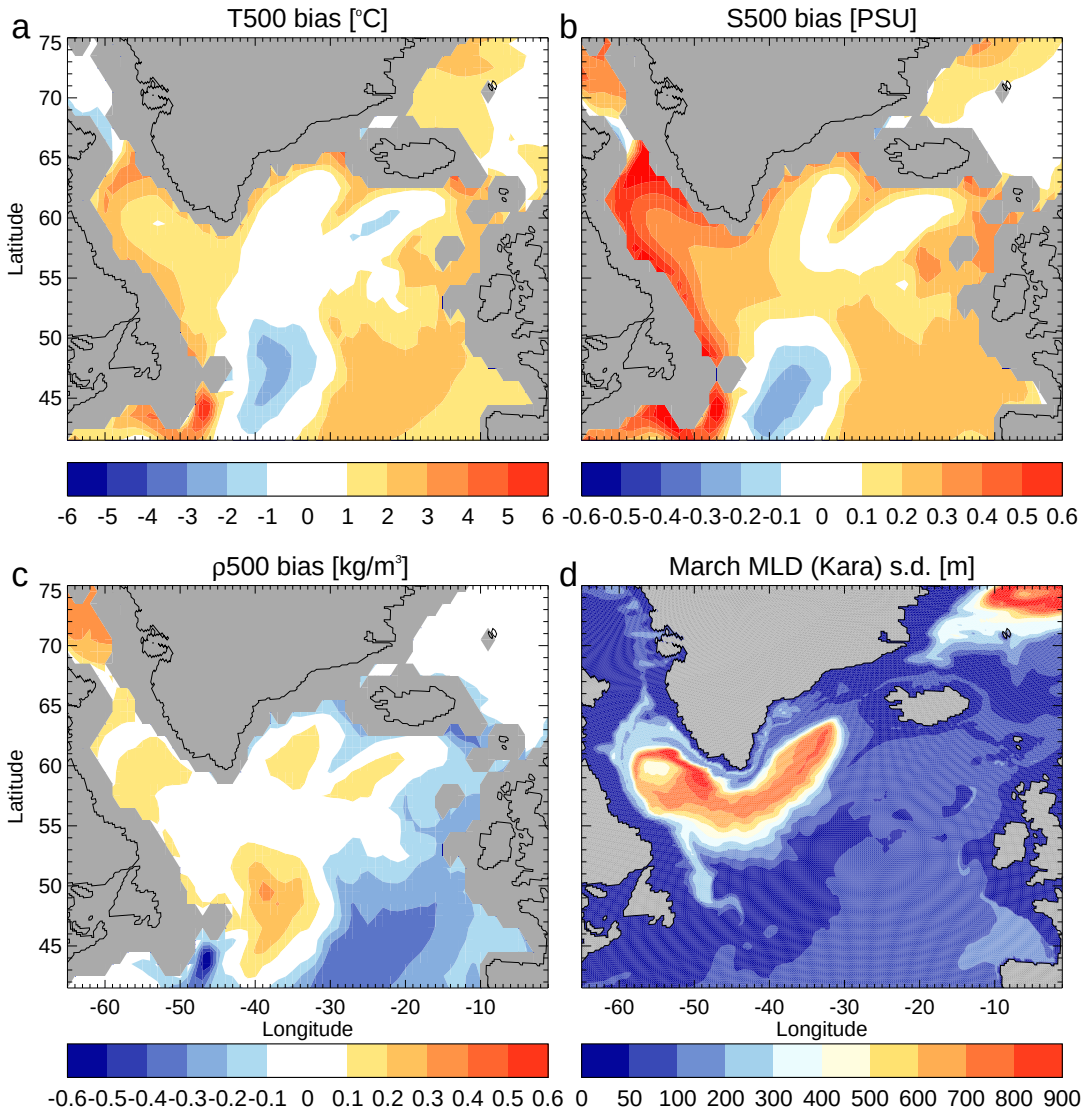

Figure 2.1: Top 500m depth averaged temperature (T500, a) salinity (S500, b), and density (ρ500, c) biases in HadGEM3 (computed from full model time series) compared to EN4. Grey shading is used for regions shallower than 500m. d) Standard deviation in March mixed layer depths (*Kara et al.*, 2000), to highlight where deep convection occurs

line in the model sits at a depth of 2–3km with the maximum overturning occurring at a depth of approximately 1km. The deeper overturning cell, representing Antarctic Bottom Water (AABW) and Lower North Atlantic Deep Water (LNADW) has a strength of around 3 Sverdrups (Sv = $10^6 m^3/s$), whereas the shallower AMOC cell, representing the western boundary current and Upper North Atlantic Deep Water (UNADW) has a mean strength of 17Sv for the last 200 years of the simulation.

At 26°N it is possible to directly compare the streamfunction in the model to the RAPID (*Cunningham et al.*, 2007) observations. The depth of the RAPID overturning maximum is marked with a cross and is approximately 200m deeper than in the simulations, which at these depths represents a single model grid cell in the vertical. The depth of the RAPID zero streamfunction line is around 4km, much deeper than simulated. This is not uncommon in models and may be partly an artefact of computing the simulated overturning using the full 3-dimensional velocities (*Roberts et al.*, 2013), although some models do represent a much deeper upper cell (*Yeager and Danabasoglu*, 2012). Indeed, using a 'RAPID-style' calculation, after *Roberts et al.* (2013) (with a depth of no motion at 4740m) yields a zero streamfunction depth approximately only 250m deeper than using the full 3-dimensional velocities; the structure and variability of the streamfunction shallower than this are essentially unchanged. Finally, the NA SPG barotropic streamfunction and associated time series are also shown (Figure 2.2, b and d) and broadly compare well to observational estimates and high resolution models (*Tréguier et al.*, 2005).

Although the depth (1000m) and strength (17Sv) of the maximum of the upper AMOC cell are consistent with observations, the simulated annual variability in this index is weaker than observed. The simulated annual mean AMOC streamfunction at 26.5°N and 1000m depth has a standard deviation of 1.2Sv (0.9Sv if first detrended), compared to an annual standard deviation of 2.3Sv from the 10 years of RAPID data available (Figure 2.2c). Additionally, the simulated index begins at a low value and then takes several centuries to spin-up to a more stable state more favourably comparable to the observed mean. Although this represents an improvement in this index of the NA circulation, the spin-up of the overturning circulation also results in an increase in northward heat and salt transport within the Atlantic Ocean, causing the NA SPG to drift away from its initialised state to a warmer and saltier state, seen in Figure 2.1

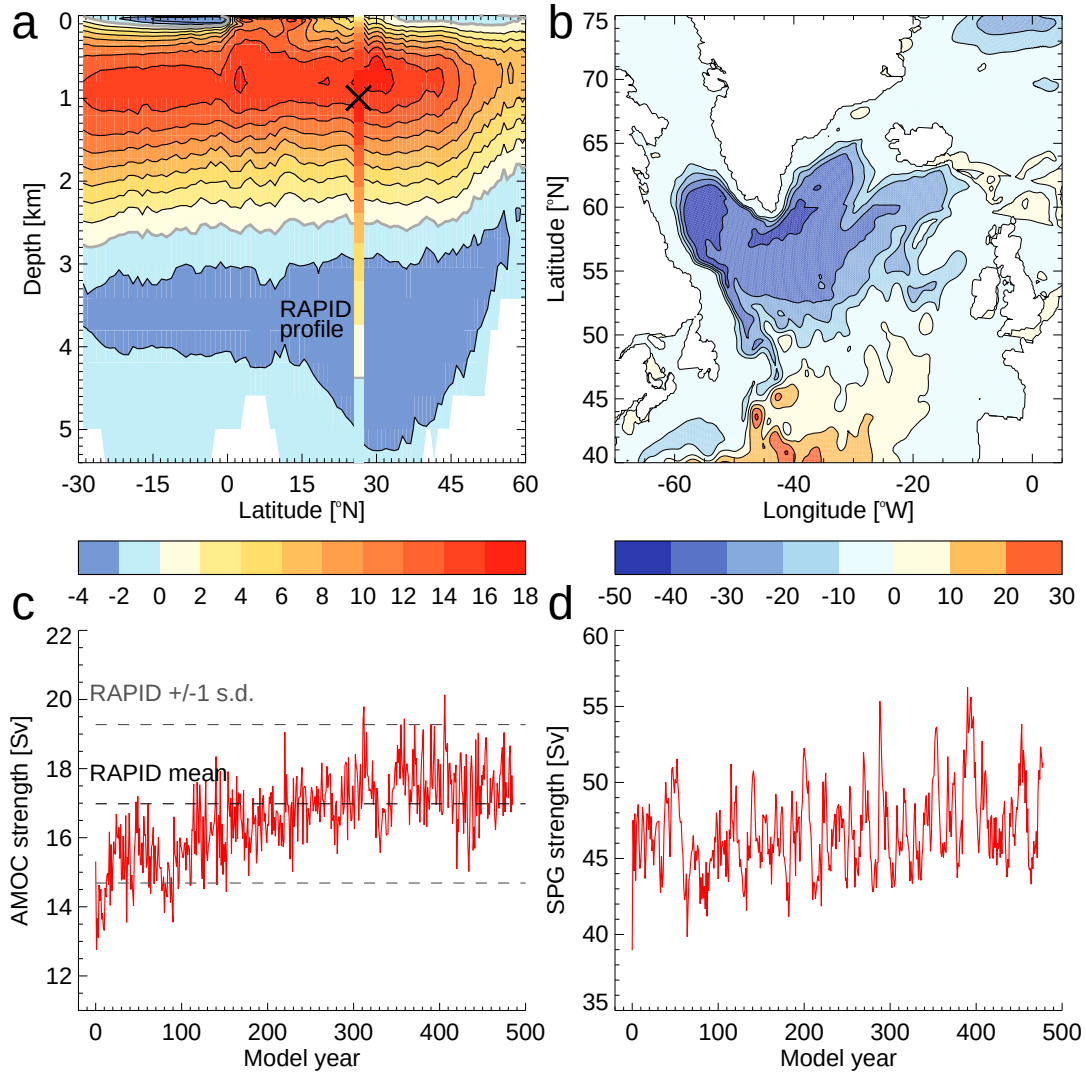

Figure 2.2: a) Time mean Atlantic overturning streamfunction in HadGEM3. The contour interval is 2Sv and the zero-line is marked with a grey contour. At 26.5°N the profile from the RAPID array (*Cunningham et al.*, 2007) is overlaid on the same colour/contour scale. The depth of the maximum in the RAPID profile is marked with a cross. Note that the latitudes north of 45°N are approximate (within 1°) due to the increasingly curved nature of the model grid towards the two northern poles. b) Time mean NA SPG barotropic streamfunction in HadGEM3. Contour interval is 10Sv. c) Time series of the overturning streamfunction at 26.5°N and 1000m in HadGEM3 (red). Also shown are the time mean and annual mean standard deviation from the 10 years of RAPID data (black). d) Time series of the minimum (multiplied by -1) of the barotropic streamfunction in the NA SPG in HadGEM3.

The simulated AMOC index also shows some evidence of multi-annual/decadal variability, particularly at the more northerly latitudes of the NA SPG (not shown) in addition to 26N, as in other models (*Zhang*, 2010). The maximum correlation between the simulated AMOC indices at 26.5°N and 50°N occurs when the 50°N index leads by 1 year (correlation of 0.63), suggesting the lower latitude variability is responding to variability further north in the NA SPG. The lag-zero correlation of the AMOC transport, calculated after *Bingham et al.* (2007), is shown in Figure 2.3 and highlights the separation between the South Atlantic, mid latitude Atlantic, and the NA SPG. Interestingly, the North Atlantic coherence pattern is more similar to that seen in the lower resolution HadCM3 model (1.25°) than higher resolution OCCAM model (0.25° *Bingham et al.*, 2007) but in all cases there is clear separation between the subtropical and subpolar gyres. We now move on to examine the decadal variability of the NA SPG in more detail.

2.3.2 Signal of decadal variability

The time-mean T500 simulated in HadGEM3 is shown in Figure 2.4a along with contours at 6 and 10 degrees to mark the general shape of the NA SPG. A comparison with observations (EN4) again shows the general warm bias of the NA SPG, particularly towards the edges of the gyre. A power spectrum for T500 averaged over the whole region reveals a significant peak at a period of 16 to 17 years (Figure 2.4b). This periodicity exists whether using the entire simulation or alternatively removing the first 200 years (not shown), suggesting it is not merely an adjustment process, and so we use the entire time series to maximise the available data. Additionally, the periodicity is not unique to any of the four individual subregions within the NA SPG (dashed regions in Figure 2.4a); all show a significant peak at 16 to 17 years, as well as the North Atlantic Current (NAC) region (Figure 2.5). Indeed, in HadGEM3 many other large scale ocean indices in the NA SPG also reveal peaks in their power spectra at periods of 16 and 17 years, such as SSTs, depth averaged salinities, the AMOC at 50°N, or the strength of the NA SPG itself (as defined by the barotropic streamfunction, *cf.* Figure 2.2).

In addition to these ocean indices, the NAO index also shows periodicity at 16 to 17 years in its otherwise much whiter spectrum (Figure 2.4c). This is suggestive of a link from

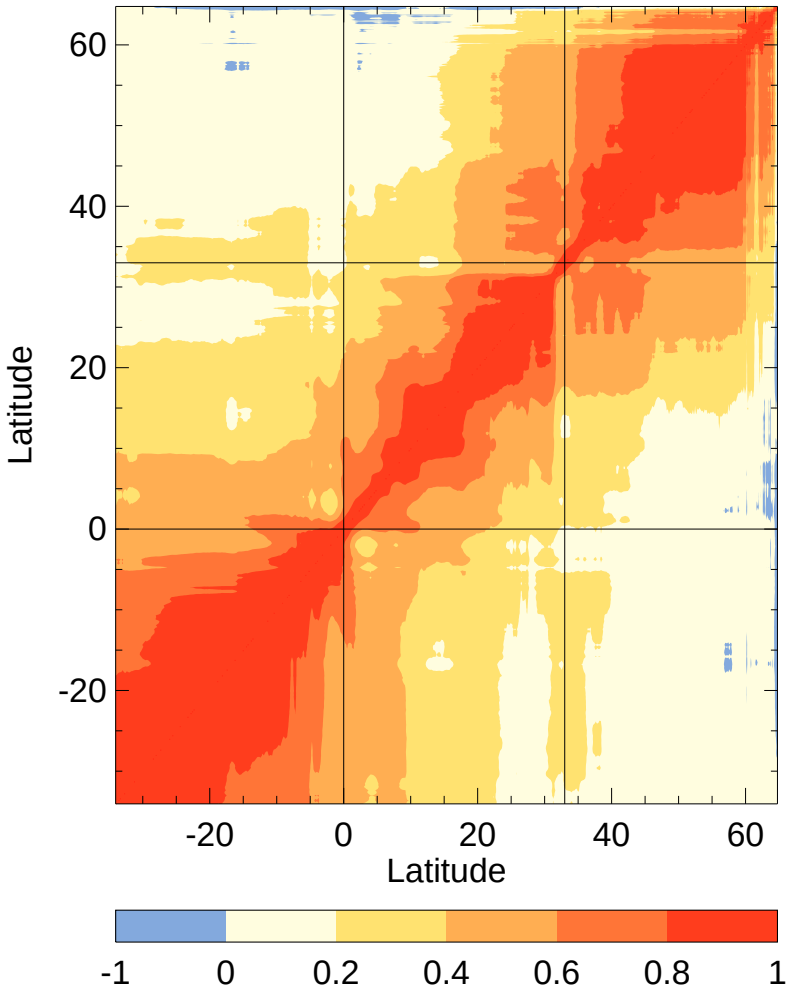

Figure 2.3: The cross-correlation between the simulated depth averaged (100–1000m) zonal mean transport in the Atlantic as a function of latitude. Data have been detrended. Black lines denote the latitude of the equator and where the Gulf Stream separates from the coast of North America.

ocean to atmosphere in the region of the NA in which the ocean can impart some of its long term memory on to the atmosphere. Such a feedback might in general be expected to be weaker than similar atmosphere to ocean processes, and related to the strength of the ocean circulation and SST gradients (*Nonaka and Xie*, 2003), and thus detection of this feedback is perhaps at least in part due to the increased signal to noise ratio resulting from the length of the control simulation (though we note this is still short compared to many previous studies with lower resolution models). The mechanistic drivers behind this 17 year mode in the ocean and atmosphere, and the reasons for the particular timescale, are investigated in the next section; initially characterising the variability in the NA SPG as a whole before targeted analysis of the processes in different regions.

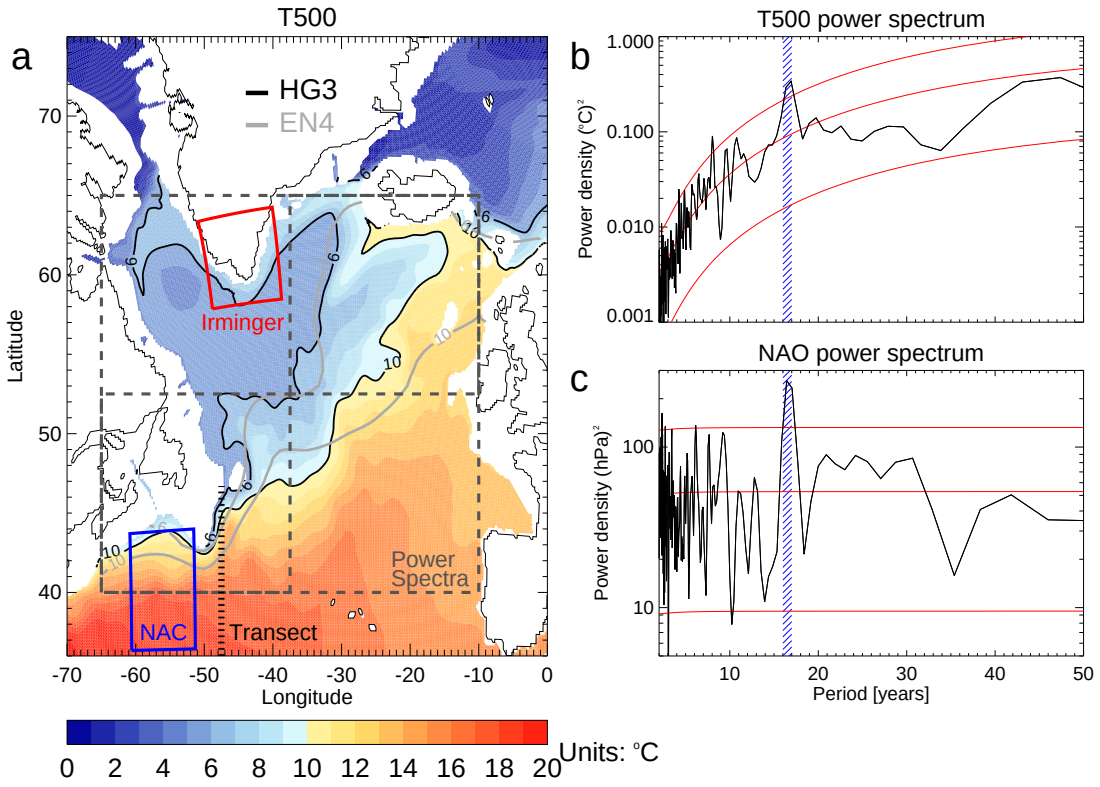

Figure 2.4: a) Time mean top 500m depth averaged temperature (T500) in HadGEM3. Contours at 6°C and 10°C are also marked (black) to show the shape of the gyre and for comparison with equivalent contours from EN4 (grey). Areas in white are shallower than 500m. The dashed grey box denotes the four quadrants and fifth overall region for which power spectra of T500 were produced. The Irminger Current region (red box) and Gulf Stream/North Atlantic Current region (blue box) analysed in the text are also marked. The dashed black line stretching south from the Grand Banks denotes the transect location for dynamic height analysis in Figure 2.12. b) The T500 power spectrum for the whole subpolar gyre region (combination of all four quadrants). An estimate of significance is given by the 5–95% confidence intervals for a red noise process with the same mean and standard deviation. Periods of 16–17 years are highlighted with the blue shading. c) As (b) but for the NAO index, defined as the difference between simulated sea level pressures over the Azores and Iceland. Time series are linearly detrended prior to calculating the power spectra.

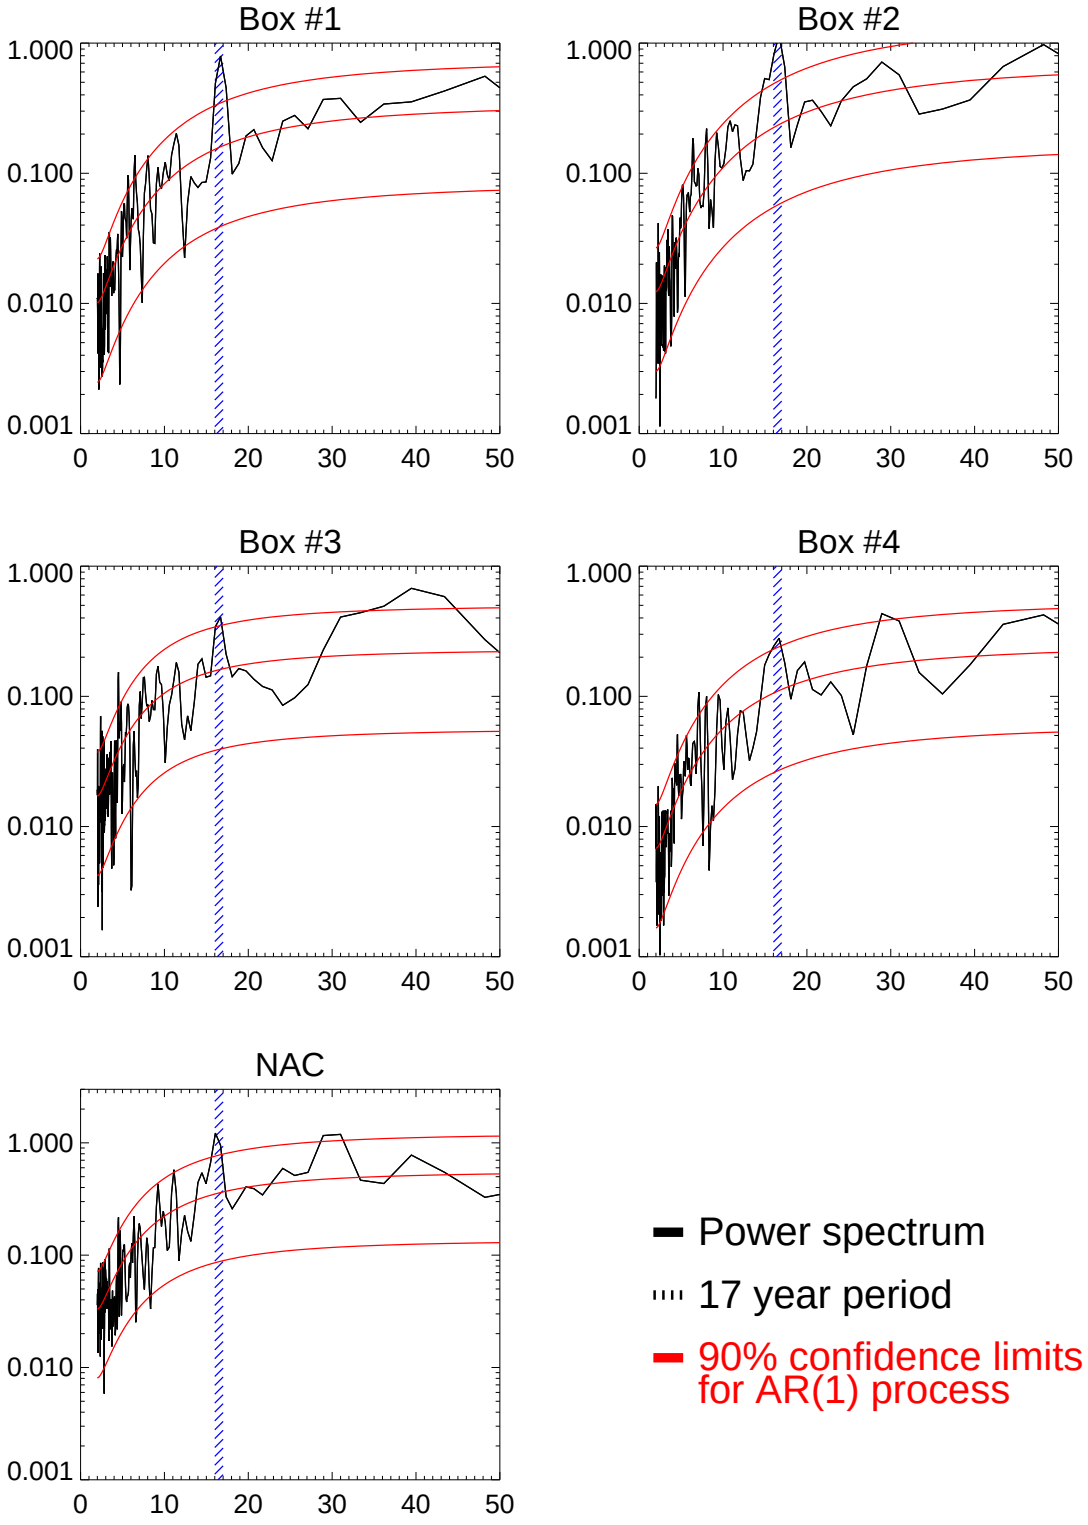

Figure 2.5: The top 500m depth averaged temperature power spectra for the 4 smaller boxes encompassed by the dashed lines in Figure 2.4, reading from left to right and top to bottom, as well as the box further south in the North Atlantic Current (NAC) region. An estimate of significance is given by the 5–95% confidence intervals for a red noise process with the same mean and standard deviation. Periods of 16–17 years are highlighted with the blue shading.

2.4 Mechanism of decadal variability in the NA SPG

We now diagnose the mechanism of decadal variability within the NA SPG, beginning with a heat budget for the region before investigating how temperature anomalies propagate around the gyre.

2.4.1 Heat budget

To begin to understand the variability of T500 in HadGEM3 a heat budget of the NA SPG is diagnosed. The basin-wide, full depth NA SPG heat budget is shown in Figure 2.6 for the latitude range 53–73°N. Due to the lack of availability of the correct ocean diagnostics at high enough output frequency (precluded by the expense of storing high resolution atmosphere and ocean data), the heat budget of the NA SPG does not close perfectly (*cf.* red and black lines in Figure 2.6a). However, the error is negligible, less than 1% of the net surface fluxes of the region. Sensitivity tests where all output diagnostics were computed online and stored revealed that horizontal isopycnal diffusion was the most important missing heat flux. The heat budget of the NEMO ocean is further complicated by the use of a linear free surface with variable volume which sits on top of the fixed volume ocean grid cells and a heat flux between the two. For further details of the precise formulation of the heat budget within the NEMO ocean model see *Madec* (2008).

The heat budget (ocean heat content (OHC) rate of change) of the NA SPG can first be broken down into advective (Q_{adv}) and surface fluxes (Q_{surf}), which add together to give the net heat flux *into* the volume. There are also additional smaller heat fluxes from the ice to ocean, between the linear free surface and fixed ocean volume, and from geothermal heating of the abyssal ocean, particularly in the vicinity of the mid Atlantic ridge, Q_{ice}, Q_{free}, and Q_{geo} respectively:

$$\frac{\mathrm{d}OHC}{\mathrm{d}t} = Q_{adv} + Q_{surf} + Q_{ice} + Q_{free} + Q_{geo} \tag{2.1}$$

The advective fluxes can be further broken down into fluxes from the north ($OHTN$) and

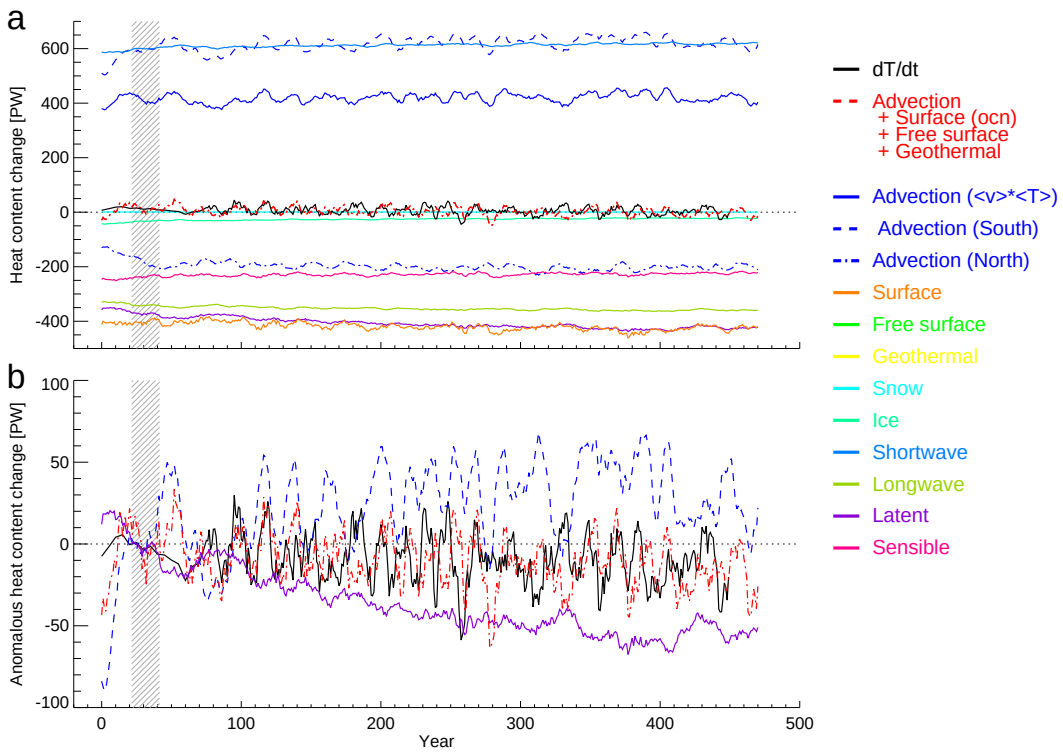

Figure 2.6: The full depth heat budget of the NA SPG (53–73°N) volume plotted using 9 year running means for clarity and to highlight decadal variability. Positive is into the specified region. a) Individual components of the heat budget as denoted in the legend. b) The anomalous heat budget (referenced against years 22–42) to highlight the trends in latent heat fluxes and advective heat fluxes through the southern boundary. Note that the heat content change ($\frac{\mathrm{d}T}{\mathrm{d}t}$, black) and sum of heat fluxes (red) do not match prior to the year 100 as instantaneous ocean temperatures (used to calculate $\frac{\mathrm{d}T}{\mathrm{d}t}$) were stored with intermittent frequency and don't necessarily represent the same time window as the other heat fluxes during this time.

south ($OHTS$, positive northward) whilst the surface fluxes can be broken down into the shortwave (solar), longwave, latent, and sensible heat fluxes:

$$Q_{adv} = OHTS - OHTN \tag{2.2}$$

$$Q_{surf} = Q_{SW} + Q_{LW} + Q_{lat} + Q_{sens} \tag{2.3}$$

Inspection of these terms reveals that $OHTS$ dominates the variability in advective heat fluxes: Using annual data, the standard deviation of $OHTS$ is 28PW, compared to 17PW for $OHTN$. The variability in $OHTS$ is split between vertical 'AMOC' (Chapter 1,

Equation 1.2) and horizontal 'gyre' (Equation 1.3) heat transport variability at these latitudes (annual correlation between $OHTS$ and OHT_{AMOC} is 0.74, and between $OHTS$ and OHT_{gyre} is 0.88). The surface fluxes (directed into the ocean) are dominated by shortwave (solar) heating of the NA SPG, whereas longwave, latent, and sensible heat fluxes represent net heat loss from the NA SPG.

To investigate the relative magnitudes of their variability, and any trends, the mean of each of the key heat fluxes over the years 22–42 is removed (Figure 2.6b). Rather than remove the full time mean, removing the mean from just the period soon after the model was initialised serves to additionally show how the heat fluxes diverge. Net advective heat fluxes into the region are increasing throughout the period, balanced largely by increasing surface heat flux loss, but with some residual heating of the NA SPG. The advective heat flux trend is dominated by the increase in heat flux from the south, which is due to the strengthening AMOC (Figure 2.2c), with much of this heat lost via latent heat loss as well as longwave emission. The rate of net warming is highest in the first century, which is also why the net heat flux appears to be below zero for the remainder of the time, *i.e.* the net warming rate is slower in the subsequent years.

There is considerable variability in the net heat flux into the NA SPG, the majority of which appears to be attributable to the advective heat fluxes from the south, which results in decadal timescale heat content changes within the NA SPG. Annual and decadal correlations between the total heat flux and net advective fluxes are 0.75 and 0.69, whereas the same for the total heat flux and net surface fluxes are 0.63 and 0.29 (the regression gradients scale similarly) suggesting that particularly on decadal timescales advective heat fluxes dominate the variability. Once within the NA SPG, how do these heat content anomalies evolve?

2.4.2 Spatial characteristics of decadal variability

In order to investigate the spatial characteristics of the heat content variability, lagged regressions were performed of NA SPG T500 on to SST spatially averaged over the NA SPG (Figure 2.7, second column). T500 anomalies can be seen propagating around the

NA SPG: eastwards along the southern boundary whilst spreading into the interior with a timescale of around 4–6 years (notably slower than implied by the mean circulation speed in this region); westwards along the northern edge but south of the Greenland, Iceland, Norwegian (GIN) Seas; into the central Labrador Sea as opposite sign anomalies form in the Gulf Stream region. A similar evolution of anomalies was also found when regressing T500 on to T500 spatial averages over the eastern SPG, NAC region, or Labrador Sea (not shown). Features such as the Reykjanes Ridge can be seen diverting the flow. Although not shown here, there is little evidence of significant amounts of the signal diverting into the GIN Seas in the far northern part of the SPG. The heat content anomalies reach the Labrador Sea from the eastern SPG within a couple of years but several more years are required for the anomalies to spread into the interior SPG. As the heat content anomalies in the Labrador Sea build up so does a cold anomaly in the Gulf Stream/NAC region. The opposite phase of the cycle now begins.

The underlying essence of the cycle is captured by regressing T500 indices in the northern and southern edges of the NA SPG against each other (Figure 2.8). This shows the southern edge of the NA SPG leading the northern edge by $\approx$6–7 years and subsequently lagging changes in the northern edge by $\approx$2 years with opposite sign, yielding a half period of 8–9 years and a full period of 16–18 years (constrained here to be even by the use of annual data). The range in periodicities is further increased by 2 years if a third location in the eastern SPG is added to the regression model (not shown), forcing the signal to go via the eastern SPG, suggesting that the spread in timescales is perhaps related to the superposition of various advective pathways. This decadal mode is generally confined to the top 500m–1km with the exception of the central Labrador Sea where it extends to around 2km (Figure 2.9). Decadal variability in the band 10–30 years, encompassing the spectral peak at 17 years, explains $>$15% of the annual variability in T500 within the NA SPG, with this value rising to $>$30% in the centre of the gyre.

The lagged regression analysis leads to two key questions: Firstly, what is controlling the apparent propagation of the heat content anomalies in both a) the Gulf Stream extension/NAC, and b) the northern boundary currents/Irminger Current? Secondly, what is the negative feedback that forms the opposite sign anomaly in the NAC, resulting in a cyclical mechanism and a spectral peak in NA SPG temperatures? To investigate these

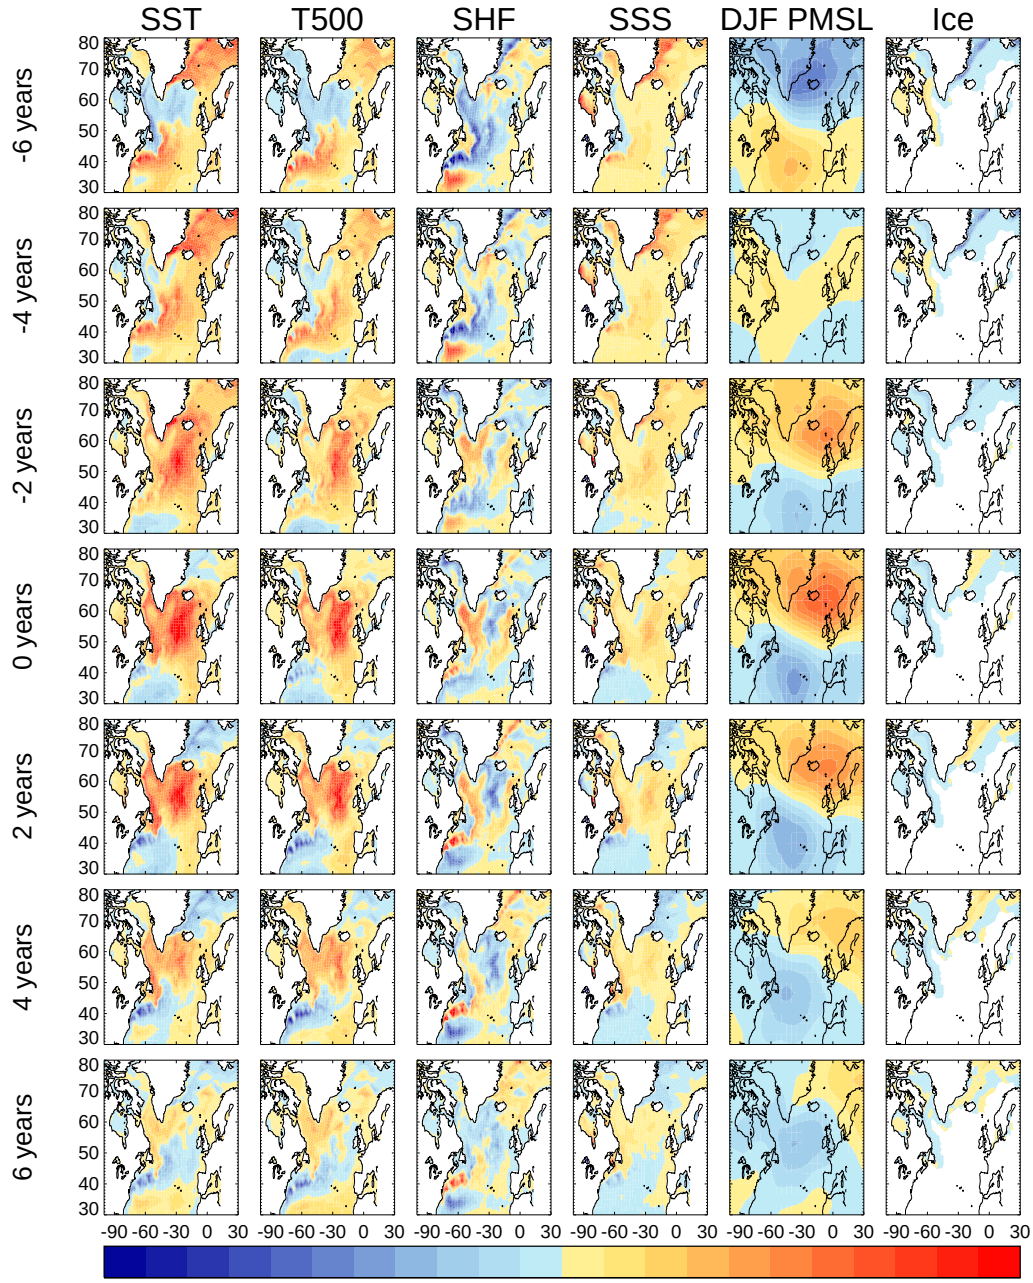

SST [°C/°C]	-1	-0.8	-0.6	-0.4	-0.2	0	0.2	0.4	0.6	0.8	1
T500 [°C/°C]	-1	-0.8	-0.6	-0.4	-0.2	0	0.2	0.4	0.6	0.8	1
SHF [W/m²/°C]	-20	-16	-12	-8	-4	0	4	8	12	16	20
SSS [PSU/°C]	-0.5	-0.4	-0.3	-0.2	-0.1	0	0.1	0.2	0.3	0.4	0.5
DJF PMSL [hPa/°C]	-5	-4	-3	-2	-1	0	1	2	3	4	5
Ice [frac/°C]	-0.25	-0.2	-0.15	-0.1	-0.05	0	0.05	0.1	0.15	0.2	0.25

Figure 2.7: Regressions between North Atlantic subpolar (45–65°N) sea surface temperatures (SSTs) and, from left to right: SST, top 500m depth averaged temperature (T500), net surface heat flux into ocean (SHF), sea surface salinity (SSS), wintertime mean sea level pressure (DJF MSLP), and ice fraction. From top to bottom, the SST index lags then leads the fields from -6 to +6 years. All data are bandpass filtered to remove interannual variability and the trend using a bandpass filter of 5–70 years.

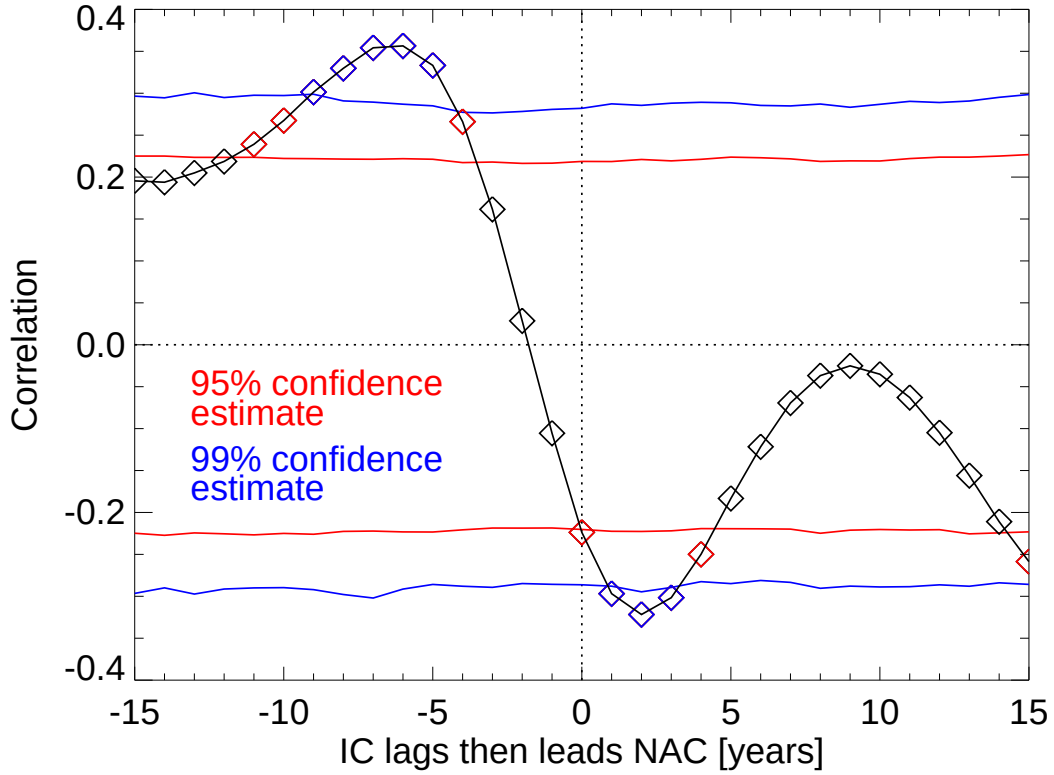

Figure 2.8: The lagged correlation between the Irminger Current and North Atlantic Current top 500m depth averaged temperatures (T500). Regions are as marked in Figure 2.4a. Time series have been detrended and smoothed with a 5 year running mean to highlight the decadal correlations by removing annual variability and the long term drift. An estimate of the significance is provided by the 95% (red) and 99% (blue) confidence intervals estimated by creating 40,000 random time series with the same mean, standard deviation, and applied filtering.

questions, we first break down the advective heat budget into components related to the mean and anomalous circulation/temperatures.

2.4.3 Decomposition of the advective heat budget

Previously *Dong and Sutton* (2001), showed the advective heat budget for a region in approximate long term equilibrium could be estimated by considering perturbations around a long term mean as:

$$q(t) = \bar{q} + q'(t) \tag{2.4}$$

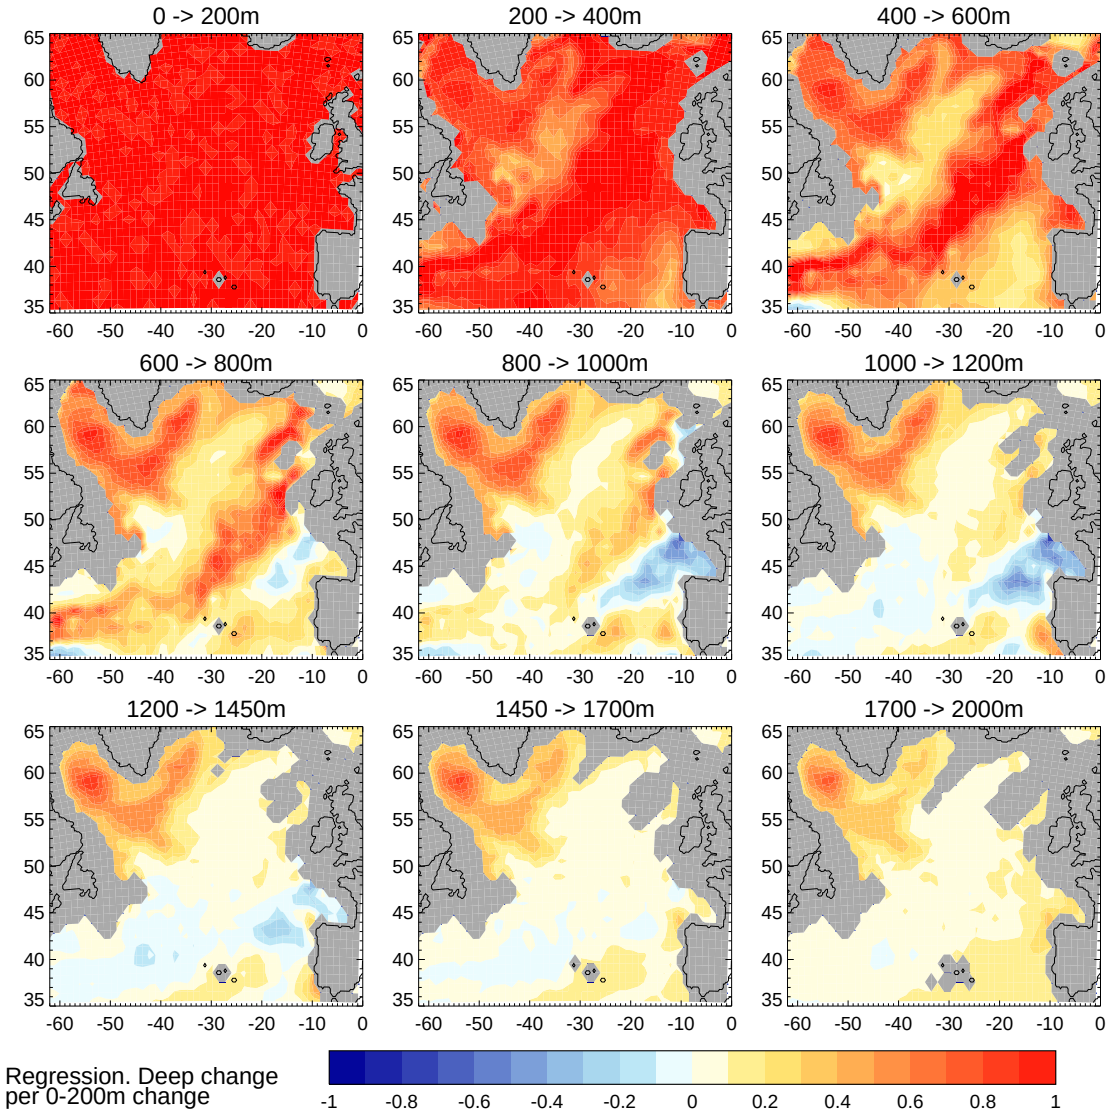

Figure 2.9: . Regression slopes of 200m depth averaged temperature against approximately 200m thick ranges in the water column. Regressions computed at every horizontal location to highlight how the depth extent of the decadal mode varies in space. Data is filtered allowing only period of 10–30 years to focus on the decadal mode.

where $q(t)$ is some quantity varying in time, $\bar{q}$ is its time mean, and $q'(t)$ is the anomaly in q at each time, t. For the case of the net advective heat transport convergence, replacing q with both v (velocity) and T (temperature) and dropping the (t) on the right hand side for clarity gives:

$$OHT(t) = \rho c_p \int \left(\bar{v}\bar{T} + v'T' + \bar{v}T' + v'\bar{T} \right) \mathrm{d}A \tag{2.5}$$

where ρ is density, c_p is the heat capacity of seawater, $\bar{v}\bar{T}$ is a constant (the mean heat transport when multiplied by ρc_p), $v'T'$ is the heat transport due to co-variances in circulation and temperature (and is usually but not always small for large enough areas), $\bar{v}T'$ is the heat transport by the mean circulation, $v'\bar{T}$ is the heat transport by the anomalous circulation, and $\mathrm{d}A$ indicates integrating over all faces (horizontal and vertical) enclosing the volume. This is the same as Equation 1.1 except we now explicitly integrate around a closed volume.

However, as previously mentioned, there is a trend in the NA SPG temperatures in HadGEM3, and so the breakdown of the heat budget is made more complicated. For the case of a known trend in one or more of these parameters (*e.g.* temperature) the q' term will not just represent the annual/decadal anomaly but will also have a component due to the trend with the relative contributions to q' varying in size depending on the magnitude of the trend compared to the magnitude of the variability. Thus q must be detrended prior to combining the terms together, *e.g.*

$$q(t) = q_0 + q_1 t + q'(t) \tag{2.6}$$

where q_0 is the intercept, q_1 is the linear trend multiplied by time, t, and q' is the perturbation from this trend. Setting $t = 0$ at the midpoint of the linearly trending time series results in q_0 also representing the mean (previously $\bar{q}$). This results in the OHT becoming an equation of nine terms (as we detrend v as well due to the trend in the AMOC, Figure 2.2c):

$$OHT(t) = \rho c_p \int \Big(v_0 T_0 + v_0 T_1 t + v_0 T' + v_1 t T_0 + v_1 T_1 t^2 + v_1 t T' + v' T_0 + v' T_1 t + v' T' \Big) \mathrm{d}A \quad (2.7)$$

where the terms inside the integral on the right hand side respectively refer to: 1) The time mean OHT, 2) the interaction between the temperature trend and the mean circulation, 3) the OHT due to anomalous temperature advected by the mean circulation, 4) the interaction between the mean temperature and the trend in circulation, 5) the interaction between the trends in both circulation and temperature, 6) the interaction between the trend in circulation and the anomalous temperatures, 7) the OHT due to mean temperature advected by the anomalous circulation, 8) the interaction between the anomalous circulation and the trend in temperatures, and 9) the OHT due to co-variances in circulation and temperature. Analysis of these components reveals a non-zero contribution from the trend-related terms to the advective heat budget variability, but these are much smaller than (less than half) the mean and anomalous circulation terms ($v_0 T'$ and $v' T_0$) and so we focus on these latter circulation terms.

2.4.4 Heat content anomalies in the NAC region

To determine what controls the heat content changes on the southern boundary of the NA SPG, the heat budget of the NAC region is examined in more detail. A region was chosen where simulated zonal currents are much stronger than meridional or vertical currents (See Figure 2.4a, blue box). This simplifies the later interpretation of the decomposition of advective heat fluxes into circulation and temperature components. As noted in Section 2.4.1, it is not possible to close the heat budget precisely, which becomes more apparent for smaller subregions. Table 2.1 shows the time mean advective components and net surface heat fluxes for the NAC top 500m. Note that the choice of reference temperature becomes irrelevant when considering the net transport through all faces combined but not when considering open sections (*Schauer and Beszczynska-Möller*, 2009). The most important advective heat fluxes are from the east and west, associated with the mean volume transport through the region from east to west. These advective heat fluxes are

	NAC	Irminger Current
East advection	-1660	493
West advection	1753	-392
North advection	33	0
South advection	-83.1	-23.7
Net vertical advection	7.5	-63.9
Net convergence	51.4	15.2
Surface	-49.8	-16.9
A: Sum of advection and surface (net sum)	1.6	-1.7
B: Ocean heat content change (net actual)	0.1	0.1
Correlation A:B (Monthly, Annual, Decadal)	0.96, 0.93, 0.98	0.96, 0.94, 0.95

Table 2.1: Time mean simulated heat fluxes into the North Atlantic Current (NAC) and Irminger Current regions (TW, referenced to 0°C).

approximately balanced by the surface heat fluxes but the sum of the two is not identical to the actual heat content change implied by the in-situ temperatures. This is likely due to missing diagnostics (See Section 2.4.1) and the use of monthly means when computing vT, rather than at each model time step. However, although the means are slightly different, the variability in both time series is well correlated on all timescales at monthly or longer sampling (Table 2.1). Thus, in the ensuing analysis of the variability, we treat the budget as sufficiently closed for our purposes.

The annual and decadal timescale correlations (regression gradients, W=Watts) between the advective heat fluxes and the net heat content changes in the NAC are 0.82 (0.92 $\mathrm{W}_{\mathrm{d}OHC}/\mathrm{W}_{adv}$) and 0.54 (0.40 $\mathrm{W}_{\mathrm{d}OHC}/\mathrm{W}_{adv}$) respectively, as compared to 0.43 (0.92 $\mathrm{W}_{\mathrm{d}OHC}/\mathrm{W}_{surf}$) and 0.20 (0.20 $\mathrm{W}_{\mathrm{d}OHC}/\mathrm{W}_{surf}$) for the correlation between surface heat fluxes and the net heat content change (for annual and decadal data the 95% significance levels, assuming a two-tailed t-test, are 0.12 and 0.37 respectively). Thus much of the annual and decadal variability in the heat content changes in the Gulf Stream is associ-

	Monthly	Annual	Decadal
NAC v_0T'	139	43	31
NAC $v'T_0$	149	44	33
NAC $v'T'$	58	16	11
Irminger Current v_0T'	13.1	4.0	3.2
Irminger Current $v'T_0$	6.7	3.6	2.7
Irminger Current $v'T'$	4.0	1.1	1.0

Table 2.2: Standard deviations of advective heat flux components in the North Atlantic Current (NAC) and Irminger Current at various timescales (TW).

ated with advective heat fluxes but there is a role for surface fluxes to modulate these changes, even on decadal timescales. Of the advective heat fluxes, the remaining question is whether these are due to the anomalous circulation or anomalous temperature. See Section 2.4.3 for the full heat transport breakdown.

For the NAC region it can be seen that slightly more of the advective heat flux variability arises from anomalous circulation advecting the mean temperature ($v'T_0$, Table 2.2) than terms involving anomalous temperatures (v_0T', $v'T'$). Although the magnitudes are similar between $v'T_0$ and v_0T' components, the relationship with the net ocean heat transport (OHT, *i.e.* vT) is not, with $v'T_0$ having a higher correlation with OHT. Correlations (regression gradients in brackets) between $v'T_0$ and OHT are 0.29, 0.36, and 0.42 (0.82, 2.4, 2.9 $\mathrm{W}_{OHT}/\mathrm{W}_{v'T_0}$) on monthly, annual, and decadal timescales respectively, compared to 0.00, -0.16, and -0.23 (0.01, -1.1, -1.6 $\mathrm{W}_{OHT}/\mathrm{W}_{v_0T'}$) for v_0T' (95% significance levels, assuming a two-tailed t-test and accounting for some missing data, are 0.03, 0.12, and 0.37 respectively). This holds throughout the western half of the southern edge of the NA SPG (not shown), and is associated with a strong background temperature gradient. Thus $v'T_0$ appears to be the dominant advective heat flux in the NAC region on all timescales.

2.4.5 Heat content anomalies in the Irminger Current region

The same breakdown of heat content changes into a particular region was applied to the Irminger Current at the entrance to the Labrador Sea (Figure 2.4a, red box). Similarly to the NAC region, this was chosen where horizontal circulation was well defined in a particular direction and much larger than all orthogonal circulations. The breakdown of heat fluxes into surface, advective, and advective subcomponents is shown in Table 2.1, right column. Similarly to the Gulf Stream region, the net surface and net advective heat fluxes approximately balance but do not fully explain the directly calculated heat content change. However, as before, the correlation between the sum of the surface and advective components and the flux implied by the actual heat content change is very good on all timescales and so we again treat the budget as sufficiently closed.

For the individual fluxes, on annual timescales, the correlation (regression gradient) between the advective heat fluxes and net heat content changes is 0.56 (0.56 $\mathrm{W}_{\mathrm{d}OHC}/\mathrm{W}_{adv}$), again marginally greater than the correlation between surface heat fluxes and net heat content changes at 0.47 (0.52 $\mathrm{W}_{\mathrm{d}OHC}/\mathrm{W}_{surf}$). On decadal timescales these drop to 0.21 (0.08) and 0.19 (0.09) for advective and surface fluxes respectively. Despite these low decadal correlations, there is still a very large correlation between their sum and the actual net heat content change (Table 2.1), suggesting that on these decadal timescales no single component of the heat budget can be considered the controlling influence. This is also indicated by the strong anti-correlation between advective and surface heat fluxes of -0.87 on decadal timescales.

In contrast to the Gulf Stream region, for the Irminger Current the most important advective heat flux is that due to the mean circulation advecting anomalous temperature (v_0T', Table 2.2). v_0T' has slightly greater variability on all timescales than $v'T_0$ and also shows larger correlations (and regression gradients) with the actual OHT changes on all timescales. Correlations between OHT and v_0T' for monthly, annual, and decadal variability are 0.83, 0.34, and 0.29 (0.83, 0.66, 0.53 $\mathrm{W}_{OHT}/\mathrm{W}_{v_0T'}$) respectively, whereas correlations between OHT and $v'T_0$ are much smaller at 0.19, -0.1, and -0.14 (0.10, -0.18, -0.24 $\mathrm{W}_{OHT}/\mathrm{W}_{v'T_0}$; for monthly, annual, and decadal data, the 95% significance levels, assuming a two-tailed t-test and accounting for some missing data, are 0.03, 0.12,

and 0.37 respectively). In our Irminger Current box the zonal currents are an order of magnitude larger than in all other directions, and so we suggest that it is the zonal mean circulation which is playing an important role in moving heat content anomalies from east to west on the northern edge of the NA SPG.

In summary, the heat budget for the NA SPG as a whole has been diagnosed and it has been seen that advective heat fluxes play an important role on decadal timescales, but that the relative contributions of circulation and temperature anomalies to the OHT are region specific. We now investigate the remaining question of what controls the negative feedback between Labrador Sea and NAC temperature anomalies.

2.4.6 Negative feedback between Labrador Sea and Gulf Stream

The anomalous temperatures in the Labrador Sea, which are related to the increased heat flux into the region, affect deep water formation in this region. As noted in Section 1.3.3, an assessment of related studies suggests an approximately even split between temperature and salinity control of the Labrador Sea density changes related to increased deep water formation on decadal timescales. Following *Delworth et al.* (1993) we decompose the simulated density changes in the Labrador Sea into those due to temperature and those due to salinity (Figure 2.10a). This analysis suggests that in HadGEM3 simulated density changes in the Labrador Sea are due to temperature induced density changes (annual correlation with actual density: 0.64), rather than salinity induced density changes (annual correlation with actual density: 0.06). A lagged correlation analysis confirms that on both annual and decadal timescales density changes are temperature-controlled (Figure 2.10b). In the wider context of the proposed mechanism of decadal variability these temperature anomalies are related to OHT anomalies through the Irminger Current region.

In the Labrador Sea, on these decadal timescales, the temperature/density signal is manifest throughout the whole water column. The deeper signal, between 1-1.5km depth, spreads southwards along the particular deep water pathways of the model. Model simulated deep water pathways show much inter-model diversity and are notoriously difficult

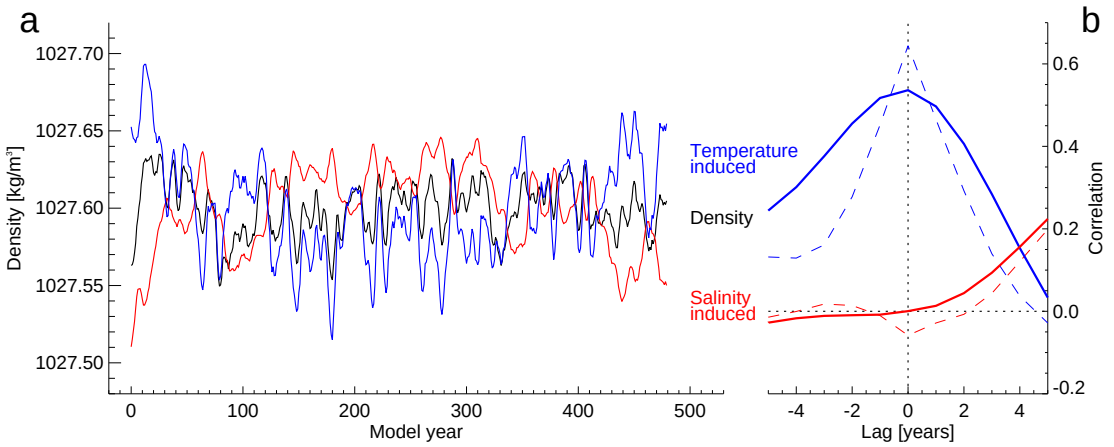

Figure 2.10: a) 5-year smoothed Labrador Sea (50–60°W, 55–62°N) top 500m mean density (calculated using annual mean temperature and salinity, black) and contributions from temperature (by keeping salinity at the time mean in the density equation of state, blue) and from salinity (by keeping temperature at the time mean in the density equation of state, red). Temperature and salinity both linearly detrended prior to computing density. b) Lagged correlation of temperature-induced (blue) and salinity-induced (red) density against actual density for detrended data, either unsmoothed (dashed) or smoothed with a 5-year running mean (solid, as in (a)). Temperature/salinity-induced density leads at negative lags.

to validate (*Spence et al.*, 2011). In HadGEM3 these pathways involve a combination of the fast deep western boundary current and slower interior pathways (Figure 2.11). Young water is formed in the Labrador Sea (Figure 2.11, a) which spreads eastwards and southwards along the deep western boundary current (Figure 2.11, b and c). The core of relatively young deep water can be seen transiting underneath the older northward flowing water (Figure 2.11, d). These deep advective pathways are relatively slow compared to the timescale of the simulated variability within the NA SPG and are more likely to modulate longer timescale variability in the North Atlantic extending outside of the SPG (*e.g. Jungclaus et al.*, 2005; *Menary et al.*, 2012). Instead, we hypothesise that simulated dense water formation in the Labrador Sea in HadGEM3 contributes to circulation anomalies in the NAC region via the creation of an anomalous north-south density gradient, and as such acts as a negative feedback on to NA SPG temperatures.

To examine this hypothesis we calculate a composite difference in the density in a cross section through the NAC which lags the density upstream in the Labrador Sea (Figure 2.12a). To the north the connection between surface and deep water is revealed with the signal sinking below the surface as it progresses southwards. The north-south density

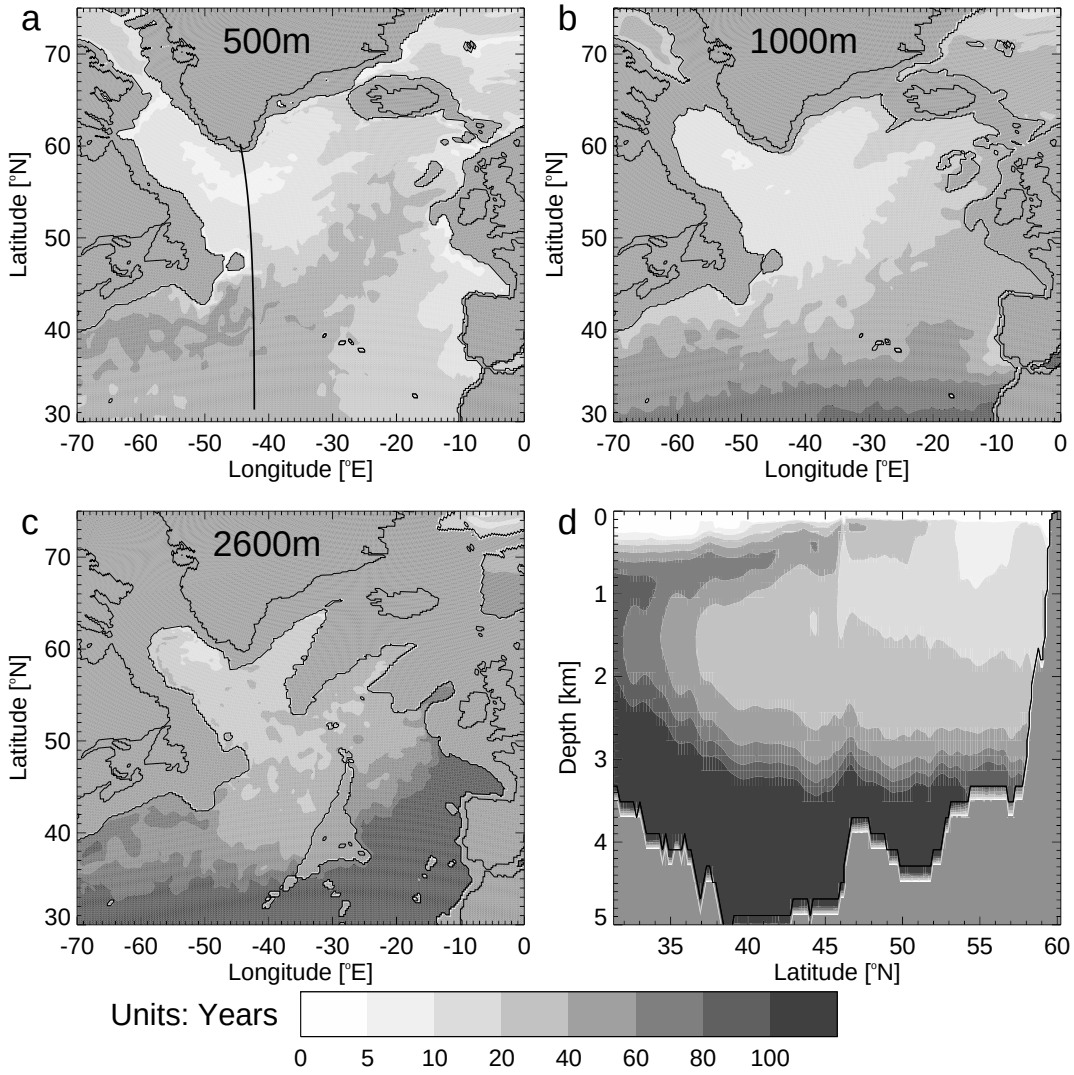

Figure 2.11: NA SPG 'age of water' tracer simulated in HadGEM3 where age of water is defined as the time since the water was last in a surface grid cell. It is implemented by resetting all surface grid cell ages to zero at the beginning of each model time step and incrementing the age of non-surface cells by the length of the time step (in this case 20 minutes). Subsequently the age tracer evolves as any other passive tracer. Thus the tracer represents the grid cell mean age of the water. The tracer was applied at model year 268 and had been running for 213 years at the time this figure was created. Panels a, b, c: The age of water at depths of 500m, 1000m, and 2600m respectively. Panel d: The age of water in a cross section at approximately 42°E (due to the curved grid lines), as marked by the line in (a).

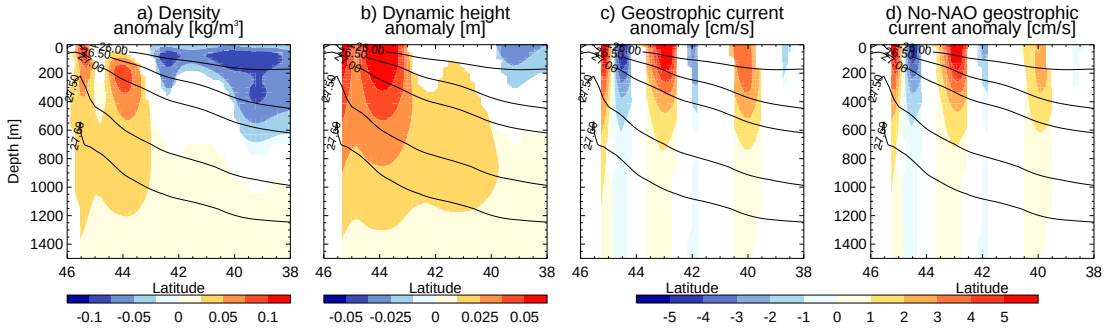

Figure 2.12: Transect south from the Grand Banks through the North Atlantic Current at 47.5°W, as shown in Figure 2.4a. Density (referenced to 0m) composite of high minus low densities in the Labrador Sea, computed by averaging all cases where Labrador Sea volume mean density (computed over the region 47–55°W, 56–61°, 0–1000m) was at least one standard deviation larger than the time mean, and subtracting the average of all cases where density was at least one standard deviation less than the time mean. Cross sections lag by 1 year the Labrador Sea index. b) As (a) but for dynamic height composites (relative to 1500m). c) As (a) but for the geostrophic circulation (relative to 1500m). d) As (c) but first removing the NAO signal from the density field after *Polo et al.* (2014), see text. Cross-sectional time mean density is indicated by the black contours. All data have been detrended and 5-year smoothed. Data insignificant at the 99% level for a two-tailed t-test has been masked.

gradient is associated with a change in the local dynamic height (Figure 2.12b). Despite the negative density anomaly in the south it can be seen that a large part of the dynamic height anomaly is controlled by the northern, positive density anomaly.

As the signal of anomalous density spills out of the Labrador Sea this dynamic height gradient increases and is balanced by anomalous shear in the geostrophic velocities (Figure 2.12c). The mean geostrophic velocity anomaly between the surface and 500m for the pictured transect is 0.9cm/s, increasing to 1.2cm/s for the top 200m only. Thus, an increase in density in the Labrador Sea, associated with a cooling in this region, is followed by a strengthening of the circulation in the NAC, and thus an increase in northward OHT into the NA SPG (with likely also some additional contribution from $v'T'$ as the anomalous circulation acts on anomalously warm, low density surface water, *cf.* Figure 2.12a). This acts as a negative feedback on the NA SPG temperatures. We now discuss the atmospheric contribution to these ocean feedbacks.

2.4.7 The role of the atmosphere

Although the proposed mechanism of decadal (17 year) variability in HadGEM3 has been described mostly in terms of ocean dynamics there are regions where the atmosphere directly forces, or acts as a positive feedback on, the ocean variability.

For example, the negative feedback dipole between Labrador Sea and NAC temperatures is reminiscent of the Ekman response to NAO forcing. To quantify the instantaneous (*i.e.* zero lag) impact of the NAO we attempt to isolate its signal similarly to the analysis of *Polo et al.* (2014). Specifically, the annual mean 3-dimensional ocean density field was regressed onto the wintertime NAO index (both unfiltered, not shown). The direct impact of the NAO was then removed from the density field by scaling the regression pattern by the NAO index and removing the pattern from the density at each time point before re-calculating the composites. Removing the instantaneous NAO-related signal weakens the density/dynamic height and thus geostrophic current response calculated in Section 2.4.6 (Figure 2.12d), hence suggesting that some of the proposed negative feedback in the ocean is forced by the atmosphere and not merely an ocean-only process. On annual timescales the magnitude of the current response, as calculated in Section 2.4.6 and depicted in Figure 2.12c, is reduced by 45% but on longer, decadal timescales the reduction is less stark (13% reduction). This analysis assumes that the instantaneous impact of the NAO is annually independent and can be linearly separated. To what extent the NAO and ocean temperatures/densities can be seen as one-way forcing from atmosphere to ocean, and to what extent it is actually a coupled feedback (*i.e.* some of the NAO signal is itself forced by the ocean, implied by the spectral peak in the NAO power spectrum Figure 2.4c), is discussed below. However, the reduction in anomalous circulation response when removing the NAO suggests that atmospheric forcing/the NAO may act to reinforce this ocean feedback.

In the northern NA SPG we have previously shown a role for ocean advection in moving heat content anomalies westwards via the mean circulation (Section 2.4.5). At the same time, surface heat fluxes were also shown to be non-negligible. In Figure 2.7 the SST, T500, SHF, Sea Surface Salinity (SSS), Mean Sea Level Pressure (MSLP), and Sea ice are regressed at various lags against NA SPG mean SSTs. The SHF is directed into

the ocean and at lag=0,+2 is having a cooling effect in the eastern SPG and a warming effect in the western SPG, i.e. it is effectively moving heat content anomalies from east to west. This is likely related to the concomitant strongly negative NAO anomaly in the MSLP field at the same lags. The actual magnitude of the SHF contribution to the Irminger Current OHC change is similar to the contribution from advective fluxes but, as noted in Section 2.4.5, both are individually quite poorly correlated with the OHC change on multi-annual timescales. This is consistent with a mechanism whereby the ocean integrates up the interannually independent forcing from the atmosphere/NAO resulting in decadal timescale variability in ocean heat content.

In the eastern SPG, the SSTs are anti-correlated with the NAO index, seen both at the lag=0 regression and with the opposite phase at lag=-6. These SSTs are likely a combination of the direct forcing of both 1) the NAO via SHFs and anomalous Ekman and gyre circulation (*Hakkinen and Rhines*, 2004; *Sarafanov et al.*, 2008) and 2) the advective heat flux associated with the diagnosed mechanism of decadal variability. However, the simulated NAO shows a spectral peak at 17 years similarly to ocean indices within the NA SPG. It would appear most likely that this atmospheric memory must come from the ocean but unfortunately long enough atmosphere-only experiments with this model are not available to further test this hypothesis.

The anomalous NAO-related SHFs show the same sign change over both the Labrador Sea and Gulf Stream/NAC but over the Gulf Stream/NAC are of the wrong sign to explain the heat content changes (both at the surface and throughout at least the top 500m of the water column). This is consistent with advective heat fluxes playing a much more dominant role in the heat budget of the NAC region (see Section 2.4.4) than the Irminger Current/Labrador Sea region (Section 2.4.5). However, as noted at the beginning of this section, in the NAC region there is a significant portion of the ocean geostrophic circulation (and associated heat transport) response which is itself related to the NAO (*cf.* Figures 2.12c and 2.12d). In short, it is impossible to completely separate the effects of either the atmosphere or ocean without further experiments.

SSS evolves similarly to SST in the NA SPG although the largest changes are associated with movement of the ice edge in the GIN Seas (Figure 2.7). In general in the NA SPG,

positive salinity anomalies co-vary with positive temperature anomalies in both space and time (mostly density compensating), again suggesting a role for advective fluxes. NAO-related surface freshwater fluxes are also proposed to be of only secondary importance due to the fact that simulated SSS anomaly magnitudes are independent of the amplitude of the NAO.

Similarly to other large scale variables within the NA SPG, ice edge changes exhibit decadal variability with a spectral peak at a period of 17 years (not shown). However, unlike in similar work with the IPSL model (*Escudier et al.*, 2013) these changes do not appear to lead variability in either the East Greenland Current or deep water formation in the Labrador Sea; we suggest that in our simulations ice edge changes are primarily a passive response to the temperature dominated decadal variability within the NA SPG, perhaps again via the NAO (*Deser et al.*, 2000), rather than a direct driver of this variability.

2.4.8 Summary of the proposed mechanism

The mechanism of decadal (17 year) variability simulated in the NA SPG T500 and SSTs is summarised in Figure 2.13. Positive circulation anomalies in the southern part of the SPG move heat eastwards and northwards into the eastern SPG with a timescale of around 5 years (orange). These heat content anomalies are then transported by the mean circulation around the northern edge of the SPG with a timescale of around 2 years (red). In the Labrador Sea these anomalies affect the stability of the water column. These negative density anomalies, associated with reduced deep water formation, spill out from the Labrador Sea into the SPG, deepening as they go. In the region north of the Gulf Stream these negative density anomalies affect the north-south density gradient and induce geostrophic circulation anomalies weakening the NAC. The weaker circulation reduces ocean heat transport and acts to cool the NA SPG (blue). The phase of the oscillation is thus reversed.

The postulated role of the atmosphere is also noted (black dashed lines in Figure 2.13): As temperature anomalies build up in the eastern SPG the atmosphere acts to strengthen

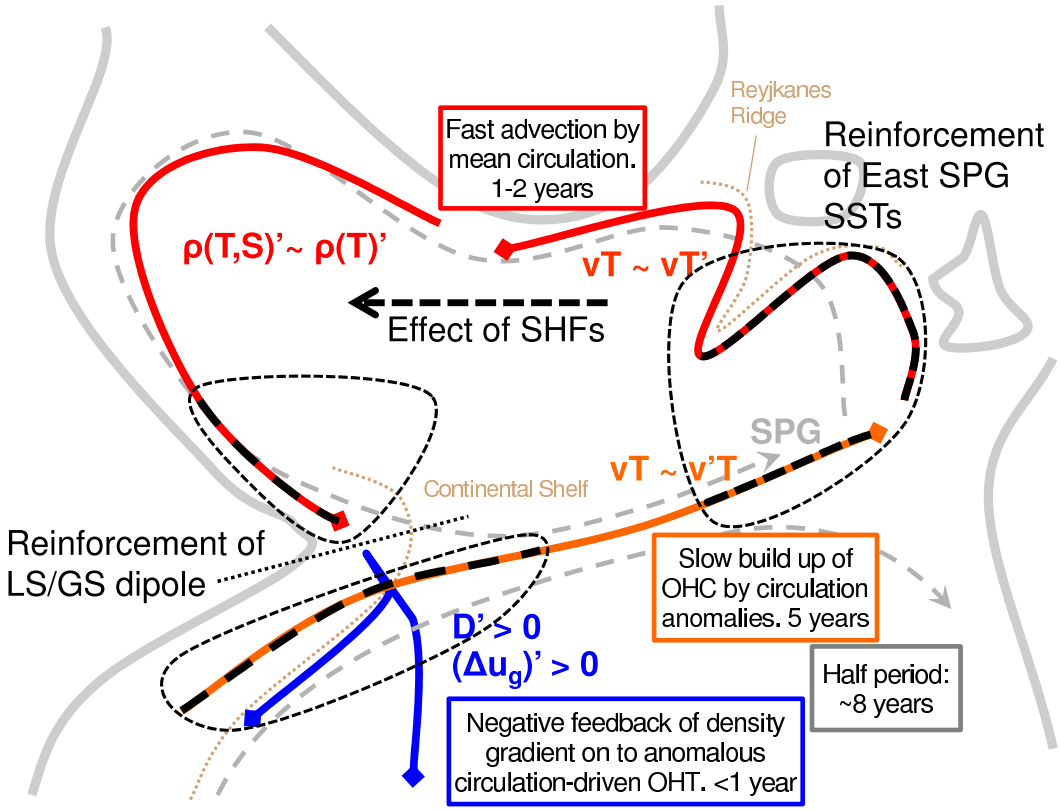

Figure 2.13: A schematic of the proposed mechanism. The various processes in different regions, the timescales, and the postulated role of the atmosphere are as described in the text. Dashed grey lines denote the approximate location of NA SPG and subtropical gyres with bathymetry of particular interest marked brown. Regions dominated by circulation anomalies (orange) and temperature anomalies (red), and where the negative feedback is suggested to occur (blue) are also marked. Additionally, black dashed lines denote regions where the atmosphere is postulated to play a role in forcing or feeding back on ocean anomalies.

these anomalies. When the east of the NA SPG is anomalously warm or cold SHFs also act to move the ocean heat content anomaly westwards. Lastly, in the region of the Labrador Sea/Gulf Stream temperature (density) dipole the NAO is associated with around 13% of the ocean-circulation feedback (*cf.* Figure 2.12d).

We now discuss the implications of our work and similarities between it and previous studies.

2.5 Discussion

In the context of the literature summary in Section 1.3.3, and the schematic illustration presented in Figure 1.5, our simulations broadly fall into a temperature-dominated regime in the Labrador Sea in which the mechanism could be described as 'Ocean*' *i.e.* A positive feedback between the NAO and SSTs may be amplifying the mode. The timescale is set in part by mean circulation speeds in the northern SPG but with a transition to anomalous circulation in the southern SPG — although it is not clear from the simulations precisely where this transition occurs.

The simulated timescales between changes in the Labrador Sea, NAC and eastern SPG have been attributed to advective processes. However, confounding this are wave processes which are also weakly detectable within the model. Analysis of the deep density field (1500–3000m) reveals signals characteristic of boundary waves propagating from the Labrador Sea to the equator; propagating along the equator to the eastern boundary; subsequently propagating north and south along the eastern boundary, all the while radiating Rossby waves westwards (Figure 2.14). The evolution is very similar to that found in the idealised model of *Johnson and Marshall* (2002) and yield a lag between the Labrador Sea and eastern SPG of 5 years, broadly similar to that due to the proposed advective feedback. Although detectable, these wave signals require heavy filtering of the deep density field whilst the proposed mechanism exists mainly in the top 1km (Figure 2.9). Additionally, recent work has shown that, while detectable and associated with decadal variability, Rossby waves may not necessarily be a significant driver of this variability (*MacMartin et al.*, 2013). We can only conclude that wave processes may play an additional role in our simulated variability but the magnitude of this is unclear. We also note that the relatively diffuse thermocline in HadGEM3 (*Megann et al.*, 2014) may act to dampen these wave processes (*Grotzner et al.*, 1998) as compared to the updated seasonal forecast model, GloSea5 (which will be similar to the new Met Office decadal prediction model).

Despite the lagged regression analysis used in this study, and its ubiquity within studies of decadal variability within climate models, there are some hints from the present work

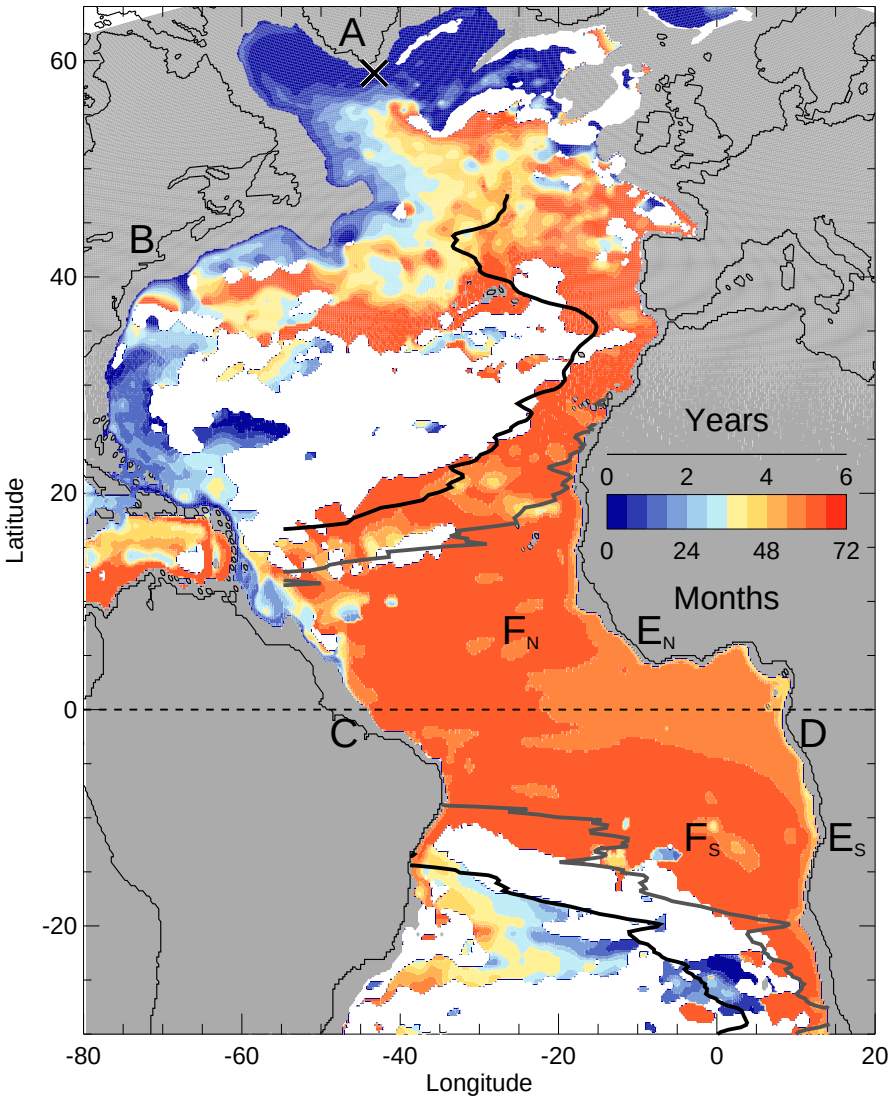

Figure 2.14: A map of the lag in months of the maximum correlation with density at the marked location in the Irminger Current. Densities are the depth average between 1500–3000m after *Hodson and Sutton* (2012). Data have been bandpass filtered to remove periods less than 3 years and greater than 40 years. Wave signals propagate around the boundary of the Labrador Sea, beginning at **A**, and down the west coast of the North Atlantic via **B** to the equator at **C**, with an arrival time at the equator of 15 months. From the west coast of the equator the signal is next seen on the east coast at **D**, arriving there after 43 months implying a pan-equatorial journey time of 28 months. Subsequently, the signal propagates poleward in the direction of $\mathbf{E}_N$ and $\mathbf{E}_S$, all the while radiating Rossby Waves into the ocean interior ($\mathbf{F}_N$ and $\mathbf{F}_S$). Grey and black contours mark the distance from the eastern boundary that a first-mode baroclinic Rossby Wave would be estimated to travel after 5 and 6 years respectively, given the simulated thermocline depth (far shallower than 3000m; these lines mark the general extent of Rossby Wave penetration) and the signal's initial arrival time on the eastern boundary of the Atlantic at that latitude.

that the proposed mechanism may be asymmetric. This asymmetry is manifest in the timescales of various phases of the cycle being also dependant on the sign of the anomaly; *i.e.* the same processes are at work in opposite phases of the mechanism but may evolve with different timescales. Some evidence for this can be seen in Figure 2.7 in which all the fields reverse sign over 6–8 years, implying a periodicity of 12–16 years, and yet the spectral peak occurs at the upper end of this at 16 to 17 years. If we construct lagged composites of the T500 (or SST) field based on the top/bottom 10% of phases of the SST index we find a reversal timescale of 9 years following a high SST phase, but a reversal time of 7 years following a low SST phase (not shown). This asymmetrical timescale doesn't appear to be directly due to the effect of heat transport by the anomalous circulation in the southern SPG (the $v'T_0$, see Section 2.4.4) as the lags between the NAC and eastern SPG result in the same timescale in high and low phases. It is important to note though that constructing composites, which only use 20% of the total data, reduces the effective number of degrees of freedom.

Additionally, atmosphere-only sensitivity experiments (see Chapter 4) suggest a stronger coupling in the NA SPG between anomalously positive NAO/negative SSTs than anomalously negative NAO/positive SSTs. This atmospheric asymmetry also appears evident in the coupled simulation when compositing MSLP based on high/low phases of the SST index. Although both MSLP patterns, composited against positive SSTs (Figure 2.15a) and negative SSTs (Figure 2.15b), show a pattern broadly similar to the NAO, the magnitude and precise structure are clearly different, with a stronger NAO signal associated with the negative SST composite (Figure 2.15b). This asymmetry is examined further in Chapter 4.

It is difficult to prove the mode of variability reported here is inconsistent with observational data due to the paucity of observational records in the NA SPG, particularly in the northern half, and the presence of confounding additional transient forcings in the observational record. However, palaeo proxies from the NA SPG suggest there is 20 year variability in some indices in the region (*Sicre et al.*, 2008; *Chylek et al.*, 2012), although it must be noted that there is disagreement on the spectral characteristics of all proxies (*Mann et al.*, 1995). The specific elements of our proposed mechanism (anomalous circulation OHT in the southern part of the NA SPG, mean circulation OHT in the northern

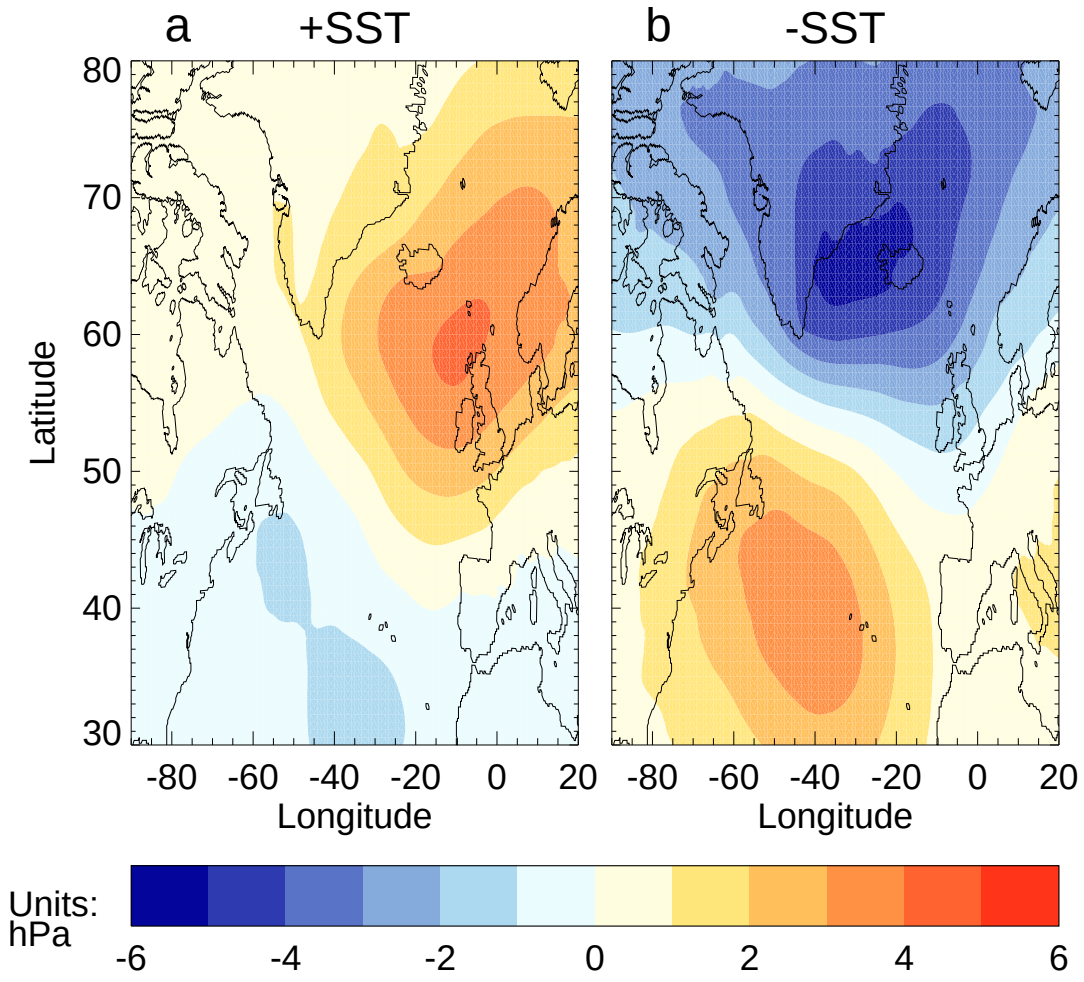

Figure 2.15: Mean sea level pressure (MSLP) composites created using the same SST index as in Figure 2.7 to highlight the asymmetry between positive and negative phases of the proposed mechanism. a) Composite created using the highest 10% of SST anomalies. b) As (a) but for the lowest 10% of SST anomalies.

part, a negative feedback between Labrador Sea and NAC temperatures) are also broadly consistent with the observational literature. For example, there are some similarities to the anti-correlated relationship between Labrador Sea and NAC temperatures/transports seen in observations (*Curry and McCartney*, 2001). This observational work also highlights the significant role of the NAO in this relationship as well as the dominant role for temperature (as opposed to salinity) in driving these changes. We note that as a result of the northern NA SPG warm bias in HadGEM3 there is less ice in the mean, which may detrimentally affect the ability of ice/freshwater fluxes to affect the decadal variability. In models where the NA SPG mean state bias is cold, feedbacks involving ice and freshwater fluxes have been shown to be crucial to the diagnosed decadal variability (*Es-*

cudier et al., 2013). To what extent the simulated decadal variability is consistent with observational data is discussed further in Chapter 3.

2.5.1 Comparison with other models

Similar to our findings, recent ultra-high resolution ($^1/_{12}^\circ$ horizontal resolution) eddy resolving ocean-only model studies show that much of the OHT into the eastern NA SPG occurs in the near surface (but below the Ekman layer) originating in the subtropics (25% of virtual floats at 500m, compared to less than 10% at 50m or 1000m, *Burkholder and Lozier*, 2011, 2014). In addition, the role of anomalous circulation transporting the mean temperature gradient in the southern part of the NA SPG is indirectly supported by these ocean-only simulations, which find that the mean circulation is unable to explain the slow timescale by which temperature anomalies move from the subtropics to the eastern SPG. Important for decadal variability in our simulations are advective heat fluxes from the southern edge of the NA SPG due to the anomalous circulation ($v'T_0$). The heat flux across the southern boundary correlates well with both the vertical 'AMOC' and horizontal 'gyre' heat transports. However, the standard deviation in the annual mean AMOC at 26.5°N from 10 years of RAPID data is approximately double the annual standard deviation in HadGEM3. Thus, if the proposed mechanism exists in reality then it could be expected to have a larger amplitude or faster timescale. A more critical comparison against observational data is provided in the next chapter (Chapter 3).

The mechanism we have presented has a timescale of 17 years, similar to the 20 years found in the IPSL-CM5A-LR model recently investigated by (*Escudier et al.*, 2013, hereafter E13). However, a similar timescale does not imply the same mechanism: see for example an identical 17 year timescale but different mechanism reported by *Born and Mignot* (2012). The present study reports a mode of variability where temperature dominates the density budget, whereas E13 report a mode in which freshwater/salinity fluxes have an important role. Indeed, salinity advection within the SPG has been proposed as a cause of bistability in the SPG (*Born et al.*, 2013), albeit on longer timescales. It is intuitive that whether the density budget is dominated by temperature or salinity would affect whether a strengthening northward circulation acted as a positive or negative feedback

Mean state	Density change for one s.d. change in temperature	Density change for one s.d. change in salinity
EN4 + HadGEM3 bias	0.027	0.014
EN4 + IPSL bias	0.010	0.014
EN4 (original)	0.023	0.014

Table 2.3: Characteristic magnitudes of density changes (kg/m^3) in different simulated/estimated T/S regimes. As there is limited raw data from EN4 to reliably estimate decadal variability in the Irminger Current, and to simplify the experimental design and interpretation, we use HadGEM3 estimates of the decadal variability in temperature and salinity in all cases. *s.d.* = Standard deviation.

— but why are NA SPG density changes differently controlled in the two models?

One hypothesis is that the nature of the biases (compared to observations) affect the variability as the non-linear equation of state for density becomes increasingly salinity dominated at cooler temperatures. To estimate this effect we compute the density change in the Irminger Current region, mechanistically important in both studies, for a one standard deviation change in temperature and salinity (whilst keeping the other of salinity or temperature at climatological values) in both HadGEM3 and the IPSL-CM5A-LR model as well as an observational estimate from EN4 (Table 2.3). In HadGEM3, such a temperature change has double the impact on density than a change in salinity. This is not the case in the IPSL model where salinity changes are found to be more important. The EN4 data suggest that the real world may be in a temperature dominated regime, similar to HadGEM3. This points to there being some relationship between the NA SPG mean state biases of a given model and the subsequently diagnosed mechanisms of decadal variability. Note that this cursory analysis merely compares mean states and variability, and does not explicitly investigate whether density variability is temperature- or salinity-controlled. Nevertheless, one implication of this would be that decadal prediction studies using anomaly-assimilation methods, in which the mean state biases are implicitly assumed to be independent of the variability, would need to re-evaluate the validity of this assumption (*Robson*, 2010). We will investigate this further in Chapter 6.

2.6 Chapter conclusions

We have analysed a decadal mode of variability in the North Atlantic subpolar gyre (NA SPG) in a 460 year control simulation with a version of the high resolution coupled climate model HadGEM3.

- The mode of variability exists primarily in the top 1km and involves the propagation of heat content anomalies around the NA SPG with a periodicity of around 17 years.

- Simulated decadal variability (between 10 to 30 years) in the NA SPG explains more than 15% of the annual mean variance in top 500m depth averaged temperatures. This rises to $>$30% of the variance within the interior NA SPG and Labrador Sea. Some of the processes/feedbacks we have identified have a low signal to noise ratio, which reinforces the case for continuing to invest in long control simulations even with high resolution, computationally expensive, coupled climate models.

- The simulated NA SPG heat budget is dominated by advective, rather than surface, heat fluxes on decadal timescales, with advection from the subtropics playing the primary role. For the specific regions of interest, namely the Irminger Current and North Atlantic Current (NAC), advective fluxes were also found to dominate. The large depth extent of the mode is also consistent with an important role for advection (*Saravanan and McWilliams*, 1998), as has also been shown for the analogous Kuroshio extension (*Qiu and Kelly*, 1993).

- The role of mean or anomalous circulation in transporting heat content anomalies was found to vary with region: Anomalous circulation dominated the variability in the NAC with mean circulation, and hence temperature anomalies, dominant in the Irminger Current region.

- A negative feedback, required for the mechanism to result in a spectral peak, occurs between the Labrador Sea and NAC. Here, density anomalies spill out of the Labrador Sea resulting in a dynamic height gradient across the NAC/Labrador Sea which induces vertical shear in the geostrophic currents. These current anomalies

result in heat transport anomalies which reverse the cycle. The density changes are temperature, rather than salinity, driven.

- Variability in the NAO directly contributes to various stages of the mechanism as well as showing signs of responding to ocean variability. Removing the North Atlantic Oscillation (NAO) signal from the negative feedback between Labrador Sea and NAC temperatures/densities (see Section 2.4.7) shows about 45% of the geostrophic current speed feedback is related to the NAO on annual timescales but that on decadal timescales the ocean feedback still dominates. The atmosphere also acts to reinforce temperature anomalies in the eastern NA SPG and aid their westward propagation in the northern SPG. The proposed mechanism is summarised in Figure 2.13.

- Whether density changes are temperature or salinity controlled effects where, and how, negative feedbacks can occur. This may also be expected to affect the particular mechanism simulated in the model. This could have important implications for decadal prediction studies that use the method of anomaly-assimilation and prediction, in which the future evolution of the model is assumed to be independent of the mean state — an assumption which we suggest may not be valid.

A modified version of the model presented here will be used as part of the Met Office decadal prediction system and analyses such as we have presented will be important in developing and evaluating such systems. Given the relationship between resolution and the improved realisation of particular processes, as well as mean state biases, further high resolution coupled model studies would be valuable in testing whether these results are model-specific.

How robust various elements of this mechanism of decadal variability are will be discussed in Chapters 4 and 5. However, in the next chapter we first critically compare the mechanism to available observations within the NA SPG.

Chapter 3

Confronting the mechanism of simulated decadal variability with real-world observations

3.1 Introduction

In the previous chapter the mechanism of decadal variability within the North Atlantic subpolar gyre (NA SPG) was diagnosed within HadGEM3. We now attempt to test the validity of this mechanism against available observations of the NA SPG. Our goal is to do more than merely test the magnitude of the biases in the model mean state, which has already been done elsewhere (*Walters et al.*, 2011). Although this is useful for understanding the base climate of a model, and provides guidance as to the 'plausibility' of the variability within a particular model, it does not speak to the specific mechanisms of that variability (although we note that this does have some utility for understanding mechanisms of variability when undertaken in a multi-model framework, see Chapter 6). As such, in Section 3.2 we begin with a brief review of the periodicity and variability observed in the NA SPG and describe the key datasets that we will use in our analysis. In Section 3.3 we investigate the surface evolution of observed variables in the NA SPG before investigating the depth structure in more detail in Section 3.4. We test key elements

of the mechanism of variability simulated in HadGEM3 in Section 3.5 before discussing our results in Section 3.6. Finally, conclusions are presented in Section 3.7.

3.2 Review of relevant observations

3.2.1 Variability in palaeo and direct observations

In Chapter 1 we discussed some of the direct and indirect observations of the mean state and variability of the North Atlantic region, as well as many of the simulations that aim to reconstruct and understand that variability. Aside from numerical models, much of the information about the periodicity of variability in and around the North Atlantic subpolar gyre (NA SPG) originates from palaeo reconstructions (*Mann et al.*, 1995) with palaeo records only recently able to resolve temporal variability on sub-decadal timescales (*Sicre et al.*, 2008; *Chylek et al.*, 2012). However, despite recent efforts (*Ahmed et al.*, 2013), the spatial patterns of this variability and any depth structure in the ocean are still not estimable from the palaeo archive.

Direct/instrumental observations of temperature in the North Atlantic have been made for many centuries, with Benjamin Franklin reported to have made measurements on voyages between Europe and America as early as 1775 (*Richardson*, 1980). A relatively well sampled record of surface temperatures in the subtropics and southern half of the subpolar gyre exists for the last 160 years (*Kennedy et al.*, 2011). However, observations in the northern edge of the NA SPG are well sampled for perhaps only the last 70 years (*Kennedy et al.*, 2011) and observations of surface salinity or subsurface temperature and salinity (in the top 2000m) in the NA SPG are poorly sampled until at least 1960 (*Good et al.*, 2013). Observations at depths below this are still sparse to this day, as are direct estimates of circulation indices (*Cunningham et al.*, 2007), which is particularly alarming given the recently detected trend in the AMOC (*Smeed et al.*, 2014) and its potential relationship to variability in the NA SPG (*Robson et al.*, 2014). Despite the observational paucity in temperature and salinity in the NA SPG compared to the rest of the Atlantic Ocean, the high signal to noise ratio in this region makes it a prime candidate for detecting

significant decadal (or even longer) variability (*Hakkinen and Rhines*, 2004) — here we approximately define decadal variability as timescales of 10–30 years and multi-decadal variability as timescales >30 years although in the literature there is considerable overlap of these terms.

Large multi-decadal variations in the mean sea surface temperature (SST) in the subtropical and subpolar North Atlantic have been observed over the last century, often referred to as the Atlantic Multidecadal Oscillation (AMO, *Schlesinger and Ramankutty*, 1994). However, it is not clear to what extent these are forced signals (*Booth et al.*, 2012; *Zhang et al.*, 2013) and whether they are simultaneously representative of the same variability in the NA subpolar and subtropical regions (*Knight et al.*, 2005; *Zhang et al.*, 2013; *Häkkinen et al.*, 2015). Within the NA SPG, using records since 1960, signals of multi-decadal variability have been observed in upper North Atlantic deep water (UNADW) whilst lower North Atlantic deep water (LNADW) has exhibited a long term trend (*Mauritzen et al.*, 2012).

On shorter, decadal timescales, multiple analyses have noted the relationship between the strength of the North Atlantic Oscillation (NAO) and a proxy for the strength of the NA SPG (*Curry and McCartney*, 2001; *Rhein et al.*, 2011; *Roessler et al.*, 2015). The NA SPG strength is postulated to have declined in the past decades (*Hakkinen and Rhines*, 2004), possibly related to decadal variability in the strength of the Atlantic meridional overturning circulation (AMOC) with opposing signals in the surface and subsurface temperatures (*Zhang*, 2008). Indeed, there is often an anticorrelation between surface and subsurface temperature anomalies, as well as between NA SPG and subtropical temperature anomalies (*Lozier et al.*, 2008). Model studies suggest these may be internally (rather than externally) forced (*Zhang and Vallis*, 2007) as well as due to advective (rather than surface) processes (*Williams et al.*, 2014).

In summary, there is clear evidence of decadal (and multi-decadal) variability in the NA SPG, particularly in ocean temperatures, and with a depth structure that has implications for the processes involved. In order to quantitatively test the simulated mechanism of decadal variability in HadGEM3 we first source and describe an observational dataset.

3.2.2 The EN4 dataset

The 'EN4' optimally-interpolated objective analyses provide gridded T/S data throughout the global oceans based on quality controlled temperature and salinity profiles (*Good et al.*, 2013). EN4 represents an improvement upon the previous version EN3 (*Ingleby and Huddleston*, 2007) owing to the provision of uncertainty estimates derived via a novel 'observation influence' approach (*Donlon et al.*, 2012). This approach involves using a recursive filter to iteratively estimate the optimal value of the interpolated data, hence 'optimal interpolation'. The number of iterations has been increased from 10 to 50 in this version of EN4 to enable closer convergence.

Data are provided as monthly means on a 1x1° grid with 42 vertical levels. The data are infilled, using statistically determined three-dimensional decorrelation lengthscales and persistence of anomalies in time. The horizontal covariances are modelled using two second-order autoregressive functions, with lengthscales of 300km and 400km. Within four degrees of the equator the first of these lengthscales increases exponentially to 1500km. Vertical covariances are estimated using lengthscales of 100m and 200m. The persistence of anomalies in time is modelled month by month and combines the present monthly climatology with the previous monthly mean anomaly (from climatology), scaled by a factor of 0.9 resulting in an e-folding time of 9.5 months. Further details can be found in *Good et al.* (2013).

The data span the period January 1900 to the present day (although there are no observations within the interior NA SPG prior to August 1908) with a lag of 2 months to real time[1]. Unless otherwise stated, in the subsequent analysis annual mean data is used up to and including the year 2014.

Before analysing the depth structure in EN4 (Section 3.4), it is prudent to check where, in the 1900–2015 record, there are subsurface observations. Figure 3.1a counts the number of subsurface (below 200m) temperature and salinity observations per month within the important Labrador Sea region. Although the infilled record stretches back to 1900, there are few years before 1925 with any subsurface observations in this region at all.

[1]Data freely available from http://www.metoffice.gov.uk/hadobs/en4/ (August 2015)

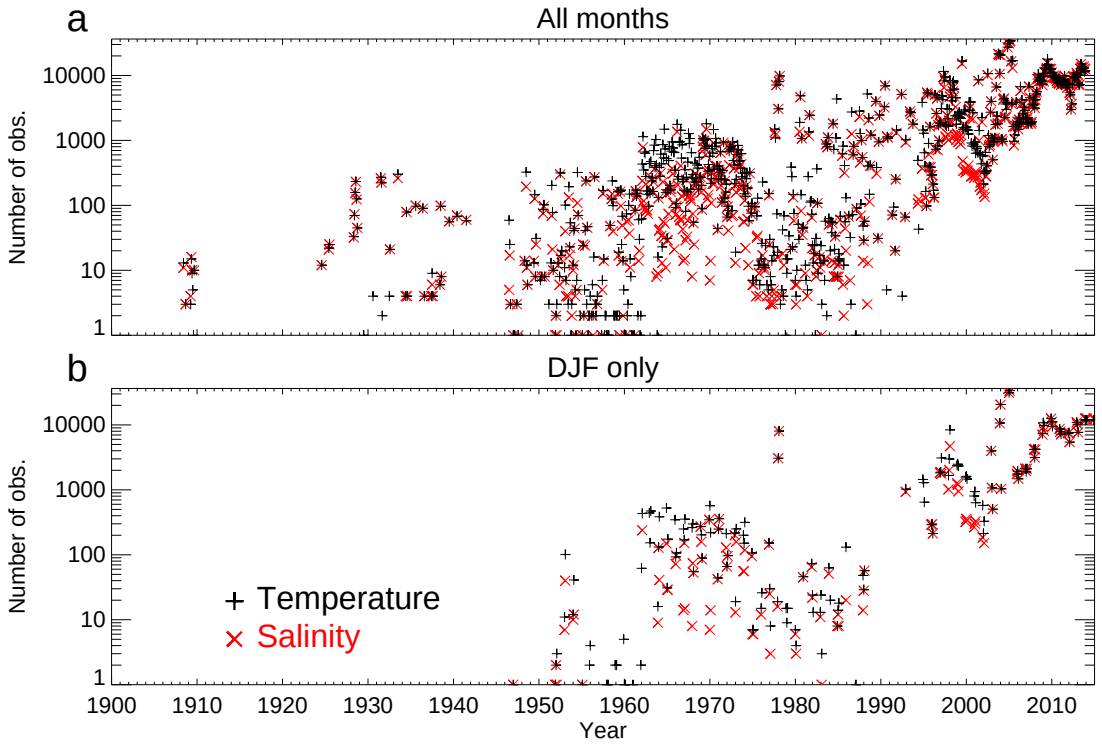

Figure 3.1: a) The number of independent (as estimated from the quality control procedure within EN4, *Good et al.*, 2013) observations per month in EN4 within the Labrador Sea, defined as the region 45–60°W, 55–65°N between 200m and 6000m. b) As (a) but just for the winter months: December, January, February.

Indeed, analysing just the winter months, when deep convection most commonly occurs (*Yashayaev*, 2007) reveals that prior to 1960 there were at best only tens of observations and even some subsequent years with no wintertime observations at all, *e.g.* the early 1990s (Figure 3.1b).

3.2.3 Other datasets

In Section 3.3.1 we initially compare the variability in near surface fields from a variety of observational datasets. For SST, sea surface salinity (SSS), and top 500m depth averaged temperature (T500) we use EN4 (*Good et al.*, 2013). For December-February inclusive (DJF) wintertime mean sea level pressure (MSLP) we use HadSLP2r, a near-real-time update of HadSLP2 (*Allan and Ansell*, 2006) that provides 5x5° gridded fields of land and marine pressure observations for the period 1850–2004 (updated to 2014 in the present analysis with use of HadSLP2r). For sea-ice fraction we use HadISST (*Rayner et al.*,

2003) that provides 1x1° gridded fields of monthly mean sea ice concentrations for the period 1870 to present day with a lag of 1 month to real time. All fields are converted to annual means (except wintertime MSLP) for the common period 1900–2014.

3.3 Comparing the surface evolution of NA SPG variability in the model and observations

3.3.1 The spatio-temporal evolution of surface fields

In this section we attempt to recreate a key figure from our HadGEM3 analysis (Chapter 2, Figure 2.7) using observational data and an index of NA SPG variability. In Chapter 2 we used a basinwide SST index for the model simulations, which comprise more than 400 years of data without annually varying external forcings. For the observations, we use a depth averaged index over the top 500m (Figure 3.2). This is because the observed SST record is made up a series of large amplitude multi-decadal changes that are potentially externally forced (*Booth et al.*, 2012; *Menary and Scaife*, 2014; *Swingedouw et al.*, 2015) and may be subject to more global influence than T500. Testing with the model (not shown) suggests this has little effect on the lagged relationships seen in the simulations.

The time series of NA SPG T500 is shown in Figure 3.2 and highlights the rapid warming in the 1990s, discussed in detail by *Robson et al.* (2012) and also later in this section. Removing the regression against global mean SST since 1900 does not have a dramatic effect on the NA SPG T500 index (Figure 3.2, red) but does reduce the linear trend in the period since 1960 (see the orange line during this period). Alternatively, removing the regression against global mean SST from 1960–2000 has the opposite effect on the linear trend in NA SPG T500 during this time (Figure 3.2, blue). The sensitivity of the subsequent regression maps to these choices will be discussed.

Figure 3.3 shows the regression slopes between an index of T500 between 45–65°N in the North Atlantic, estimated using EN4 data, and various other fields at a variety of lags. At a lag of zero years (fourth row) the T500 index is consistent with anomalously

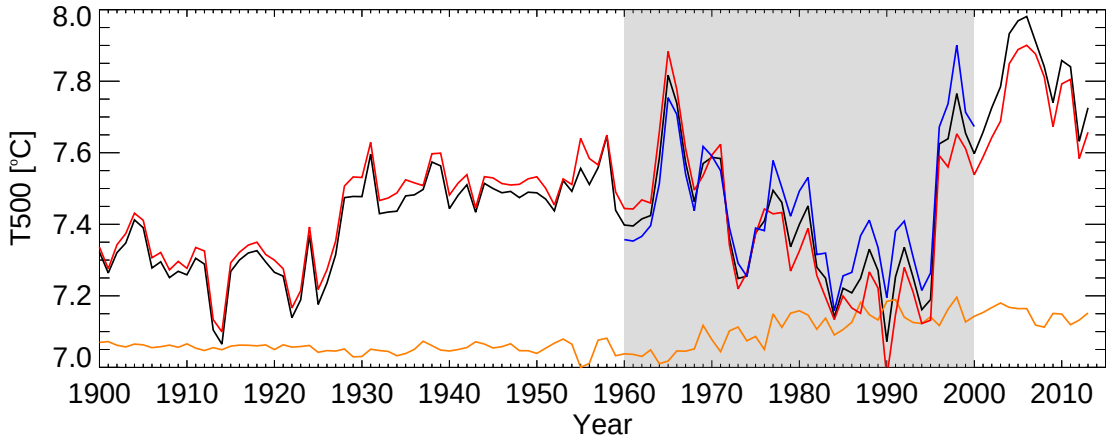

Figure 3.2: Time series of North Atlantic subpolar (45–65°N) top 500m depth averaged temperature (T500) from EN4 (black). An estimate of the component associated with global mean climate change is made by regressing T500 against the global mean SST (from EN4) and this component (orange, displayed with arbitrary mean) is subsequently removed from the T500 index as well as all the fields used in Figures 3.3 and 3.4. The final T500 index (red) is subsequently regressed against a variety of similarly processed spatial fields in Figure 3.3. The full period 1900–2014 is used in Figure 3.3 and the shorter period 1960–2000 (blue, highlighted with grey shading) is used in Figure 3.4, with the climate change signal recomputed and removed for this shorter period.

warm conditions at the surface and throughout at least the top 500m of the water column (Figure 3.3, first two columns). Concomitant with this warm anomaly is a salinity anomaly (third column) that does not appear to be due to ice formation (brine rejection) as ice concentrations are also anomalously negative at this time (fifth column). These warm/saline anomalies are consistent with a joint (advective) origin, though also consistent with merely temperature-induced evaporation or reduced mixing with fresher, subsurface water. The warm conditions are also consistent with the atmospheric pressure pattern (*Flatau et al.*, 2003), which resembles a negative North Atlantic Oscillation (NAO) pattern and implies a reduction in wind-driven cooling.

In the years leading up to this maximal warm anomaly in the NA SPG (Figure 3.3, first three rows) surface warming begins in the Nordic Seas before becoming focussed most heavily on the Labrador Sea (Figure 3.3, first column) whilst the T500 warming appears to grow *in situ* in the NA SPG (Figure 3.3, second column). Similar to the SST anomalies, the SSS anomalies are largest in the Nordic Seas and Labrador Sea, although unlike the SST anomalies they remain larger in the Nordic Seas than Labrador Sea. Throughout this time, the atmosphere inverts from an initially neutral NAO phase (Figure 3.3, fourth

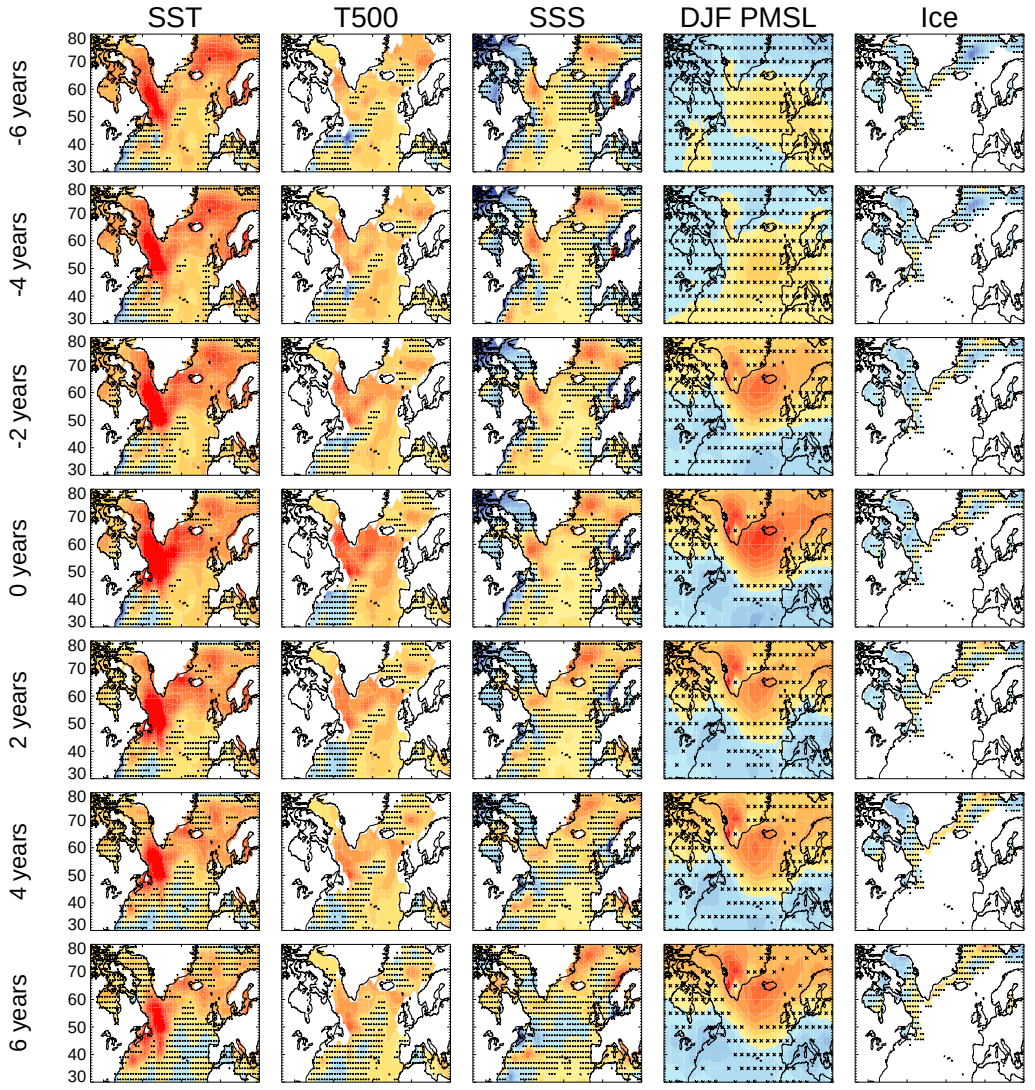

SST [°C/°C]	-2	-1.6	-1.2	-0.8	-0.4	0	0.4	0.8	1.2	1.6	2
T500 [°C/°C]	-2	-1.6	-1.2	-0.8	-0.4	0	0.4	0.8	1.2	1.6	2
SSS [PSU/°C]	-1	-0.8	-0.6	-0.4	-0.2	0	0.2	0.4	0.6	0.8	1
DJF PMSL [hPa/°C]	-10	-8	-6	-4	-2	0	2	4	6	8	10
Ice [frac/°C]	-0.5	-0.4	-0.3	-0.2	-0.1	0	0.1	0.2	0.3	0.4	0.5

Figure 3.3: Regressions over the period 1900–2014 between North Atlantic subpolar (45–65°N) top 500m depth averaged temperature (T500, from EN4) and, from left to right: sea surface temperature (SST, from EN4), T500, sea surface salinity (SSS, from EN4), wintertime mean sea level pressure (DJF MSLP, from HadSLP2r), and ice fraction (from HadISST). From top to bottom, the T500 index lags then leads the fields from -6 to +6 years. To attempt to remove the climate change signal, all fields are first detrended by removing the local regression against a global mean SST index. Stippling denotes regions insignificant at the 90% level using a two-tailed t-test for correlations between random variables with the same mean, standard deviation, and lag=1 autocorrelation as the real data *i.e.* testing the null hypothesis that both the T500 index and each of the fields merely exhibit damped persistence.

column, first two rows) to its lag=0 negative NAO phase although both the MSLP and ice cover at these lags are not significantly correlated with the T500 index.

Figure 3.3 has been designed to be comparable to the same figure created using more than 400 years of simulated data with HadGEM3 (Figure 2.7) in Chapter 2 with some exceptions: Firstly, in the present figure, the index used is depth averaged temperature (T500) as opposed to SST, for the reasons given at the beginning of this section but we note that creating the simulated version of the figure using T500 data does not appreciably affect the patterns for HadGEM3 (not shown). Secondly, surface heat flux observations were far too sparse in time or space to be meaningfully included in our observational comparison. We avoid using reanalysis products that use dynamical ocean/atmosphere models in order to avoid potentially contaminating the observed signal with the model dynamics that we are trying to test. Finally, note that in all cases the scale used for the observed regression slopes in Figure 3.3 is double that for the simulated variables in Figure 2.7, which is related to the use of a T500 index rather than SST index (consistent with the scaling between T500 and SST indices in HadGEM3, not shown).

There are some key similarities and differences between the two sets of figures. The NAO signal reverses from broadly positive to negative in both simulations and observations between lag=-6 and lag=0. Similarly, ice extents in the Labrador Sea also reverse from anomalously positive to negative during this time. However, surface and near surface temperature anomalies, which appear to propagate into the NA SPG and reverse the sign of the anomaly in HadGEM3 (Figure 2.7, first two columns) arise *in situ*, without appearing to propagate, in the observations (EN4) and do not switch sign between lag=-6 and lag=0. Similarly, following the maximum at lag=0, the observed anomalies merely appear to gradually decay (Figure 3.3, first two columns, fifth and sixth rows) whereas in HadGEM3 opposing anomalies can once again be seen propagating into the region (Figure 2.7, first two columns, fifth and sixth rows). Finally, the pattern of all anomalies at lag=-6 and lag=+6 has gone full circle and is becoming similar in HadGEM3 (Figure 2.7, first and last rows), indicative of the decadal periodicity, whereas this is not the case in EN4 (Figure 3.3, first and last rows).

Despite removing the global climate change signal from the observational indices (by

removing the local regression against the global mean SST) the lagged relationships between T500 and the other indices do show some sensitivity to the precise period used. To attempt to compare our analyses of the available observations with previous work (*Robson et al.*, 2012) we recreate Figure 3.3 for the shorter period 1960–2000 (Figure 3.4). As noted in *Robson et al.* (2012), the direct correlation between T500 and the NAO appears to break down after the year 2000 and, as previously noted here and in *Robson et al.* (2012), observations of North Atlantic ocean variables are somewhat sparse prior to 1960 — indeed many reanalysis products do not extend much further back in time than this (*e.g. Uppala et al.*, 2005). Similar to the full period, at lag=0 the NA SPG is warm and under negative NAO conditions (Figure 3.4, fourth row). However, unlike for the full period, this state appears to have arisen from a cool NA SPG under positive NAO conditions at lags of -4 to -6 years (Figure 3.4, top two rows), with some signs that the warm anomalies propagate into the NA SPG from the south.

The phase reversal in Figure 3.4 (top to bottom) is consistent with the analysis of *Robson et al.* (2012) who found a shift from a positive NAO/cool NA SPG in the decade 1986–1995 to a negative NAO/warm NA SPG in the decade 1996–2005. In addition, the anomaly propagation from the southern edge of the NA SPG (Figure 3.4, first two columns) is again consistent with the analysis of *Robson et al.* (2012) who found a possible role for the circulation (AMOC) in producing these warm anomalies. Subsampling the observations this way also gives better agreement with the simulations: both the lag=-6 to lag=0 phase reversal and the possible role for heat content anomaly propagation can be seen in the simulated case (Figure 2.7). As such, it is not clear whether the apparent disparity between simulations and observations — covering the full period 1900–2014 — is related to the effects of severe data paucity (which relax to climatology where there is missing data) or whether this reflects a more subtle distinction between ‘average’ decadal variability and specific decadal events in the late twentieth century.

Additionally, regressing the observed MSLP against an SST index (rather than T500 index, not shown) over the period 1960–2000 yields instead a positive NAO associated with a warm NA SPG with the negative NAO shifted to a lag of -2 years (*i.e.* preceding the positive NAO). This is again consistent with the analysis of *Robson et al.* (2012) in which an extended period of an anomalously positive phase of the NAO, followed by a

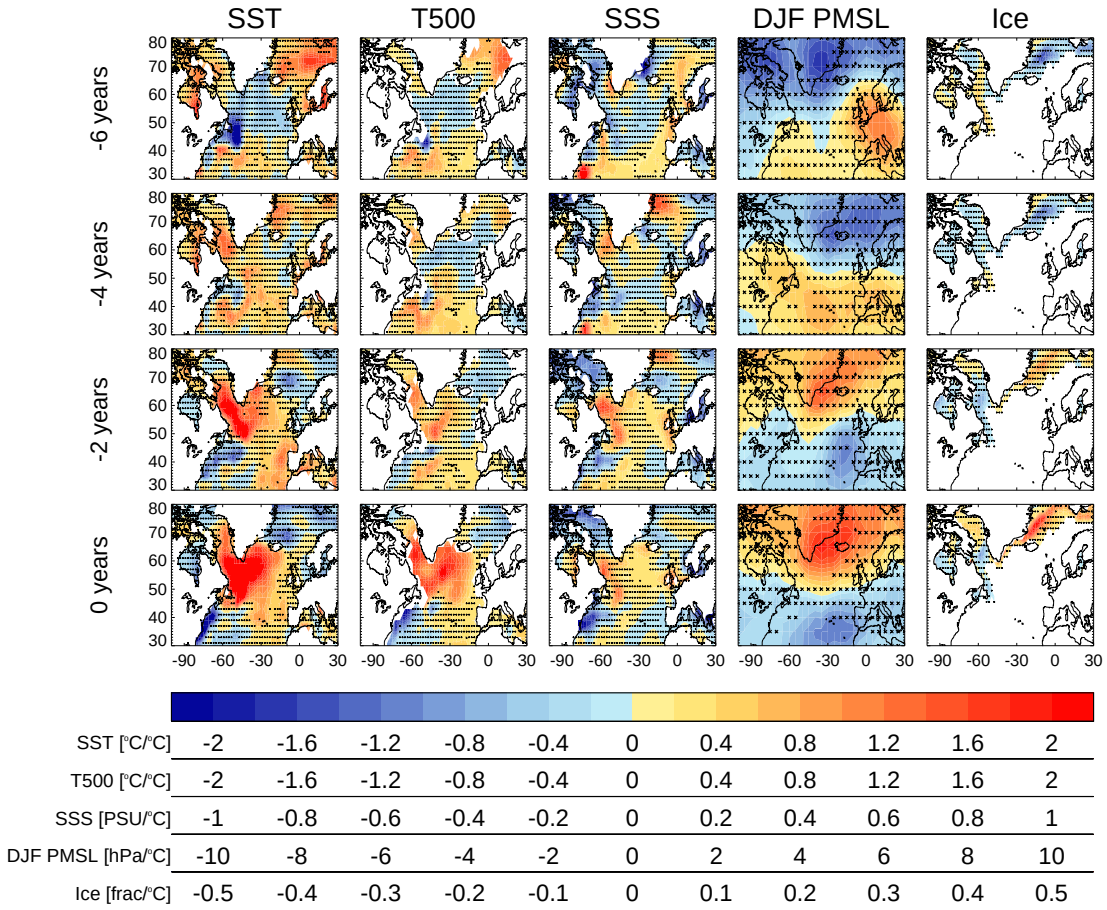

Figure 3.4: As in Figure 3.3 but only computing regressions for the period 1960–2000, for comparison with *Robson et al.* (2012). Regressions between North Atlantic subpolar (45–65°N) top 500m depth averaged temperature (T500, from EN4) and, from left to right: sea surface temperature (SST, from EN4), T500, sea surface salinity (SSS, from EN4), wintertime mean sea level pressure (DJF MSLP, from HadSLP2r), and ice fraction (from HadISST). From top to bottom, the T500 index lags the fields from -6 to 0 years. To attempt to remove the climate change signal, all fields are first detrended by removing the local regression against a global mean SST index. Stippling denotes regions insignificant at the 90% level using a two-tailed t-test for correlations between random variables with the same mean, standard deviation, and lag=1 autocorrelation as the real data *i.e.* testing the null hypothesis that both the T500 index and each of the fields merely exhibit damped persistence.

strongly negative phase, was associated with a rapid warming of the NA SPG. This result does call in to question our use of a T500 index (rather than SST index) in Figure 3.3 but, as previously noted, the model (to which we compare) is insensitive to this choice. Once again this may reflect a distinction between average and specific decadal variability.

In summary, the key similarities between simulations and observations are the lag=0 relationships between the T500 index and various surface fields as well as the possible role for advection of heat content anomalies and the phase reversal of the NAO when using only observations for the period 1960–2000. However, there are also potentially important differences including the unclear origins of heat content anomalies when using the full period 1900–2014 and the sensitivity of the observed relationships to the use of either surface or depth averaged temperature indices when using the shorter period 1960–2000. It is not clear to what extent many of these differences could be explained by the confounding influences of data paucity, the representativeness of particular decadal events, and the presence of transient external forcing in reality (which masks the internal variability). In the next section we begin to address this by considering only the most well observed dataset: SSTs.

3.3.2 Testing the simulated variability in regions of high observation density

The previous analysis comparing the surface evolution of various fields attempted to maximise the number of available years by allowing the use of data infilling, which is approached in different ways for the different data sets and variables but essentially amounts to reverting to climatology over a given length and/or timescale. As such, the regressions shown in Figure 3.3 may only represent a reduced subset of years, compared to the potentially 114 total years used. We now attempt to account for incomplete spatial and temporal coverage by 1) inquiring where there are the most observations of SST within the NA SPG, subsequently 2) characterising the simulated decadal variability by relationships between those locations before 3) applying the same analysis to the (now more well sampled) observations.

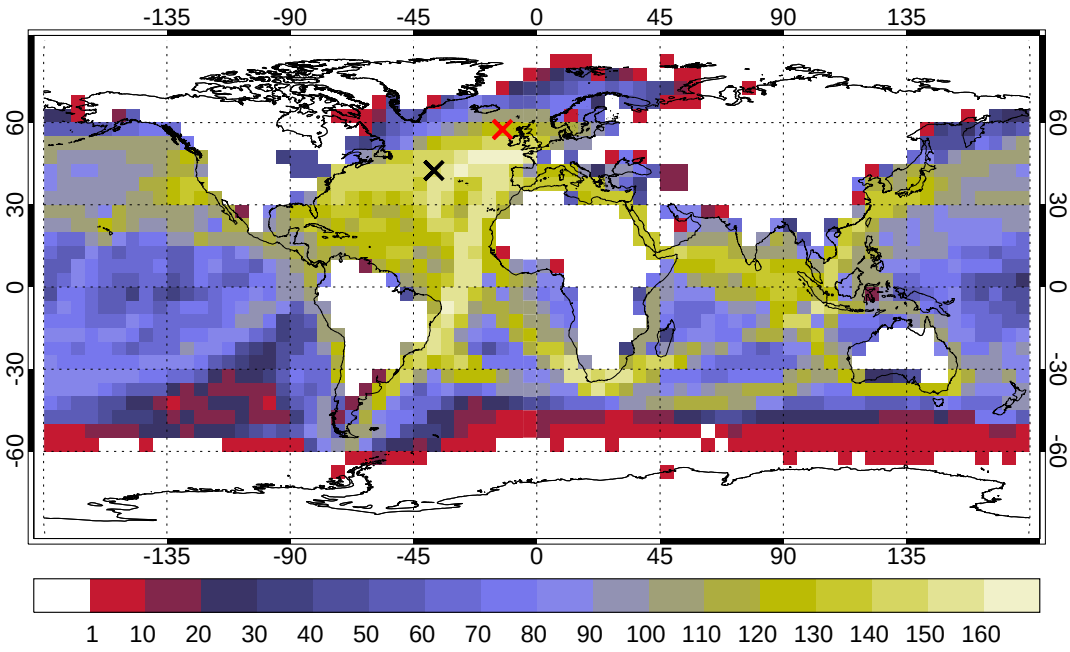

Figure 3.5: The number of individual years within the HadSST3 sea surface temperature dataset (*Kennedy et al.*, 2011) that have at least 10 months with which to create an annual mean. Data are provided on a regular 5x5° grid and span the period 1850–2014 inclusive.

The spatial coverage of SST data within the HadSST3 dataset (*Kennedy et al.*, 2011) is shown in Figure 3.5 and highlights the North Atlantic as a relatively well observed basin. The data are provided pre-binned on to a 5x5° regular grid (*Kennedy et al.*, 2011) and, within the North Atlantic, most of the subtropical and southern edge of the subpolar gyre have more than a century of annual means (where we have defined an annual mean as containing a minimum of 10 sampled months). However, there is a large disparity between the southern and northern edges of the subpolar gyre, with the northern edge approximately half as well sampled as the southern edge. This is of particular concern for attempting to directly test the mechanism of simulated decadal variability in HadGEM3, which relies on ocean advection and feedbacks within this poorly sampled region (see Chapter 2 and Section 2.4.5).

In the analysis of the mechanism of decadal variability in HadGEM3 (Chapter 2) we noted the long timescales (several years) for near surface (T500) signals to propagate around the NA SPG, as well as the shorter timescales (around 1 year) between anomalously positive northern NA SPG signals and anomalously negative southern NA SPG signals (Figure 2.8). We now recreate that figure using observed SSTs from HadSST3,

and slightly modifying our index locations to be the most well sampled regions (crosses in Figure 3.5). Changing the index locations also affects the simulated timescales and so we recompute these using the observationally constrained locations.

The time series of the observed SSTs in the two index locations are shown in Figure 3.6a, and Figure 3.6b and highlight the large multi-decadal variability in observed SSTs in the North Atlantic region. The subsequent lagged correlation between the two indices is shown in Figure 3.6c. The nature of the lagged relationship does not appear too sensitive to the precise grid point locations chosen. The timescale for signals to transfer between the NAC region and north east subpolar gyre (NE SPG) region is somewhat shorter than in the original estimate with HadGEM3, which the model suggests is due to the shorter path used here as well as the use of SSTs rather than depth averaged temperatures. Estimating the same lagged correlation using the simulated data yields a timescale of 2 years, which is broadly similar to the 2–3 years implied by the peak correlations in Figure 3.6c.

The timescale for simulated inverse anomalies to form in the NAC lagging the NE SPG is not clear from the observed correlations, and shows no indication of a minimum around lag=-6 years, compared to HadGEM3 (filtering to remove periods $>$30 years yields qualitatively similar results, not shown). The observed correlations do show a minimum at around lag=-15 years, which would imply an overall timescale of 34 years, consistent with the multi-decadal variability seen in the SST indices (Figure 3.6a, b). This would also imply different processes at play in the observed record, possibly externally forced, to enable a longer timescale relationship and could suggest a reduced role for the simulated negative feedback between the Labrador Sea and NAC, which is a relatively quick process in the model taking at most only a few years. A reduced role for the simulated negative feedback in reality is also consistent with later analysis of the key processes of variability in the NA SPG (see Section 3.5.1). The timescales of signal propagation are investigated further in Section 3.5.2.

Given the data paucity, even in the relatively well observed SST record, it clearly remains difficult to apply the same lagged regression analysis previously used successfully with simulated data. That is, in the observed record, the degrees of freedom in the combined horizontal and temporal axes are still low. To address this, we now extend our analysis

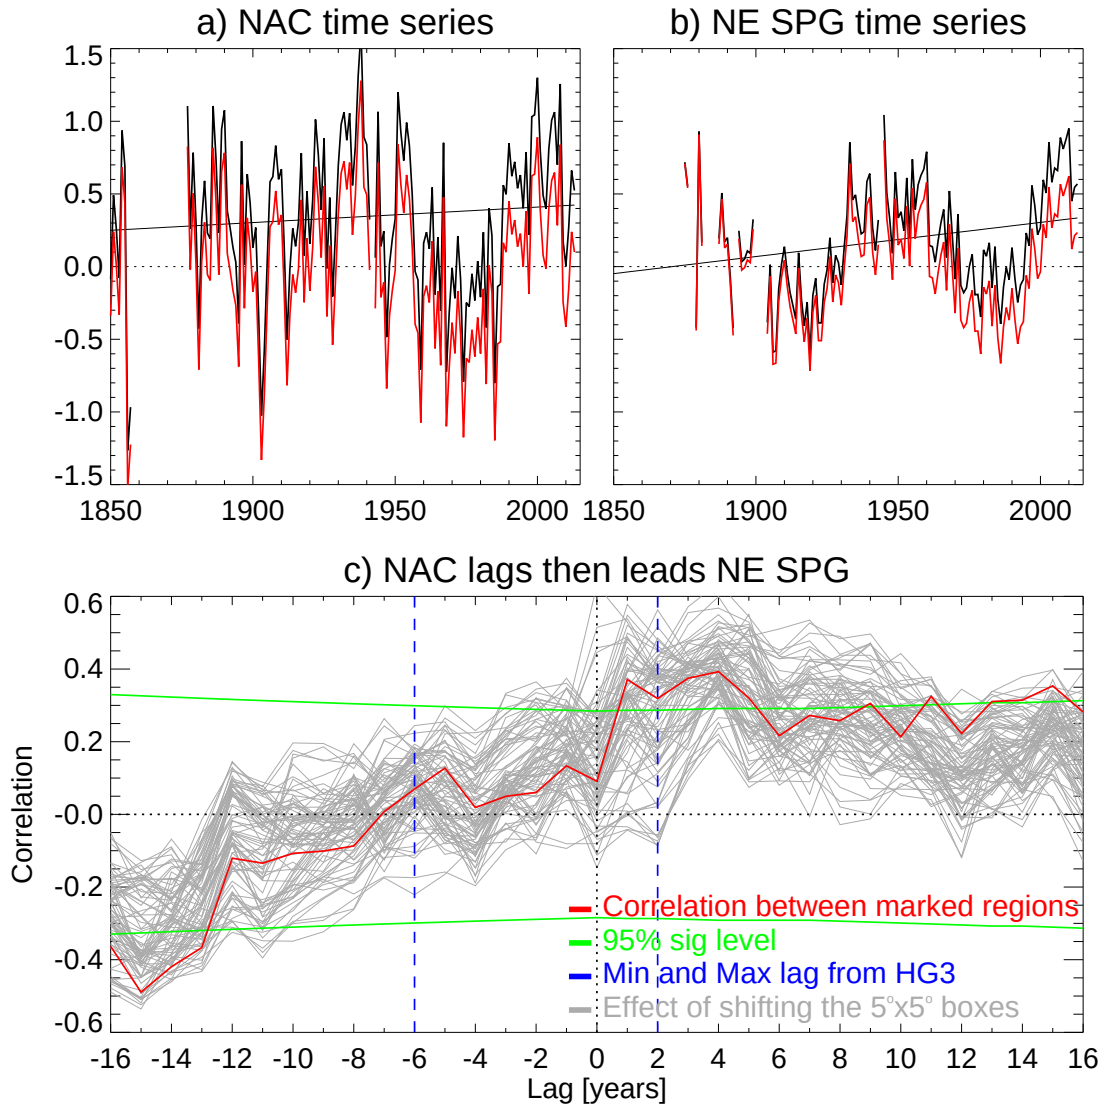

Figure 3.6: a) HadSST3 sea surface temperature (SST) time series in the North Atlantic Current (NAC) region, as indicated by the black cross in Figure 3.5 for both the raw data (black) and detrended data (red). Data are anomalies relative to the period 1961–1990 (*Kennedy et al.*, 2011). b) as (a) but for the North East subpolar gyre (NE SPG) region, as indicated by the red cross in Figure 3.5. c) The lagged correlation between detrended SSTs in the NAC and NE SPG regions, with NAC leading at positive lags (red). Significance levels (green) are estimated using a two-tailed t-test and assuming individual years are independent. Similarly computed times of maximum and minimum lagged correlation simulated in HadGEM3 are also highlighted (blue lines). The effect of shifting one or both of the north-eastern or south-western box by one five degree grid point meridionally or zonally is estimated by recomputing the correlation with every possible combination of paired locations (81 possibilities, grey lines).

to the depth structure of variability within the NA SPG and investigate to what extent surface and deeper signals are coherent in space and time.

3.4 Comparing the depth evolution of NA SPG variability in the model and observations

To analyse the depth structure of variability within the NA SPG, we show the temporal correlations between a near surface layer between 100-200m (such that it is below the Ekman layer) and approximately 250m thick layers in both EN4 and HadGEM3 (where we have sub-sampled HadGEM3 into 54 year sections comparable to EN4 and subsequently computed the ensemble mean values, Figure 3.7). This aims to elucidate the depth extent of the near surface (but not directly wind-driven) variability that is more likely to be due to advective ocean processes. In this respect these estimates of the depth coherence are different to merely estimating the mixed layer depths, the latter being more likely to reveal a signal of surface forcing (though in regions of particularly strong/frequent convection such as the Labrador Sea they are likely to be well correlated). For example, the wintertime mixed layer depths outside of the Labrador Sea are deepest in the eastern NA SPG in both EN4 and HadGEM3 (Figure 2.1d for HadGEM3), whereas the coherence analysis suggests the strongest depth coherence in the western NA SPG. In order to minimise the effects of the infilling methodology we use only the period 1960–2014 in the subsequent analysis with EN4 data.

As can be seen, both EN4 and HadGEM3 show similar coherence in the top 2 layers (Figure 3.7, top row) with high correlations throughout the NA SPG. At layers 3 and 4 (between approximately 500-1000m, Figure 3.7, second row) there are some differences, with EN4 highlighting stronger coherence south of the Denmark Straits, whereas HadGEM3 finds both this region and the Labrador Sea to be particularly depth coherent. Additionally, in HadGEM3, the depth extent of the NAC can be seen as the track extending from the south west to north east, which is not as readily visible in EN4. This may be linked to the relative importance of oceanic advection in this region in HadGEM3 (see Chapter 2 and Section 2.4.4) and perhaps suggests this effect may be less prominent in

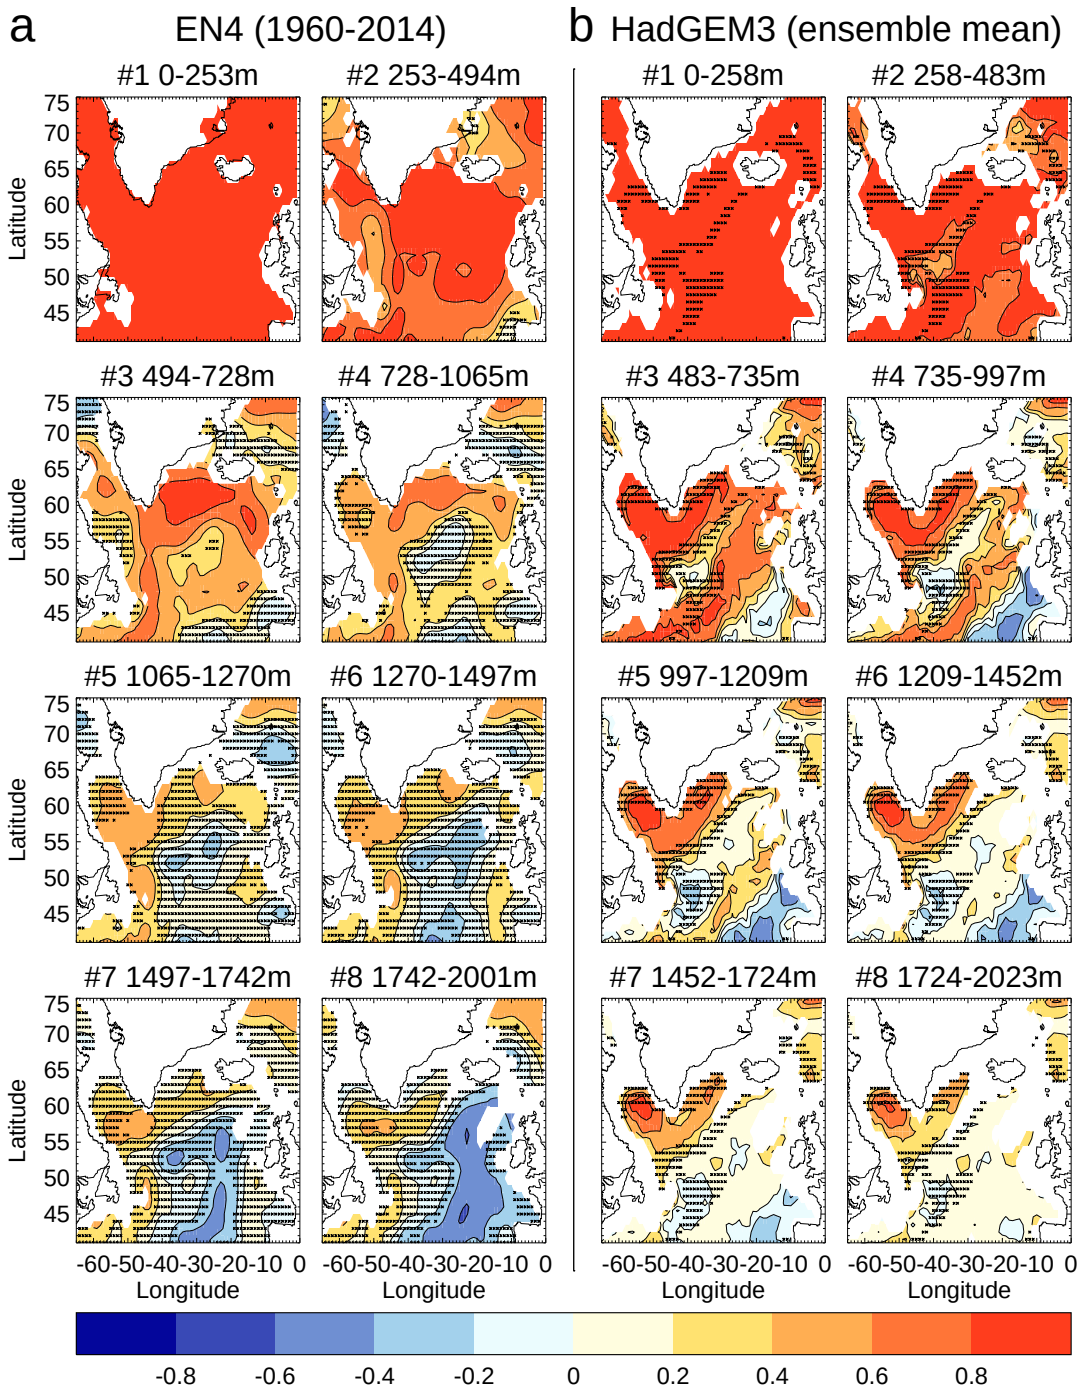

Figure 3.7: a) The correlation between layer average temperature data from EN4 (layer bounds noted above panels) and a depth layer between 100–200m. Layer averages are approximately 250m thick but the original grid cell boundaries are used instead of interpolating in the vertical. b) As (a) but for HadGEM3. EN4 is linearly detrended over the 54 year period and HadGEM3 high-pass filtered to remove periods greater than 108 years *i.e.* assuming the linear trend removed from EN4 represents a half period of at least 54 years. Stippling denotes regions insignificant at the 90% level using a two-tailed t-test for correlations between random variables with the same mean, standard deviation, and lag=1 autocorrelation as the real data *i.e.* testing the null hypothesis that both the 100-200m layer index and each of the lower 250m thick layers merely exhibit damped persistence.

reality. At even greater depths, between 1000–2000m (Figure 3.7, third and fourth rows) the NAC track is still visible in HadGEM3, but in EN4 this region is now increasingly anticorrelated with the near surface variability. Further, the east-west dipole that can be seen in EN4 (Figure 3.7, fourth row) is skewed in favour of positive western correlations in HadGEM3, perhaps suggesting that the spatial extent/influence of Labrador Sea water appears to be reduced in EN4 compared to HadGEM3 (see Section 3.6 and Appendix A).

Figure 3.7 uses detrended annual data. Smoothing the data with a running mean of, for example, 10 years to highlight multi-decadal variability, yields a qualitatively similar picture in both EN4 and HadGEM3 (not shown). This is because the annual variability is dominated by the multi-decadal signal (*cf.* the time series from 1960 onwards in Figure 3.2). Removing the multi-decadal variability (to analyse the interannual signal) has the effect of reducing the depth extent of the Labrador Sea signal in both EN4 and HadGEM3 and generally reducing the depth extent of all variability in EN4 (not shown). In general, the patterns of interannual and multi-decadal variability are broadly similar in HadGEM3, which is not the case in EN4. This difference in EN4 may be indicative of different processes existing at these different timescales, perhaps also related to whether the forcing is internal or external (note that HadGEM3 has no interannually varying external forcings), but could also reflect the short time series used in EN4 that clearly cannot resolve many individual multi-decadal variations.

The depth at which the correlation with near surface variability falls below a defined threshold is shown in Figure 3.8, which represents a condensed version of Figure 3.7. Although this removes information about the spatial extent of negative correlations at deeper levels, it does highlight several other key features. For example, the Labrador Sea and Irminger Current regions can be seen to have similar depth coherence in both HadGEM3 and EN4 despite the actual correlation values generally being lower in EN4 (Figure 3.7). In addition, the NAC track can be clearly seen in HadGEM3 as a region of strong depth coherence, down to around 1250m, compared to 750m in EN4. In HadGEM3 this region is separated from the northern SPG except at the east and west coasts. This is consistent with subsurface signals propagating around the NA SPG in HadGEM3 and subsequently feeding back on the circulation in the NAC region in the west (see Chapter 2). In EN4, although the connections in the west and east coasts of the

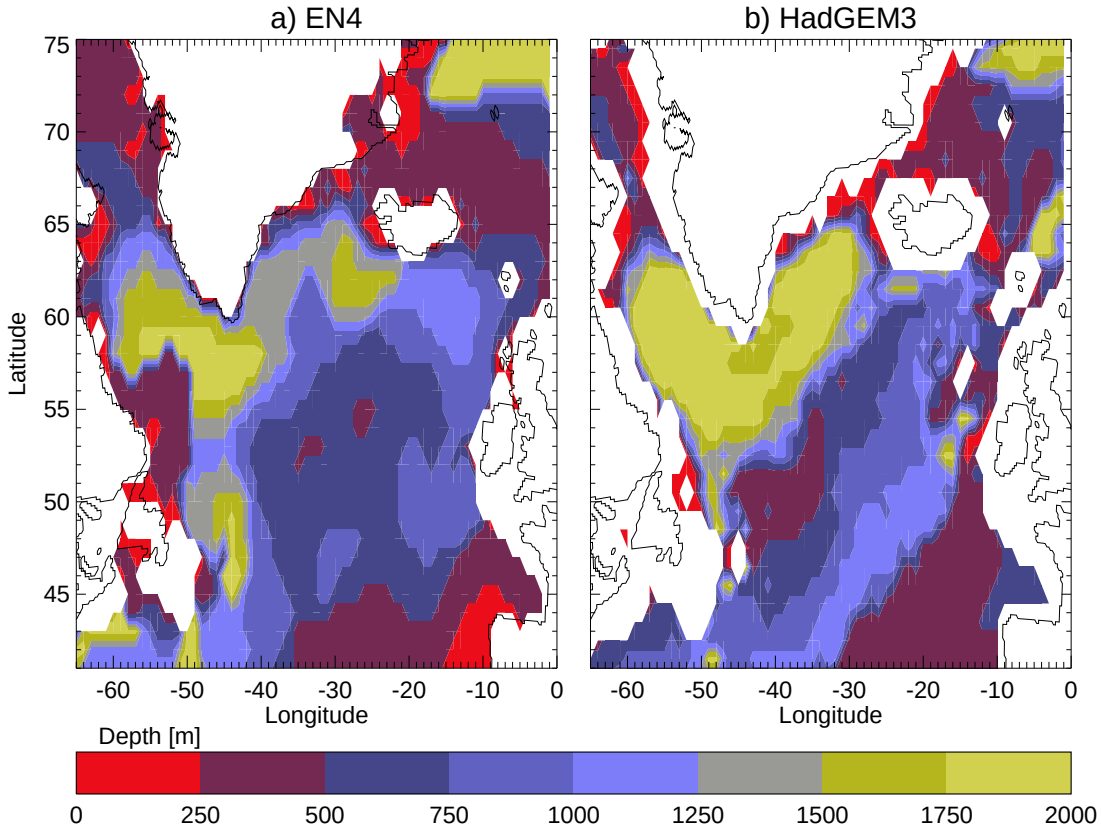

Figure 3.8: a) The depth at which the correlation between the 100–200m layer and other 250m thick layers (shown in Figure 3.7) falls below the arbitrary level of r=0.3 in EN4 (a), and HadGEM3 (b)

NA SPG exist, there is a much weaker depth coherence within the NAC, possibly pointing to a reduced role for advection of heat content anomalies in this region. Once again, removing the multi-decadal variability has the effect of reducing the maximum depths of significant correlations in both EN4 and HadGEM3.

In summary, a comparison of the depth coherence of variability within the NA SPG between EN4 and HadGEM3 suggests a broadly similar structure in terms of the prominence of the Labrador Sea and northern NA SPG. However, there are some differences, such as the reduced depth extent of the NAC coherence in EN4, which may be key to the processes and amplitude of variability in HadGEM3 as compared to EN4. We now examine these key processes in some more detail, beginning with the simulated negative feedback between the Labrador Sea and NAC (Section 3.5.1) before investigating the timescales of propagation around the NA SPG (Section 3.5.2).

3.5 Observational analysis of key simulated processes

3.5.1 Negative feedback between Labrador Sea and NAC

To investigate the negative feedback between Labrador Sea and NAC temperatures that was simulated in HadGEM3 we compute spatial maps of the driver of interannual density changes (Figure 3.9). The data we use are horizontally varying top 500m depth averaged temperature and salinity in EN4, from the relatively well observed period of 1960–2014. This analysis follows the methodology of *Delworth et al.* (1993) and decomposes density changes into those due to temperature and those due to salinity by time-meaning salinity or temperature in the density equation of state respectively (Equation 3.1).

$$\rho_T = \rho(T, \bar{S}, p), \quad \rho_S = \rho(\bar{T}, S, p) \tag{3.1}$$

Having computed temperature-induced density changes, we then regress density against these temperature-induced density changes to estimate the regression slope between the two *i.e.* for a given density change, how much of this change is due to temperature variability (Equation 3.2)?

$$\rho_{Tcontrol} = regr[\rho, \rho_T], \quad \rho_{Scontrol} = regr[\rho, \rho_S] \tag{3.2}$$

We do the same for salinity-induced density changes, and subsequently subtract the two regression slopes to estimate the magnitude of temperature or salinity control of density changes (Equation 3.3).

$$\rho_{TorScontrol} = \rho_{Tcontrol} - \rho_{Scontrol} \tag{3.3}$$

A value of $\rho_{TorScontrol}>0$ indicates density changes that are temperature controlled, with $\rho_{TorScontrol}<0$ indicating salinity-controlled density changes. The linearisation of the non-linear density equation of state, using depth averaged temperature and salinity, explains >99% of the variance in depth averaged density in the NA SPG in both EN4 and

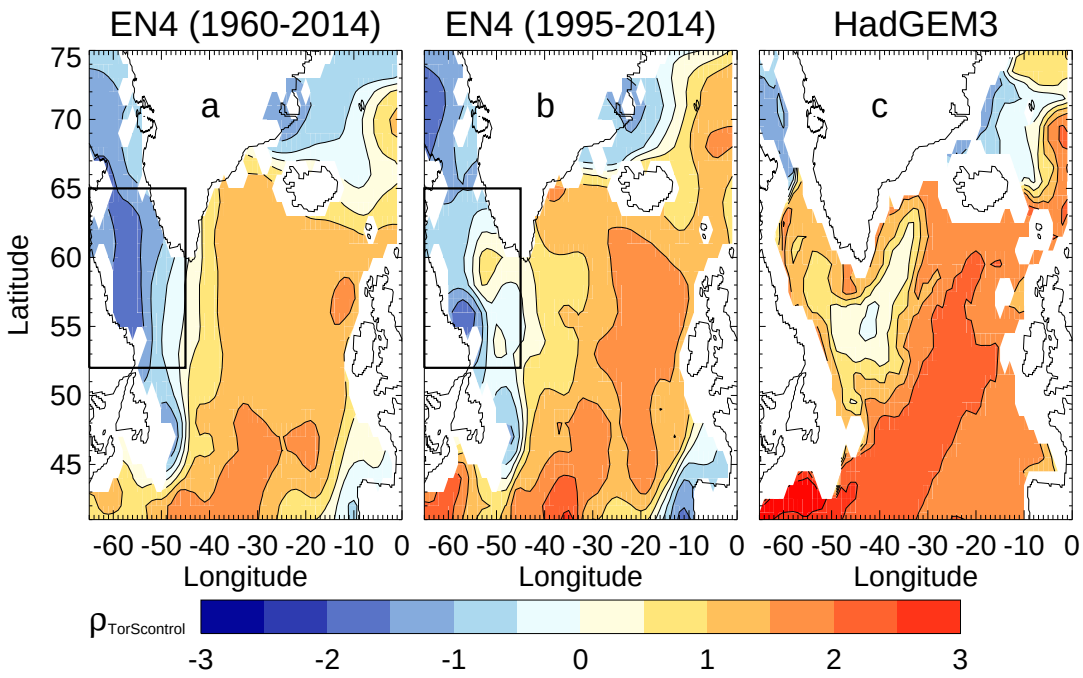

Figure 3.9: $\rho_{TorScontrol}$, as described in the text, for the top 500m in EN4 for the period 1960–2014 (a), for the period 1995–2014 (b), and HadGEM3 (c). The Labrador Sea region in EN4 is marked and analysed in Figure 3.10. Both EN4 and HadGEM3 are linearly detrended prior to computing $\rho_{TorScontrol}$.

HadGEM3 (not shown) and is thus able to provide a framework to investigate the drivers of interannual density variability.

Within the NA SPG, in both EN4 and HadGEM3, top 500m interannual density variability in the eastern half of the basin can be considered to be temperature controlled, although this is more strongly the case in HadGEM3 (Figure 3.9c) than in EN4 (Figure 3.9a). In contrast, within the Labrador Sea region, the density variability is driven by salinity variability in EN4 (Figure 3.9a) but remains temperature-controlled in HadGEM3 (Figure 3.9c). Given that, in both EN4 and HadGEM3, interannual anomalies in temperature and salinity co-vary in the NA SPG such that warm anomalies are generally also more saline (not shown), this suggests that a given anomaly in the Labrador Sea region would be expected to have opposing effects on density in either EN4 or HadGEM3. *e.g.* A warm/saline anomaly would be expected to decrease density in HadGEM3 (because density changes are controlled by temperature, and increasing temperature reduces density) but increase density in EN4 (because density changes are controlled by salinity, and increasing salinity increases density). However, we suggest some caution in interpreting the limited EN4 variability and note that assimilation of observed anomalies into a

dynamical ocean model suggests a prominent role for temperature in controlling density variability (*Robson*, 2010) and that hydrographic analyses suggest recent Labrador Sea density variability may be more complex than merely whether it is temperature or salinity controlled (*Yashayaev*, 2007). The sensitivity to the exact period is also implied by the near-zero values of $\rho_{TorScontrol}$ in the Labrador Sea for the most well observed period 1995–2014 (Figure 3.9b).

To investigate the specific contributions of temperature and salinity to observed density changes in the Labrador Sea, in Figure 3.10a we plot time series of density (ρ), density due to temperature changes (ρ_T), and density due to salinity changes (ρ_S). The linear sum of the components can be seen to be a good estimate of the actual non-linear density anomaly. Throughout the period 1960–2014, the relative contributions of temperature and salinity to density variability are not the same, highlighting the potential difficulty in choosing a representative period. For example, there are large density anomalies that are sometimes salinity driven with little temperature compensation (*e.g.* 1962), and other times large ρ_T and ρ_S anomalies that are broadly compensating (*e.g.* 1984).

To quantify the contributions of temperature and salinity induced density changes through time, we plot the correlation of 20 year subsections of the data (Figure 3.10b). Similar to the regression maps presented in Figure 3.9a, for most of the period the largest correlations are between density and ρ_S, contrary to the relationship simulated in HadGEM3. However, for the most recent 20 years (*i.e.* 1995–2014, as in Figure 3.9b) interannual density variability appears to be more driven by temperature changes, with higher correlations between density and ρ_T. The simulated relationships (indicated by the shading in Figure 3.10b) are more consistent with this recent period. It is not clear to what extent the observed density variability, and drivers thereof, are here exhibiting internal variability/noise (the correlation window is only 20 years long so could severely alias longer period variability, should it exist — see Section 3.3.2) or to what extent the apparent switch from salinity to temperature-controlled density variability represents a secular/climate change. Further analysis on this topic is outside the scope of this investigation into the mechanisms of simulated decadal variability, but would be worthwhile in the future. Additional analysis in a multi-model context is presented in Chapter 6.

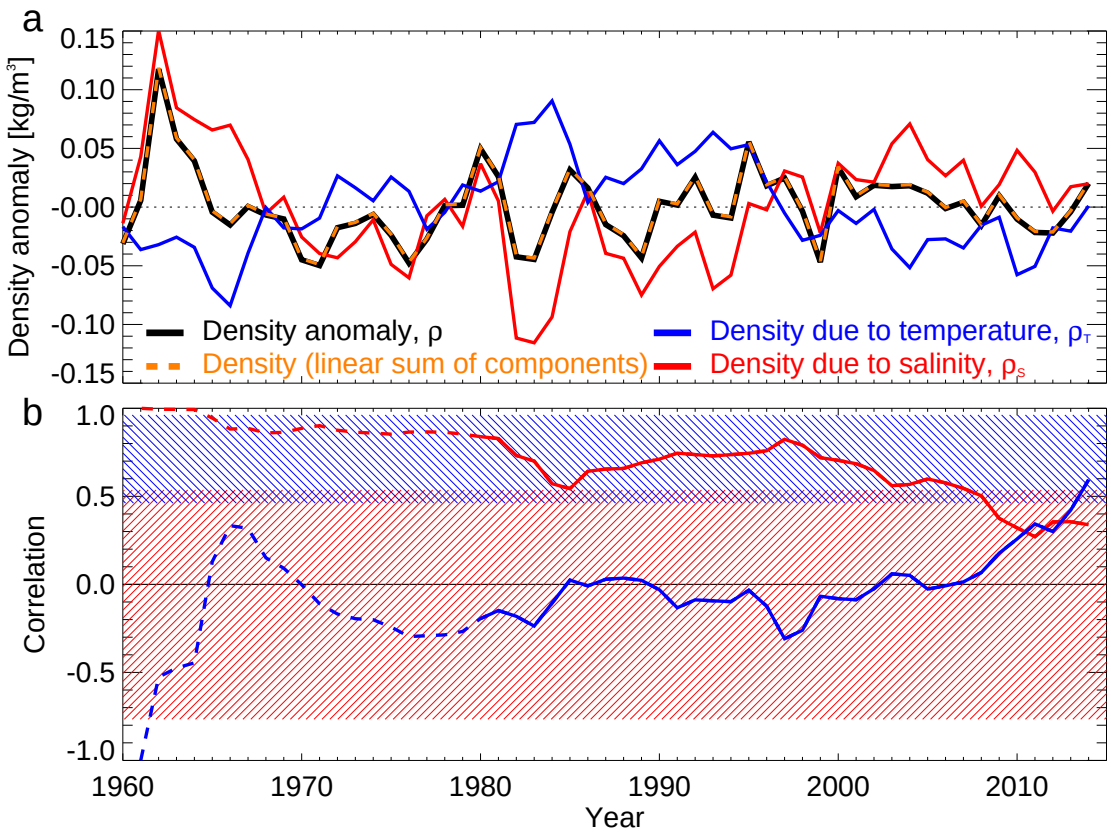

Figure 3.10: a) Time series of the volume averaged density anomaly (computed from linearly detrended volume averaged temperature and salinity) in the Labrador Sea (52–65°N, 45–65°W, top 500m, as marked in Figure 3.9) for the actual density (ρ, black), for density due to temperature changes (ρ_T, blue), and for density due to salinity changes (ρ_S, red). Also plotted is the sum of ρ_T and ρ_S (orange). b) Correlations between density and ρ_T (blue) and ρ_S (red) using a moving 20 year window, plotted at the final year. Prior to 1980 a reduced window length is used (dashed). The 95% ranges of the same correlations in HadGEM3 (using an identical 20 year moving window) are highlighted with red and blue shading.

In both HadGEM3 and EN4, T500s in the Labrador Sea and NAC region are anticorrelated (not shown), seemingly consistent with a situation whereby warm anomalies in the Labrador Sea could induce cool anomalies in the NAC (Chapter 2, Section 2.4.6). However, the disparity in the driver of Labrador Sea density changes between HadGEM3 and observations (the last 20 years notwithstanding) has implications for the feedback identified in HadGEM3 and whether it acts as a positive or negative influence. Consequently, the anticorrelation of Labrador Sea and NAC T500s in EN4 are perhaps more likely due to a local response to the NAO (*Visbeck et al.*, 1998) rather than an ocean feedback. In HadGEM3, positive (negative) temperature anomalies in the Labrador Sea induce negative (positive) temperature anomalies in the NAC via a weakening (strengthening) of the

meridional density gradient (which is temperature controlled) and associated geostrophic balance. In EN4, if the same process exists, it would result in positive (negative) temperature anomalies in the Labrador Sea inducing positive (negative) temperature anomalies in the NAC, as the associated Labrador Sea salinity anomalies would have the effect of strengthening (weakening) the meridional temperature gradient (which is salinity controlled), *i.e.* there would be a positive feedback. We analyse this feedback further in a multi-model context in Chapter 6.

3.5.2 Propagation timescales in the southern half of the NA SPG

In Section 3.3.2 we attempted to compute observed lagged relationships between locations in the NA SPG using SSTs. Although a similar methodology showed utility when analysing the multi-century simulation with HadGEM3, it appeared to show limited scope when given the shorter time series—and increased set of forcings present—in the observed record. To address this, we instead calculate the implied propagation timescales in both EN4 and HadGEM3 by essentially calculating the extent of the spatial coherence of depth averaged temperature anomalies (Figure 3.11).

For example, taking the EN4 dataset for the period 1960–2014, we begin by defining an initial index region, which is the area averaged top 250m depth averaged temperature in the south western corner of the NA SPG (41–45°W, 41–45°N, marked region in Figure 3.11). We then correlate the (linearly detrended) time series of this index with time series of top 250m temperature in all locations in the NA SPG at a lag of one year (*i.e.* one year later), and find regions where the correlation is greater than r=0.3 (the correlation required for significance at the 95% level for a two-tailed t-test with 54 years of data and assuming independence). Having done this, we calculate the mean longitude and latitude of the NA SPG points that meet this criteria and define a new box centred at this location, with zonal and meridional extents of $\pm 2^\circ$ from the centre. We then proceed to use this new location as our index region and begin the process again. Each iteration or 'pass' is marked on the map. We do this both for the layer 0–250m (Figure 3.11a) and the layer 250–500m (Figure 3.11b), once for the 54 year time series from EN4 and 75 times using 54 year long sections of the HadGEM3 simulation to estimate the simulated variability

around the mean pathway. This method has the joint benefits of utilising much more of the available data to constrain our result than merely finding the lagged correlation between two distant points (*cf.* Figure 3.6) as well as minimising the loss of degrees of freedom due to the lead time (fixed at 1 year). As the timescales we infer from this analysis are a function of the correlation cutoff value they cannot be directly compared to the timescales estimated from correlations between the southern and northern edges of the gyre (*cf.* Figure 2.8 and Figure 3.6) but, as we have used time series of identical lengths in EN4 and HadGEM3, these results can be meaningfully intercompared, with the spread in the HadGEM3 ensemble also providing an estimate of the sampling uncertainty.

Within the 0–250m layer, in both EN4 and HadGEM3, the propagation pathway follows the mean circulation pathways within the NA SPG. The pathway derived from the EN4 data (Figure 3.11a, black line) is slightly outside the range of pathways simulated by HadGEM3, and to the south and east of the ensemble mean from HadGEM3. In addition, in HadGEM3 the signal arrives in the eastern NA SPG, around 33°W, 55°N within a single iteration whilst this takes approximately 4 iterations in EN4. This suggests the near surface propagation timescales in HadGEM3 are faster than in reality. Given that the strength of the circulation, measured by the barotropic streamfunction, is broadly consistent with available observations (see model validation in Chapter 2, Section 2.3.1) this implies that processes other than merely the current speed are important for the slower signal propagation in EN4, which the model cannot properly simulate (*e.g.* eddies), or does not properly simulate. For example, the signal propagation in reality could be inhibited by damping either from surface fluxes (*e.g.* related to more/less NAO variability in reality), due to interaction with deeper flows in the ocean, or from a greater/lesser role for heat storage within the ocean.

To investigate the role of the deeper circulation in aiding signal propagation we also plot the pathways for a layer between 250–500m (Figure 3.11b). In HadGEM3 the ensemble mean pathway is the same as for the shallower layer and the timescales are also broadly consistent, though there is increased spread across the ensemble members. This is somewhat dissimilar to EN4, in which the signal initially travels slowly eastwards, before halting around 37°W, 44°N.

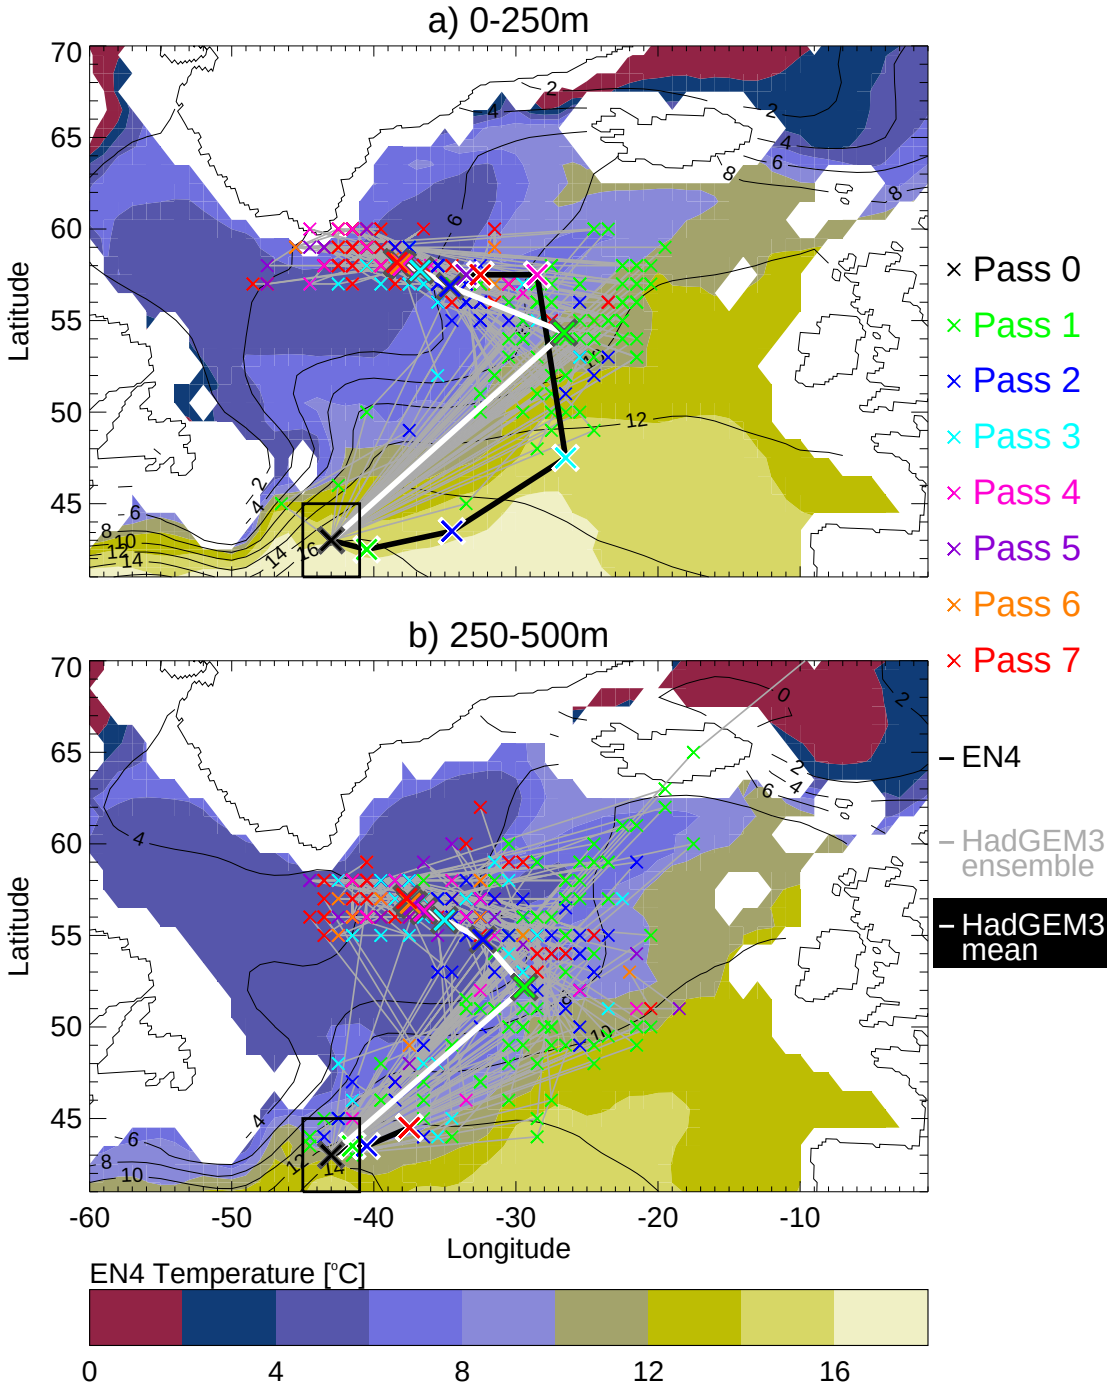

Figure 3.11: The inferred signal propagation timescales in 250m thick temperature layers between 0–250m (a) and 250–500m (b), estimated from one year lagged correlations, as described in the text. The pathways are shown for EN4 (black), the HadGEM3 ensemble members (grey), and the HadGEM3 ensemble mean (white). The first pass is at a box centred at 43°W, 43°N with a zonal and meridional extent of ±2°. Subsequent iterations are as marked in the figure legend up to a maximum of 7 iterations. In the background are the time mean depth averaged temperature in EN4 (colours) and HadGEM3 (contours) for the layer 0–250m (a) and the layer 250–500m (b).

Although there are differences in the top 500m depth averaged temperature profiles between HadGEM3 and EN4 (See Figure 2.1a for the difference) the locations of the temperature gradients are broadly similar between the two (Figure 3.11, background colours/contours) and don't appear to explain either the different timescales or pathways by which signals propagate from the NAC region into and around the NA SPG. However, we suggest that the timescale for signal propagation in EN4 in the 0–250m layer and the subsequent lack of an advective pathway in the deeper 250–500m layer are self consistent, with effective damping (or lack of reinforcement, compared to HadGEM3) from this deeper layer slowing down the signal propagation in the upper layer of EN4. This is also consistent with the analysis of the depth coherence of subsurface signals in EN4 and HadGEM3, in which EN4 had a much shallower coherence than HadGEM3 (Figure 3.8, note that the specific depths are a function of the arbitrary choice of a correlation cutoff). To bring together our assessment of the veracity of the key processes in HadGEM3 we now investigate the evolution of the depth structure.

3.5.3 Evolution of the depth structure

In the previous sections we analysed whether it was plausible that the negative feedback between Labrador Sea and NAC T500s existed in reality (Section 3.5.1), and whether the timescales of signal propagation around the NA SPG were likely to be similar (Section 3.5.2). In this final section, we attempt to bring these together by investigating depth profile composites in one part of the NA SPG based on significant departures from a particular index in another part of the NA SPG and how these profiles change in time (Figure 3.12).

Figure 3.12b shows the difference in temperature profiles (located in the blue marked region) between profiles that lag by three years or are in phase with a temperature index (based on volume averaged temperatures over the top 500m in the red marked region, see figure caption), in both EN4 and HadGEM3. That is, they show how the structure of a temperature profile in the NAC changes following an increase in the volume mean temperature in the eastern NA SPG. As can be seen, following a warming of the eastern NA SPG, the temperature profile throughout the top 500m in HadGEM3 becomes cooler,

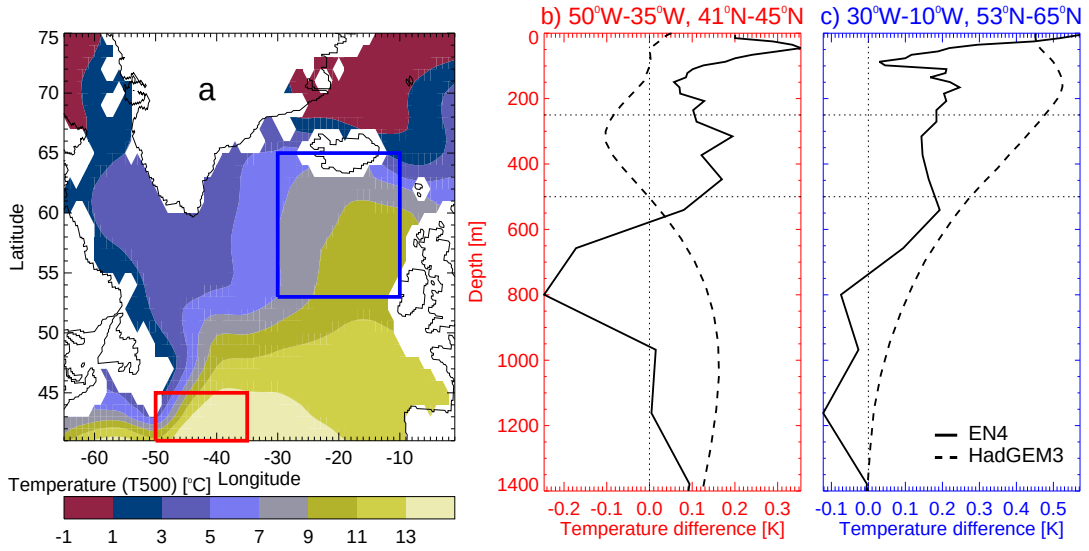

Figure 3.12: a) Time mean (period 1960–2014) top 500m depth averaged temperatures (T500) in the North Atlantic subpolar gyre (NA SPG) region from EN4. b) and c) Depth profile lagged composites in HadGEM3 and EN4, area averaged at the titled locations. Composites profiles are created by taking the difference between high and low instances (greater than one standard deviation from the mean) of a T500 index, which is volume averaged over the alternate region, *e.g.* Panel b (red) is based on an index in the blue marked region. The difference between lag=3 and lag=0 composites are then taken to construct the final lagged composite. As such, the lagged composites highlight the effect of the 3 years of temporal evolution of the temperature anomalies and are symmetric about high and low phases (due to the initial 'high' minus 'low' differencing). EN4 is linearly detrended over the 54 year period and HadGEM3 high-pass filtered to remove periods greater than 108 years *i.e.* assuming the linear trend removed from EN4 represents a half period of at least 54 years. Lags of 3 years are used to balance highlighting the temporal evolution of the signal with increasing noise at greater lags, particularly given the short EN4 time series.

with subsequent warming at depth (Figure 3.12b, dashed). This cooling of the upper part of the water column is indicative of the negative feedback between temperatures in the Labrador Sea (and upstream in the eastern NA SPG) and temperatures in the NAC (see Section 2.4.6). However, such a signal is not seen in EN4, in which warming of the eastern NA SPG is followed by even greater warming of the NAC region throughout the water column (Figure 3.12b, solid). This result is consistent with our previous analysis of the negative feedback in EN4 (Section 3.5.1) in which, due to interannual density changes being driven by salinity rather than temperature, a warm (and saline) anomaly in the Labrador Sea would be expected to exert a positive feedback on temperatures in the NAC.

Conversely, we next construct profiles in the eastern NA SPG, based upon an index of temperatures in the NAC region (Figure 3.12c). These highlight how the structure of a temperature profile in the eastern NA SPG becomes warmer or cooler three years after an increase in the mean temperature in the NAC region. In HadGEM3, the eastern NA SPG becomes increasingly warm following warming in the NAC, with the warming exhibiting a maximum at around 200m depth (Figure 3.12c, dashed). In EN4, this warming is also evident, though much reduced (Figure 3.12c, solid). Once again, this is consistent with our previous analysis, in which we found that temperature anomalies could be more clearly seen propagating around the NA SPG in HadGEM3 than in EN4 (Section 3.5.2). The greater subsurface extent evident in HadGEM3 may help the signal to remain undamped for longer.

In summary, the negative feedback simulated in HadGEM3 appears unlikely to have existed in exactly the same form in reality due to the differing drivers of density variability in the Labrador Sea, with the possible exception of the most recent decades. In addition, the propagation timescales of near surface signals, which are an important part of the mechanism of simulated decadal variability in HadGEM3, are also likely to be modified in reality, possibly due to the slightly different depth structure in HadGEM3 than in EN4. In light of these findings we now discuss our results, beginning with a discussion of the rationale behind our approaches and further discussion of the role of the Labrador Sea.

3.6 Discussion

We begin this section by briefly highlighting the difficulties in applying the same analysis techniques to observations as to models before discussing our results for the simulated and observed Labrador Sea region in some more detail. We discuss the relative strengths of statistical (*e.g.* EN4) and dynamical analysis products and the potential benefits of increased model resolution before concluding with some recommendations for future observational networks.

In Section 3.3 we attempted to analyse a suite of observational datasets using a similar framework to that which we applied to the simulations (Chapter 2). Although these

datasets use sophisticated methods to attempt to deal with spatially and temporally sparse observations, without prior knowledge of the relative contributions of internal and external forcings to the observed variability it is not possible to fully determine to what extent model-observation disparities are due to model deficiencies (when our model simulations do not include any estimate of transient external forcings). The strength of these simulations however is the statistical power that results from more than 400 years of simulation. As noted by the comparison of Figures 3.3 and 3.4, it may also be the case that average decadal variability (observed or simulated) may not necessarily evolve similarly to specific, large decadal events. Further work on this topic could involve a case study approach to find simulated analogues to observed events and assess their nature and importance/frequency in more detail, similar to the approach used by *Robson et al.* (2012) using a model reanalysis.

In many investigations into simulated decadal variability in the NA SPG, the Labrador Sea has been found to play an active role in this variability (see Chapter 1 and Figure 1.5) including in HadGEM3 (Chapter 2) where the Labrador Sea is a region of particularly strong mixed layer depth variability (and by inference, convection, Figure 2.1d). To investigate the relative importance of deep water formed in the Labrador Sea and in the Nordic Seas (another important deep water (precursor) formation site) we have performed additional tracer release experiments that investigate the downstream evolution of water formed in these two regions, detailed in Appendix A. These experiments suggest that North Atlantic Deep Water (NADW) formed in the Labrador Sea (Upper NADW) and Nordic Seas (Lower NADW) is indistinguishable in the model, whereas in reality the depth maxima of these two water masses are separated by approximately 1km (*Toole et al.*, 2011). As such, it is possible that the model overemphasises the role of the Labrador Sea in its decadal variability.

In addition to the simulated prominence of Labrador Sea water, the model also suggests a deeper extent of coherent signals within the NAC (Section 3.4), which may be the cause of the larger depth range in which advective signals are seen to propagate around the NA SPG in HadGEM3 compared to EN4 (Section 3.5.2). This may be due to the overly diffuse thermocline in HadGEM3 (*Megann et al.*, 2014) that results in a more permeable barrier between near surface and deeper waters in the NAC. Further increases

in resolution to 'eddy resolving' scales may also reduce the depth extent of the coherence in the model by reducing the average depth of high latitude mixed layers (*Oschlies*, 2002, see discussion below of resolution effects). These are two examples of how model mean state biases can potentially affect the simulation of variability, and so cannot be merely "subtracted off" the final solution or removed by conducting analyses in anomaly-space. This will be discussed in more detail in a multi-model context in Chapter 6.

In addition, also relevant to the discussion of the relationship between mean state biases and modes of variability are the implied dominant drivers of density variability and their non-stationary nature in EN4. One of the key differences between the simulations and observations is the time mean difference in the driver of interannual density variability (Figure 3.9), in which the simulations suggest temperature drives density variability throughout the NA SPG whereas in EN4 this switches to be salinity driven in the Labrador Sea region. As previously discussed, neglecting other differences, this would imply a positive (rather than negative) feedback between Labrador Sea and NAC temperatures in EN4. Despite this, further analysis suggests that the dominant driver of Labrador Sea density variability in EN4 may not be stationary, with recent decades implying a temperature dominated regime. However, this result should be interpreted with some caution, partly because the correlation window width is necessarily narrow (20 years) and may alias longer term variability. Nonetheless, if the present relationship is maintained, it would suggest that the variability simulated by HadGEM3 may become more likely/prominent in the real world in the future — in the absence of confounding variability driven by external forcings, the relative magnitude of which is unclear, as noted above. Further model simulations with realistic external forcings would help to address this issue but we note that, even if forced simulations in the same model framework were available, it would not be trivial to quantify the externally forced signal (*Frankcombe et al.*, 2015).

In our analysis of the real world we have made significant use of the latest Objective Analysis EN4 dataset (*Good et al.*, 2013, described in Section 3.2.2) from the Met Office Hadley Centre rather than alternative reanalysis datasets, such as, for example, the Simple Ocean Data Assimilation reanalysis (SODA, *Carton and Giese*, 2008), or the latest Operational Reanalysis System (ORAS4, *Balmaseda et al.*, 2013). The fundamental difference between the Objective Analysis of EN4 and these reanalyses is that the Objec-

tive Analyses provides its analysis using fixed spatial and temporal decorrelation scales (statistical), whereas the model reanalyses use a dynamical ocean model (physical). The question is then to what extent the physical approach can be considered more reliable than the statistical one, given the imperfect nature of these reanalyses (*Carton and Giese*, 2008; *Balmaseda et al.*, 2013), and which of the different physical models should be used (*Kröger et al.*, 2012). To sidestep these issues our approach has been to test the model against the Objective Analysis in the regions (in time and space) where there are enough direct observations to be most confident in the observed signal. In addition, we note that boundary currents are important for the NA SPG decadal variability in HadGEM3 (See Chapter 2 and Section 2.4.5), which are not well resolved by the resolution ($\geq 1^\circ$) of current ocean reanalyses. However, this will also be the case with EN4, which uses decorrelation length scales of $>$300km.

Although it is difficult to isolate the precise mechanisms by which increased ocean or atmosphere resolution may have altered our results — without a parallel set of low resolution simulations within the same model framework — there are specific features of the simulated decadal variability that are likely to be affected by enhanced resolution. For example, our proposed mechanism of NA SPG decadal variability suggests a prominent role for boundary currents, which may be improved by higher resolution (*Grotzner et al.*, 1998; *Gelderloos et al.*, 2011). Additionally, the increased atmospheric resolution (which represents the main computational burden for the coupled model) may affect the innate atmospheric variability over the North Atlantic (*Matsueda et al.*, 2009), while the role of the atmosphere may also be modulated by the improved ocean resolution (*Scaife et al.*, 2011). Recent work comparing 1°, 0.25°, and $^1/_{12}^\circ$ resolution simulations with the same underlying model suggest that, in the NA SPG, 0.25° is a significant improvement over 1°, but that there are still further improvements to be had at even higher resolution (*Marzocchi et al.*, 2015). In short, although the variability simulated in HadGEM3 does not appear to be identical to that which exists in reality, there are reasons — such as its representation of the location of key dense water formation sites and the boundary currents that supply these — to suppose it may be closer to emulating the real world than previous modelling work using lower resolution (in both the ocean and atmosphere) models.

Finally, we note that, while the observational density of temperature and salinity obser-

vations in the North Atlantic is growing (*Good et al.*, 2013), there are still significant gaps in the observation network. Perhaps the most significant of these are high quality observations of surface heat fluxes between the ocean and atmosphere (notably missing in Figure 3.3). Reanalysis products are not yet able to reliably simulate surface heat fluxes (*Josey*, 2001) and yet the variability in these heat fluxes is crucial in understanding the coupling between the atmosphere and ocean (*Gulev et al.*, 2013). Aside from direct heat flux observations, greater knowledge of the transports of mass, heat, and freshwater in the northern edge of NA SPG would be helpful in order to characterise the magnitude of variability, particularly in the heat content anomalies that are important in the simulations with HadGEM3. The recently begun Overturning in the Subpolar North Atlantic Program (OSNAP[2]) may help to address these issues. Lastly, merely sustaining the present day observational density is of critical importance and we note that, even in the late twentieth century, this was not guaranteed (*cf.* Figure 3.1b).

3.7 Chapter conclusions

We have compared the simulated mode of decadal variability in the North Atlantic subpolar gyre (NA SPG) in HadGEM3 (described in Chapter 2) against observational analyses of surface fields and subsurface ocean fields in this region.

- The interannual evolution of sea surface temperature (SST) and mean sea level pressure (MSLP) fields is broadly similar between HadGEM3 and EN4, with warm temperatures in the NA SPG associated with negative North Atlantic Oscillation (NAO) anomalies (Section 3.3.1). However, using the full period 1900–2014, the implied propagation of anomalies in either SST or top 500m depth averaged temperature (T500) fields is not as clear in EN4 as in HadGEM3.
- Using the shorter period 1960–2000 (to be consistent with the analysis of *Robson et al.*, 2012) yields qualitatively different evolution of observed ocean/atmosphere anomalies (Section 3.3.1), compared to the (significantly more infilled) period

[2]http://www.o-snap.org/ (October 2015)

1900–2014. This evolution is more consistent with that exhibited by the simulations with HadGEM3. It is not clear whether this represents either 1) the relative increase in noise when using a much shorter time series, 2) a related reduction in the damping effect of heavily infilled data, or 3) a more nuanced distinction between 'average' decadal variability (which may only explain a small amount of the variance, and could be internally forced) and particular decadal 'events' (which may episodically explain large amounts of variance and may be externally forced *Robson et al.*, 2012).

- Data paucity, even in the most well observed variables such as SSTs, inhibits direct comparison of lagged relationships between different locations in the NA SPG (Section 3.3.2). This is further complicated by the unknown magnitude of the contribution of externally forced variability in the observed record, and to what extent this will result in similar or different ocean feedbacks.

- Analysis of the depth coherence of interannual variability within the NA SPG suggests a shallower extent in reality than in HadGEM3, particularly in the North Atlantic Current (NAC) region (Section 3.4). The larger extent in HadGEM3 may help to explain the increased efficacy of signal propagation around the NA SPG.

- Consistent with the above, the propagation pathways of temperature anomalies around the NA SPG appear to occur over a shallower depth range in EN4 than in HadGEM3 (Section 3.5.2). The timescales of this signal propagation are also longer in EN4 (though specific lags cannot be attributed, see Section 3.5.2).

- The feedback *process* between Labrador Sea and NAC temperatures (as simulated in HadGEM3, Chapter 2) may exist in reality but does not appear to be negative (Section 3.5.1), *i.e.* it acts as a positive feedback in reality but a negative feedback in HadGEM3. Specifically, in the Labrador Sea, the dominant driver of interannual density variability is different in EN4 (salinity) than in HadGEM3 (temperature) consistent with the sign of lagged relationships between depth averaged temperatures in the eastern NA SPG and NAC region (negative in HadGEM3, positive in EN4, see Section 3.5.3).

- Despite the above statement, it is not clear that in reality the interannual driver of

Labrador Sea density variability is stationary (Section 3.5.1). In recent decades, the relative contributions of temperature and salinity to this density variability appear to have inverted, with the most recent two decades implying a larger role for temperature-driven variability in EN4, similar to the simulations. It is not clear whether this represents a transient or permanent role-reversal and warrants further investigation.

- Some of the disparities between model and observations may be attributable to the overly prominent role of Labrador Sea water in the NA SPG in HadGEM3, which is overly extensive and homogenises upper and lower North Atlantic deep water (Section 3.6).

The most crucial difference between HadGEM3 and observations, in terms of the veracity of the simulated decadal variability in HadGEM3, is the absence in observations of the simulated *negative* feedback between Labrador Sea and NAC temperatures in the observed record. This feedback, and the conditions in which it may exist, is investigated further as part of the multi-model analysis in Chapter 6. Despite this crucial difference, various other elements of the simulated variability, such as the evolution of surface fields and the propagation of anomalies around the NA SPG, are detectable in reality, albeit with different temporal or spatial structures. Having understood these differences it may still be possible to combine simulated variability in the model (or a version thereof: *MacLachlan et al.*, 2015) with this knowledge to make useful decadal predictions by, for example, post-processing prediction simulations to account for model-observation discrepancies, or by giving certain simulated predictions more/less confidence depending on the processes by which those predictions arise. This will be discussed further in the context of an initial condition ensemble in Chapter 5. However, in the next chapter we further investigate the asymmetry between phases of the simulated mechanism of decadal variability.

Chapter 4

Asymmetry in the simulated variability

4.1 Introduction

In Chapter 2 it was noted that the timescales of opposing phases of the proposed mechanism of bidecadal variability in the North Atlantic subpolar gyre (NA SPG) may be different. The timescale of reversal following anomalously high NA SPG sea surface temperatures (SSTs) is around 9 years, somewhat longer than the timescale of reversal following anomalously low NA SPG SSTs at around 7 years. Additionally, Figure 2.15 highlights the different mean sea level pressure (MSLP) patterns associated with anomalously positive or negative SSTs in the NA SPG in the coupled simulation, suggesting possibly different characteristics of atmosphere-ocean interactions in opposing phases of the variability. In this chapter, we present further evidence of asymmetry in the coupled simulation in Section 4.2 before describing targeted decoupled experiments to test the possible causes of this asymmetry in Section 4.3. We discuss the similarities between these coupled and uncoupled results in Section 4.4 before chapter conclusions in Section 4.5.

4.2 Evidence of asymmetry in the coupled control

In this section we describe the method we use to diagnose asymmetry in the coupled model (Section 4.2.1) before presenting evidence of spatial and temporal asymmetries in MSLP composites (Section 4.2.2) and other fields (Section 4.2.3) and discussing the possible causes of this in the coupled control simulation (Section 4.2.4).

4.2.1 Method: composite analysis

In Chapter 2 we presented (lagged) linear regression analysis, similar to many previous studies of climate variability (*cf.* Chapter 1 and schematic Figure 1.5). However, this method of analysis, which uses all of the available data to maximise the number of degrees of freedom (DoF), forces the subsequent results to be symmetric, *i.e.* the negative phase of any oscillation is, by construction, the inverse of the positive phase. For the case where a system evolves differently (that is, asymmetrically) in positive or negative phases, these differences/asymmetries will be averaged out, possibly resulting in spurious diagnoses of the spatial or temporal evolution of the variability. This could occur if either the dependent or explanatory variables are non-normally distributed, or if the relationship between the two is non-linear (*Allison et al.*, 2015).

An alternative to linear regression is to construct composites of one field/time series based on a particular index representing a subset of the data. For example, constructing a composite field of MSLP based on the top 10% of the NA SPG SST index (Figure 2.15a) and comparing to a composite field based on the bottom 10% of the same index (Figure 2.15b) to independently determine the relationships in high/low SST phases. This can also be extended to lagged composites — similarly to lagged regressions. The composite approach has the advantage of no longer assuming *a priori* that the two phases are symmetric in space or that they evolve symmetrically in time (although it does assume that the data within each composite is representative of the same process). However, the major disadvantage is the reduction in the effective degrees of freedom by a large factor: in linear regression 100% of the data is used to create a single regression slope, whereas with composites a small percentage (for example, 10%) of the data is used for each of the

two composites. Table 4.1 summarises the advantages and disadvantages of regression and composite analysis.

In this chapter we use composites based on the NA SPG SST index used in Chapter 2 (Figure 4.1 and against which lagged regressions are shown in Figure 2.7), which is the bandpass filtered (5–70 years) basinwide average SST in the North Atlantic between 45–65°N from the coupled control simulation. The highest and lowest 10% of index values are used to construct anomaly composites of various fields at a variety of lags, with anomalies referenced to the time mean of the specific field from the coupled control simulation. The SST index does not appear significantly different from a normal distribution and is mostly symmetric (skewness of -0.06).

4.2.2 Asymmetry in MSLP composites

MSLP anomaly composites based on the NA SPG SST index are shown in Figures 4.2 (for positive SST anomalies) and 4.3 (for negative SST anomalies) and highlight the asymmetrical relationship between SSTs and MSLP in opposing phases of the decadal variability. Note that the composited MSLP anomaly associated with the lowest 10% of SSTs in Figure 4.3 is shown on an inverted scale to aid visual comparison with the composited MSLP anomaly associated with the highest 10% of SSTs (Figure 4.2). In both sets of composites, lags of -1 and 0 are associated with strong North Atlantic Oscillations (NAOs, defined as the difference in wintertime MSLP over the Azores and Iceland) — a negative NAO associated with positive SSTs (Figure 4.2, lag=-1, 0) and a positive NAO associated with negative SSTs (Figure 4.3, lag=-1, 0, note inverted scale). However, the amplitude of these two in-phase NAOs are quite different, with an in-phase NAO anomaly associated with positive SSTs of -1.7hPa (-3.6hPa/°C) and an in-phase NAO anomaly associated with negative SSTs of 3.1hPa (-6.1hPa/°C). The amplitude of these anomalies is significantly different at the 99% level for a two-tailed t-test.

In addition to the difference in strength of the in-phase NAO anomalies and overall North Atlantic MSLP field, there are also differences in the temporal evolution of these anomalies. The maximum opposite-sign NAO anomaly preceding the lag=0 anomaly occurs at

	Linear regression	High/low Composites
DoF assuming *i.i.d.* data	$(n-2)$	$(n \times 0.1)$ (for 10% composites)
Implicit assumptions about processes	Data towards the middle of the distribution (near zero anomaly) is a low-weighted version of data towards the extrema, *i.e.* all data represents the same single process but at different scalings	Data within each composite represents the same process
High and low phases	Symmetric by design	Can be asymmetric
Effect of outliers	Outliers have little effect with a large enough sample and can be easily removed	Outliers can have a large effect that can be reduced by increasing the composite size at the expense of assuming increasing symmetry
Units and interpretation of amplitude	[dependant]/[explanatory], *e.g.* hPa/K, can be scaled given any value of the explanatory variable	[dependent] *e.g.* hPa, amplitude is a function of the composite size, affected by the relative weight of outliers

Table 4.1: Advantages and disadvantages of linear regression and composite analysis. DoF: Degrees of freedom, *i.i.d*: independent and identically distributed, n: Number of data points/years.

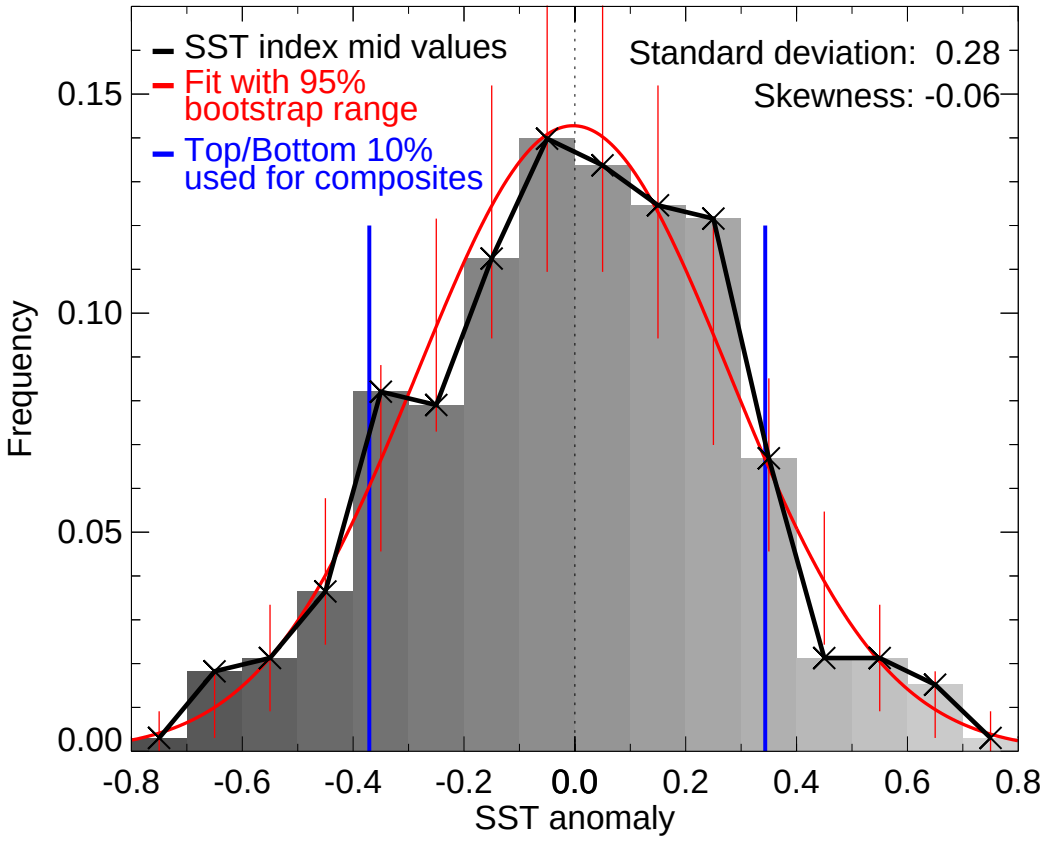

Figure 4.1: Histogram showing the distribution of the SST index (bandpass filtered (5–70 years) basinwide average SST in the North Atlantic between 45–65°N from the coupled control simulation). A normal distribution with the same mean and standard deviation (normalised to have the same area under the curve) is shown in red. Given that the simulated SST index uses a finite number of data points, a bootstrap approach (*Wilks*, 1997) is used to estimate the 95% confidence intervals, indicated by the vertical red bars. Blue lines are drawn at the SST anomaly values that occur at 10% and 90% of the total data *i.e.* the top and bottom 10%.

a lag of -7 years for the composite using positive SSTs (Figure 4.2, lag=-7). That is, prior to the in-phase relationship between high SSTs and an anomalously negative NAO, is an anomalously positive NAO 7 years previously. This reversal timescale, from positive NAO to negative NAO, deduced from composites created using high SSTs, is also consistent with the timescale for the same transition instead deduced from composites created using low SSTs (Figure 4.3, lag=+7): Noting the reversed scale in Figure 4.3, it can be seen that following a positive NAO in phase with a negative SST anomaly (Figure 4.3, lag=0) is a negative NAO (Figure 4.3, lag=+7). Thus, using either 1) MSLP composited on to the highest 10% of SST index anomalies, or 2) MSLP composited on to the lowest 10% of SST index anomalies, yields a reversal timescale from positive NAO to negative

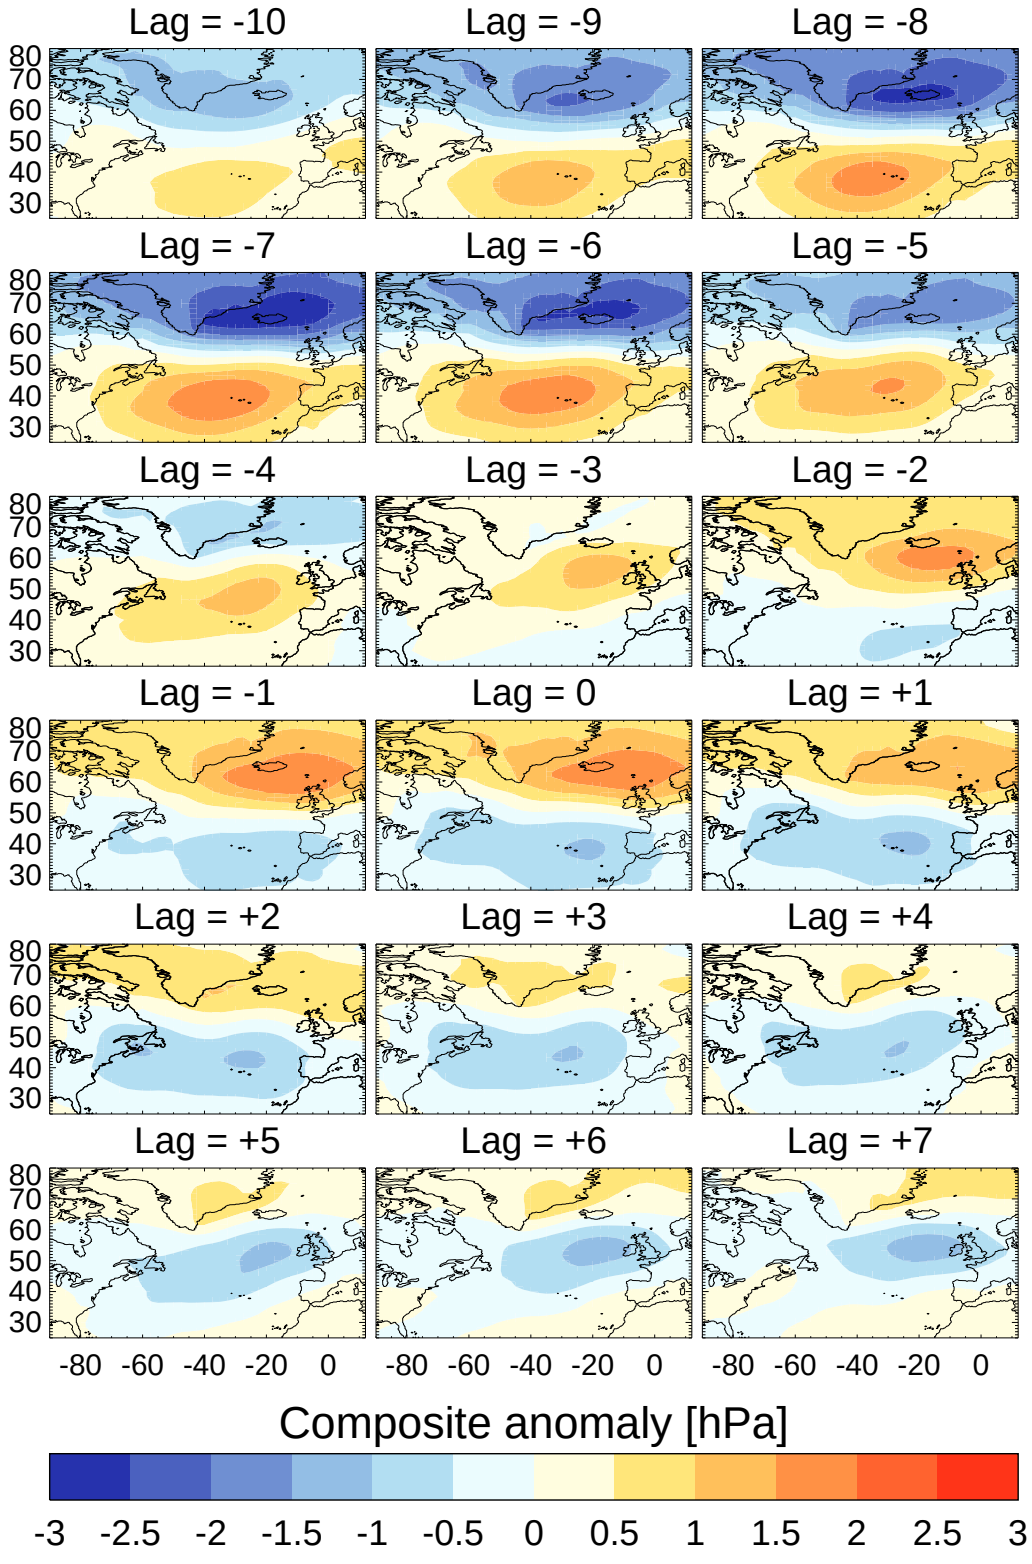

Figure 4.2: Lagged composite plot of wintertime (seasonal mean over December, January and February) MSLP composited from the highest 10% of NA SPG SST index values (5–70 year bandpass filtered basin-wide average SSTs in the North Atlantic between 45–65°N) at various lags. Anomaly with respect to the time mean wintertime MSLP. SST index lags then leads the MSLP field.

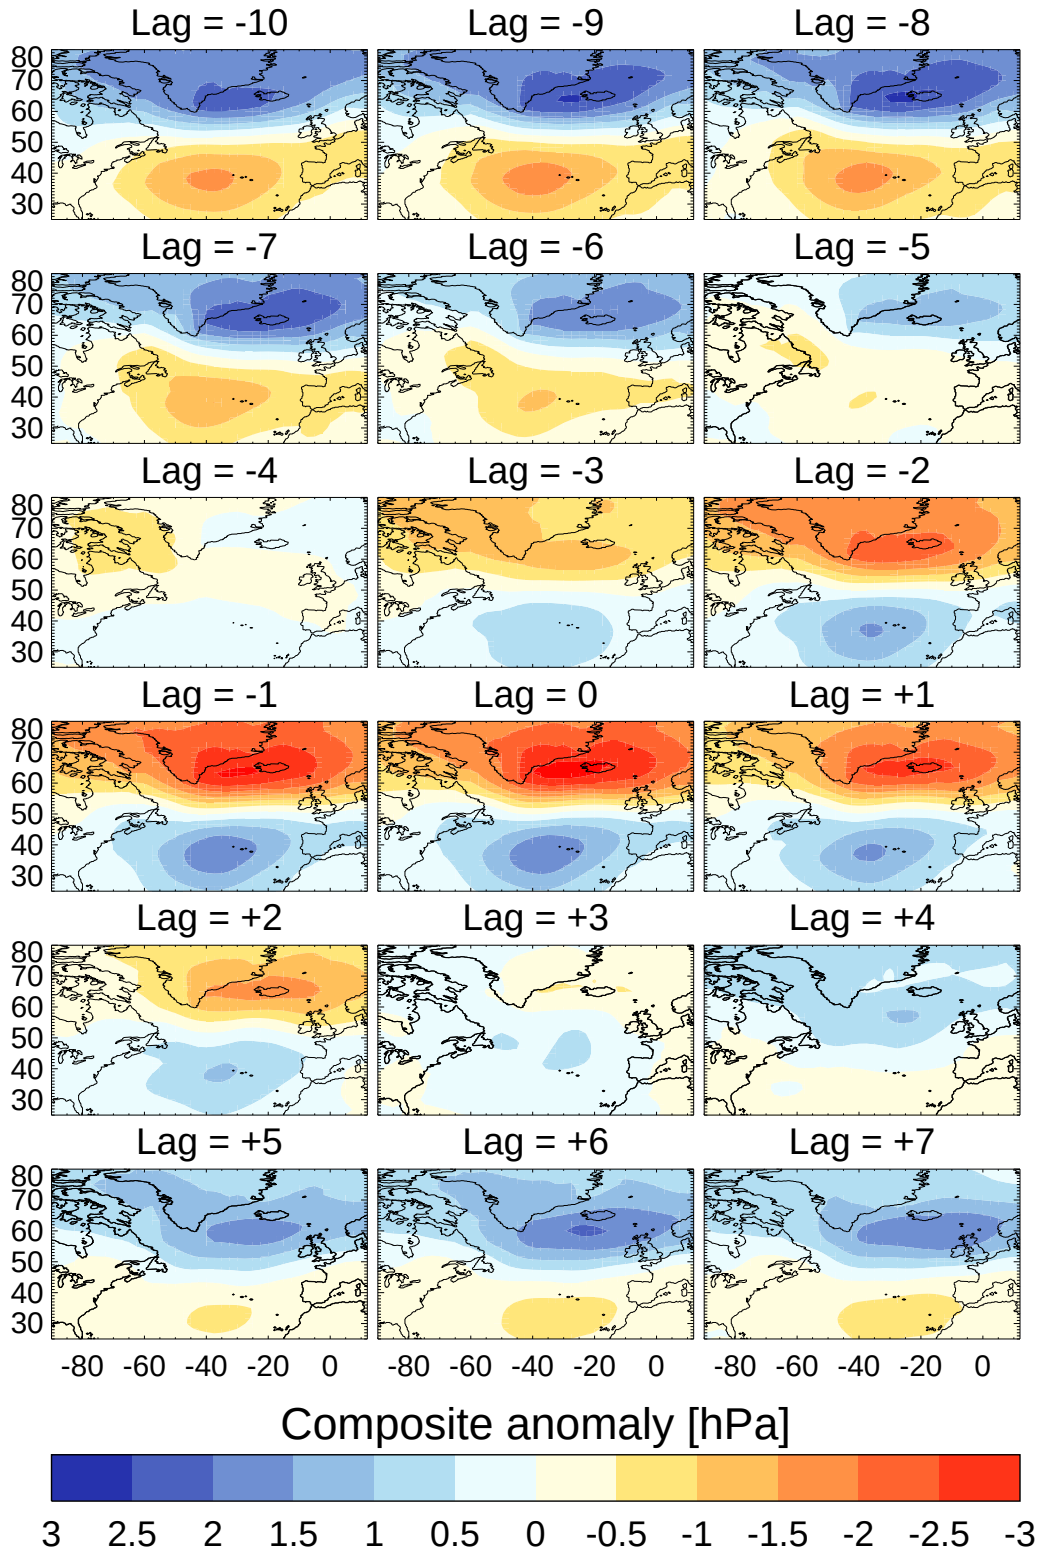

Figure 4.3: Lagged composite plot of wintertime (seasonal mean over December, January and February) MSLP composited from the lowest 10% of NA SPG SST index values (5–70 year bandpass filtered basin-wide average SSTs in the North Atlantic between 45–65°N) at various lags. Anomaly with respect to the time mean wintertime MSLP. SST index lags then leads the MSLP field. Scale has been inverted to aid comparison with Figure 4.2.

NAO of 7 years.

The reversal timescale from negative NAO to positive NAO can be estimated in a similar sense to the above. However, instead of transitioning from a negative to a positive NAO anomaly, the MSLP signal in Figure 4.2 merely tends to zero at increasing lags (Figure 4.2, all positive lags), and is still small at lags greater than +7 (not shown). As such, it is not possible to estimate the timescale of reversal from negative NAO to positive NAO using the lagged composites created using the highest 10% of SSTs. Nevertheless, analysing instead the composites created using the lowest 10% of SSTs (Figure 4.3), we find that the maximum negative NAO preceding the in-phase (positive NAO) relationship occurs at a lag of 9 years (Figure 4.3, lag=-9). Thus, from MSLP composited on to the lowest 10% of SST index anomalies yields a reversal timescale from negative to positive NAO of 9 years.

In summary, the analysis of MSLP composited on to the NA SPG SST index implies a total cycle timescale of around $7+9 = 16$ years, consistent with the 16/17 year periodicity in many NA SPG indices, *cf.* the power spectra for ocean and atmospheric variables in Figure 2.4. We now briefly investigate the manifestation of this asymmetry in related North Atlantic variables before describing experiments to diagnose its origin.

4.2.3 Asymmetry in other fields

Given the asymmetry between MSLP and SSTs in terms of both 1) the magnitude of the in-phase NAO and 2) the timescale of transitions between anomalous NAO states, it is reasonable to expect further asymmetrical relationships in related fields. Figure 4.4 shows net surface heat flux (SHF; positive values are directed into the ocean) composites, constructed as described in Section 4.2.1, following anomalously high values of the SST index (left column) and anomalously low values of the SST index (right column, note again the inverted scale to aid visual comparison). The SHF anomaly over the NA SPG following warm SSTs shows a weakening. At lags of +5 to +7 years the largest remaining anomalies are related to a shift in the position of the NAC, the mean position of which is highlighted (see also Section 4.2.4). This is in contrast to the SHF following cool SSTs

(right column, inverted scale) where inverse SHF anomalies build up in the NA SPG (initially in the eastern SPG at lags +3 to +5 and in the Labrador Sea at lags +5 to +7), implying subsequent warming by the SHFs. These SHF responses are consistent with the MSLP responses shown in Figures 4.2 and 4.3 in which the relationship between a positive NAO and cool SSTs is stronger than the negative NAO/warm SSTs and shows a more rapid phase reversal.

Similarly to the SHF fields, there is also asymmetrical evolution in top 500m depth averaged temperatures (T500) in the NA SPG (Figure 4.5). The T500 field was analysed extensively as part of the mechanism of decadal variability diagnosed in HadGEM3 (Chapter 2) and the symmetric lagged linear regression against the same SST index is shown in Figure 2.7 (column 2), along with SHF (column 3) and MSLP (column 5). Using the same time window as for SHFs, the T500 composites show the reversal from anomalously warm temperatures (left column) to cool temperatures, and from anomalously cool temperatures (right column, note inverted scale) to anomalously warm temperatures. At lags of +3 and +4 years (*i.e.* 3 or 4 years after maximal values of the SST index) the high and low composites are very similar. However, at increasing lags it can be seen that the temperature anomaly in the eastern NA SPG is stronger and more extensive in the low composites. That is, after anomalously cool temperatures in the NA SPG SST index, a subsequent opposite-sign anomaly builds up more strongly than in the inverse phase of the mechanism. As above, this appears consistent with larger SHF anomalies in the eastern SPG and a larger MSLP signal following the lowest 10% of SST index values. Conversely, the T500 signal in the Labrador Sea is not so different in opposite phases, despite the differences in the SHFs seen there, perhaps suggesting a role for ocean processes to remove the additional SHF, either horizontally or vertically. In general, despite the differences in opposite phases of the T500 composites, the magnitude of the difference between the evolution of T500 high and low composites is not as clear as for MSLP and SHF. This likely reflects the competing contributions to T500 of atmospheric and oceanic processes. Indeed, as was shown in Chapter 2, advection by both the mean and anomalous ocean circulation is important for the build-up of NA SPG T500 anomalies.

Given the diagnosed asymmetry in the relationship between SSTs and MSLP, as well as other variables, we next ask how that asymmetry is formed.

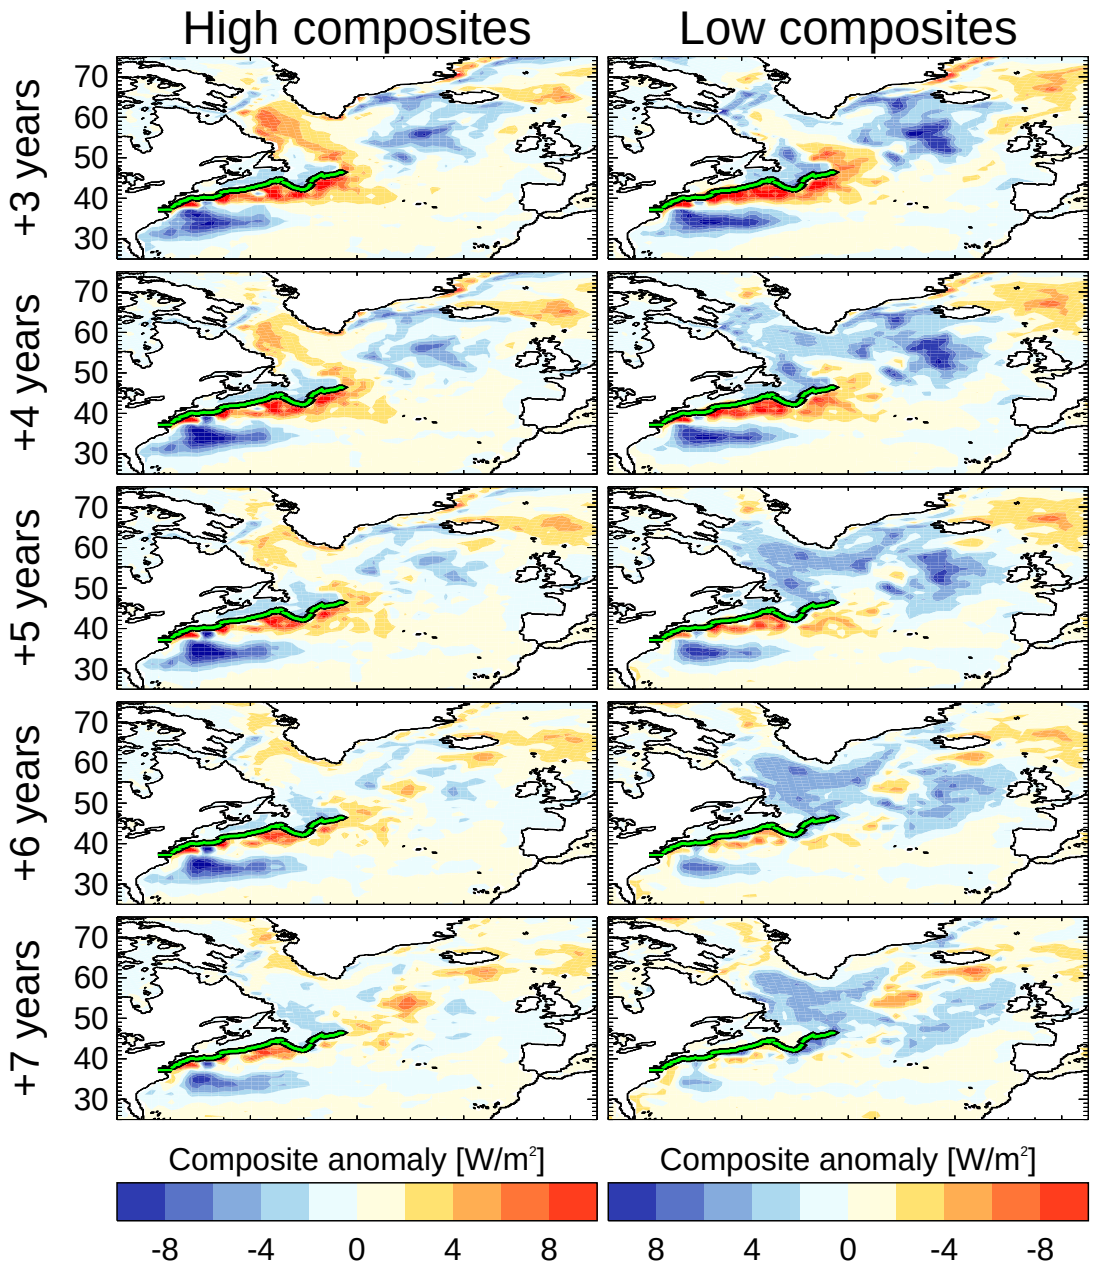

Figure 4.4: Lagged composite plot of net surface heat flux (SHF, into ocean) composited from the highest 10% of NA SPG SST index values (5–70 year bandpass filtered basinwide average SSTs in the North Atlantic between 45–65°N, left) and lowest 10% (right, scale inverted) at various lags. Anomaly with respect to the time mean net SHF. SST index leads the SHF field. The mean location of the North Atlantic Current is estimated by the location of the maximum time mean SST gradient (green).

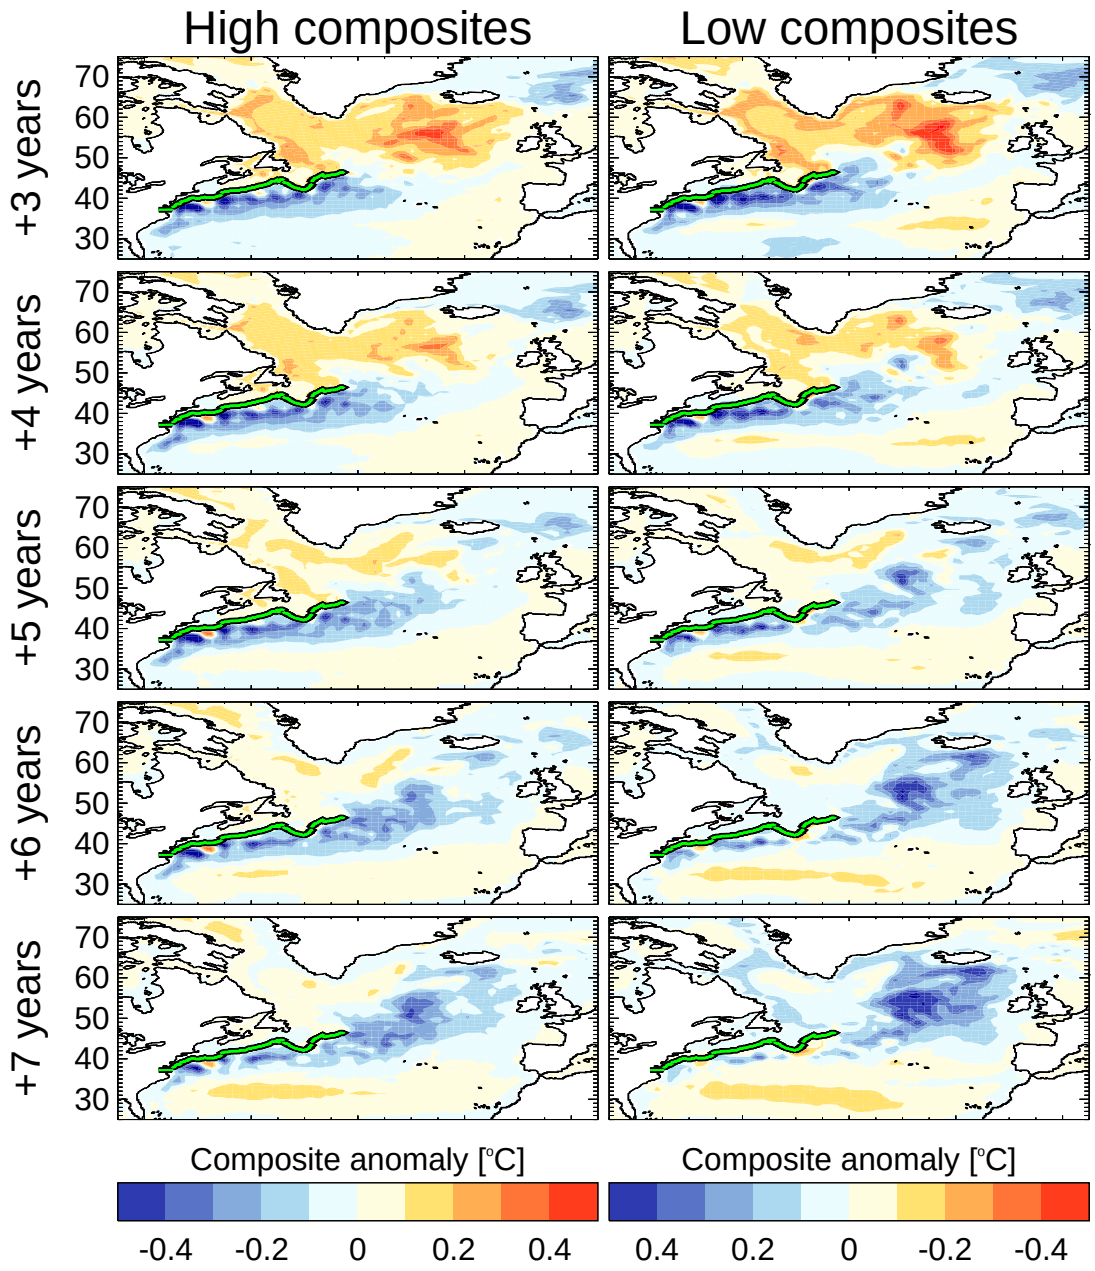

Figure 4.5: Lagged composite plot of top 500m depth averaged temperature (T500) composited from the highest 10% of NA SPG SST index values (5–70 year bandpass filtered basinwide average SSTs in the North Atlantic between 45–65°N, left) and lowest 10% (right, scale inverted) at various lags. Anomaly with respect to the time mean T500. SST index leads the T500 field. The mean location of the North Atlantic Current is estimated by the location of the maximum time mean SST gradient (green).

4.2.4 Origin of the asymmetry

The asymmetry between opposing phases of the decadal variability in the NA SPG is evident in sea level pressure, associated surface heat fluxes, and their integral in depth averaged temperatures. In the proposed mechanism of decadal variability (see Figure 2.13 for a schematic) there are two important non-linear processes that are prime candidates for the origins of this asymmetry. These are: 1) the role of convection, a highly non-linear process, in mixing down near surface density anomalies in and around the Labrador Sea, and 2) the importance of circulation anomalies (*e.g.* $v'(T_0 + T')$) for the advective heat transport in the southern edge of the NA SPG.

The mixed layer depth (MLD) composites, constructed using the same method as in previous figures, are shown in Figure 4.6. Consistent with the net surface heat flux composites (Figure 4.4) they show a stronger relationship with negative SST anomalies (right column) with the initial MLD anomaly in the Labrador Sea/Irminger Current (the location of the climatological mixed layer depth variability, Figure 2.1) able to reverse sign between lags +3 to +7. This is in contrast to the MLD anomaly following positive SST anomalies (left column) in which the initial MLD anomaly merely tends to zero throughout this time. This implies that some of the asymmetry in the timescales between phases of the mechanism of decadal variability may be related to asymmetries in the MLDs, perhaps related to the physical limits of mixed layer depth variability (*i.e.* the bathymetry and the ocean surface). However, there is a symmetric relationship between the MLDs (or net surface heat flux) and composites based instead on an NAO index (not shown). This suggests that the NAO is related to similar strength heat fluxes and subsequent mixed layer depth/convection anomalies and would imply that the asymmetry must come from the oceanic response to these symmetric anomalies, perhaps through the circulation.

The time mean top 500m depth averaged zonal current speed (U500) in the North Atlantic is shown in Figure 4.7a. The North Atlantic Current (NAC) can be clearly seen after the Gulf Stream detaches from the coast between 35–40°N. An index of the NAC is defined as the top 500m volume averaged zonal mean current in the marked box, which has a mean value of 5.85cm/s. Composites of this index, as described in Section 4.2.1, are then created based on the NA SPG basinwide SST index and are shown in Figure 4.7b.

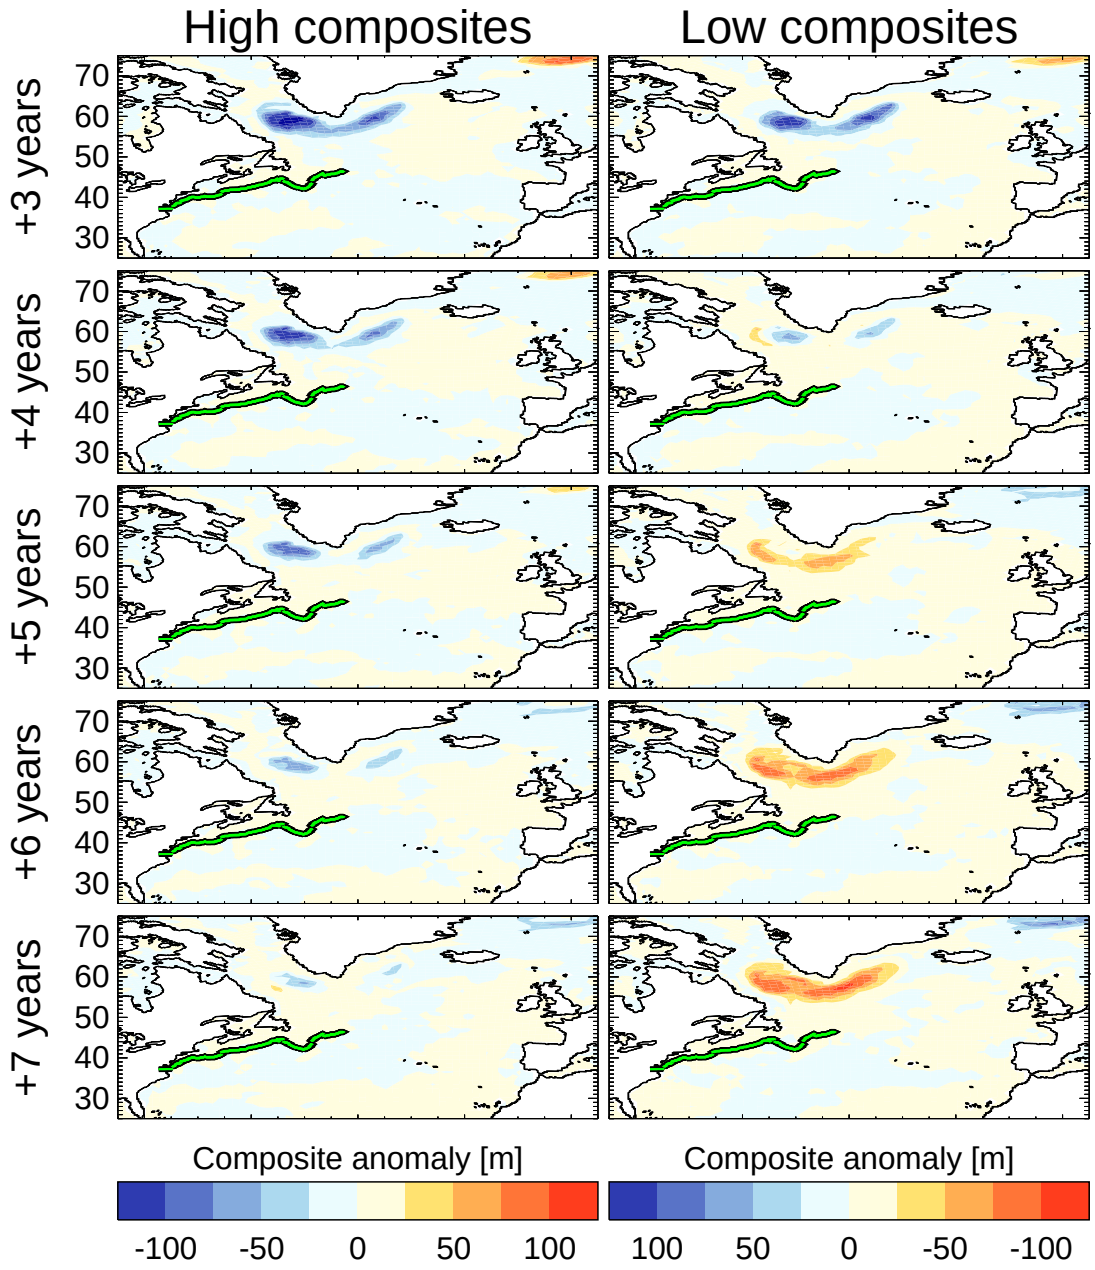

Figure 4.6: Lagged composite plot of March-time mixed layer depths (MLDs) computed using the methodology of *Kara et al.* (2000) composited from the highest 10% of NA SPG SST index values (5–70 year bandpass filtered basinwide average SSTs in the North Atlantic between 45–65°N, left) and lowest 10% (right, scale inverted) at various lags. Anomaly with respect to the time mean MLD. SST index leads the MLD field. The mean location of the North Atlantic Current is estimated by the location of the maximum time mean SST gradient (green).

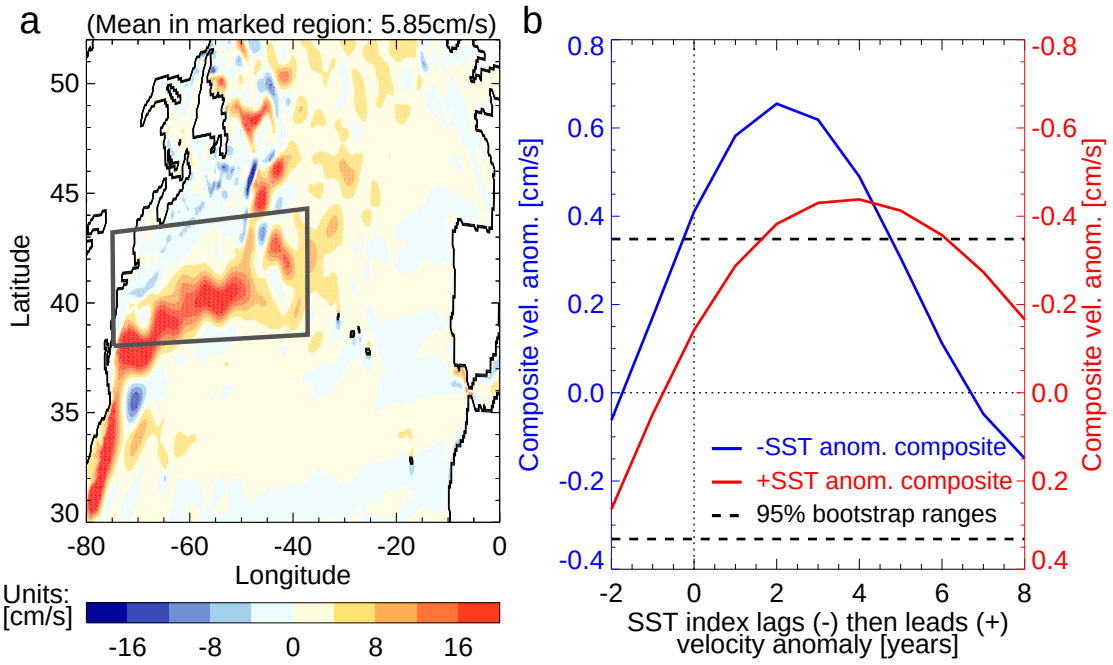

Figure 4.7: a) Time mean top 500m depth averaged zonal velocity (*i*-direction, which is approximately zonal at these non-polar latitudes, see Chapter 1, Section 2.3 for details of the curvilinear grid) in the coupled control simulation. The dark grey box demarks the North Atlantic Current (NAC) region for which composites are shown. b) Lagged composite NAC top 500m volume averaged zonal current speeds in the marked region, composited from the SST index used in Figure 2.7 and Figures 4.2–4.6 using the highest 10% of SST index values (+SST, red) and the lowest 10% of SST index values (-SST, blue). Climatological mean current speed in this region is 5.85cm/s. Note the inverted scale for the velocity anomaly following +SST. Confidence intervals (black, dashed) are estimated use the moving blocks bootstrap technique of *Wilks* (1997), using a block length of 4, resampling the original data 40000 times.

Consistent with the proposed negative feedback (see Chapter 2 and Section 2.4.6), anomalously warm temperatures in the NA SPG lead to a negative circulation anomaly in the NA SPG region whilst anomalously cool temperatures lead to a positive circulation anomaly (Figure 4.7b, note the inverted scale for the high SST index composite). The spatial *patterns* of the circulation anomalies (not shown) are symmetrical and depict an anomaly located on the southern edge of the NAC but the magnitude and temporal evolution are not symmetric (Figure 4.7b). Even if these circulation (and subsequent heat content) anomalies were of equal (but opposite) magnitude it would imply different timescales of anomaly propagation in the NAC region due to the inverse relationship between speed and time, *e.g.*

$$OHC'_{\pm} \propto (v_0 \pm v')T' \times t_{\pm} \tag{4.1}$$

where notation is as in Section 2.4.1. OHC_+ is only equal to OHC_- in the trivial case of v'=0 or when t_+ is not equal to t_-. In such a situation, it could be said that symmetric forcing (of the circulation anomaly, v') led to an asymmetric response. However, as the circulation anomalies in either composite are not of comparable magnitude we conclude that an asymmetrical forcing, of unknown origin, leads to an asymmetrical response. The velocity anomaly per degree of the SST high/low composite is ≈1.2cm/s/°C for the negative SST phase, and ≈0.8cm/s/°C for the positive SST phase (using the peak anomalous velocities from Figure 4.7b), suggesting that some of the asymmetry may indeed arise from an asymmetrical ocean response to the SST changes. However, the difference in timescales between the two composites, estimated by assuming a length scale that is the width of the box (3000km) and a velocity anomaly scale that is the peak anomalous velocity (0.6cm/s for the high composite, -0.4cm/s for the low composite), is 4 months. As noted in Chapter 2, the location where heat transport anomalies transition from being dominated by the anomalous to the mean circulation is not clear. In addition, we have neglected the potential role for the temperature anomalies (which, in the absence of compensating salinity anomalies, will also be potential vorticity anomalies) to interact with the circulation. As such, this likely represents an upper estimate on the asymmetry due to the anomalous circulation and is much less than the ≈2 year difference in phase-reversal time suggested by the MSLP composites (*cf.* Figures 4.2 and 4.3), suggesting other processes must also exist to explain the asymmetry.

In addition to the larger magnitude of the anomalous velocity following cool SSTs this velocity anomaly also subsequently subsides more quickly. The anomaly following cool SSTs becomes 'insignificant', as defined by the moving blocks bootstrap estimate of significance at the 95% level, within 5 years, compared to the 6 year timescale for the weaker but longer lasting anomaly following warm SSTs. This can also be seen in Figure 4.4, in which the surface heat flux anomalies in the NAC region (marked in green), indicative of a slight shift in the latitude of the NAC, are longer lasting following anomalously warm SSTs (left column) than following anomalously cool SSTs (right column).

In summary, there is evidence of asymmetry within the proposed mechanism of simulated decadal variability in the NA SPG. The asymmetrical timescales for reversal between positive and negative phases (phases characterised by basinwide SSTs in the NA SPG) cannot solely be attributed to the non-linear processes within the ocean that we have investigated (*e.g.* convection or the role of circulation anomalies). However, given the lag=0 asymmetry in MSLP composites based on this SST index (*cf.* Figures 4.2 and Figure 4.3, lag=0), which to some degree characterise the strength of the relationship between SSTs and MSLP/the NAO, it is possible that the overall asymmetry between the opposing phases of the variability is related to asymmetries in the strength of ocean-atmosphere coupling. That is, one phase of the mechanism can elicit a larger MSLP response in the atmosphere, which amplifies or accelerates that part of oscillation. We examine this hypothesis in the next section.

4.3 Atmosphere-only experiments

We begin this section with the details of some forced atmosphere-only experiments (Section 4.3.1) and briefly discuss the choice of a control baseline (Section 4.3.2) before presenting the results in Section 4.3.3.

4.3.1 Experimental design

To investigate possible asymmetries in the strength of ocean-atmosphere coupling we design and run a set of atmosphere-only ensembles, *i.e.* using just the atmosphere component of the coupled model, summarised in Table 4.2 and described next. The atmosphere model is forced with annually-repeating daily mean SST and sea ice fields, linearly interpolating the daily mean values on to the atmosphere time step of 15 minutes. Note that there is no representation of the diurnal cycle in either the forcing fields or in the atmosphere simulations. The individual daily mean climatological forcing fields are calculated from the coupled model by separately averaging each day over a 20 year period towards the end of the simulation, *i.e.* averaging all firsts of January to create a mean

Experiment name	+NAO	-NAO
Forcing fields	Daily mean SSTs and sea ice, repeated annual cycle (no diurnal cycle)	
Forcing fields construction	Composite wintertime SST and sea ice fields on to NAO index using:	
	highest 10% of wintertime NAOs from coupled control	lowest 10% of wintertime NAOs from coupled control
Ensemble members	7	7
Length (years)	30, 32, 40, 40, 41, 44, 45	26, 35, 35, 35, 38, 39, 44
Total years/winters	272	252

Table 4.2: Design of the atmosphere-only ensembles.

January 1st forcing field for both SST and sea ice. We use daily values from this shorter subset of the full simulation due to the availability of the correct diagnostics and computational overheads. Although it is difficult to validate the daily variability against the full simulation, the time mean NAO during this period is in a neutral state compared to the full simulation (not shown). Subsequently, the additional anomalous NAO-related SST and sea ice forcing are added to the climatological forcing fields to create the forcing fields for the '+NAO' and '-NAO' experiments. These NAO-related anomalies are estimated by creating composites of wintertime (December, January, February time mean) SSTs and sea ice using the highest (+NAO) and lowest (-NAO) 10% of wintertime NAOs from the whole time series, similarly to the composite construction detailed in Section 4.2.1. Wintertime values are used as the strength of the relationship between the ocean and the NAO is largest at this time (*Rodwell et al.*, 1999) but the anomalous forcing is applied constantly throughout the annual cycle. As such, the forcing fields are designed to represent an ocean state that is perpetually associated with a positive or negative NAO anomaly.

The SST and sea ice anomaly fields, which were added to the daily mean SST and sea ice climatology, are shown in Figure 4.8. Note that the fields are global rather than regional. This was to avoid making any *a priori* assumptions that the NAO is forced from anywhere in particular as evidence of bidecadal variability exists globally in proxy records (*Mann*

et al., 1995) and in our model (*e.g.* Arctic sea ice, not shown). That is, we did not want to damp any ocean to atmosphere forcing we might have seen in the atmosphere-only simulations by removing additional — but related — non-local forcing (see discussion in Section 4.4). Additionally, the use of global forcing fields significantly simplified the experimental design by avoiding the use of buffer regions between regions of climatological forcing and anomalous forcing. However, we note that global forcing may also have limitations (see Section 4.4). Some of the key features of the forcing fields are the tripole pattern in Atlantic SSTs (top two panels), consistent with the NAO, and a much stronger signal in the northern hemisphere than southern hemisphere. There is also a pattern reminiscent of El Niño/La Niña in the Pacific Ocean, the significance of which we discuss in Section 4.4. The sea ice fields, only shown for the northern hemisphere, yield a pattern of increased/decreased sea ice consistent with local decreases/increases in SST (bottom two panels). As expected, SSTs in the NA SPG are anomalously cool associated with a positive NAO (top left), and anomalously warm associated with a negative NAO (top right). The North Atlantic SST gradient between the subtropical and subpolar gyres (defined as the area mean SST in 2x2° boxes centred at 57°W, 29°N and 31°W, 53°N respectively) is similar in both sets of forcing fields at -1.03K in the +NAO ensemble and 1.00K in the -NAO ensemble.

4.3.2 Control baseline choice

Prior to analysing the atmosphere-only ensembles, it is first prudent to consider the choice of baseline against which to compare the output from the ensembles. There are several possible candidates for a control baseline, which are summarised in Table 4.3. These represent a choice between using the coupled control simulation, from which the ensembles are initialised, or separate atmosphere-only simulations, which should have a similar background climate. Additionally, there are choices over the time window to use, balancing an improved signal to noise ratio from a long time series against the computational expense and/or issues with drifts in the mean state. We compared our analysis of the atmosphere-only ensembles against baseline options 1–4 in Table 4.3 and found that the results were sensitive to the baseline choice (not shown). Given this sensitivity,

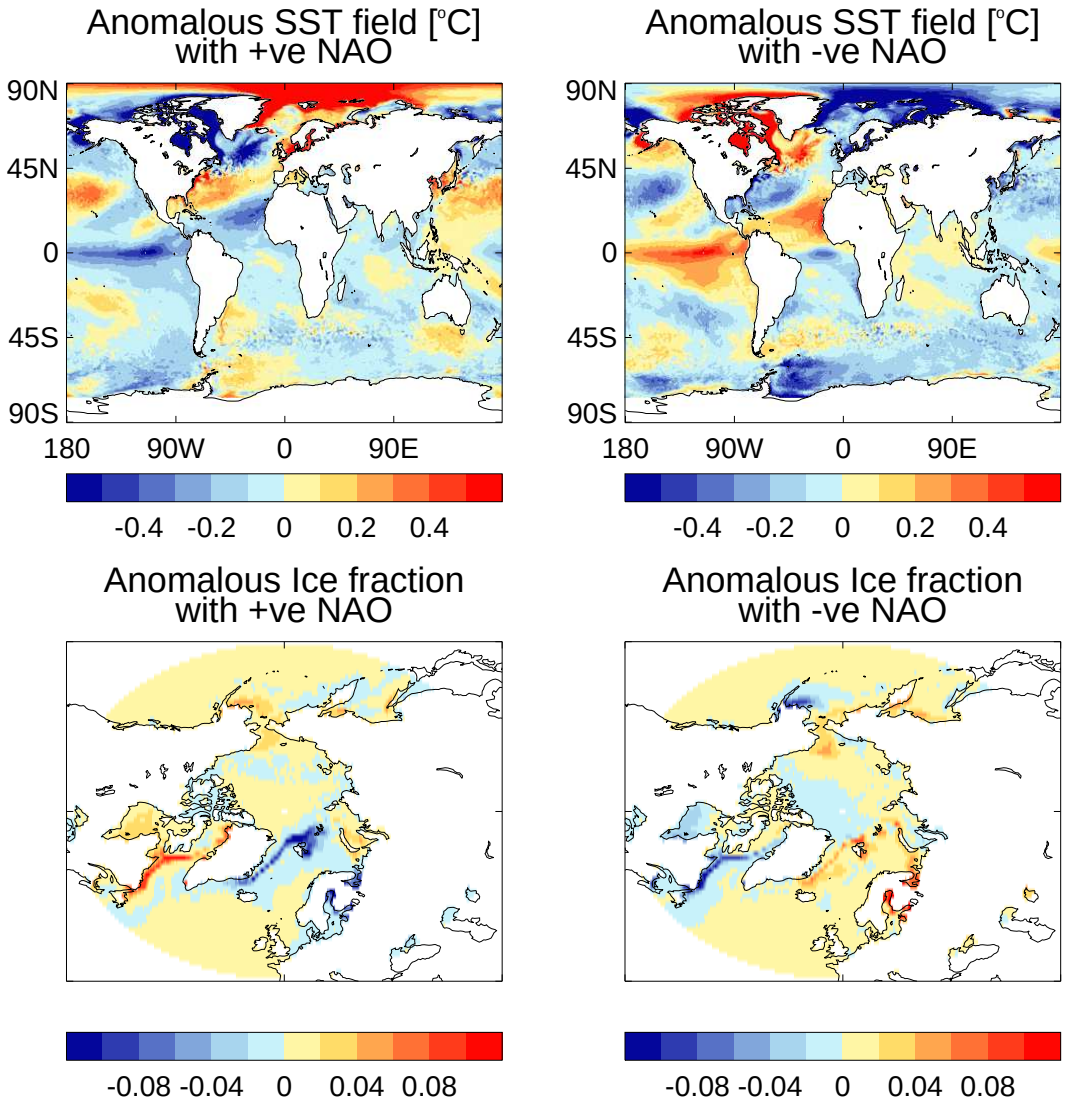

Figure 4.8: The anomalous SST (top) and sea ice (bottom, only shown for the northern hemisphere) forcing applied to the atmosphere-only ensemble, in addition to the daily-varying climatology as described in Section 4.3.1 and Table 4.2. Left column: Forcing fields associated with the highest 10% of NAO anomalies (+NAO). Right column: Forcing fields associated with the lowest 10% of NAO anomalies (-NAO).

we chose to use the computationally expensive, but most robust, baseline of a parallel set of atmosphere-only ensembles, identically forced to the +NAO and -NAO ensembles but forced with only climatological forcings (Table 4.3, number 5). Seven atmosphere-only control ensemble members were run, with simulation lengths of 26, 27, 27, 28, 28, 29, and 29 years, totalling 194 individual years/winters (*cf.* the 272 and 252 individual years/winters in the +NAO and -NAO ensembles respectively). This ensemble is hereafter referred to as the control ensemble.

Having defined the atmosphere-only ensemble experimental design, and our choice of control baseline, we now analyse the results of these ensembles. We ask: How does the atmosphere respond to the SSTs/sea ice associated with a positive or negative NAO anomaly? Does the atmosphere return a similar NAO anomaly, implying ocean to atmosphere forcing, or return no significant NAO anomaly, implying the relationship between simulated SSTs/sea ice and the NAO is due to atmosphere to ocean forcing? Is the strength of the atmospheric response similar for both +NAO and -NAO forcing, or is the response asymmetric, as suggested by the coupled analysis in Section 4.2?

4.3.3 Asymmetry in atmosphere-only ensembles

4.3.3.1 North Atlantic sea level pressure response

The time mean intra-ensemble wintertime MSLP response to the anomalous SST/sea ice forcing is shown in Figure 4.9 for the +NAO (left) and -NAO (right) experiments. It can be seen that the response to SSTs/sea ice associated with a positive NAO is to return an MSLP pattern that is also suggestive of a positive NAO (Figure 4.9, left). Conversely, the response of the atmosphere to SSTs/sea ice associated with a negative NAO is an MSLP pattern that is more reminiscent of the East Atlantic Pattern (EAP, the second most dominant mode of MSLP variability in the North Atlantic (*Barnston and Livezey*, 1987) after the NAO, Figure 4.9, right). In both cases, the signal in the northern NA SPG is not significantly different from the control ensemble, related to the large interannual/inter-ensemble variability in/across the control ensemble in this region. However, defining the NAO index as the difference in MSLP between the Azores and Iceland does reveal a

Control baseline	Advantages	Disadvantages
1. Full ≈500 years from coupled control	Long enough to average out other modes of variability. Coupled control uses 3-hourly coupling so does represent the diurnal cycle	Simulation shows some drift in SSTs/sea ice that will affect the MSLP. Background state in coupled and uncoupled modes cannot be assumed the same
2. 20 year parallel portion of coupled control	Short enough to remove drift effects and attempts to represent the same background climate as in the forced ensembles. Represents the diurnal cycle	Other modes of variability/noise become amplified relative to the signal. Background state in coupled and uncoupled modes cannot be assumed the same
3. 20 year atmosphere-only simulation forced by time varying SSTs/sea ice from parallel coupled control	Background climate state should be more similar to the forced ensembles. The SST and sea ice forcing will contain the effects of the diurnal cycle, which exists in the coupled control	Interannually varying SSTs/sea ice from the coupled simulation may contain the imprint of interannual NAO variability, which may project on to the time-mean response
4. Inter-ensemble mean	The background climate state is identical to the forced ensembles by design	Forces the results to be symmetric
5. Control ensemble forced with periodic SSTs/sea ice but without the addition of NAO-related SST/sea ice anomalies	Background state is the same as in the forced ensembles and effects of drift and/or noise affecting the signal are the same in both the control and forced ensembles	Many ensembles needed to get statistically significant results, which are computationally expensive to run. No representation of diurnal cycle - though this is also the case in the anomaly-forced simulations

Table 4.3: Advantages and disadvantages of various control baseline choices for the atmosphere-only forced ensembles, as described in Table 4.2.

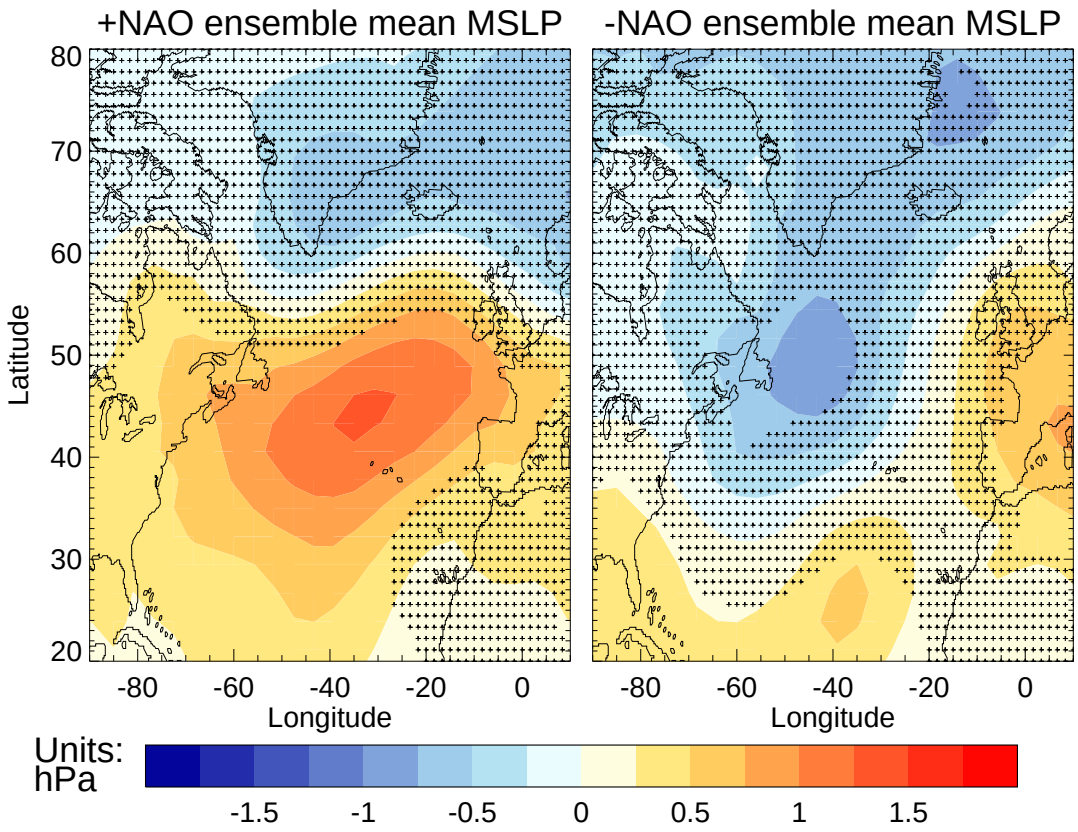

Figure 4.9: Ensemble mean wintertime MSLP response [hPa] to anomalous SST/sea ice forcing in the +NAO ensemble (left) and -NAO ensemble (right, scale inverted). Stippling denotes anomalies indistinguishable from the control ensemble at the 90% level for a two-tailed t-test.

significant difference between the control and the +NAO ensemble. The ensemble mean NAO is 1.35hPa stronger in the +NAO ensemble than the control ensemble, with a t-statistic of 1.86 that is significant at the 90% level. In the -NAO ensemble, the ensemble mean NAO is *not* significantly different from the control ensemble at the same level.

These results imply a stronger coupling between the wintertime NAO and wintertime SSTs in the positive NAO phase than in the negative NAO phase. That is, anomalously cool SSTs in the NA SPG (and associated sea ice changes) are more able to force anomalous atmospheric circulation in this region than can anomalously warm SSTs. This is consistent with the composite analysis of asymmetry within the coupled control simulation in Section 4.2 in which a stronger in-phase relationship between cool SSTs and a positive NAO was found than between warm SSTs and a negative NAO. The atmosphere-only ensembles suggest this stronger coupled relationship could be due to the ocean-to-atmosphere forcing being stronger during a positive NAO state.

It is not clear why the coupling should be stronger in one phase compared to the other. Using model output from the coupled control simulation, we find the net surface heat flux into the NA SPG is a linear function of the atmosphere-ocean temperature difference (not shown) using both annual and wintertime data. This is consistent with the bulk formula (*e.g. Large and Yeager*, 2004) in which the relationship is approximately linear for small temperature differences. It is possible that the asymmetries are introduced by the interaction between the atmosphere-ocean temperature difference and the specific humidity in the resultant latent heat flux. In addition, non-local forcing from the tropical Atlantic (*Sutton et al.*, 2000) or elsewhere (*Hoerling et al.*, 2001) could be important (allowed by our experimental design) both for the manifestation of the NAO signal and perhaps the non-linear/asymmetrical response (*Sutton et al.*, 2000). The possible role of non-local forcing is discussed in the next section and further in Section 4.4.

4.3.3.2 Global and upper troposphere responses

The forcing fields shown in Figure 4.8 reveal a Pacific SST anomaly of the same sign as the NA SPG SST anomaly. These SST anomalies co-vary such that the anomaly related to a positive NAO is concomitant with an SST anomaly consistent with a La Niña episode, and similarly for a negative NAO and El Niño. To highlight the global response, Figure 4.10 shows the ensemble mean MSLP anomalies in the atmosphere-only ensembles (top row) along with the relative magnitudes of these anomalies compared to the control ensemble variability (bottom row). It can be seen that there are also significant MSLP anomalies in the tropical Pacific in both ensembles. These are much smaller than the MSLP anomalies in the North Atlantic and North Pacific regions but exist in a region where the annual variability in MSLP is also much smaller, resulting in relatively large excursions from the variability described by the control ensemble (Figure 4.10, bottom). Thus, an important open question is to what extent the resultant anomalies in the NA SPG are due to forcing from the tropical Pacific.

To investigate the relationship between the tropical Pacific and North Atlantic we analyse the 500hPa geopotential height anomalies (Figure 4.11). In the North Atlantic region these are suggestive of a broadly barotropic response in the -NAO ensemble, whereas

there is increased baroclinicity in the +NAO ensemble, as evidenced by the large negative height anomaly over the Labrador Sea that has little surface expression (*cf.* Figure 4.10, top right). The tropical Pacific annual variability in both ensembles is of comparable magnitude to the annual variability from the control ensemble (Figure 4.11, bottom). However, an important difference is the fraction of variability in the tropical Atlantic region, which is much larger for the +NAO ensemble than the -NAO ensemble. This may be related to the increased surface response in the tropical and midlatitude Atlantic in the +NAO ensemble, which ultimately drives the NAO anomaly (see discussion section, next).

4.4 Discussion

It is important to clarify the actual relationship that shows asymmetry in the above analysis: In both the composites created from the coupled control simulation (Section 4.2) and the atmosphere-only ensembles (Section 4.3.3) we have found indications of a stronger relationship between SSTs/sea ice and MSLP in the positive NAO/cool NA SPG phase. However, when compositing MSLP based on the NAO index (using the highest and lowest 10% of values, as in Section 4.2.1) we find a similar in-phase relationship in both positive and negative NAO states (not shown). That is, the coupled control simulation suggests it is the strength of the relationship between SST/sea ice and NAO that is asymmetric, and not the strength of the NAO itself. In other words, in the coupled control the absolute magnitude of the positive NAO anomaly using the highest 10% of NAO values is similar to the absolute magnitude of the negative NAO anomaly using the lowest 10% of values. Indeed, it is the existence of the postulated *asymmetrical coupling strength* that the atmosphere-only ensembles are testing (and have found) and not merely the skewness in the distribution of NAO anomalies in the coupled control simulation.

Some of the asymmetry in the coupling strength may arise due to the relationship between the North Atlantic eddy-driven jet/storm tracks and the NAO (*Hoskins and Valdes*, 1990). Recent work has shown that whether it is the position or strength of the storm tracks that affect the NAO may depend on the timescale of the variability, being either multi-annual

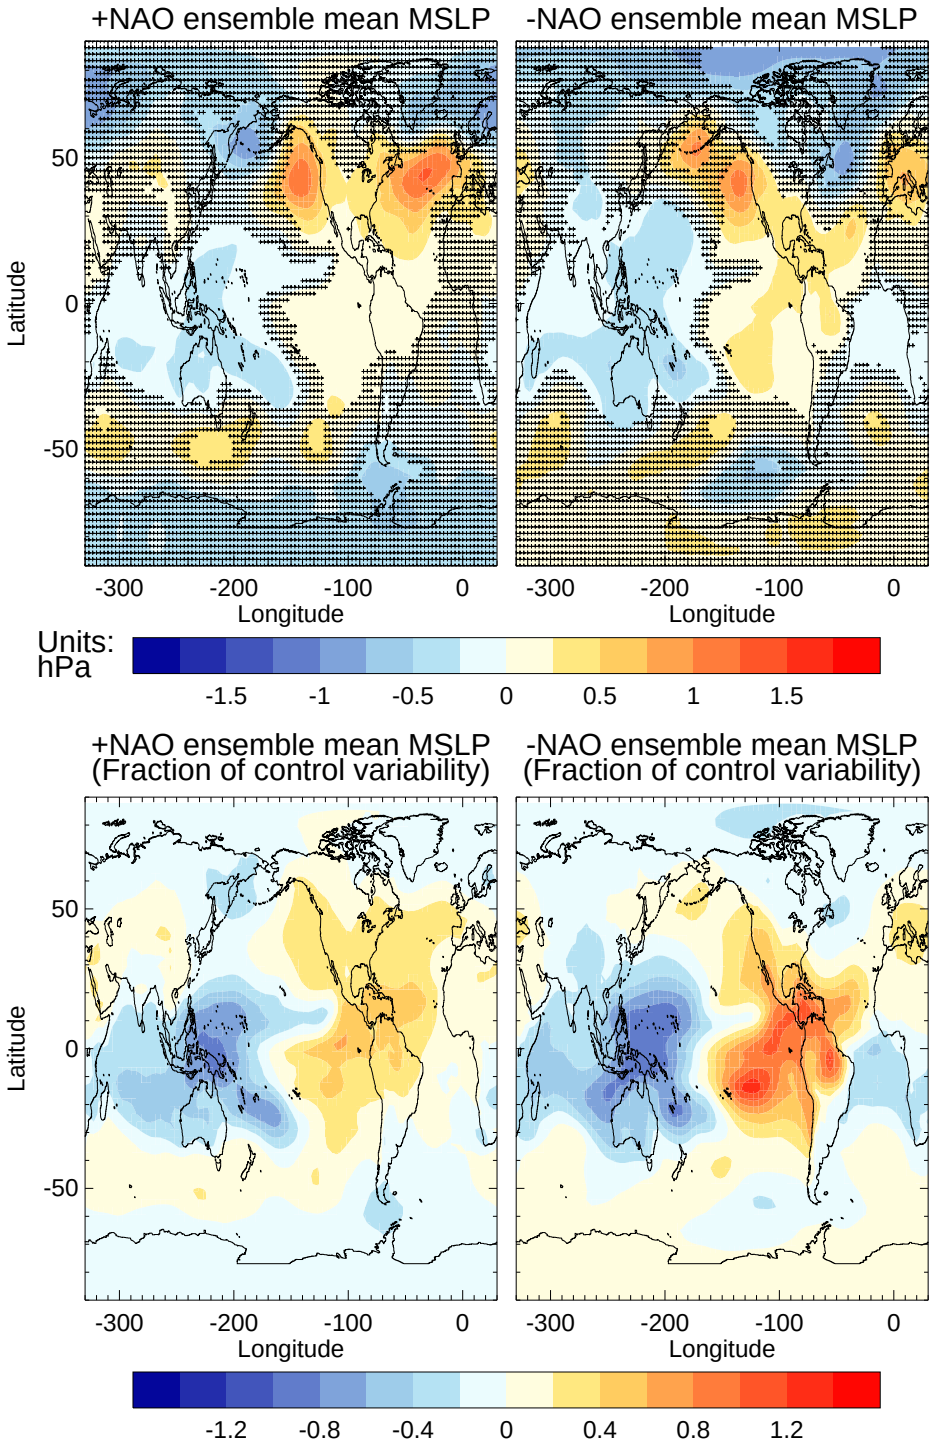

Figure 4.10: Top row: Global ensemble mean MSLP anomalies, with respect to the control ensemble, in the +NAO (left) and -NAO (right, scale inverted) ensembles, as in Figure 4.9. Stippling denotes anomalies indistinguishable from the control ensemble at the 90% level for a two-tailed t-test. Bottom row: As for the top row but dividing by the annual standard deviation in MSLP from the control ensemble, to estimate the relative magnitude of the anomalous responses.

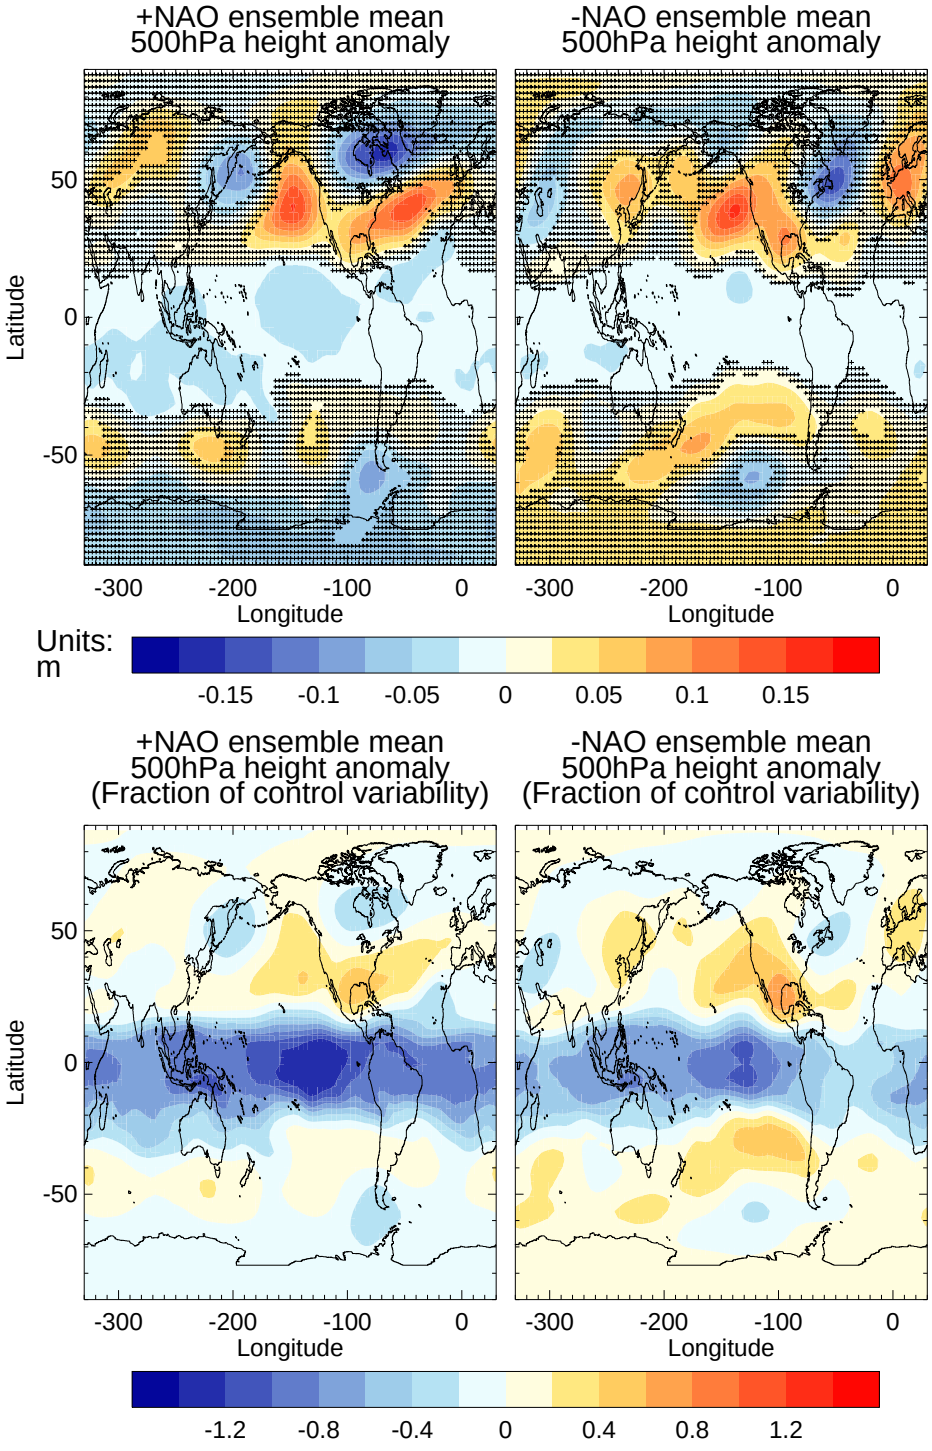

Figure 4.11: As in Figure 4.10 but for geopotential height anomalies at 500hPa. Units: metres.

or multi-decadal (*Woollings et al.*, 2014). Our atmosphere-only ensembles are forced with an anomalous SST gradient that is essentially equal and opposite between forcing sets (see Section 4.3.1) but this does not preclude differences in the precise patterns of the resulting static stability. Indeed, it has been shown that the extratropical atmospheric response of the North Atlantic can be forced by both the local (extratropical) and non-local (tropical/subtropical) Atlantic SSTs (*Sutton et al.*, 2000). Some evidence of this in the present uncoupled experiments can be seen in the subtropical 500hPa height anomalies in both the +NAO and -NAO ensembles (Figure 4.10). In addition, the combination of these forcings has been postulated to add non-linearly to give the resulting extratropical NAO response (*Sutton et al.*, 2000) and thus represents another possible driver of the asymmetrical coupling strength we have diagnosed.

In Section 4.3.3.2 we highlighted the asymmetrical global response to the global SST/sea ice forcing. Previous work has also highlighted an asymmetrical response of the NAO to forcing from the tropical Pacific in an observational framework, with a stronger relationship in the +NAO/La Niña phase (*Pozo-Vázquez et al.*, 2005), similar to the response we find. Indeed, further work has linked Pacific Rossby wave breaking to the +NAO phase, though not necessarily related to El Niño/La Niña (*Strong and Magnusdottir*, 2008). The favouring of +NAO over -NAO arises due to the preference for anticyclonic wave breaking (*Drouard et al.*, 2013). As such, it is possible that the MSLP anomalies (and the asymmetry) over the NA SPG are driven by non-local forcing, allowed by our experimental design, which provides a tantalising link between decadal variability in the North Atlantic and tropical Pacific, although we also note that there is no evidence of bidecadal periodicity in SSTs in the tropical Pacific (not shown).

Ultimately, the root cause of the asymmetry in the coupling strength between opposing phases of the simulated decadal variability in the NA SPG, and how this translates to different timescales for phase reversal, has yet to be diagnosed. Within the coupled system, there are key regions in which the atmosphere is postulated to play a role in reinforcing or driving ocean anomalies (Figure 2.13). If the timescales of phase reversal in the coupled system are, on average, related to the anomaly magnitudes (*e.g.* in T500) then it is possible that stronger coupling in the +NAO phase would, on average, hasten the +NAO (cool T500) part of the cycle.

Given the analysis in Section 4.3.3.2, future analysis in the current framework of investigation of the NA SPG could involve further atmosphere-only simulations performed with anomalous SST/sea ice forcing only in the NA SPG region, rather than globally. These would help to confirm that the asymmetry in coupling strength between SSTs/sea ice and the atmosphere does indeed occur *in situ* in the NA SPG, or whether other, larger scale, processes are involved. As noted, there is a tropical Pacific signal reminiscent of El Niño/La Niña in the forcing fields we have used (Figure 4.8). Using an atmosphere-only model that was a precursor to the model used in this analysis *Ineson and Scaife* (2009) found that, under some circumstances, El Niño-related SSTs could impact wintertime temperatures in northern Europe. As such, further analysis could probe the relative roles of both Pacific and Atlantic forcing as well as tropical Atlantic and subpolar Atlantic forcing — the latter two having already been shown to add non-linearly to give an extra-tropical atmospheric response over the North Atlantic (*Sutton et al.*, 2000).

Recent work has shown that NA SPG ocean variability on timescales greater than 15 years can be driven by NAO-related forcings (*Mecking et al.*, 2014). To further investigate whether the asymmetry we have detected arises due to atmosphere or ocean processes, a parallel set of forced ocean-only simulations, given surface forcing associated with anomalously high and low NAO states, could be performed. These were not performed here due to complications involving the most optimal way to restore the ocean surface salinities, which show a tendency to drift significantly when in uncoupled mode (*Behrens et al.*, 2013). Additionally, as noted in Section 4.2.4, there is a stronger response in ocean circulation in the NAC region associated with cool SSTs than warm SSTs (Figure 4.7) — with the anomalous circulation in this region previously shown to be related to the geostrophic response to temperature anomalies (Section 2.4.6). Ocean-only simulations would help to elucidate whether the asymmetrical response seen in Figure 4.7 is due to the asymmetrical coupling strength differently affecting the ocean temperatures (and hence geostrophic circulation), or whether it is due to processes internal to the ocean.

More broadly, the analysis we have presented in this chapter emphasizes the complementary approaches of composite analysis (low DoF but no assumption of symmetry) and regression analysis (high DoF but assumes symmetry, see Table 4.1 for further examples). Indeed, it highlights the potential role of asymmetry in the mechanism of decadal vari-

ability presented in Chapter 2 and why linear regression (symmetric) analysis may not always reveal the precise timescales involved. Such a combination of analysis techniques could be applied to other processes in the climate system where there is enough data to allow substantial sub-sampling, such as variability associated with El Niño/La Niña.

4.5 Chapter conclusions

We have investigated the existence of asymmetry between opposing phases of the mechanism of decadal variability in HadGEM3 that was diagnosed in Chapter 2. This asymmetry is highlighted through the use of composite analysis, as opposed to linear regression analysis, where the latter implicitly assumes symmetry between opposing phases of an oscillation.

- The coupled control simulation shows asymmetry in a variety of climate-relevant fields in the North Atlantic subpolar gyre (NA SPG), such as mean sea level pressure (MSLP), surface heat fluxes, and top 500m depth averaged temperature.
- The asymmetry is manifest in space as a stronger MSLP anomaly associated with an anomalously *cool* NA SPG rather than with an anomalously *warm* NA SPG. The asymmetry is also manifest in time as a faster reversal timescale (by approximately 2 years) from positive North Atlantic Oscillation (+NAO) to negative NAO (-NAO) conditions than from -NAO to +NAO conditions.
- An ensemble of atmosphere-only experiments suggest the asymmetry in *space* is related to an asymmetrical coupling strength between the ocean and atmosphere, *i.e.* that anomalously cool NA SPG sea surface temperatures (SSTs) exert a stronger influence on the atmosphere than anomalously warm NA SPG SSTs.
- Although still unclear, the asymmetry in *timescales* could be linked to the asymmetry in coupling strength, given the previously diagnosed role for the atmosphere in reinforcing ocean anomalies. The timescale asymmetry cannot be fully explained by the role of anomalous ocean circulation in the North Atlantic Current (NAC) region, which explains only a 4 month offset.

- Non-local forcing from the tropical Pacific and tropical Atlantic cannot be discounted as playing a significant role in the asymmetry. Further atmosphere-only ensembles, with non-local SSTs/sea ice relaxed to climatology, would help to resolve this issue.

This analysis has highlighted the potential role of the asymmetrical evolution of anomalies in the NA SPG in HadGEM3. The asymmetry between the evolution of, for example, MSLP anomalies, can lead to significantly different timescales (up to 2 years difference) in opposing phases of the variability. As such, it further highlights the need for multi-century control simulations when investigating simulated decadal variability to allow for the low-DoF asymmetrical analysis we have presented.

We have previously diagnosed a range of processes that are important to the NA SPG decadal variability in HadGEM3 and tested them against observations revealing varying degrees of agreement (Chapter 3). In the next chapter we discuss how robust and predictable are various elements of the proposed mechanism (Chapter 5) before moving on to try and draw broader conclusions across many more models in Chapter 6.

Chapter 5

Examining initial condition ensembles as a means of testing the mechanism of decadal variability

5.1 Introduction

In Chapter 2 we diagnosed the mechanism of internal decadal variability within the North Atlantic subpolar gyre (NA SPG) in HadGEM3. In this chapter we investigate whether specific phases of the decadal variability are more or less robust than others, using as our experimental apparatus a set of initial condition ensembles. Additionally, given that this model, or a version thereof, will provide the basis for seasonal to decadal prediction at the Met Office Hadley Centre (*MacLachlan et al.*, 2015) for the forthcoming years it is also appropriate to begin to test whether the diagnosed variability can contribute to increased predictive skill. As such, we begin with a brief review of decadal prediction and the motivation for our initial condition ensembles in section 5.2. In Section 5.3 we describe the initial condition ensemble experimental design before presenting some results for the NA SPG in Section 5.4, both generally and for some of the specific processes noted in Chapter 2. We discuss our findings in the context of the simulated mechanism of decadal variability (Chapter 2) and comparisons against observations (Chapter 3) in Section 5.5

before presenting chapter conclusions in Section 5.6.

5.2 Brief review of decadal prediction

Decadal climate prediction aims to provide useful information for society on, for example, near-term atmospheric temperature and precipitation changes on regional scales (*Meehl et al.*, 2009). The skill and/or predictability (defined next) of these predictions may well arise from oceanic processes due to the long memory in the ocean, particularly the deep ocean (*Delworth et al.*, 2007). Therefore, understanding whether simulated ocean variability is realistic (see Chapter 3) is important in order to assess the reliability of decadal predictions. When evaluating the performance of a decadal prediction system, one can test it against real-world observations (with ensuing difficulties relating to the procedure for assimilating the real-world state (*Magnusson et al.*, 2013) and initialisation shock) or against itself in a 'perfect model' framework (*Hawkins et al.*, 2011). Given its relative simplicity, and that we also aim to test the robustness of various parts of the decadal variability described in Chapter 2, we use the latter approach.

The potential 'skill' of a prediction of a climate variable (see Chapter 1, Section 1.2.2.1)), such as European surface air temperature, can be assessed in a model framework by conducting multiple 'Initial Condition' ensembles, each with multiple ensemble members (*Collins*, 2002). Here, a set of parallel simulations begun at multiple start dates initialised from some control simulation are performed and their evolution at a given lead time is compared to the control evolution. The control simulation represents the target, or 'truth' — in this analysis the simulation described in Chapter 2 is the control simulation. The better the anomalies (from some time mean) in the parallel simulations match the control simulation anomalies, the higher the anomaly correlation coefficient (ACC, skill). In addition, each start date is made up of multiple ensemble members to attempt to average out the noise and retain only the predictable signal. Ensemble member generation can be complex (*Karspeck et al.*, 2013; *Ham et al.*, 2014) but in the present experiment we choose to simply apply a bit-level perturbation to the atmosphere model restart files after *Dunstone et al.* (2011). The ACC, which is a function of lead time, is given by the

following formula (modified from *Collins*, 2002), where a value of 1 represents perfect skill, and a value of 0 represents no skill:

$$ACC\left(t\right) = \frac{\sum_{j=1}^{N}\sum_{i=1}^{M}\sum_{k\neq i}\left(x_{kj}\left(t\right)-\overline{x}(t)\right)\left(x_{ij}\left(t\right)-\overline{x}(t)\right)}{\sum_{j=1}^{N}\sum_{i=1}^{M}\sum_{k\neq i}\left(x_{kj}\left(t\right)-\overline{x}(t)\right)^{2}} \tag{5.1}$$

where x is some climate variable, j represents the individual start dates (N total), i represents each ensemble member (M total), and k represents the 'truth', which is taken to be each ensemble member in turn (*Collins*, 2002). The grand ensemble mean, $\overline{x}$, is given by:

$$\overline{x}(t) = \frac{1}{N\times M}\sum_{j=1}^{N}\sum_{i=1}^{M}x_{ij}\left(t\right) \tag{5.2}$$

In addition to the ACC, another useful measure is the prognostic potential predictability (PPP, *Pohlmann et al.*, 2004). This relates the amount of variability across the ensembles and ensemble members at a given lead time (for example, 1 year) to the interannual variability of the control simulation (*i.e.* the annual standard deviation in the detrended control simulation). As such, it diagnoses lead times where the ensembles remain more tightly constrained than the control simulation, suggesting that there is some useful constraint provided by the initial conditions in these periods. For the PPP, a value of 1 represents an ensemble with no intra-ensemble variance, and a value of 0 represents an ensemble with the same intra-ensemble variance as exists through time in the control simulation. The PPP, as a function of lead time, is given by the following equation (notation modified from *Pohlmann et al.* (2004) to be consistent with the above):

$$PPP(t) = 1 - \frac{\frac{1}{N(M-1)}\sum_{j=1}^{N}\sum_{i=1}^{M}\left(x_{ij}(t)-\overline{x}_{j}(t)\right)^{2}}{\sigma^{2}} \tag{5.3}$$

where $\overline{x}_j$ is the jth ensemble mean and σ^2 is the variance of the particular climate variable throughout the control simulation (detrended as appropriate). The PPP is complementary to the ACC. For example, one can imagine a situation where the anomalies in the ensembles are of the same sign as the control simulation but with considerable spread across

the ensemble members — a high ACC but low PPP, possibly indicating a signal-to-noise problem. Similarly, one can also imagine a situation where the anomalies in both the ensembles and control are spread tightly around zero anomaly, yielding a low ACC but high PPP (if the variability is normally much larger).

Finally, in addition to the ACC and PPP, we also define a new index: the concatenated anomaly correlation coefficient (CACC), to allow us to investigate the relative skill in different *subsets* of start dates/initial conditions. This is defined similarly to the ACC but to offset the effect of reducing the number of start dates it uses multiple lead times to build the correlation. As such, the CACC is not calculated for each lead time but is sensitive to the lead times (P) incorporated:

$$CACC = \frac{\sum_{l=1}^{P}\sum_{j=1}^{N}\sum_{i=1}^{M}\sum_{k\neq i}\left(x_{kjl}-\overline{x}_{cacc}\right)\left(x_{ijl}-\overline{x}_{cacc}\right)}{\sum_{l=1}^{P}\sum_{j=1}^{N}\sum_{i=1}^{M}\sum_{k\neq i}\left(x_{kjl}-\overline{x}_{cacc}\right)^{2}} \tag{5.4}$$

where the notation is as previously and l represents the (not independent) lead times (P total) whilst N is reduced to investigate the skill for groups of start dates separately. The grand ensemble mean, $\overline{x}_{cacc}$, is given by:

$$\overline{x}_{cacc} = \frac{1}{P\times N\times M}\sum_{l=1}^{P}\sum_{j=1}^{N}\sum_{i=1}^{M} x_{ijl} \tag{5.5}$$

The CACC is a measure of where the ensembles and ensemble members evolve similarly in time. However, whereas multiple start dates can be assumed independent (if well separated), multiple lead times cannot, which reduces the statistical power of this index. Nonetheless, given the few start dates available to us (see Section 5.3) this metric at least provides an indication of how various sets of initial conditions compare. Finally, we limit our analysis to these relatively simple and intuitive metrics but note that there are many alternative methods of quantifying skill. These are often more appropriate for verification against real observations (rather than in a perfect model framework) or for analysis of discrete rather than continuous variables (*Weigel et al.*, 2007).

5.2.1 Motivation for an initial condition ensemble

Our motivation for conducting investigations with initial condition ensembles is twofold.

Firstly, we aim to investigate in more detail which of the elements of the mechanism of simulated decadal variability in HadGEM3 are more robust than others by choosing initial conditions with large anomalies in a variety of subregions within the NA SPG (**Goal 1**). That is, to what extent are different phases of the variability more or less susceptible to random fluctuations that can create or destroy apparently low frequency modes. This question is not easily answered with lagged regression analysis (as presented in Chapter 2) that inherently post-rationalises the simulated variability. For example, *Wittenberg et al.* (2014) described how simulated, unforced, large decadal variability in El Niño, which one might expect to be predictable in a perfect model framework and to provide skill to — or drive variability in — other climate indices, showed very little (if any) predictability at all.

Secondly, we aim to begin to analyse the predictability characteristics of this model, noting its potential future use in seasonal to decadal forecasting (*MacLachlan et al.*, 2015, **Goal 2**). Although this analysis is in a 'perfect model' framework (where the target 'truth' is not reality but the model control simulation), we aim to combine it with our analysis of observed constraints on the simulated decadal variability (Chapter 3) to discuss the potential real-world predictability. In the next section we describe the design of our initial condition experiments.

5.3 Initial condition ensemble — experimental design

To determine the instances from the control simulation to use as initial conditions for our ensemble experiments we first define some broad indices that aim to capture particular phases of the simulated decadal variability. Given the large vertical extent of the decadal mode in the NA SPG (Figure 2.9) we use depth average temperatures over the top 1000m (T1000). Additionally, given the differing processes/timescales associated with different stages of the variability (see Figure 2.13) we define indices in three locations: the

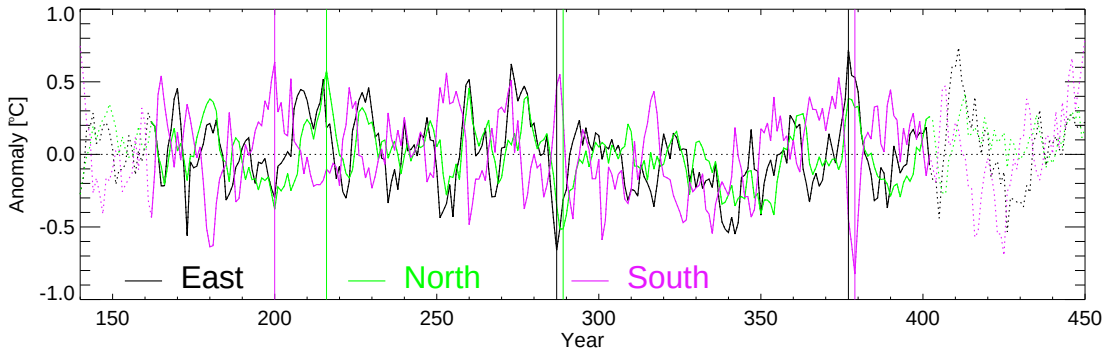

Figure 5.1: Time series of detrended (high pass filtered to remove periods greater than 70 years) top 1000m depth averaged temperature (T1000) in three regions within the North Atlantic subpolar gyre. These are a region in the East (10–25°W, 50–57°N, black), a region in the North at the entrance to the Labrador Sea (37–52°W, 57–63°N, green), and a region in the South in the North Atlantic Current (50–65°W, 35–40°N, pink). Regions are marked on Figure 5.2. The model year of the maximum and minimum values of the detrended time series are denoted by the vertical lines. Years towards the beginning of the time series are truncated due to the unavailability of model restart files during this time. Data at either end of the time series are ignored due to the end-effects of the detrending.

East (10–25°W, 50–57°N), the North (37–52°W, 57–63°N), and the South (50–65°W, 35–40°N). Time series of the T1000 interannual variability in these regions are shown in Figure 5.1. We find both the maximally positive anomaly and maximally negative anomaly for each index, resulting in a total of six sets of initial conditions/start dates. Each of these is comprised of five ensemble members (plus the control — see below) and run for five years, amounting to a total of 150 years of model integration in addition to the control simulation. Ensembles are initialised on December 1st (note that the model year runs from December 1st to November 30th by convention). Bit-level perturbations are applied to all fields in the atmospheric component of the restart files for each ensemble member to provide the initial seed for the subsequent chaotic evolution. We present annual means unless otherwise stated.

In the present analysis we use an idealised framework in which the model control simulation provides the ‘truth’ that a set of short simulations with perturbed initial conditions aim to recreate. Therefore, for a given start date, the period of the control simulation that is parallel to this ensemble can be considered to be merely another ensemble member. Thus, when computing the ACC and PPP, we include this additional ensemble member to increase the robustness of our statistical estimates. For the ACC, this is achieved by us-

ing a resampling technique, after *Collins* (2002), which treats each ensemble member in turn as the 'truth' (Equation 5.1). For the PPP, we merely increase the effective ensemble size by one member, after *Pohlmann et al.* (2004, Equation 5.3).

It can be seen from Figure 5.1 that some of the initial conditions are not well separated in time, due to the co-variability of the East, North, and South indices. We chose not to manually interfere with our algorithm for selecting the initial conditions but note that the effective number of independent start dates may be closer to four than to six. However, it should also be noted that it is not always the same pairs of indices that are adjacent (*cf.* the East and North indices around model year 290, and the East and South indices around model year 380, both of which are separated by only 2 years). In addition, the spatial structures of these nearby start dates are not obviously more similar than other well separated start dates (Figure 5.2). This is discussed further in Section 5.5.

The annual mean T1000 anomalies from the control simulation for the year that ends at the initialisation of the ensembles are shown in Figure 5.2. We use annual means for simplicity when searching for anomalous phases of the multi-annual/decadal variability but note that, as the model is initialised from a single point in time (or more precisely at two adjacent time-steps due to the specifics of the numerical scheme), these are not the precise initial conditions. It can be seen that all the initial conditions show to some degree a dipole between the northern and southern halves of the NA SPG (discussed in Section 5.5) and that the sign of the anomaly in the eastern and northern regions is generally the same (consistent with the short timescales linking the simulated decadal variability in these regions, Section 2.4.2). The apparent similarity of many of the initial conditions is related to our first goal of investigating which elements of the decadal variability are most robust but is not optimal for general assessment of the predictability characteristics of the NA SPG, for which a more complete sampling of parameter space would be preferable (though we note this is also limited by available computing resources). Nonetheless, within the limits imposed by the available resources and the dual aims of the experimental design, we next present our analysis of the initial condition ensembles.

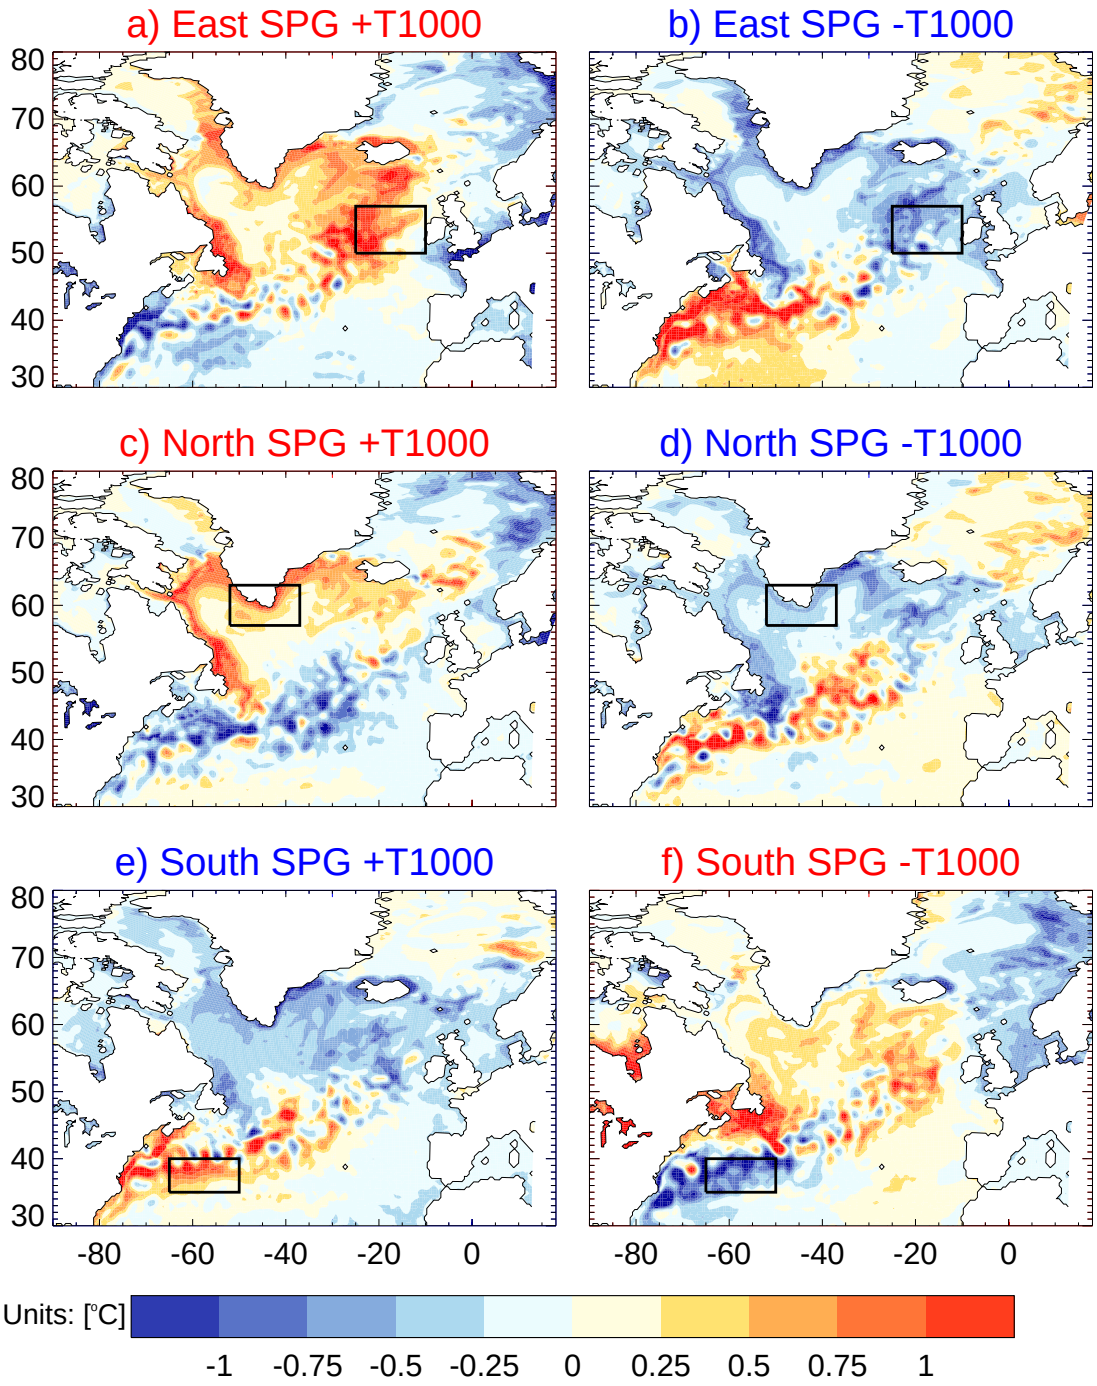

Figure 5.2: Maps of the control simulation annual mean top 1000m depth averaged temperature (T1000) anomalies prior to the initialisation of the model ensembles. Anomalies are with respect to the parallel control low frequency variability (low pass filtered to remove periods shorter than 70 years) to highlight their magnitude relative to neighbouring years. The start dates are chosen using filtered time series as shown in Figure 5.1 (vertical lines). There are six initial condition states, representing anomalously warm conditions in the East subpolar gyre (SPG, a), North SPG (c), and South SPG (e), and anomalously cool conditions in the East subpolar gyre (b), North SPG (d), and South SPG (f). Also highlighted is whether the majority of the central subpolar gyre is warm (a, c, f, red titles) or cool (b, d, e, blue titles).

5.4 Results

In this section we present the results of our initial condition ensemble experiments, beginning with the general evolution of the NA SPG (Section 5.4.1) before analysing the skill and predictability in different regions (Section 5.4.2) and finally in specific processes related to the simulated decadal variability (Section 5.4.3).

5.4.1 General evolution

To characterise the general evolution of anomalies initialised in each of the eastern, northern, and southern NA SPG regions we show the normalised ensemble mean T500 anomalies relative to the low frequency variability in the long control simulation (Figure 5.3). Each of the panels comprises two initial conditions (positive and negative anomalies) that each consist of five ensemble members as well as the parallel portion of the control simulation, making a total of twelve fields over which we compute the mean anomaly (multiplying the negative anomalies by minus one). At t=1 (which denotes the annual mean anomaly for the first full year), the location of the region in which we maximised the anomaly values in the control simulation (Figure 5.3) can still be seen, with large anomalies in the eastern NA SPG when initialising based upon an East index (Figure 5.3, left column) and similarly for northern and southern regions. In addition, for the East index, there are already large anomalies in the north and west of the NA SPG, which are consistent with the fast timescales linking the decadal variability in these two regions (note though that these timescales will both influence the initial condition anomaly (Figure 5.2) as well as the subsequent evolution of the initialised anomalies — see further discussion of this in Section 5.5).

For increasing lead times, the location of the maximum positive anomaly (normalised by the interannual variability) for the East initial conditions spreads from the east into the Labrador Sea region (Figure 5.3, left column) whilst the location of the maximum negative anomaly spreads from the NAC region into the eastern NA SPG. The relationship between positive anomalies in the Labrador Sea and negative anomalies in the NAC region (and subsequently the eastern NA SPG) is further highlighted by the North initial

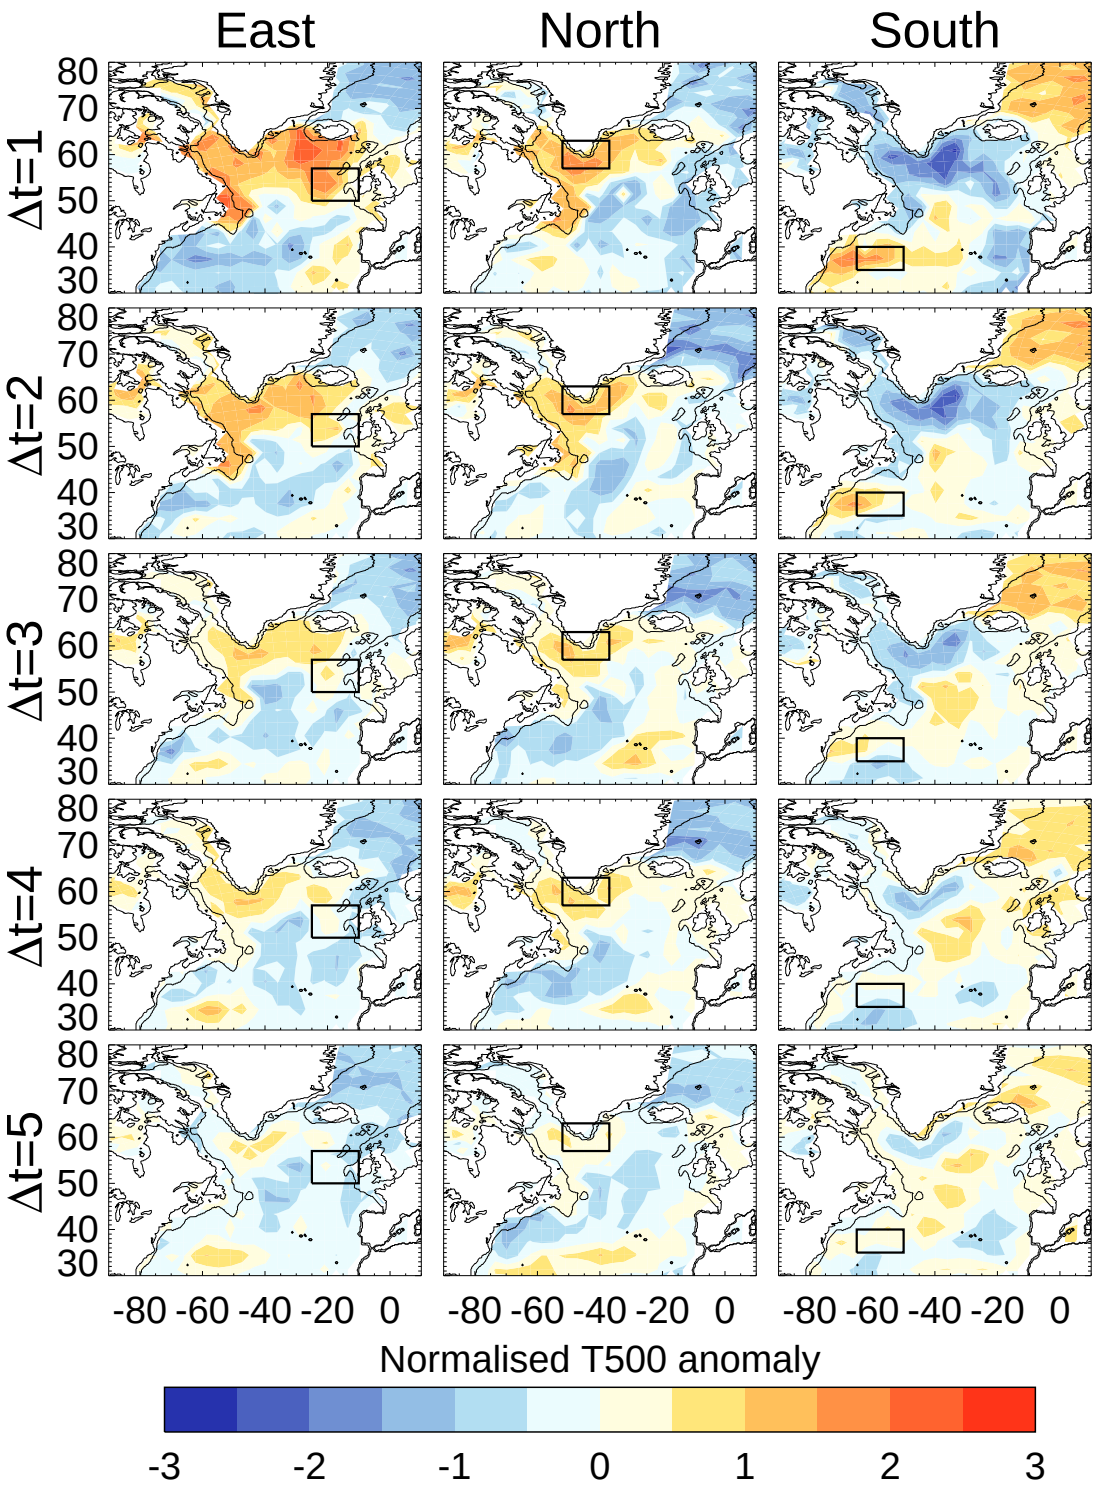

Figure 5.3: Annual mean top 500m depth averaged temperature (T500) ensemble mean anomalies with respect to the long-period variability in the parallel control simulation (low pass filtered to remove periods less than 70 years) as a function of lead time (*e.g.* Δt=1 is the annual mean over the first full year) for the East (left column), North (middle column), and South (right column) initial conditions. Pairs of anomalies for each region are combined by subtracting the cool anomaly from the warm anomaly and dividing by 2. Anomalies are further normalised by dividing by the detrended control annual standard deviation to highlight the relative magnitude of the ensemble mean signal compared to the annual variability.

conditions (Figure 5.3, middle column). For this set of initial conditions negative anomalies in the NAC region appear to follow positive anomalies in the Labrador Sea region by a few years, whilst negative anomalies located more towards the centre of the NA SPG appear in-phase with the same Labrador Sea anomalies. This difference between the locations of negative anomaly responses may suggest that the nature of the negative feedback between Labrador Sea and NAC anomalies (described in Section 2.4.6) is sensitive to the precise location of the anomaly in and around the Labrador Sea, which would likely further affect the timescales for phase reversal.

Finally, the South initial conditions (Figure 5.1, right column) show a strong NAC anomaly that appears to gradually propagate into the eastern NA SPG. Notably, at t=1, the location of the maximum negative anomaly for the South initial conditions is not centred in the Labrador Sea but located approximately evenly between the locations of maximal positive anomalies at t=1 for the East and North initial conditions, further suggesting some nuance as to the precise centres of action linking northern and southern NA SPG variability via the diagnosed negative feedback (Section 2.4.6).

The ensemble mean evolution of anomalies is an instructive first-order method of analysing the initial condition ensembles and the possible differences relating to the location of the largest anomalies in the initialisation. However, it is not clear from the ensemble mean evolution to what extent individual ensemble members evolve similarly and to what extent the ensemble mean evolution merely represents a damped response to large anomalies and/or regression to the mean (zero anomaly). In Figure 5.4 we plot an index of T500 variability throughout the NA SPG and the associated ACC and PPP scores for this index. Consistent with the large signals seen in the ensemble mean evolution (Figure 5.1) the initial condition ensemble members (Figure 5.4a, black lines) often appear to evolve similarly to the control simulation (Figure 5.4a, red line) — for example, the warming around year 290 and the cooling around year 380. Despite this, given the autoregressive characteristics of the control simulation, the ensemble mean evolution (averaged over all start dates and all ensemble members) is not more skilful than damped persistence (calculated after *Boer* (2000)) for lead times greater than 2 years (Figure 5.4, b). However, there are some important caveats to this estimate of the ACC: Firstly, even including all start dates and ensemble members together, the uncertainty on this ACC score is likely

still quite high. Secondly, note that our NA SPG index is between 45–65°N, to try and capture the general skill associated with the NA SPG. However, as can be seen from Figure 5.1, at lead times of t=3–4 years the NA SPG is in a somewhat transitional phase, which results in the basin-wide ensemble mean T500 anomaly being near to zero and thus reduces the correlations (assuming that the ensemble variance at these lead times continues grow, as indeed implied by Figure 5.4, c). This may contribute to the apparent recovery of skill at lead times of five years.

From the ocean, any transmission of skill to — or impacts on — atmospheric variability, such as the NAO or precipitation over Europe (*Sutton and Hodson*, 2005), must take place at the ocean surface. As such, we also show the skill and predictability in NA SPG SSTs in Figure 5.5. The variability (Figure 5.5, a), and lead time evolution of skill (Figure 5.5, b) and potential predictability (Figure 5.5, c) are very similar to the depth averaged temperatures. Nonetheless, for all lead times the absolute values of skill and potential predictability are slightly lower than for T500, consistent with a role for the subsurface ocean in providing some of the skill/predictability in T500. Related to this, there is less autocorrelation in the SSTs and so the damped persistence model becomes easier to improve upon, resulting in skill and predictability in SSTs that is better than damped persistence at broadly similar lead times as for T500. However, we note that the dip and subsequent improvement in skill, which results in poor skill in SPG T500 at lead times of 3–4 years but an improvement thereafter (Figure 5.4, b), is shifted forward by a year for SSTs (Figure 5.5, b), occurring at lead times of 2–3 years. This could imply that 1) some of the skill in the subsurface ocean originates at the surface, possibly driven by the atmosphere (see Section 5.4.3.3), or 2) that the propagating T500 anomalies exhibit some shear such that any transitional phase of the variability (noted above) occurs at the surface before the subsurface. However, as also noted above, small differences between ACC scores are likely sensitive to the limited number of start dates and ensemble members.

Given that there appears to be some (limited) skill in predicting the simulated evolution of basin-wide NA SPG T500 we now present the spatial structure of this skill to investigate in which regions this skill arises.

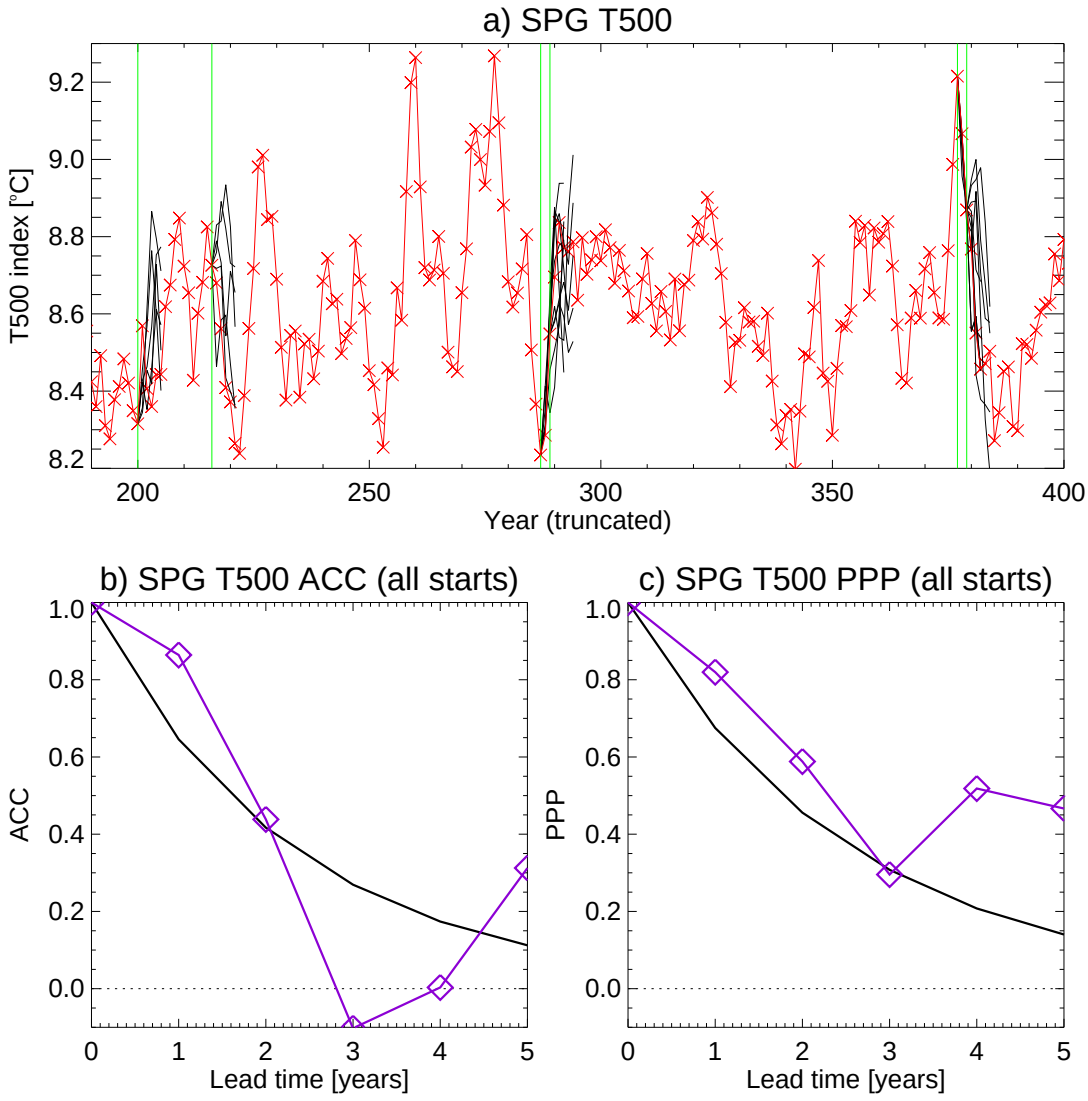

Figure 5.4: a) Time series of annual mean detrended (high pass filtered to remove periods longer than 70 years) top 500m depth averaged temperature (T500) in the North Atlantic subpolar gyre (SPG, 45–65°N, red) truncated to highlight the initialisation dates (green lines) and ensemble members (black). Ensemble members (and control) detrended by removing the low frequency (periods longer than 70 years) part of the parallel control time series and all predictability scores calculated using detrended data. b) Resampled anomaly correlation coefficient (ACC, purple with squares) calculated after *Collins* (2002) and an estimate of the ACC for damped persistence (black). c) Prognostic potential predictability (PPP, purple with squares) calculated after *Pohlmann et al.* (2004) and an estimate of the PPP for damped persistence (black). Damped persistence calculated after *Boer* (2000).

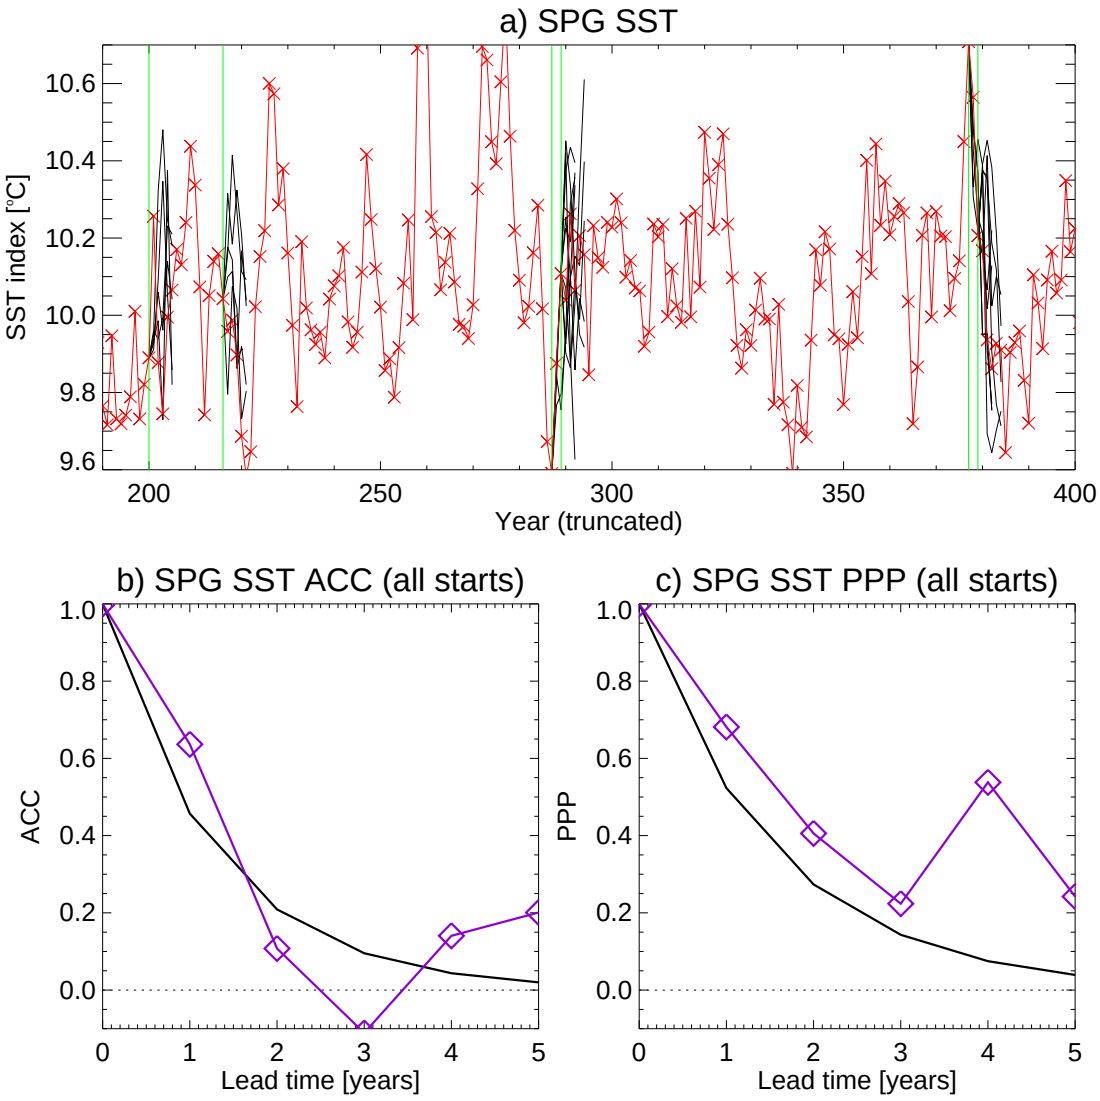

Figure 5.5: As Figure 5.4 but for sea surface temperature (SST).

5.4.2 Regional skill and potential predictability

In Figure 5.6 we show the resampled anomaly correlation over all start dates and ensemble members for each annual mean following the initialisation for T500. The skill falls away rapidly in the eastern and southern subpolar gyre from t=2 onwards but in the western subpolar gyre, and in particular the central Labrador Sea and Irminger Sea (south of Greenland), the skill remains above r=0.6 up to and including a lead time of 4 years. There is also considerable skill north of Iceland in the Nordic Seas. It is notable that both of these locations are where the model mixes shallow and deep water during convection, and so the skill may arise from the mixing up of slowly varying deep water

(*Alexander and Deser*, 1995). In addition, between t=2–5 it is possible to see an area of relatively high skill (compared to its surroundings) apparently propagating from the NAC region eastwards and northwards, consistent with the diagnosed mechanism of decadal variability in which heat content anomalies propagate along this path (see Section 2.4.4). Unfortunately, it is not possible to separately assess the skill provided by initialising based on the three regions as the ACC score relies on multiple start dates. The theoretical ACC score increases with the number of ensemble members, up to a maximum level given an infinite number of ensemble members (*Murphy*, 1990), with the number of start dates increasing the confidence in that ACC score. Nonetheless, it is possible to estimate the PPP for individual (or pairs of) start dates.

Although the PPP cannot explain which, if any, start dates provide more skilful predictions (our second goal) it can give some clues as to whether particular sets of start dates provide more coherent (across the ensemble) variability. This could then be interpreted as evidence of whether particular elements of the mechanism of simulated decadal variability are more robust than others (our first goal). In Figure 5.7 we show the evolution of PPP for pairs of initial conditions based on the East index (Figure 5.7, left column), the North index (Figure 5.7, middle column), and the South index (Figure 5.7, right column). It can be seen that initialisation based on large anomalies in the East provides high potential predictability in the northern NA SPG after 2 years, and more so than initialisation based on large anomalies in the North. Indeed, at a lead time of two years, the largest PPP signal based on the North index is in the subtropical gyre, which is also true for initialisation based on the South index. However, this similarity is not necessarily unexpected, given that anomalies in the North index and South index (of opposite sign) occur at short lags (Figure 2.8). One possible reason that the potentially predictable signal spreads into the subtropical gyre is that the anomalies affect the mean circulation here and may thus remain large (but for only a few years) whereas the link between the NAC and the eastern subpolar gyre is in part due to the anomalous circulation, which involves longer timescales and traversing along the inter-gyre boundary. Indeed, the low predictability inter-gyre boundary can be clearly seen for the South index at t=1.

Although the spatial maps suggest limited predictability at timescales longer than two years, this does not preclude longer timescale potential predictability for broader regions

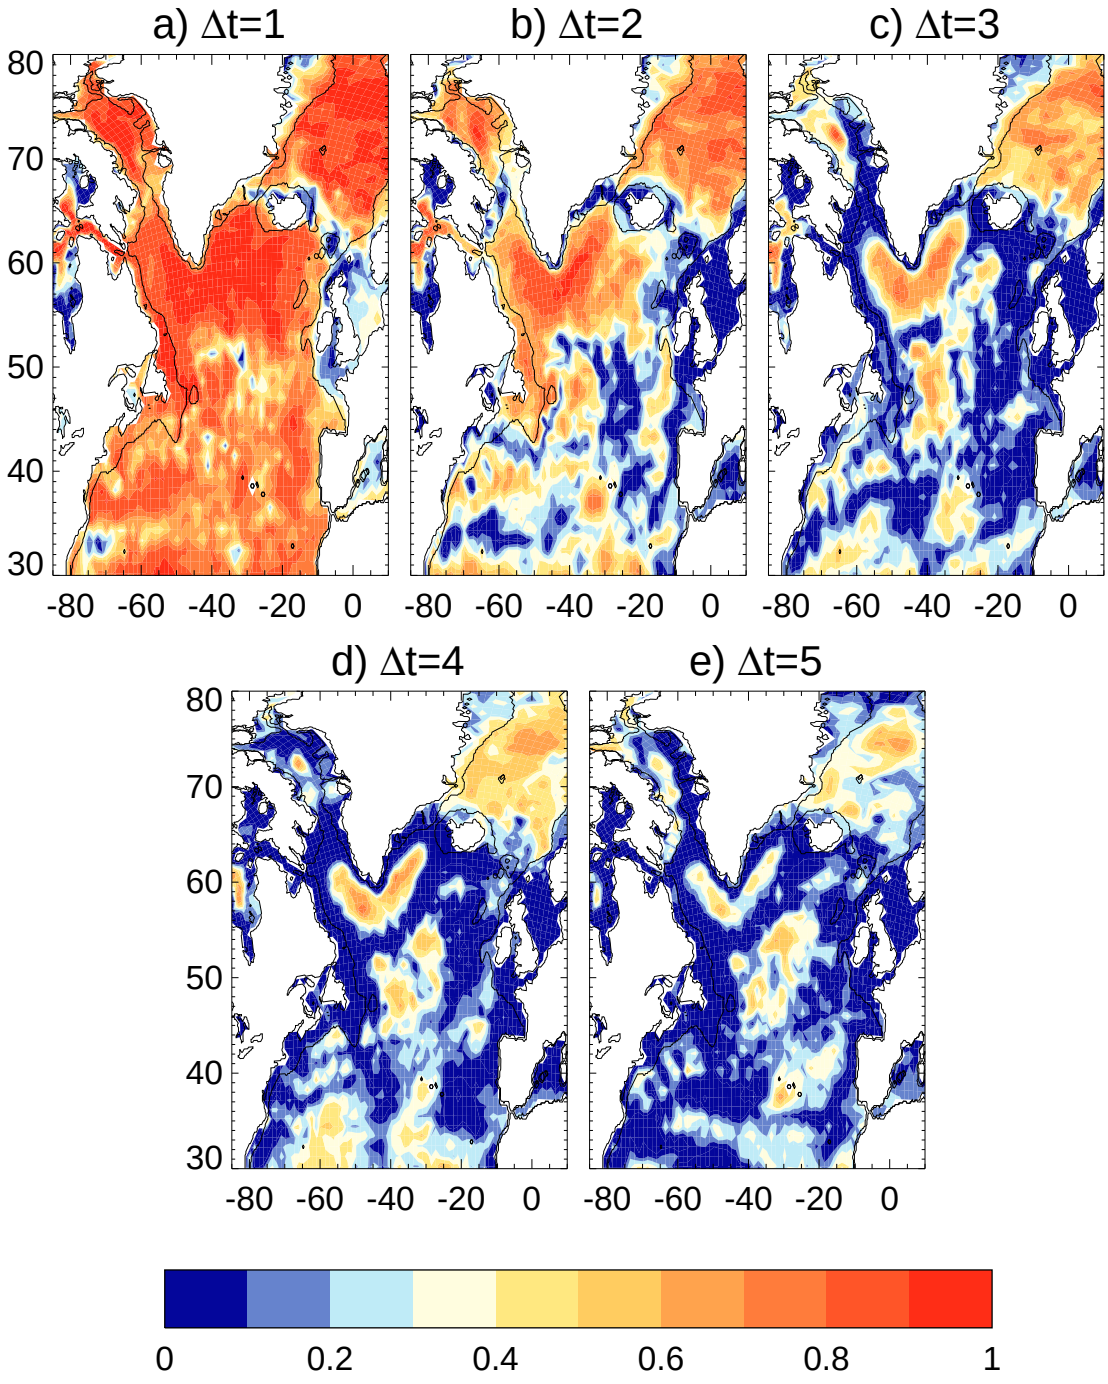

Figure 5.6: Annual mean top 500m depth averaged temperature (T500) resampled anomaly correlation coefficient (ACC, *Collins*, 2002) for all start dates as a function of lead time (years).

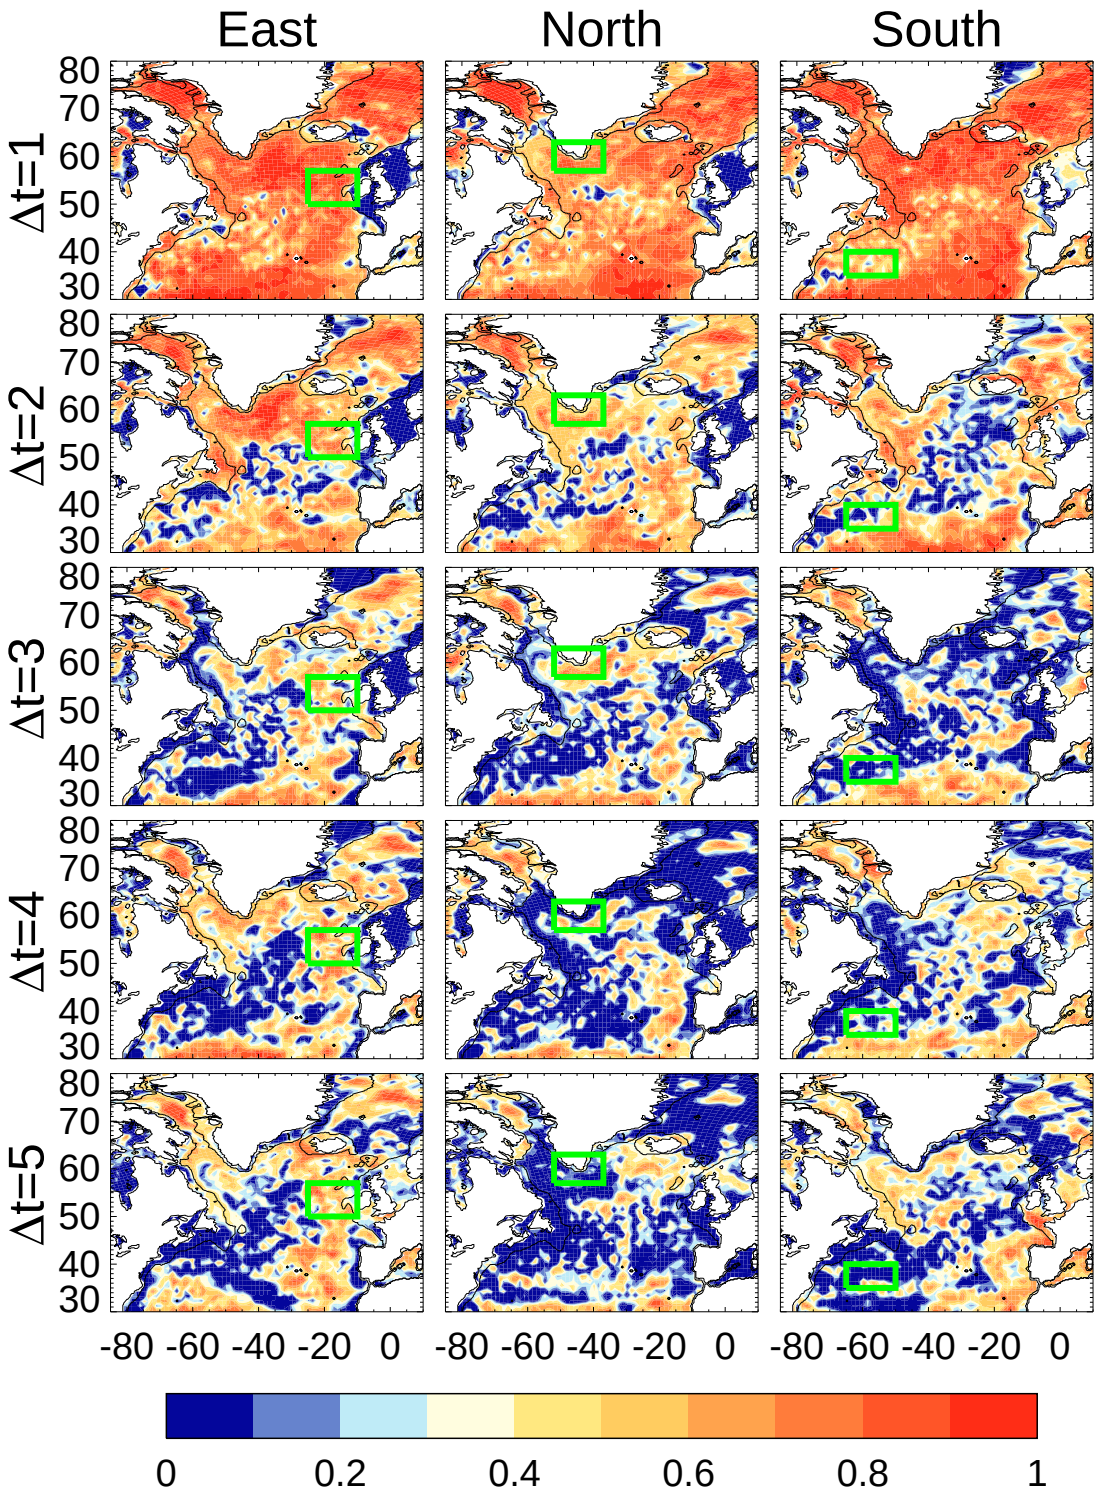

Figure 5.7: Prognostic potential predictability (PPP) in annual mean top 500m depth averaged temperature (T500) for initial condition ensembles based on: an eastern subpolar gyre index (left column), a northern subpolar gyre index (middle column), and a southern subpolar gyre index (right column), as a function of lead time (years). Each column represents the combination of a pair of ensembles describing anomalously warm and anomalously cool states (in each region) from the total of 6 initial condition ensembles. Green boxes highlight the region from which the initial conditions were chosen.

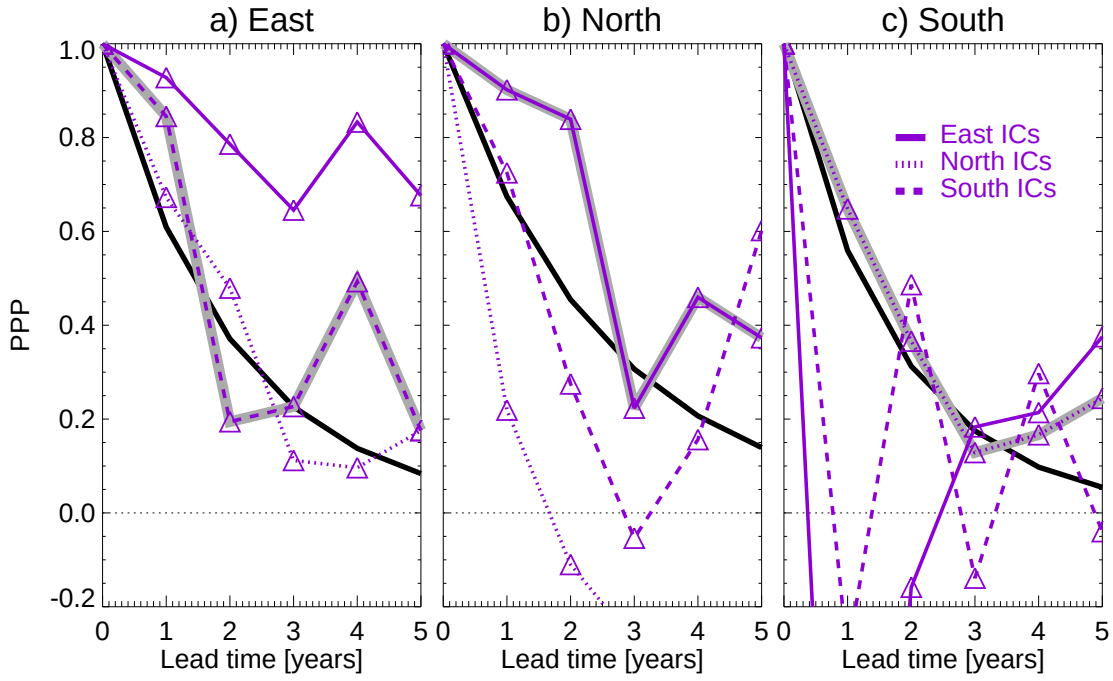

Figure 5.8: Prognostic potential predictability (PPP, purple) in annual mean top 500m volume averaged temperature measured in the eastern subpolar gyre (10–25°W, 50–57°N, a), northern subpolar gyre (37–52°W, 57–63°N, b) and southern subpolar gyre (50–65°W, 35–40°N, c). The PPP is calculated for pairs of ensembles describing anomalously warm and anomalously cool states (in each region) from the total of 6 initial condition ensembles, as in Figure 5.7. These are a pair of ensembles based on large eastern anomalies (solid), northern anomalies (dotted), and southern anomalies (dashed). In addition, the ensemble pair 'upstream' of the measurement region (according to the simulated mechanism of decadal variability in the North Atlantic subpolar gyre, see Figure 2.13) is highlighted with grey shading. The PPP for damped persistence (calculated separately for each region) is shown in black.

— for example, to average across the circulation-driven heat transport anomalies in the NAC region. For simplicity, we choose to measure the potential predictability of top 500m volume averaged anomalies in each of the regions that we also used to create the initial conditions (Figure 5.8). Over the broader eastern region (Figure 5.8, a), we find that the PPP generally remains greater than that expected from damped persistence, *i.e.* the ensembles remain better constrained than the innate interannual variability (estimated from the control simulation). Initialisation based on an index in the same East region gives the greatest predictability, but initialising 'upstream' in terms of the mechanism of simulated decadal variability (*i.e.* in the South), also gives high predictability.

Similar to the above, measuring the PPP over the North region (Figure 5.8, b) also suggests that initialising based on an index upstream in the East NA SPG gives high pre-

dictability. Interestingly, the initialisation based on an *in situ* North index gives predictability worse than damped persistence. It is possible that this is an artefact of the heat content anomalies in the measurement (and initialisation) index being strongly dominated by the northern boundary currents (see Section 2.4.5) and thus the initial condition ensembles based on large anomalies here may also capture large variability (for example, if the temperature/potential vorticity anomalies interact strongly with the circulation), resulting in large ensemble spread, compared to more quiescent times when the temperature anomalies are smaller and so the intra-ensemble spread is also small (*cf.* initialisation based on East or South temperature anomalies, Figure 5.8b). In contrast to the North measurement region, the South region shows initialisation based on large North anomalies gives the highest potential predictability (Figure 5.8, c). Once again this highlights a preference for the upstream region. However, we caution that the limited size of these ensembles mean that it is difficult to quantify the statistical difference between these upstream or downstream PPP scores.

In summary, across all three measurement regions, initialisation based on large anomalies upstream of the measurement region gives the highest (twice) or second highest (once) potential predictability. Overall, this suggests that anomalies seem to spread around the NA SPG consistent with the mechanism of decadal variability described in Chapter 2. However, it should also be noted that this is also consistent with just the mean circulation pathways, given that the NA SPG is a cyclonic gyre. We discuss this further in Section 5.5.

5.4.3 Skill and potential predictability in important processes

Given that there is some skill (ACC) and potential predictability within the NA SPG on both broad and regional scales we now investigate some of the processes previously shown to be important for the negative feedback in the decadal mode, as well as whether the ocean predictability shows any link to the North Atlantic Oscillation (NAO).

5.4.3.1 Labrador Sea deep water formation

In Chapter 2 we suggested that the negative feedback between Labrador Sea temperatures and those in the NAC was also linked to deep water formation (DWF) in the Labrador Sea, which mixes down the surface signal. Additionally, DWF in the Labrador Sea could project on to longer timescale variability via its effect on the Atlantic Meridional Overturning Circulation (AMOC, *Pohlmann et al.*, 2013). Figure 5.9 shows the time series and ACC and PPP scores for wintertime mixed layers in the Labrador Sea. The interannual autocorrelation is low, resulting in low ACC and PPP from damped persistence (Figure 5.9, b, c, black lines), which the initial condition ensembles are able to beat. Despite this, at lead times longer than two years the ACC and PPP are very low, although the ACC increases again at lead times greater than four years (similarly to the ACC for T500 and SST across the NA SPG, Figures 5.4 and 5.5). We also find that the previous winter provides greater skill (the autocorrelation between winters in the control simulation is 0.42) than the previous autumn season (r=0.30). This may suggest that some of the interannual memory is related to the subsurface storage and re-emergence of anomalies (*Alexander and Deser*, 1995).

5.4.3.2 Dynamic height gradient between Labrador Sea and North Atlantic Current

Similarly to convection in the Labrador Sea, the dynamic height gradient between the Labrador Sea and North Atlantic Current (NAC) was postulated to be important for the phase reversal of decadal variability. In Figure 5.10a we show the interannual variability in the control simulation for a dynamic height index defined identically to that used previously (Chapter 2, Section 2.4.6 and Figure 2.12), as well as the evolution of the initial condition ensembles. Although we use an annual index of the dynamic height there is little autocorrelation in the control simulation, which results in low estimates of the skill or predictability expected through damped persistence. As such, both the ACC (Figure 5.10, b) and PPP (Figure 5.10, c) remain higher than damped persistence for the first four years. The skill shown by the initial condition ensembles is larger for the dynamic height index than wintertime convection (mixed layers, Figure 5.9b) at com-

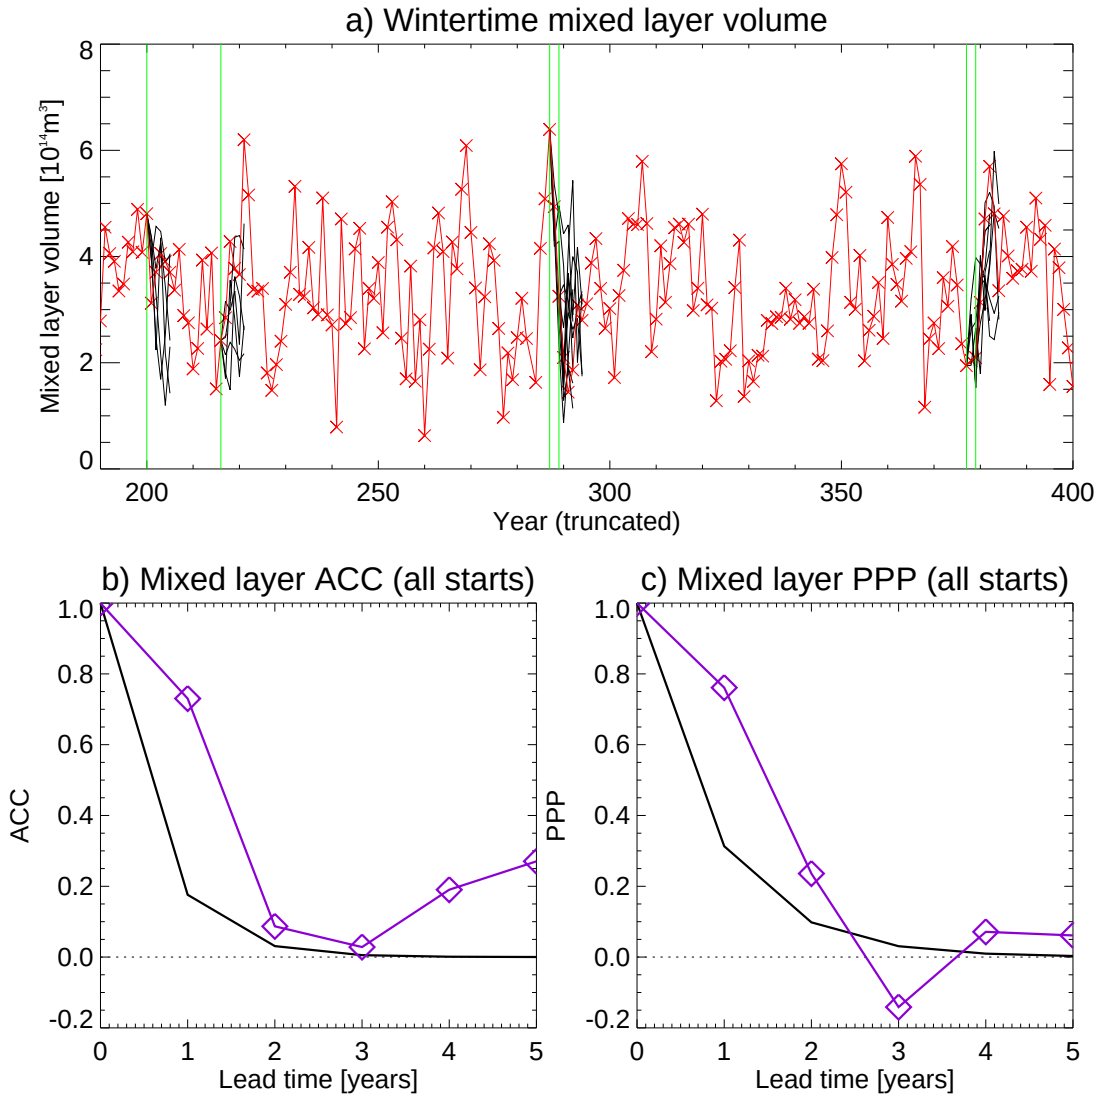

Figure 5.9: As Figure 5.4 but for wintertime (December to February inclusive) mixed layer volume in the Labrador Sea (50–60°W, 55–62°N). The damped persistence model is based on wintertime mixed layer volumes (*i.e.* the first lead time is 9–12 months rather than the 0–3 months for the initialised ensembles) as the previous winter provides greater skill (r=0.42) than the previous season (autumn, r=0.30).

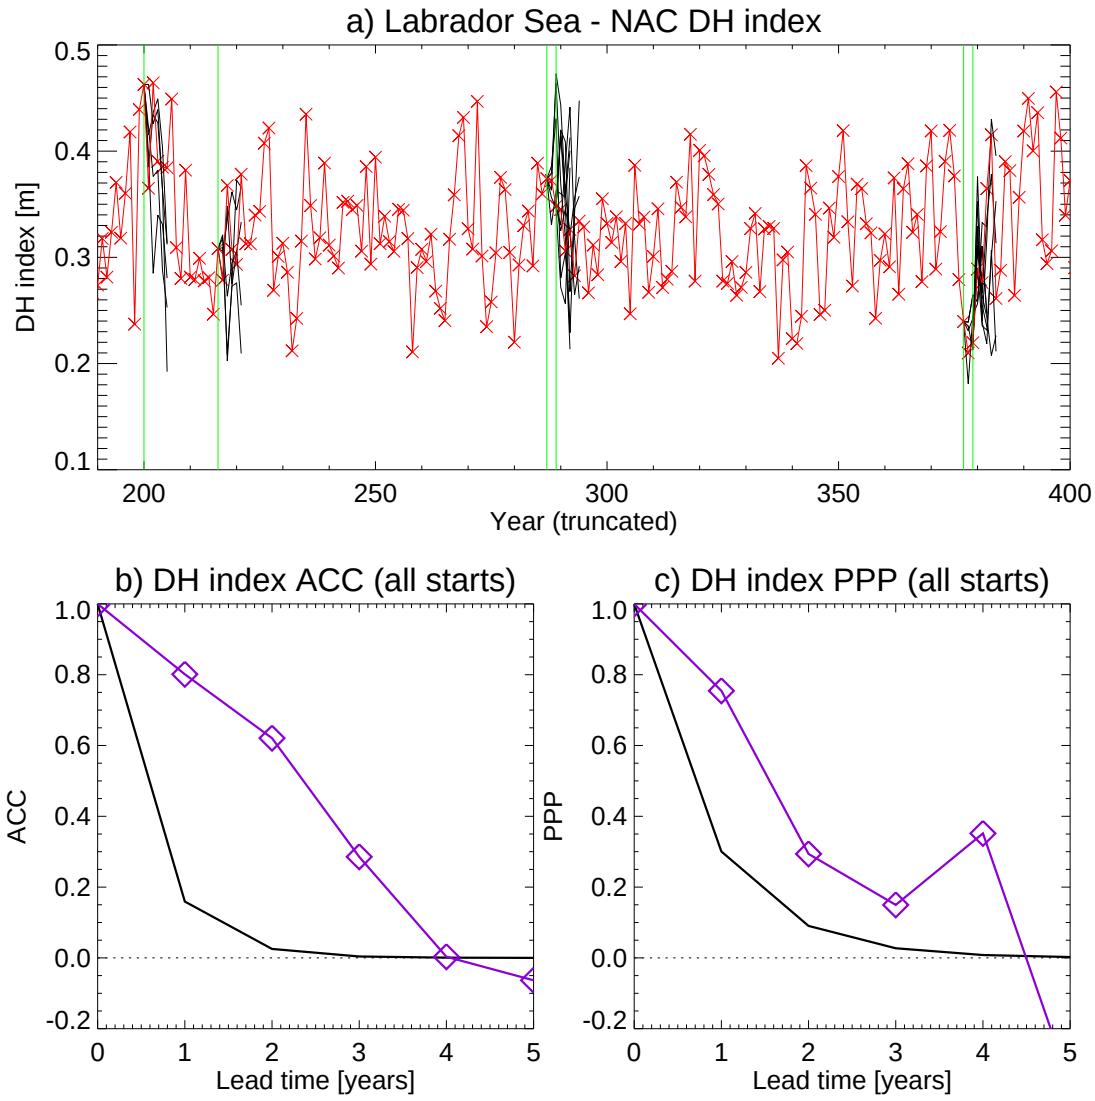

Figure 5.10: As Figure 5.4 but for the dynamic height (DH, relative to 500m) difference between the southern edge of the extended Labrador Sea (44°N) and the North Atlantic Current (NAC, 39°N) at 47.5°W, *i.e.* the same index as used in Chapter 2, Figure 2.12.

parable lead times, perhaps suggesting that the dynamic height index also encompasses skill from other sources. For example, in Section 2.4.7 we noted the potential role of the atmosphere in reinforcing the dynamic height gradient. For this reason, as well as for its climate/societal relevance, we conclude our analysis of specific processes by looking at the skill and predictability in the wintertime NAO index.

5.4.3.3 The North Atlantic Oscillation (NAO)

The NAO index computed from the control simulation shows very little autocorrelation (Figure 5.11, a), which can also be seen in the mostly white spectrum in Figure 2.4c (despite the spectral peak at 16/17 years). In addition, over the six sets of initial conditions and all ensemble members there is little evidence of skill or potential predictability at lead times longer than one season. This is also true for an annual NAO index (not shown) and indeed there is little skill or predictability in wintertime or annual mean sea level pressure anywhere in the North Atlantic (not shown). Nonetheless, the skill at a lead time of one season (r=0.34, Figure 5.11, b) is suggestive of a link between ocean and atmosphere, although the PPP is very low even at this lead time (Figure 5.11, c), indicating that the spread across the ensemble is already approaching the climatological annual variability in the wintertime NAO.

In addition to the PPP for the full set of six initial conditions, we also group the ensembles by whether they are initialised with an anomalously warm or cool northern NA SPG (see Figure 5.2). The PPP for the full ensemble is the linear sum of the PPP for the warm and cool subsets (Figure 5.11, c). It can be seen that the poor PPP for the full ensemble is primarily due to the warm set of initial conditions, whereas the cool initial conditions give higher potential predictability, although there is considerable spread across the ensembles. This result is consistent with the analysis in Chapter 4 in which we suggest a stronger coupling between ocean-atmosphere in the cool SPG/positive NAO state than in the warm SPG/negative NAO state. This increased coupling strength appears to result in increased predictability in the NAO index when the subpolar gyre is anomalously cool (*i.e.* positive NAO), as opposed to when it is anomalously warm (*i.e.* negative NAO). However, as noted previously, it is not possible to reliably assess the skill (ACC) of subsets of ensembles, given the limited number of start dates.

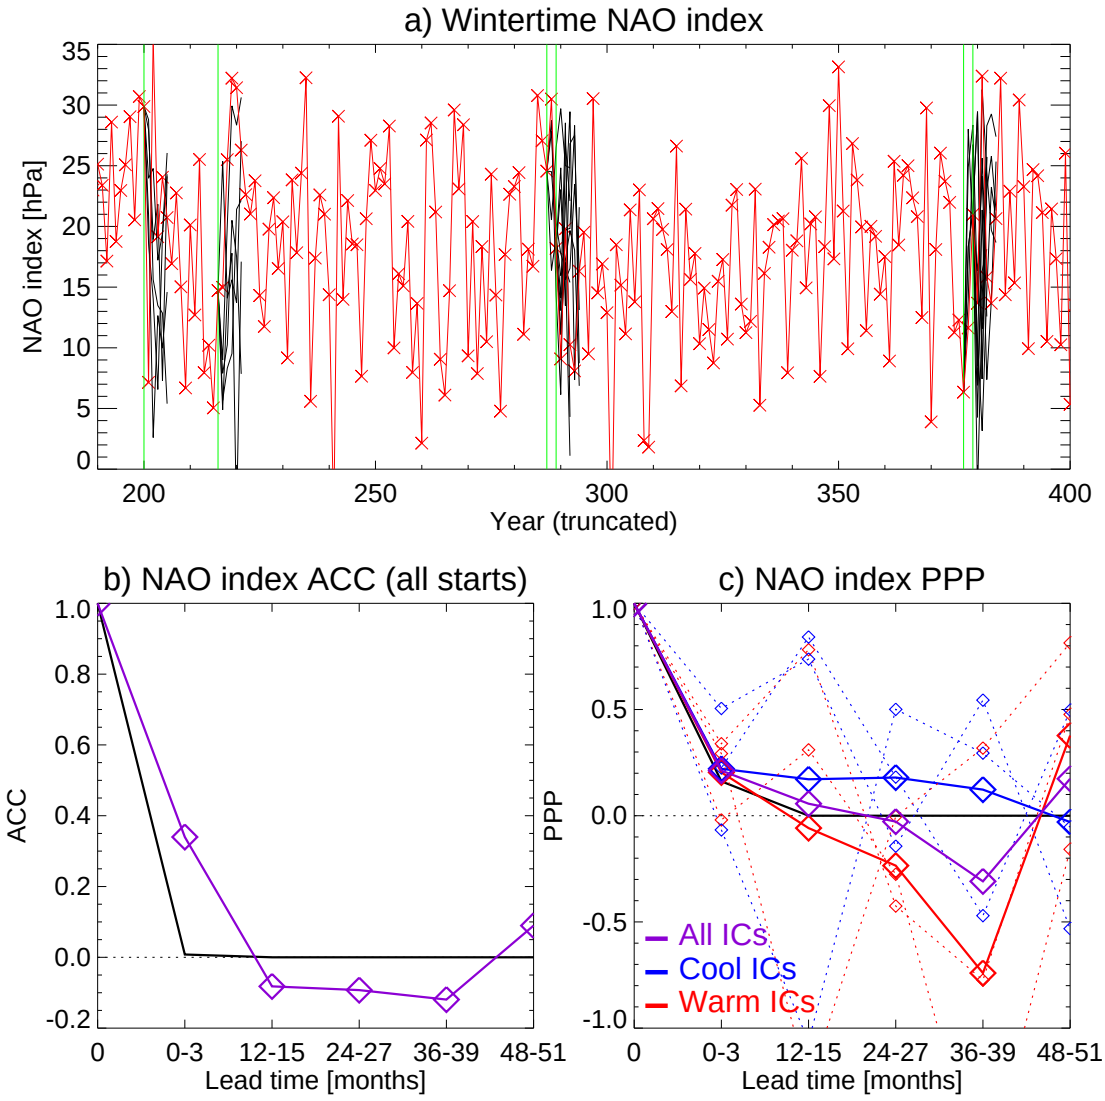

Figure 5.11: As Figure 5.4 but for the wintertime (December to February inclusive) North Atlantic Oscillation (NAO) index (defined as the difference between mean sea level pressure over the Azores and Iceland). The damped persistence model is based on seasonal (autumn to winter) autocorrelation. In addition, the prognostic potential predictability (PPP, c) is further separated into sets of initial conditions where the majority of the northern North Atlantic subpolar gyre is cool (blue) or warm (red), with individual initial conditions also shown (dotted). With respect to Figure 5.2, warm initial conditions correspond to spatial maps (a), (c), and (f), and cool initial conditions to spatial maps (b), (d), and (e).

5.5 Discussion

In this section we begin by summarising and discussing the origins of the predictability diagnosed in our initial condition ensemble experiments. We then discuss the relationship between these experiments and our previous analysis of the mechanism of internal decadal variability in the NA SPG (Chapter 2) and the plausibility of various elements of that variability given the available observations (Chapter 3). We conclude with some suggestions for future work in this context.

The limits on — and sources of — predictability in the NA SPG appear to vary with the process that supplies the potentially predictable variability in a given region. For example, between the Labrador Sea and NAC the dynamic height gradient suggests skill at lead times of up to three years (Figure 5.10), explaining some of the ACC in the southern NA SPG (Figure 5.6). However, in the eastern NA SPG, the ACC and PPP fall away more rapidly (Figure 5.6). This could be related to the importance of circulation variability in this region, which has much more interannual variability than temperature (not shown) and perhaps requires a larger ensemble set to capture. Alternatively, this region could represent one area in which atmospheric variability drives the NA SPG decadal variability (via the provision of white noise forcing), although this is far from certain.

In terms of the three measurement regions we have defined (East, North, South), the remaining question concerns the origin of the high ACC (Figure 5.6) and PPP (Figure 5.8, b) in the North region. To investigate this, we estimate the concatenated anomaly correlation coefficient (CACC, see Section 5.2 and Equation 5.4) for pairs of start dates based on large anomalies in each of the East, North, and South regions (Figure 5.12). When initialising based on large eastern temperature anomalies (Figure 5.12, a), the CACC suggests coherent evolution *through time* of the eastern subpolar gyre as well as the northern boundary, consistent with an important role for the mean circulation advecting anomalous temperature anomalies. That is, the skill in the North region appears to arise via initialisation of upstream anomalies that are then transported by the northern boundary currents.

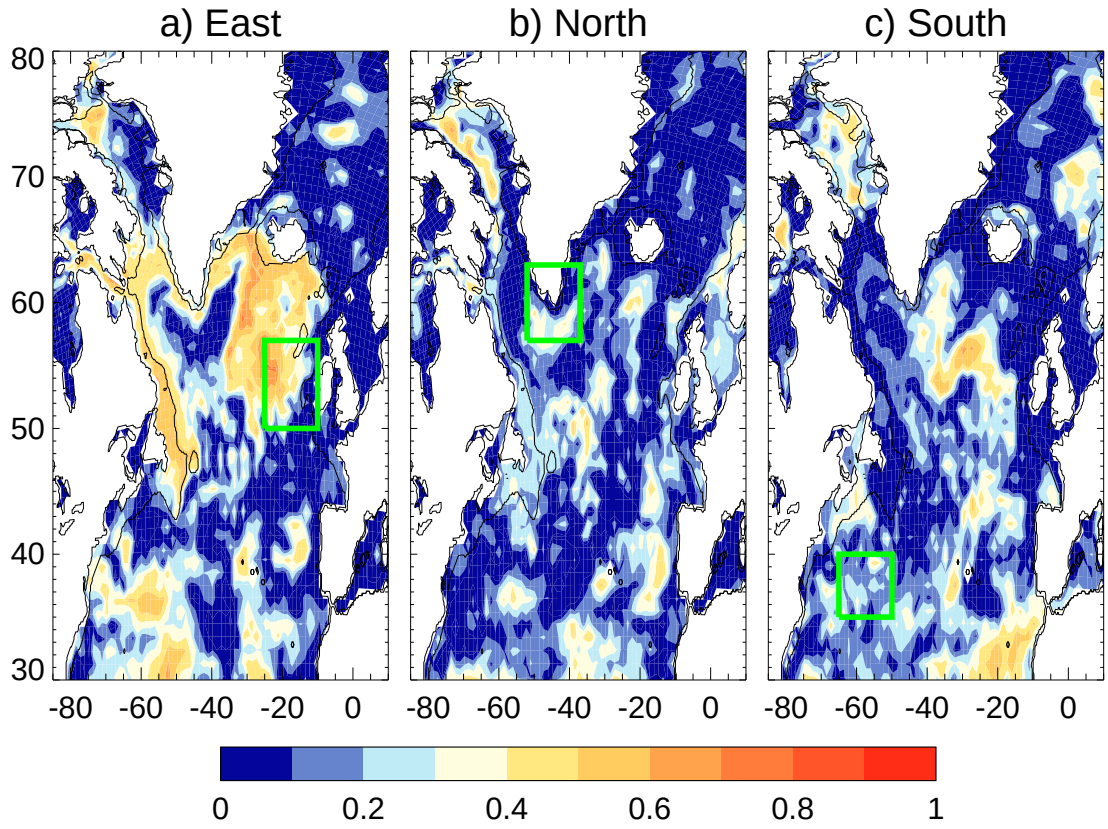

Figure 5.12: Concatenated anomaly correlation coefficient (CACC, as described in the text) in annual mean top 500m depth averaged temperature (T500) for initial condition ensembles based on: an eastern subpolar gyre index (a), a northern subpolar gyre index (b), and a southern subpolar gyre index (c). Each map represents the combination of a pair of ensembles describing anomalously warm and anomalously cool states (in each region) from the total of 6 initial condition ensembles. Green boxes highlight the region from which the initial conditions were chosen.

A similar approach can be taken for the North (Figure 5.12, b) or South (Figure 5.12, c) regions. For the North region, there appears to be little coherent evolution of temperature anomalies through time. Despite this, for both the central Labrador Sea (just south of Greenland) and northern edge of the NAC (east of Newfoundland) the temporal evolution (CACC) is larger than when initialising based on anomalies in the East and South, consistent with initialisation of the northern edge of the dynamic height gradient and a subsequent response at the southern edge (Figure 5.10). Nonetheless, the low CACC (compared to values of $r>0.6$ in for the East initial conditions) suggests the ensemble members become increasingly decorrelated within the 5 year window (recall that the CACC score is a function of the total lead time). Similarly, there are also limited regions of coherent evolution in the South ensemble within the eastern NA SPG (Figure 5.12, c),

although these are once again notably downstream of the region from which the initial conditions were chosen.

In general, initialisation based on large anomalies upstream of the region of interest results in the highest potential predictability (above analysis and Figure 5.8). In this sense, 'upstream' both describes the spatial evolution of variability associated with the diagnosed mechanism of decadal variability (Chapter 2) as well as the cyclonic nature of the NA SPG. Ultimately, the question is to what extent any given propagating anomaly within the NA SPG is representative of either 1) the diagnosed mode of decadal variability (that involves propagation along the mean circulation pathways), 2) another mode of variability at shorter or longer timescales, or 3) noise (or atmospheric forcing unrelated to the diagnosed decadal mode). The PPP shown in Figure 5.8 suggests at least some of the signal is not merely due to noise, and the analysis within Section 5.4.3 supports some role for the previously diagnosed mode of decadal variability, but what fraction remains unclear.

In our 'perfect-model' analysis, the NAO shows skill at lead times of 0–3 months, when initialised on December 1st, but this skill is much lower than that for the real-world wintertime NAO initialised at the beginning of November for a later version of this model (*Scaife et al.*, 2014). However, *Scaife et al.* (2014) use 20 start dates, each with 24 ensemble members, compared to the 6 start dates with 5 (plus the control) ensemble members in this analysis, which may explain some of the difference (*cf.* the effect of ensemble size on skill in their Figure 3). Related to this, another intriguing possibility for the limited skill in our analysis is that the ratio of predictable to unpredictable variability in the model and in reality may not be the same (*Scaife et al.*, 2014; *Eade et al.*, 2014), with the model exhibiting less predictability than reality. If this remains true for North Atlantic ocean indices then it is possible that the skill we see in the oceanic evolution of the NA SPG could be vastly improved in the real world by increasing the number of start dates and ensemble members. This may also explain why the sometimes limited skill in these initial condition ensembles exists in a system that clearly shows a significant decadal mode of variability (*cf.* Figure 2.4, b).

As noted above, our analysis has been within a 'perfect-model' framework, and as such

specific ACC scores are not necessarily indicative of skill in predicting real-world variability. However, we can combine this analysis with our extensive comparison against observations (Chapter 3) to speculate on the regions where the ACC scores are more/less likely to be representative of real skill. We do this using the three measurement regions (East, North, and South) defined and used previously.

In the eastern NA SPG there is low skill in the model initial condition ensembles, which may well be even less so in reality. This is because the model skill likely arises at least partly via the advection of anomalies along the southern edge of the NA SPG, which is a process with a much greater depth extent (and thus memory) in the model than in observations (Figure 3.11). Conversely, in the North measurement region, there is high skill in the model (Figure 5.6), likely related to the initialisation of boundary currents (Figure 5.12), which we previously suggested the model was far better able to capture than previous, lower resolution, models (Section 3.6). Finally, in the South measurement region, although the skill/predictability is moderate, arising from the initialisation and simulation of the dynamic height gradient between the Labrador Sea and NAC (Figure 5.10), we suggest the skill at predicting this region in reality will be low (at least when hindcasting recent decades — see next). This is because the negative feedback simulated in the model relies on the density variability in the Labrador Sea being dominated by temperature variability (rather than salinity, Figure 3.9, c), whereas in reality it appears that salinity is the dominant driver (Figure 3.9, a and b). Therefore, initialising with warm/saline anomalies in the Labrador Sea may give skilful predictions of the downstream model evolution but are unlikely to give skilful predictions of reality. However, it should be noted that the observed driver of Labrador Sea density variability appears to be non-stationary (Figure 3.10, b) and so it is possible that decadal predictions in this region made with this model may have more skill in future predictions than hindcasts. This represents an important caveat when using the hindcast skill of a model to assess the reliability of future predictions. However, one should also note that we have assessed the skill of the model without external forcings, which could be the driver of the non-stationarity mentioned above and would also provide increasing amounts of the longer term predictability (*Branstator and Teng*, 2012).

As outlined in Section 5.2.1, our two goals when conducting these initial condition en-

semble experiments were to 1) test the robustness of processes found to be important for the simulated decadal variability, and to 2) assess the predictability characteristics of the model in the NA SPG. In these regards, the analysis in this chapter represent a far more in-depth investigation into the plausibility of simulated decadal variability, and the predictability it implies, than is normally found in the literature on model mechanisms. Nonetheless, the experimental design could be improved if used as a framework for future studies. For example, of the six initial condition sets, there are closer to just four independent start dates as both the East -T1000 and North -T1000 start dates and the East +T1000 and South -T1000 start dates are separated by only two years (Figure 5.1). Although this is not ideal, it is consistent with the mechanism whereby anomalies propagate around the NA SPG before inducing opposite-sign anomalies in the southern NA SPG and we note at least that it is not always the same pairs of start dates that are close together. Further work could improve the algorithm to choose the start dates, by perhaps ensuring they are well separated in time, and also by defining indices that are more well separated through the temporal evolution of the variability. For example, the East and North indices are separated by a much shorter timescale than the South and East indices. Using the periodicity of the region (17 years, Chapter 2) suggests indices separated by lags of around 3 years would be optimal (noting that there are six indices over three regions if choosing both positive and negative anomalies). Finally, initialising with anomalies confined to each region (and zero elsewhere) would help to categorically confirm that the predictability arises from initialising upstream, though this would require artificially creating dynamically consistent anomalies. This would be a considerably more complex procedure than merely providing anomalies taken straight from the control simulation, and would also require the redefining of the model 'truth' as well as reinterpretation of the climatological variability.

5.6 Chapter conclusions

We have conducted initial condition ensemble experiments to jointly assess the robustness of various features of the simulated decadal variability and assess the predictability characteristics of the model

- There is skill at predicting subpolar gyre top 500m depth averaged temperatures at lead times up to and including 2 years (despite the small ensemble size) and potential predictability at all lead times up to and including 5 years.

- For depth averaged temperatures, the skill and potential predictability are largest in the Labrador Sea region as well as on the northern edge of the NAC (possibility related to the negative feedback between Labrador Sea and NAC temperatures).

- Analysis of regional PPP suggests predictability in a given region often arises from initialising large temperature anomalies upstream of that region, although it is not clear whether the predictability arises via the diagnosed mechanism of decadal variability or other means.

- Labrador Sea deep water formation and the negative feedback between Labrador Sea and NAC temperatures appear to provide the skill in the South measurement region (on the southern edge of the NA SPG). Northern boundary currents that advect temperature anomalies westwards on the northern edge of the NA SPG appear to provide the skill in the North measurement region.

- There is little skill in the NAO index, possibly related to the small number of ensembles and ensemble members, but some suggestion that cool NA SPG initial conditions provide greater potential predictability of positive NAO anomalies than warm NA SPG initial conditions provide of negative NAO anomalies, consistent with the analysis in Chapter 4.

- Comparison of the regional skill in the model and the plausibility of various aspects of the model variability (Chapter 3) suggests that the model may be expected to overestimate the real-world skill in the eastern NA SPG and southern NA SPG (related to the driver of density variability in the Labrador Sea) but may present an improvement in real-world skill levels related to the northern NA SPG due to its improved representation of boundary currents compared to previous generations of models.

As noted here and in all previous chapters, a version of this model will be used for seasonal to decadal prediction at the Met Office (*MacLachlan et al.*, 2015) and such analyses

as these (that provide comparisons of the model skill and the model reliability, with a focus on the processes that provide that skill/reliability) will be important in understanding predictions made with that model.

We have shown that a negative feedback involving the Labrador Sea is important for the NA SPG decadal variability in HadGEM3 (Chapter 2), and that this feedback may be different in reality (Chapter 3), with implications for the real-world predictability with this model (this chapter). As such, in the next chapter (Chapter 6) we widen the focus from a single model (HadGEM3) and attempt to understand the wider relationships between biases in climate models and the many documented mechanisms of decadal variability.

Chapter 6

Exploring the impact of CMIP5 model biases on the simulation of North Atlantic decadal variability

6.1 Introduction

Initial conditions within the North Atlantic subpolar gyre (NA SPG) have been shown to be important in making skilful decadal forecasts (*Dunstone et al.*, 2011). However, even when given similar initial conditions, decadal predictions of the North Atlantic between different models can be quite different (*Pohlmann et al.*, 2013). Indeed, decadal variability in the North Atlantic Ocean, although extensively investigated in both coupled and uncoupled models (e.g. *Eden and Willebrand*, 2001; *Dai et al.*, 2005; *Dong and Sutton*, 2005; *Cabanes et al.*, 2008; *Alvarez-Garcia et al.*, 2008; *Biastoch et al.*, 2008a), is still poorly understood, in part due to the paucity of constraining observational data (Chapter 3 and *Good et al.*, 2013). In this chapter we conduct one of the first, large multi-model analyses of variability in the NA SPG region to investigate whether there are relationships between broad characteristics of the models and their manifestation of decadal variability. We begin by briefly recapping the evidence for bidecadal variability in the NA SPG and the disparate model representations in Section 6.2 before describing

the model data we use in Section 6.3. In Section 6.4 we analyse the results including the relationships between model biases and variability before putting these into a wider context in Section 6.5. Chapter conclusions are presented in Section 6.6. This chapter is reproduced in a similar form in *Menary et al.* (2015b).

6.2 The Labrador Sea and its role in North Atlantic decadal variability

An array of coupled, partially coupled, and ocean-only mechanisms have been proposed to describe simulated variability in the NA SPG (see above references and *Liu* (2012) for a review) whilst the pacemaker of this variability has been attributed to a variety of processes such as Rossby Wave propagation (*Sévellec and Fedorov*, 2013), mean advection timescales (*Delworth et al.*, 1993), and interaction with the deep flow (*Eden and Willebrand*, 2001). Despite the large number of postulated mechanisms and key processes, a periodicity of around ≈20 years has begun to emerge as the common timescale of simulated multi-annual/decadal variability in the NA SPG (*Frankcombe et al.*, 2010), consistent with some high resolution palaeo records in this region (*Sicre et al.*, 2008; *Chylek et al.*, 2012). This timescale describes variability generally confined to the subpolar gyre, with feedbacks to other regions in the Arctic and/or subtropical North Atlantic usually involving longer timescales (*Jungclaus et al.*, 2005; *Menary et al.*, 2012).

Within the NA SPG, a key region in most mechanisms of decadal (or longer) variability is the deep water formation (DWF) region of the Labrador Sea, where surface signals can spread to depth and impact the large-scale dynamics of the region (*Medhaug et al.*, 2012), though we note that some models can locate their main DWF regions elsewhere (*Ba et al.*, 2014). Similar to the disparate mechanisms of variability, although models generally agree that the Labrador Sea (or model equivalent) is important, they are split roughly evenly on whether decadal timescale changes in density in this region are controlled by either temperature or salinity (Chapter 1).

In summary, there remain many systematic differences across the present generation of climate models in their representation of North Atlantic decadal variability. In Chapter 2 we hypothesised that these differences may in part be related to the relationship between mean state biases in the NA SPG and whether temperature or salinity control density changes. This would have ramifications for situations in which mean state biases and the evolution of the system are assumed to be independent, such as decadal forecasts using 'anomaly-assimilation'. In this chapter, we test the validity of this simple hypothesis linking mean state temperature and salinity biases with the controller of density variability and investigate whether this has implications for the manifestation of decadal variability in the NA SPG.

6.3 Methods/Models

Phase 5 of the Coupled Model Intercomparison Project (CMIP5) represents a coordinated approach to simulating global climate under a variety of scenarios (*Taylor et al.*, 2012). We examine 40 pre-industrial control simulations from 40 individual models from the CMIP5 archive (Tables 6.1 and 6.2). We use control simulations in all cases to isolate the internal variability whilst their length allows us to maximise the signal to noise ratio for each model. In addition, we examine 2 iterations of the latest high resolution coupled climate model from the Met Office Hadley Centre: 'HadGEM3' (which comprises Global Atmosphere (GA) version 3.0 (*Walters et al.*, 2011) and the NEMO ocean model version 3.2 (*Madec*, 2008), see Chapter 2 and *Duchez et al.* (2014) and references therein) and 'GC2' (which comprises GA6.0 and NEMO version 3.4, *Williams et al.*, 2015). The versions of HadGEM3 and GC2 we analyse are control simulations run with interannually constant forcings appropriate for the years 2000 and 1850 respectively. We compare the simulations to optimally interpolated observations from the EN4 dataset (*Good et al.*, 2013). We use EN4 data from the most well observed period 1960–2014 but note that, due to under-sampling, we have far more confidence in the estimation of the time mean values than the decadal variability (see Chapter 3).

All models are regridded on to a regular $1\times1^{\circ}$ horizontal grid to aid analysis. Testing

Institute (CMIP5 name)	Model (CMIP5 name)	Length (years)	Effective ocean resolution (degrees)
BCC	bcc-csm1-1	500	0.99
BCC	bcc-csm1-1-m	400	0.99
BNU	BNU-ESM	559	1.0
CCCma	CanESM2	996	1.1
CMCC	CMCC-CESM	277	1.6
CMCC	CMCC-CM	330	1.6
CMCC	CMCC-CMS	500	1.6
CNRM-CERFACS	CNRM-CM5	850	0.85
CNRM-CERFACS	CNRM-CM5-2	410	0.85
CSIRO-BOM	ACCESS1-0	500	0.75
CSIRO-BOM	ACCESS1-3	500	0.75
CSIRO-QCCCE	CSIRO-Mk3-6-0	500	1.35
FIO	FIO-ESM	765	0.62
ICHEC	EC-EARTH	452	0.85
IPSL	IPSL-CM5A-LR	1000	1.6
IPSL	IPSL-CM5A-MR	300	1.6
IPSL	IPSL-CM5B-LR	300	1.6
LASG-CESS	FGOALS-g2	700	1.0
LASG-IAP	FGOALS-s2	501	1.0
MOHC	HadGEM2-CC	241	1.0
MOHC	HadGEM2-ES	578	1.0

Table 6.1: A summary of the first 21 models used in this analysis. All CMIP5 models simulate pre-industrial climates. For further details of the CMIP5 models and institutions see Table 9.A.1 of *Flato et al.* (2013) and references therein. For the remaining models see Table 6.2.

Institute (CMIP5 name)	Model (CMIP5 name)	Length (years)	Effective ocean resolution (degrees)
MPI-M	MPI-ESM-LR	1000	0.47
MPI-M	MPI-ESM-MR	1000	0.53
MPI-M	MPI-ESM-P	1156	0.47
MRI	MRI-CGCM3	500	0.67
MRI	MRI-ESM1	251	0.67
NASA-GISS	GISS-E2-H	540	1.0
NASA-GISS	GISS-E2-H-CC	251	1.0
NASA-GISS	GISS-E2-R	550	1.1
NASA-GISS	GISS-E2-R-CC	251	1.1
NCAR	CCSM4	1051	0.62
NCC	NorESM1-M	501	0.62
NCC	NorESM1-ME	252	0.62
NOAA-GFDL	GFDL-CM3	500	0.99
NOAA-GFDL	GFDL-ESM2G	500	0.96
NOAA-GFDL	GFDL-ESM2M	500	0.99
NSF-DOE-NCAR	CESM1-BGC	500	0.62
NSF-DOE-NCAR	CESM1-CAM5	319	0.62
NSF-DOE-NCAR	CESM1-FASTCHEM	222	0.62
NSF-DOE-NCAR	CESM1-WACCM	200	0.62
MOHC (non-CMIP5)	HadGEM3	473	0.21
MOHC (non-CMIP5)	GC2	269	0.21

Table 6.2: A summary of the second 21 models used in this analysis. All models simulate pre-industrial climates except HadGEM3, which uses forcings appropriate for the year 2000 (GC2 uses forcing appropriate for the year 1850).

with HadGEM3 and GC2 (not shown), using both original and regridded data, suggests regridding has very little impact on our subsequent results. The original vertical discretisations are left unaltered. Additionally, all models are linearly detrended prior to analysis but this again has little effect on our results (not shown). Unless otherwise stated, top 500m depth averaged annual mean data is used.

6.4 Results

6.4.1 Biases

Models exhibit a wide variety of temperature and salinity (T/S) biases in the northern North Atlantic (Figure 6.1). Both the magnitude and spatial structure of the bias varies greatly, up to $\pm$4K and $\pm$1PSU for top 500m depth averaged temperature and salinity respectively. Models from the same institution often share similar biases, perhaps associated with a reduction in the effective number of degrees of freedom of our sample (*Knutti et al.*, 2013), but this is not always the case (e.g. IPSL-CM5A-LR and IPSL-CM5B-LR, or GISS-E2-H and GISS-E2-H-CC).

In addition to the mean state biases, interannual variability in the northern North Atlantic (as diagnosed by the annual standard deviation) varies from model to model. Figure 6.2 highlights the combination of mean state biases (compared to EN4) and interannual variability for the volume averaged Labrador Sea (55–65°N, 45–65°W, top 500m) across the entire 42-model ensemble. For each model, simulated annual standard deviations in temperature and salinity are approximately 0.5K and 0.1PSU respectively, significantly smaller than the inter-model spread in mean state temperature and salinity. In general there is a positive correlation between temperature and salinity biases (r=0.85) to the extent that biases are largely density compensated, which holds throughout the NA SPG (not shown). However, a small density bias remains, which is typically due to the salinity biases, *i.e.* models which are warmer and saltier (Figure 6.2, top right) are also generally denser than models which are cooler and fresher (Figure 6.2, bottom left). The origin of this co-variability is unclear but it is consistent with important roles for either evapora-

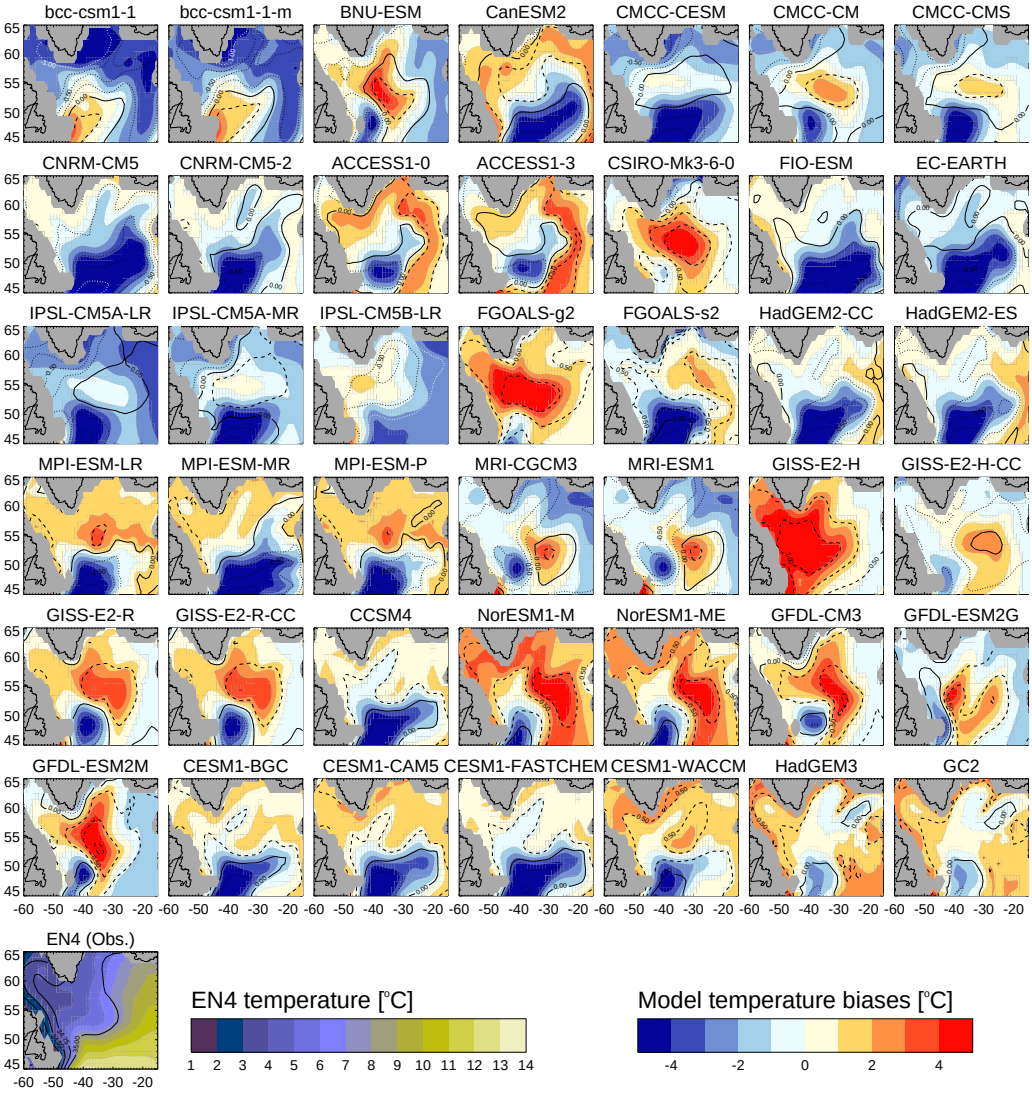

Figure 6.1: Top 500m depth averaged time mean temperature (colours) and salinity (contours) biases in the full model ensemble, relative to EN4. Salinity contours are solid for zero bias, dotted for negative bias, and dashed for positive bias, with a contour interval of 0.25PSU. Time mean absolute values in EN4 are also shown for comparison.

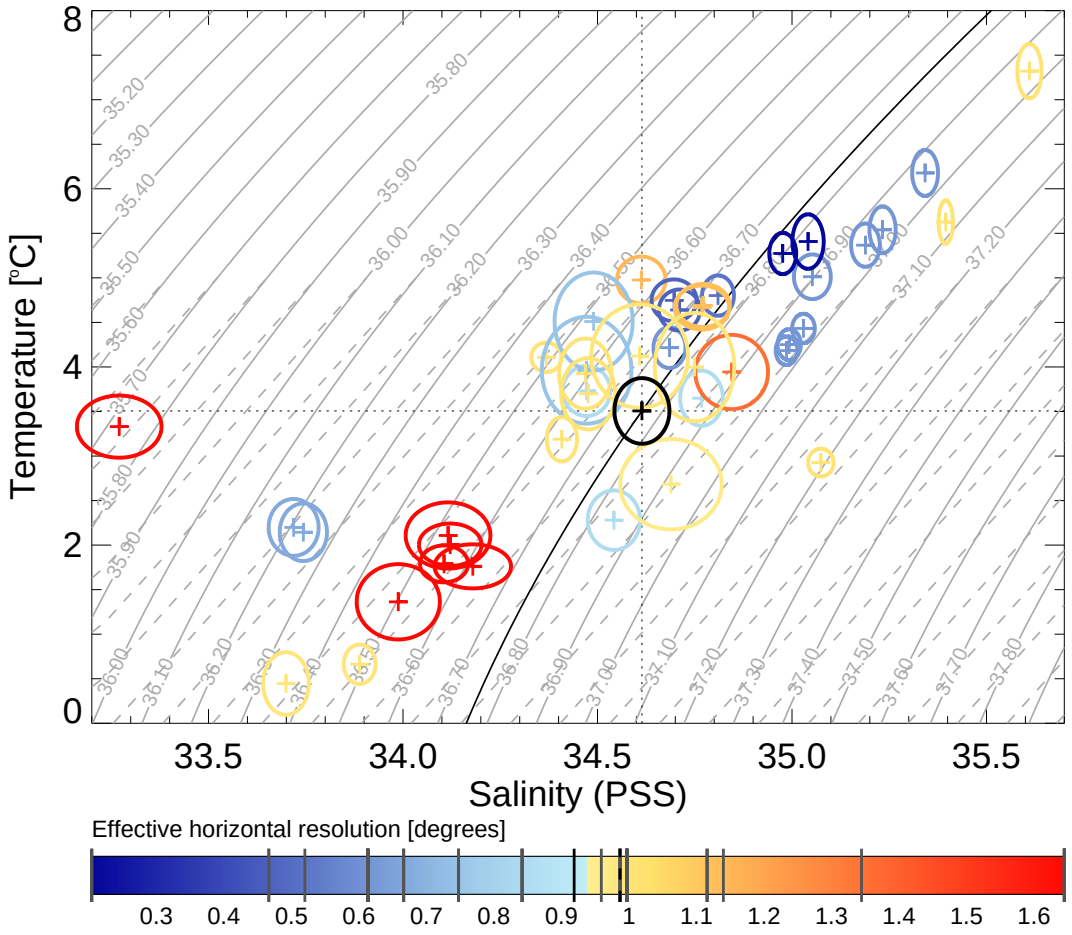

Figure 6.2: Top 500m depth averaged temperature and salinity mean (crosses) and annual ± 1 standard deviation range (ellipses) in the Labrador Sea for the model control simulations (colours) and observations (EN4, black). All time series linearly detrended prior to computing the standard deviation. Contours of constant potential density (relative to 2000m) are also plotted (grey, solid). A linear fit to these contours, taken when T=5°C and extrapolated for cooler temperatures, is shown to highlight the non-linear nature of the seawater equation of state (grey, dashed). The black contour highlights the mean density in EN4. The models have been coloured by their effective horizontal resolution in the NA SPG, the exact values of which have been indicated in grey on top of the colour bar (see Section 6.4.4, note that some resolutions occur more than once) along with the mean resolution (black, solid) and median resolution (black, dashed).

tion, deep convection (as the subsurface Labrador Sea is warmer/saltier than the surface), or ocean dynamics (see Section 6.5).

6.4.2 Density control

One of the key uncertainties in climate model simulations of the North Atlantic region is whether these simulations imply that the density changes associated with DWF are temperature or salinity controlled (Chapter 1). Although individual convective events are likely due to rapid wind-induced cooling (rather than salinifying) of surface waters during winter time, the frequency and/or intensity of these events depends on both the background mean and interannual variability of temperature and salinity profiles. As such, both temperature and/or salinity could plausibly be said to control interannual variability in density changes associated with DWF in the Labrador Sea.

In order to quantify whether temperature or salinity are controlling density changes in the Labrador Sea in our 42-model ensemble we follow the decomposition of *Delworth et al.* (1993) and decompose density changes into those due to temperature and those due to salinity as in Chapter 3. This assumes a linear decomposition of the non-linear equation of state, but we suggest this is not a bad assumption given the small size of interannual T/S variability for any given model (*cf.* Figure 6.2) and indeed the linear reconstruction explains >96% of the annual variance in density in all models. This results in the familiar equation for $\rho_{TorScontrol}$, the degree of temperature or salinity control of Labrador Sea density changes:

$$\rho_{TorScontrol} = \rho_{TControl} - \rho_{SControl} \tag{6.1}$$

This is identical to Equation 3.3 but repeated here for clarity. Figure 6.3 compares $\rho_{TorScontrol}$ with the biases in the Labrador Sea. There are strong correlations between the biases and what controls density changes with regression slopes significantly different from zero. Density changes in warm and salty models appear to be dominated by variability in temperature, whereas density changes in cool and fresh models appear to be dominated by variability in salinity. *i.e.* T/S biases appear to explain what controls interannual density variability in models and explains to a large extent the spread in the literature on this topic. This can be understood by comparison to Figure 6.2 in which the potential density contours have been overlaid. For temperatures less than 5°C a linear fit

to the density equation (extrapolated from $\rho = \rho(T{=}5, S)$) has also been added to highlight the non-linear nature of the seawater equation of state. A given salinity change has a larger effect on density when the temperature is cool than when it is warm. Similarly, when the background mean temperature is warm, a given temperature change is now more likely/able to play an important role than when it is cool. Note also that, although *mean* density biases are due to salinity, it is the combination of both temperature and salinity biases which systematically affect the controller of density *variability*. We find this result holds for both unfiltered (but detrended) time series representing interannual relationships as well as filtered time series (removing periods less than 5 years) investigating decadal relationships. Indeed, even estimating the relationships on 'multi-decadal' timescales (by filtering to remove periods less than 30 years) has little effect (not shown), *i.e.* it is not the case that temperature variability dominates on interannual timescales and salinity variability dominates on decadal timescales.

6.4.2.1 Scaling analysis

Despite the above analysis, it is also possible that the magnitude of temperature or salinity interannual variability across models varies with the biases and thus affects what we estimate controls density variability. Indeed, the magnitude of T/S interannual variability across models shows some correlation with the biases. As such, the magnitude of T/S variability could also affect what we estimate controls density variability. This would lead to the alternative interpretation that models which, for example, have larger amplitude interannual salinity *variability* (rather than particular mean state biases) would also have salinity-controlled density changes, and similarly for temperature. We account for this by scaling the T/S variability in all the models to have the same magnitude of temperature or salinity variability set to the multi-model mean. *i.e.* we effectively scale $\frac{\mathrm{d}T}{\mathrm{d}t}$ and $\frac{\mathrm{d}S}{\mathrm{d}t}$ in Equation 6.2.

$$\frac{\mathrm{d}\rho}{\mathrm{d}t} = \frac{\mathrm{d}T}{\mathrm{d}t} \cdot \frac{\mathrm{d}\rho}{\mathrm{d}T} + \frac{\mathrm{d}S}{\mathrm{d}t} \cdot \frac{\mathrm{d}\rho}{\mathrm{d}S} \tag{6.2}$$

Where $\frac{\mathrm{d}\rho}{\mathrm{d}T}$ and $\frac{\mathrm{d}\rho}{\mathrm{d}S}$ measure the density change for a given temperature or salinity change

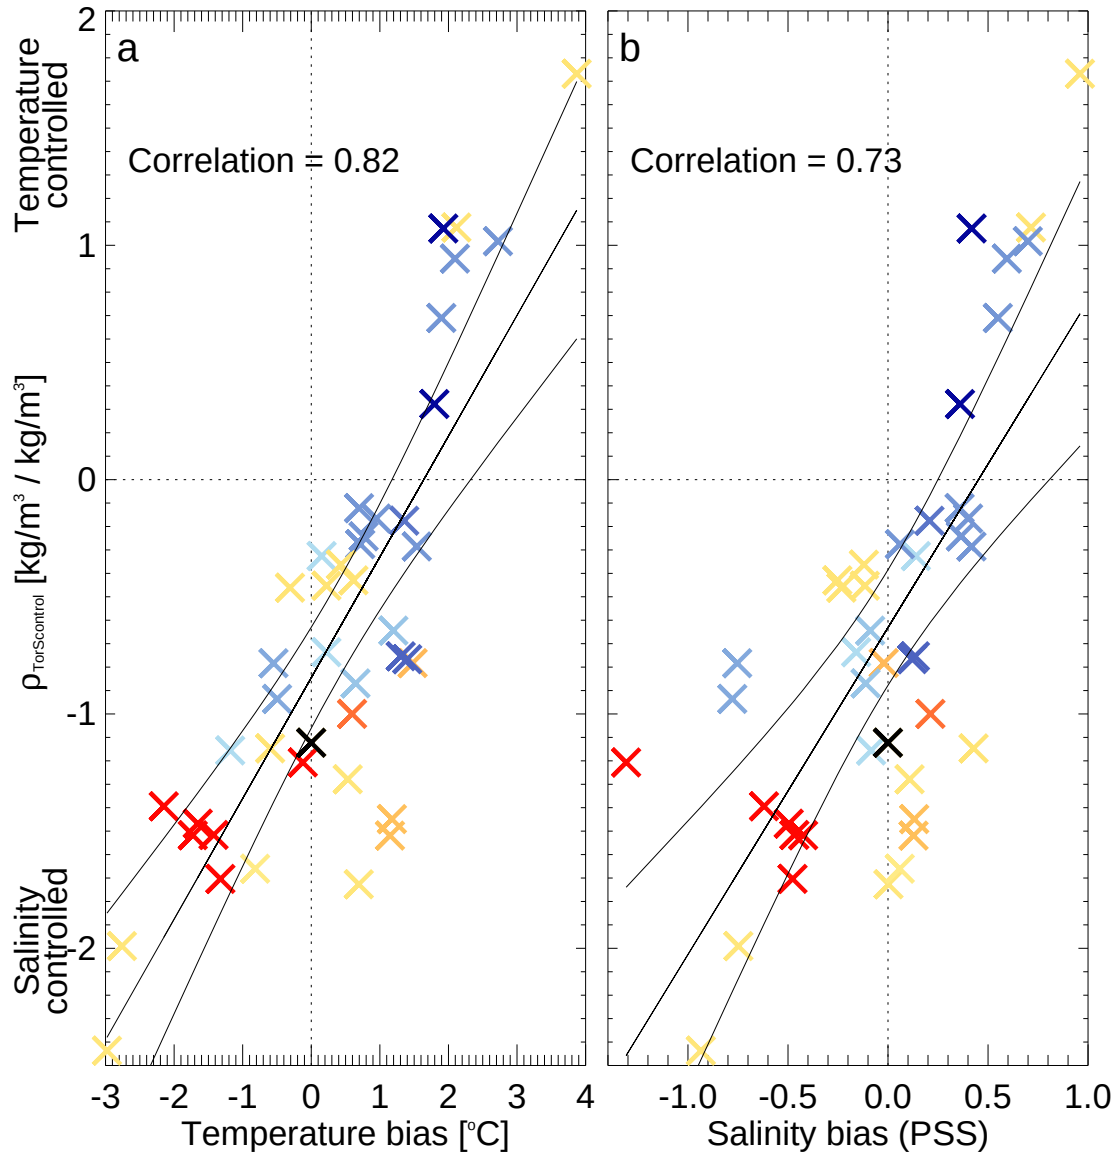

Figure 6.3: Top 500m depth averaged temperature (left) and salinity (right) biases in the Labrador Sea against a measure of the controller of interannual density changes in the Labrador Sea ($\rho_{TorScontrol}$). The colours represent each model's effective horizontal resolution in the NA SPG, as in Figure 6.2. An estimate of $\rho_{TorScontrol}$ for EN4 has also been added (with zero bias, black). Regression slopes are computed by ordinary least-squares regression with the envelope representing an estimate of the 95% confidence interval on the slope.

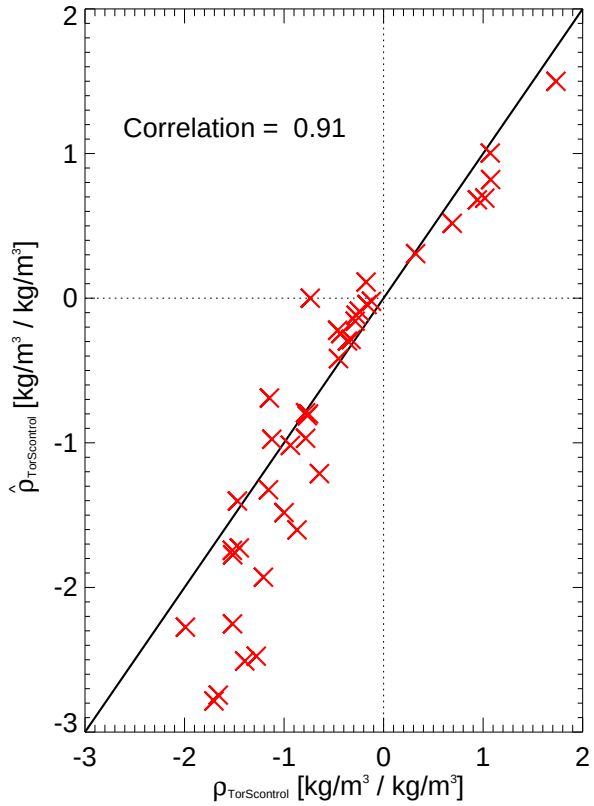

Figure 6.4: The effect of scaling temperature and salinity variability. On the x-axis is $\rho_{TorScontrol}$, computed as described in the text. On the y-axis is $\widehat{\rho}_{TorScontrol}$ which is as $\rho_{TorScontrol}$ but calculated using T/S which have been scaled to have the same standard deviation in all models (set to the multi-model mean standard deviation). The solid line represents a simple 1:1 relationship.

and are thus a function of the mean state biases in both temperature and salinity (*i.e.* for cooler temperatures, a given temperature change has a lesser effect on density than at warmer temperatures, *cf.* Figure 6.2). $\frac{\mathrm{d}T}{\mathrm{d}t}$ and $\frac{\mathrm{d}S}{\mathrm{d}t}$ represent the magnitude of T/S variability through time (t) and are scaled in this sensitivity test. We then recompute $\rho_{TorScontrol}$ with these scaled variables to give $\widehat{\rho}_{TorScontrol}$. If the magnitude of variability in T/S does affect what controls density variability then there should be large differences between $\rho_{TorScontrol}$ calculated with/without scaling. However, we find that scaling has little effect, yielding very similar values for $\widehat{\rho}_{TorScontrol}$ as $\rho_{TorScontrol}$ (Figure 6.4). As such, although T/S variability correlates with $\rho_{TorScontrol}$ we find that the magnitude of T/S variability is unlikely to causally affect $\rho_{TorScontrol}$ and suggest that the causation is more likely to be in the opposite direction.

In summary, it appears that simulated Labrador Sea mean state T/S biases do appear to

be related to whether interannual/decadal density changes are temperature- or salinity-controlled. The final question we investigate is whether these mean state biases — and their apparent relationship to what controls density variability — have a systematic impact on the mechanisms of variability within the North Atlantic.

6.4.3 Labrador Sea feedbacks

The Labrador Sea has been shown to be an important region in differing mechanisms of decadal variability in the North Atlantic (Chapter 2 and *Escudier et al.*, 2013) — here we examine whether these differences can be simply understood by investigating negative feedbacks associated with depth-averaged temperature anomalies (denoted T500) in the Labrador Sea. That is, when it is anomalously warm/cold in the Labrador Sea, where (if anywhere) in the ocean do the anomalies that reverse the state of the Labrador Sea originate. We use depth averages over the top 500m for comparison with our previous analysis and because these have a high signal to noise ratio. We focus on negative feedbacks as these are required in order to create significant periodicity in the absence of periodic forcing but limit our analysis to the ocean. Indeed, we make no *a priori* assumptions about the precise details of the feedback mechanism, be it coupled or ocean-only. We use temperature (rather than, for example, salinity) for simplicity but note that, whatever the mechanism, Labrador Sea DWF is likely to leave signals in both temperature and salinity.

Figure 6.5a shows the location of maximal negative correlation of T500 with the marked Labrador Sea region at a lag of 1 year (*i.e.* the map lags the Labrador Sea index by 1 year). A lag of 1 year is used to reduce the confounding effect of the North Atlantic Oscillation (NAO) forcing both the Labrador Sea and rest of the NA SPG simultaneously. The location of this negative correlation follows a curve broadly consistent with the shape of the NA SPG. These anomalies propagate into the Labrador Sea and reverse the sign of the initial anomaly (not shown) over the course of a few years to a decade with the timescales varying from model to model (see also Section 6.5). Many models exhibit their maximal negative correlation in the North Atlantic Current region (consistent with the feedback in HadGEM3, described in Chapter 2, Section 2.4.6) whereas some show a

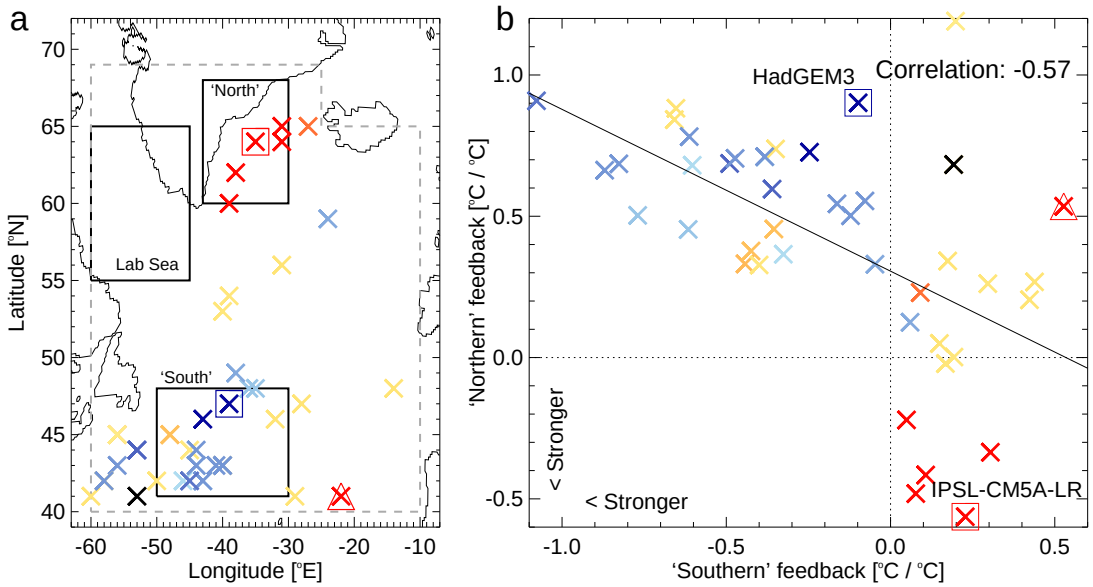

Figure 6.5: a) The location, constrained to be within the grey dashed boundary, of the maximum negative correlation with the Labrador Sea depth averaged temperature (T500, marked region) in each model. The map lags the Labrador Sea index by 1 year; annual, detrended data is used. b) The location and strength of feedbacks involving the Labrador Sea, estimated by regressing T500 in the Labrador Sea against a 'northern' and 'southern' index, as described in the text and marked on the map, for the models (colours) and EN4 (black). The trend line and correlation are estimated using a Bayesian approach to total least-squares (TLS) regression (*Kelly*, 2007) with the uncertainty on the individual regression slopes for each model used as the measurement errors within the TLS calculation. The HadGEM3 and IPSL-CM5A-LR models have been highlighted with squares and the IPSL-CM5B-LR model with a triangle (see text). The colours represent each model's effective horizontal resolution in the NA SPG, as in Figure 6.2.

preference for the East Greenland Current (consistent with elements of the mechanism, though not the specific timescales, of *Escudier et al.* (2013)).

To quantify the differences highlighted in Figure 6.5a, we design two metrics that aim to characterise these negative feedbacks in ocean temperatures: the magnitude of the (negative) regression gradient between the Labrador Sea index and an index of the North Atlantic Current region (30–50°W, 41–48°N: 'South'), and the magnitude of the (negative) regression gradient between the Labrador Sea index and an index of the East Greenland Current region (30–45°W, 60–70°N: 'North'). It can be seen from Figure 6.5b that in general models have a preference for one feedback or the other, with a stronger 'northern' feedback implying a weaker 'southern' one, and vice versa. The correlation between the two feedbacks is -0.56, which rises to -0.71 when using zero lag (Figure 6.6). This

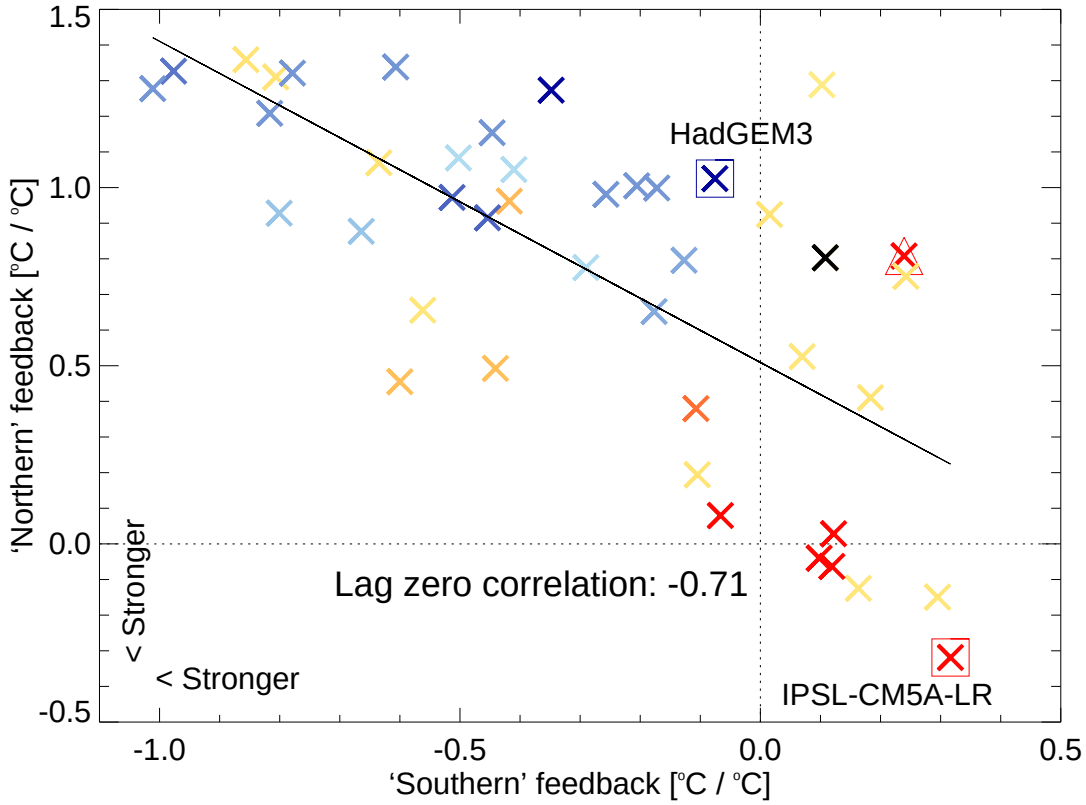

Figure 6.6: As Figure 6.5b but for zero lag. The location and strength of feedbacks involving the Labrador Sea, estimated by regressing T500 in the Labrador Sea against a 'northern' and 'southern' index with zero lag, for the models (colours) and EN4 (black). The trend line and correlation are estimated using total least-squares regression as in Figure 6.5. The HadGEM3 and IPSL-CM5A-LR models have been highlighted with squares and the IPSL-CM5B-LR model with a triangle (see text). The colours represent each model's effective horizontal resolution in the NA SPG, as in Figure 6.2.

method of analysis is also consistent with the individual simulations investigated in previous work where the southern and northern portions of the NA SPG were important in HadGEM3 (Chapter 2) and IPSL-CM5A-LR (*Escudier et al.*, 2013) respectively (highlighted on the figure). Note also the low resolution outlier IPSL-CM5B-LR (triangle), which has improved tropical atmospheric dynamics compared to IPSL-CM5A-LR but a severely worsened representation of the North Atlantic Ocean with a control Atlantic overturning strength of 4 Sverdrups (*Dufresne et al.*, 2013).

To compare against reality we add an observational estimate of the feedback using the years 1960–2014 from EN4. The observations occupy a zone where neither a northern nor southern feedback dominates, according to our analysis. This may be because in-

terannual/decadal variability in the observations during this period is also related to a number of other factors, including: Great Salinity Anomalies (*Dickson et al.*, 1988), recent rapid warmings of the subpolar gyre (*Robson et al.*, 2012), and possibly a larger role for the NAO (*Scaife et al.*, 2011), to name but a few. Unfortunately the confounding influence of transient climate change, along with a much shorter record (there are only sporadic observations in the northern North Atlantic prior to 1960), inhibits detailed comparison (see Section 6.5). Nevertheless, the observational estimate lies within the simulated ranges implying that accurate representations of the observed North Atlantic variability may be sampled by some of the model mechanisms.

6.4.4 Resolution

The refinement of horizontal resolution in the ocean to permit and even resolve eddies/Rossby waves may be particularly important for regions such as the Labrador Sea (*Gelderloos et al.*, 2011; *Marzocchi et al.*, 2015). To account for non-regular grids we estimate the effective resolution of each model in the North Atlantic region (40–76°N, 0–65°W) by simply counting the number of grid cells with centres within this domain on each model's native grid. These range from 1.6° for the IPSL and CMCC models to 0.21° for the new Hadley Centre models (though note that the new Hadley Centre models are post-CMIP5 models).

The effective resolution of the models does appear to show a broad relationship to their biases, with warm and salty models also being of higher resolution than cool and fresh models (Figure 6.2). Interestingly, the effective resolution does not appear to correlate with the absolute value of the bias, *i.e.* higher resolution models cannot be said to be 'better', in terms of their depth averaged T/S biases in the Labrador Sea, than lower resolution models. In addition, we find that higher resolution models are more likely to show temperature-controlled interannual/decadal density changes, whereas these density changes are more likely to be salinity-controlled in lower resolution models (Figure 6.3). Lastly, models which have a lower effective resolution in the North Atlantic are much more likely to exhibit signals indicative of a northern feedback than models of higher resolution (Figure 6.5), though this selection between northern and southern feedbacks

is not as clear when colouring by $\rho_{TorScontrol}$ (not shown). Possible reasons for the relationship between biases and resolution are discussed next (Section 6.5).

6.5 Discussion

We have demonstrated, across a 42-member coupled-model ensemble, that mean state T/S biases in the North Atlantic co-vary and are almost density compensating. We have shown that these biases also appear to affect what controls interannual/decadal density variability in the sinking regions, a result that perhaps helps to explain the spread in the literature on this point. Furthermore, we have highlighted how these models appear to favour Labrador Sea feedbacks either to the north or the south. Lastly, the relationships between all of these metrics show some separation by the models' effective horizontal resolution.

A key remaining question is whether there is a further systematic link between the mechanisms of simulated variability and the inherent timescales of this variability. Of the 42 models tested, 26 show a peak in their power spectrum (above the 95% level for that estimated for a similar first order autoregressive red noise process) for top 500m volume averaged temperatures in the NA SPG (45–62°N) at short 'decadal' timescales (periods of 10–40 years). This rises to 29 models when considering just the Labrador Sea region. Despite this, we find no relationships between the aforementioned mechanisms/density-control and either the preferred timescales or relative magnitudes of decadal variability. Thus, within the statistical power of our reduced sample, there appears to be no systematic relationship between the mechanisms that we have investigated and the subsequent manifestation of decadal *periodicity* in the North Atlantic subpolar gyre. This is further evidenced by the similar periodicities in HadGEM3 (17 years) and IPSL-CM5A-LR (20 years (*Escudier et al.*, 2013)) despite very different resolutions/biases/density-control. However, it is still possible that there exist systematic relationships on longer timescales, which we have not investigated, perhaps involving the AMOC. Finally, recent work has suggested that observed and simulated (in historical simulations) bi-decadal variability in the NA SPG may be pace-set by volcanic forcing (*Swingedouw et al.*, 2015). Al-

though the authors find similar timescales of variability this appears to arise by different mechanisms across the models and further highlights the apparent independence of the mechanisms and overall timescales in the NA SPG.

The relationship between the mean T/S (biases) and density control suggests that the real world may be in a state where Labrador Sea annual density variability is controlled by salinity (rather than temperature) changes, similarly to lower resolution models (Figure 6.3). However, as noted in Chapter 3, this relationship may have changed in recent decades. In either case, there is limited spatial and temporal sampling within the Labrador Sea, particularly during the important wintertime period (defined here as December to February inclusive). For example, there are multiple years in the 1990s in EN4 where there are less than two combined T/S observations below the surface anywhere in the Labrador Sea during winter (see Chapter 3 and Figure 3.1b). Additionally, it is likely that there has been significant external forcing within this time period, noted in Section 6.4.3. Therefore, we have more confidence in the first order moments (*i.e.* the mean, from which we estimate the model biases) than second order moments (*i.e.* the observed interannual/decadal variability) and as such we suggest some caution in interpreting measures relating to the observed Labrador Sea variability (*cf.* black crosses in Figures 6.3 and 6.5).

Although it is clear that the effective horizontal resolution of the models in the NA SPG is related to their mean state biases, it is not as apparent why this is. We speculate that this relationship may arise due to higher resolution allowing better representation of the strength of boundary currents (*Grotzner et al.*, 1998; *Gelderloos et al.*, 2011), which are an important component of the AMOC (*McCarthy et al.*, 2015), as well as heat transport in the NA SPG. The computational overhead of calculating the AMOC streamfunction on the various model grids, and the mixed availability of the streamfunction diagnostic on the CMIP5 archive, precluded direct comparison to the AMOC. However, the zonal mean northward geostrophic circulation at 45°N (depth averaged between 100–1000m) was calculated using the T/S data, which shows an interannual correlation with the AMOC time series (maximum at 45°N) of r=0.7 in HadGEM3. This geostrophic estimate was then calculated for all models and used as a proxy for the AMOC. There exist weak correlations between this AMOC proxy and the T/S biases in the NA SPG (45–62°N) of

r=0.46 (for both temperature and salinity biases), with stronger northward circulation implying warmer and saltier conditions in the North Atlantic across models, consistent with recent multi-model work investigating larger-scale biases (*Zhang and Zhao*, 2015). It has also been shown that the timescales of the relationship between increased Labrador Sea (or elsewhere) convection and the response of the AMOC can vary greatly between models (*Huang et al.*, 2014), which may also explain the lack of a systematic link between the mechanisms and periodicity of simulated decadal variability, but further analysis is required.

The models we have analysed are intended to represent a stable climate not undergoing transient climate change. Due to the large role of AMOC-related northward heat transport, combined with uncertainties in AMOC projections, it is not clear whether the North Atlantic subpolar gyre region will actually become warmer or cooler under future climate change (*Collins et al.*, 2013). Whatever the sign, given that according to our analysis a change in mean T/S could affect the dominant mechanisms of decadal variability (*e.g.* Figure 6.3), this implies the prevalent mechanisms of decadal variability in the NA SPG could be different under future climate change (as indeed also tantalisingly implied by Figure 3.10b).

These results have clear implications for decadal predictions using the common method of 'anomaly-assimilation' in which observed oceanic T/S anomalies are assimilated in to a climate model's (biased) mean state. The simulated evolution of these anomalies within the climate model is then assumed to be independent of the mean state biases — an assumption that we have shown appears not to be valid when considering the mechanisms of variability. Although a process based understanding of the mechanisms of North Atlantic decadal change — gained by confronting climate predictions with recent observations — could validate and constrain models (Chapter 3), an in-depth understanding of control simulations will still be invaluable to put those results into a wider context (*e.g.* Chapter 2). Until such a time as real world NA SPG decadal variability is fully sampled we recommend a twin strand approach of analysing multi-century integrations with stable coupled climate models, and assessing the hindcast skill of decadal prediction systems.

6.6 Chapter conclusions

Climate models and direct and indirect observations have highlighted evidence of significant decadal variability within the North Atlantic subpolar gyre (NA SPG). Given the paucity of direct observations it is not clear how the signature of this variability evolves, in either space or time, whilst model simulations suggest a wide range of disparate mechanisms, many of which involve an important role for the Labrador Sea. To investigate these issues, we have analysed the systematic relationships across an exceptionally large ensemble of present-generation coupled climate models. We find that:

- Mean state biases in near surface (top 500m depth averaged) temperature and salinity in the Labrador Sea co-vary, with warm models also being saltier. Ensuing density biases are mostly due to salinity (rather than temperature).

- There exists a systematic relationship between whether density changes associated with variability in the Labrador Sea are temperature- or salinity-controlled and the mean state biases. Models which are too cool/fresh tend to have salinity-controlled density variability, whereas models which are too warm/salty tend to show a greater degree of temperature control. This relationship is seen for both interannual and decadal timescale variability.

- Negative feedbacks with the Labrador Sea tend to fall into two groups that exist either to the north (around the East Greenland Current) or to the south (around the North Atlantic current). These feedback locations are consistent with analysis of IPSL-CM5A-LR (*Escudier et al.*, 2013) and with our analysis of HadGEM3 (Chapter 2) and suggest that some of the inter-model spread in decadal variability in the NA SPG is related to differences in the preferred feedback location and mechanism.

- The effective horizontal resolution in the North Atlantic (ranging from 0.21° to 1.6° in our sample) shows some relationship to the mean state biases, density-control, and dominant feedbacks. Higher resolution models are generally too warm and salty, whereas lower resolution models are too cool and fresh in the Labrador Sea.

However, there is no relationship between the effective resolution and the absolute magnitude of the biases, suggesting higher resolution models cannot be considered to be 'better' than lower resolution models, in terms of their depth averaged biases in the Labrador Sea.

- Although there are systematic relationships between biases, density-control, and feedbacks, these do not appear to result in systematic relationships with the spectral characteristics of the variability in either the Labrador Sea or wider NA SPG. This is consistent with analysis of HadGEM3 (Chapter 2) and IPSL-CM5A-LR (*Escudier et al.*, 2013) that found different mechanisms but similar periodicities.

Although the mechanisms of NA SPG variability cannot be shown to be a systematic predictor of the periodicity, within the statistical power of our ensemble, our results nevertheless suggest that mean state biases influence the characteristics of decadal variability. This implies potential problems for decadal prediction systems that use the methodology of 'anomaly-assimilation', in which the mean state and evolution of the system are assumed to be independent (*Robson*, 2010). Indeed, the range in individual model biases and mechanisms suggests caution when making decadal predictions using any given model, whether initialised with full fields or anomalies.

In the next chapter we present a final discussion of our combination of in-depth analysis of a single model (Chapters 2–5), critical comparison against observations (Chapter 3), and its extension to scores of models (this chapter).

Chapter 7

Conclusions

7.1 Introduction

This thesis has focussed on understanding the simulated North Atlantic decadal variability in the high resolution coupled climate model HadGEM3. We have diagnosed the mechanism controlling the simulated decadal variability in HadGEM3 that gives rise to a spectral peak in many North Atlantic indices (such as depth averaged ocean temperatures) of 17 years. We have critically compared this variability against available observations, modifying our analysis methods as appropriate to account for sparse observational data. We have conducted sensitivity studies to investigate the role of the atmosphere in more detail and to assess how robust are various elements of the mechanism of decadal variability. Finally, based on these analyses, we have drawn hypotheses about the fundamental relationships between model biases and variability, and tested these in a multi-model framework. Understanding the origins of model variability, and its possible link to the underlying model setup/biases, could help give appropriate confidence to decadal predictions or longer projections made with these models.

In this chapter we describe how our analyses have addressed the key research questions that we outlined in Chapter 1. In brief, these were:

1. What decadal variability exists in the NA SPG in HadGEM3, and how does this

evolve in both space and time?

2. To what extent is the decadal variability consistent with available observations?

3. Does this variability provide potential skill for decadal predictions?

4. Is there any systematic explanation for the diversity of simulated variability within the NA SPG, as shown in Figure 1.5?

Key conclusions relating to analysis of the HadGEM3 model are presented in Section 7.2, whilst in Section 7.3 we present wider conclusions relating to the plausibility of the simulated variability relative to the real world as well as the veracity of decadal variability in other climate models. In Section 7.4 we present recommendations for future work on a number of fronts, relating to 1) decadal prediction systems, 2) observational networks, and 3) further progress in understanding the fundamental mechanisms of decadal variability in the North Atlantic before some last remarks in Section 7.5.

7.2 Conclusions part 1 — investigating the mechanism of decadal variability in HadGEM3

In this section we present the major conclusions relating to the analysis of the HadGEM3 coupled climate model, beginning with the analysis methods themselves (Section 7.2.1) before discussing the mechanism of simulated variability (Section 7.2.2) and the role of the atmosphere (Section 7.2.3).

7.2.1 Analysis methods

The investigations we have performed and the results we have presented have relied on a variety of complementary analysis methods, in terms of 1) the mathematical/statistical procedures used (*e.g.* lagged regression versus composite analysis, Table 4.1), 2) the critical testing of model mechanisms by comparison against observations (Chapter 3) and

targeted sensitivity experiments (Chapters 4 and 5), and 3) the expansion of single-model hypotheses to a multi-model context (Chapter 6). In addition, many of these methods rely on the statistical power provided by a long control simulation, which at high ocean and atmosphere resolution is a computationally expensive undertaking.

The following sections describe the key conclusions attained in this thesis, but it is worthwhile to note that another important conclusion from this work is the analytical power provided by combining analysis methods and sensitivity/multi-model studies under a single, coordinated project (such as a thesis). For example, lagged regression analysis (as used in Chapter 2) was a powerful way to maximise the amount of data used in a given analysis, but precluded the discovery of asymmetries in the variable of interest. Composite analysis, founded upon the mechanistic understanding provided in Chapter 2, led to the investigation of an asymmetrical ocean-atmosphere coupling strength (Chapter 4), which subsequently provided possible reasons for differing predictability characteristics between warm/cool initial conditions in a set of initial condition ensembles (Chapter 5).

In addition, the real-world skill of the model was estimated by comparing the model mechanisms (Chapter 2) with observed variability/processes (Chapter 3). This allowed us to gain insight as to *why* the model may have skill, which is different to merely quantifying the real-world skill using decadal hindcasts (although this would also be useful). Understanding the origins of skill would also be valuable if future iterations of a prediction system 'lost' skill, as has happened before, for example with the skill in predicting Arctic sea ice concentration at some lead times (*R. Wood*, pers. comm.). Further, these initial condition ensemble experiments provided details of the more robust elements of the mechanism described in Chapter 2 but it was not until they were combined with in-depth comparison between the model and observations (Chapters 2 and 3) that we were able to begin to assess which elements were more/less likely to be applicable to the real world and form simple hypotheses to test in a multi-model framework (Chapter 6).

In summary, a combination of techniques and experiments (using the same model) can lead to powerful understanding which, we contend, is greater than merely adding another model study of the North Atlantic to the large number that already exist (*cf.* Figure 1.5).

Conclusion: Combining analysis techniques under a single project can lead to a more

holistic understanding.

7.2.2 Periodic variability

HadGEM3 exhibits periodic decadal variability in the North Atlantic subpolar gyre (NA SPG, Chapter 2 and Figure 2.4b for top 500m depth averaged temperatures), similarly to many other climate models (Chapter 1 and Figure 1.5) and palaeo-proxies of the surrounding region (Chapter 1, *Mann et al.*, 1995; *Sicre et al.*, 2008; *Chylek et al.*, 2012). Nonetheless, the precise details of this variability are different to many previous studies (e.g. *Escudier et al.*, 2013, see also Figure 1.5). The key features of this mechanism of internal decadal variability in the NA SPG are:

- The timescale of the periodic variability is 17 years, explaining $>$30% of the inter-annual variance in top 500m depth averaged temperatures (T500) in the Labrador Sea and interior NA SPG.
- The mechanism describes heat content anomaly propagation around the NA SPG.
- Salinity content anomalies co-vary with heat content and act to dampen associated density variability.
- Advective heat fluxes rather than surface heat fluxes dominate the decadal variability.
- Anomalous circulation is most important for the heat transport variability in the North Atlantic Current (NAC) region, with anomalous temperature (mean circulation) most important for the heat transport variability in the northern NA SPG.
- The atmosphere acts to amplify the mode
- A negative feedback between Labrador Sea temperatures and those in the NAC, via a dynamic height gradient that induces circulation anomalies, switches the sign of the NA SPG temperature anomalies. This feedback appears to be a robust feature of the model (*e.g.* Chapter 5 and Figure 5.10).

The mode of variability we have diagnosed shows an important role for ocean advection in moving heat content anomalies cyclonically around the NA SPG (Chapter 2), with the cyclonic nature of the mode aiding predictability of downstream regions when initialised with large upstream anomalies (Chapter 5 and Figure 5.8). The relative importance of mean circulation (Chapter 2, Section 2.4.5) and anomalous circulation (Chapter 2, Section 2.4.4) in the advective heat budget may also relate to the differing predictability throughout the NA SPG with anomalies related to mean circulation easier to predict, for example the higher predictability in the northern than southern NA SPG, Figure 5.7. The ocean variability appears to force atmospheric variability (Chapter 4) of the same periodicity (Figure 2.4c), despite an otherwise white spectrum in the North Atlantic Oscillation (NAO) index (see next).

Conclusion: HadGEM3 shows strong decadal periodicity in the NA SPG relating to the advection of heat content anomalies around the gyre with a timescale of 17 years and explaining 30% of the interannual variance.

7.2.3 Atmospheric forcing and feedbacks

The atmosphere is involved in driving/reinforcing ocean variability in the NA SPG, with NAO-related surface heat fluxes (SHFs) in the eastern NA SPG aiding the westward propagation of temperature anomalies (Figure 2.7, third column). In addition, the NAO contributes to the dipole in anomalous ocean temperatures between the Labrador Sea and NAC, with NAO-related density changes explaining about 45% of the NAC geostrophic current response on annual timescales but less (13%) on decadal timescales. One dimensional models have shown that atmospheric standing wave patterns (such as the NAO) may interact with an advective ocean to provide decadal periodicity (*Saravanan and McWilliams*, 1998) but experiments with a lower resolution version of the ocean component of HadGEM3 suggest this may not happen in a more complex model (*Mecking et al.*, 2014). Indeed, in HadGEM3, the precise nature of any driving force from atmosphere to ocean, and whether the ocean integrates up atmospheric forcing or preferentially amplifies longer periods, remains unclear.

In terms of feedbacks, the NAO in HadGEM3 appears to respond to oceanic forcing on long timescales, showing a spectral peak with the same period (17 years) as the underlying ocean (Chapter 2 and Figure 2.4c). In HadGEM3, many climate-relevant fields in the NA SPG, such as mean sea level pressure (MSLP, Chapter 2, Figure 2.15), SHFs (Chapter 4, Figure 4.4), and near surface temperatures (T500, Figure 4.5), show asymmetry between positive and negative anomaly states (when composited against NA SPG sea surface temperatures). That is, there are stronger MSLP anomalies associated with an anomalously cool NA SPG (Figure 4.2) than with an anomalously warm NA SPG (Figure 4.3) as well as a faster (by 2 years) reversal timescale from positive NAO (+NAO) to negative NAO (-NAO) conditions than from -NAO to +NAO conditions, also associated with different lagged evolutions in MSLP, SHFs, and T500. These asymmetries in the strength of ocean-atmosphere coupling in either 1) cool NA SPG/+NAO or 2) warm NA SPG/-NAO states are supported by an ensemble of decoupled atmosphere-only experiments with the same climate model (Chapter 4). However, the location of primary ocean to atmosphere coupling that results in asymmetrical NAO responses is not clear: possibly being co-located in the NA SPG or further afield in the tropical Atlantic or tropical Pacific.

Conclusion: Atmospheric forcing of the ocean may play an important role in the diagnosed North Atlantic variability, particularly on annual timescales.

Conclusion: Asymmetries in the reversal timescales are related to an asymmetrical coupling strength between the ocean and atmosphere.

7.3 Conclusions part 2 — the mechanism of variability in HadGEM3 and relationships to reality/other models

In this section we broaden our focus to draw wider conclusions relating to the variability simulated by HadGEM3 and other climate models. We begin with conclusions relating to the similarities between HadGEM3 and observed variability (Section 7.3.1) which leads

us to the origins of the 'perfect-model' predictability in HadGEM3 and the potential real world skill (Section 7.3.2). We then broaden our focus a final time and summarise the diagnosed systematic relationships between climate model biases and simulated variability (Section 7.3.3).

7.3.1 Observational constraints

A lagged regression analysis of the simulated mode of variability (described above) against observational data was limited by severe data paucity in the NA SPG region (Chapter 3), either at depth (Figure 3.1) or at the surface (Figure 3.5), as well as increasing uncertainty in the representativeness of observations further back in time (comparison of Figures 3.3 and 3.4). Nonetheless, the model and observations show similar multi-annual joint evolution of upper ocean heat content and NAO anomalies, with warm conditions associated with a negative NAO (Figure 3.4), similar also to *Robson et al.* (2012). In addition, the lagged relationships between northern and southern edges of the NA SPG suggest variability in the southern NA SPG leads variability to the north on multi-annual timescales (Figure 3.6).

In addition, we undertook critical analysis of more bespoke features, such as the depth coherence of annual variability (Figure 3.8) and the drivers of Labrador Sea density changes (Figure 3.9).

The first of these suggests the model may overestimate the depth extent of coherent annual variability in the North Atlantic Current (NAC) region by several hundred metres (Figure 3.8), which may then account for the clearer propagation pathways around the NA SPG (Figure 3.11) and may also imply the skill in eastern NA SPG temperatures is overestimated (Chapter 5 and subsequent section 7.3.2).

Secondly, the driver of real-world upper ocean Labrador Sea multi-annual density variability appears to have been mostly salinity for the decades since 1960 (Figure 3.10) compared to temperature in the HadGEM3 simulations (Chapter 2 and Figure 2.3). However, the observations imply that the most recent decades have become increasingly temperature dominated, in agreement with the model results (discussed further in Sections 7.3.3

and 7.4). The model cannot explain the non-stationarity of the observed relationships, which could be externally forced (see also Section 7.4.3).

Conclusion: The model, observations, and recent investigation of reanalysis systems show consistent evolution of upper ocean heat content and NAO anomalies in the NA SPG.

Conclusion: Observational data paucity means critical comparison with the simulations must be undertaken by using regions/methods that make best use of these observations. Upon doing this, subtle but specific differences between the model and observations can be found (such as the depth coherence of annual variability), and their potential effect on the overall variability estimated.

7.3.2 Potential skill

A suite of initial condition ensembles, totalling 150 years of model integration (Chapter 5), suggested the real-world skill in the NA SPG provided by the model mechanism of decadal variability was dependent on both the robustness (*i.e.* predictability) of the simulated process providing the model skill and whether that process was likely to exist in a similar form in reality, the latter estimated by comparison with Chapter 3.

Perfect-model skill (*i.e.* skill in predicting the control simulation evolution) lasts up to 5 years in parts of the central and northern NA SPG. Potential predictability often arises due to initialisation of large temperature anomalies upstream. For example, potential predictability of the northern boundary current is highest when initialising based on large temperature anomalies in the eastern NA SPG. Despite this long-lasting potential skill, some regions of the NA SPG are likely to exhibit more reliable predictions than others, summarised next.

In the southern part of the NA SPG, the skill in T500 arises from the initialisation of the negative feedback/dynamic height gradient between the Labrador Sea and NAC, with skill in the dynamic height gradient lasting for 4 years (Figure 5.10). However, in the model this negative feedback relies on annual upper ocean density variability being

temperature-controlled (Chapter 2 and Section 2.4.6) whereas this appears to be salinity controlled in reality (Chapter 3 and Figure 3.9a). As such, the combination of understanding model mechanisms and the implied mechanisms in reality leads us to the conclusion that, if using a version of this model, any potential skill in predicting real-world T500 variability in the southern NA SPG *that arises from the simulated feedback with the Labrador Sea* is likely to be severely overestimated. This may be the case for both anomaly and full-field initialisation (see Section 7.4.1).

We find a smaller but similar implied reduction in potential skill when considering the eastern part of the NA SPG. In the model, the decadal variability in this region is related to the integration of circulation-driven heat content anomalies (Chapter 2 and Section 2.4.4). Along the NAC path, we find the depth coherence of these heat content anomalies to be much greater in the model than in reality (Chapter 3 and Figure 3.8), which is also consistent with a larger depth coherence/extent of signal propagation in the model than in reality (Figure 3.11). As such, the potential skill in predicting real-world T500 variability in the eastern NA SPG may also be somewhat overestimated.

Despite the possible reduction in potential skill in predicting T500 in the southern and eastern NA SPG, we suggest the model likely represents an improvement, compared to other climate models, in the northern NA SPG. This is because the high resolution of the model likely improves the representation of boundary currents (*Gelderloos et al.*, 2011), which we find are the driver of model skill in the northern NA SPG (Figure 5.12). As such, it is plausible that much of the model skill in northern NA SPG T500 (Figure 5.6) may carry over to real-world predictions, particularly if initialised during times of large anomalies in the eastern (*i.e.* upstream) NA SPG.

Conclusion: The potential skill of the model, and the reasons for reductions in skill, vary by region/process but suggest that hindcasts with this model may exhibit the most skill in the northern NA SPG.

7.3.3 Drivers of Labrador Sea density variability

There are several key findings (noted previously in this chapter and summarised below) that lead us to our final conclusion on the relationships between model biases and simulated variability. These are:

- The importance of a temperature-controlled negative feedback from the Labrador Sea to NAC (Chapter 2 and Section 2.4.6) for the periodic nature of the decadal variability (Figure 2.4b). For this feedback to be negative, density variability in the Labrador Sea must be temperature — rather than salinity — controlled.
- The unclear existence of such a feedback in reality, which appears to be a salinity-controlled positive feedback for much of the observed record but may have recently transitioned to a temperature-controlled negative feedback, Chapter 3 and Figure 3.10b.
- The importance of entirely different negative feedbacks in other models in providing periodicity within the NA SPG (e.g. *Escudier et al.*, 2013).

Given these findings (as well as the non-linear seawater equation of state), we hypothesised that:

1. There is a fundamental link between the mean state biases of a climate model (in terms of the depth averaged temperature or salinity in the NA SPG) and the processes by which negative feedbacks can occur (Chapter 2, Section 2.5.1 and also stated in the Discussion section of *Menary et al.*, 2015a).
2. A link between the mean state biases of a climate model and the mechanisms of variability leads to differing timescales/spectral characteristics of the variability.

These hypotheses were tested in Chapter 6. We found the first hypothesis to be verified, but the second one remained outstanding. Specifically, temperature and salinity biases in the Labrador Sea were found to be correlated with the drivers of density variability in that

region, as well as with the resolution (in the North Atlantic Ocean) of the ocean models, but there was no detectable systematic link between biases/resolution and the timescales of the variability. We found that high resolution models were too warm and salty and had temperature-controlled density variability, while low resolution models were too cool and fresh and had salinity-controlled variability (Figure 6.3). Furthermore, the model resolution separated the models into those that suggested negative feedbacks between the Labrador Sea and NAC (such as HadGEM3, Chapter 2) and those that suggested negative feedbacks between the Labrador Sea and northern NA SPG (Figure 6.5). The relationships between biases and variability suggest that assuming the mean state and variability are independent in both models and observations, for example when attempting to simulate actual events in the real world via 'anomaly assimilation', may not be appropriate (see Section 7.4.1).

Conclusion: The link between mean state biases and variability suggests that in order to correctly simulate decadal variability models must also exhibit a good representation of the mean state.

7.4 Recommendations and future work

In this section we discuss the implications of our work for the future development of both decadal prediction systems (Section 7.4.1) and observational networks (Section 7.4.2). We conclude with discussion of possible future multi-model analyses to make further use of the powerful resource that is the Phase 5 of the Coupled Model Intercomparison Project (CMIP5) archive of control simulations (Section 7.4.3).

7.4.1 Implications for the development of decadal prediction systems

Given the aforementioned links between NA SPG biases and multi-annual/decadal variability (Chapter 6 and Section 7.3.3) it appears likely that the technique of anomaly-assimilation, in which observed anomalies are combined with a climate model's biased mean state, is introducing systematic errors into the resulting predictions. These errors

are likely to grow with lead time, as the total effect of simulated (potentially incorrect) feedbacks aggregates. However, we note that it is also possible that full-field assimilation exhibits a similar interdependence of signal and background (in this case, drift), which may be related to the reduced skill compared to anomaly methods at long lead times (*Smith et al.*, 2013). In addition, the rapid return of a climate model to a biased mean state, after initialisation using full fields (around 2 years, *Smith et al.*, 2013) means that full-field initialisation could also result in the background mean state affecting the evolution of the signal just as for anomaly-assimilation.

Recommendation: For short lead times, use full-field ocean assimilation rather than anomaly-assimilation, especially if concerned with processes/feedbacks known to be dependent on the background state.

Future work: In addition, investigate in which climate variables/regions signal and model drift interact in predictions using full-field assimilation, and whether erroneous influences of the mean state on to the signal increase at longer lead times (after the model has drifted back to its preferred state).

Related to the above, we have shown that in reality the driver of Labrador Sea density variability may not be stationary (Chapter 3, Figure 3.10 and Section 7.3.1). Thus, it is possible that the feedbacks simulated in HadGEM3 may give rise to greater/lesser skill during different epochs. To quantify whether the real world skill of the model is non-stationary related to these feedbacks, hindcast predictions could be compared for groups of start dates, 1) during the period in observations where density variability in the Labrador Sea was salinity-controlled (*e.g.* pre-1980s), and 2) during the more recent period where this density variability was temperature-controlled. The 'Decadal Climate Prediction Project' (DCPP[1]) plans to conduct hindcast simulations initialised at every year from 1960 onwards, which would provide the necessary resource to test this hypothesis. Additional work to better understand this non-stationarity is presented in Section 7.4.3.

Future work: Compare the skill downstream of the Labrador Sea in hindcast simulations

[1] http://wcrp-climate.org/dcp-overview (October 2015)

for start dates prior to 1980 and start dates around 2010 to estimate whether prediction systems can reproduce the apparent transient changes in Labrador Sea density-drivers seen in observations.

We have shown an important role for northern boundary currents in the mechanism of decadal variability in HadGEM3 (Chapter 2, Section 2.4.5 and Chapter 5, Figure 5.12), which is likely in part made possible by the increased resolution of the model (*Gelderloos et al.*, 2011). In addition, we have shown a link between model biases, resolution, and the implied location of negative feedbacks from the Labrador Sea (Chapter 6 and Figure 6.5). Although the model resolution is higher than in control simulations with most other climate models (Tables 6.1 and 6.2) there are potentially still further improvements to be had by moving from eddy-permitting resolutions to eddy resolving resolutions (*Marzocchi et al.*, 2015). Recent versions of the ocean component of HadGEM3 are also capable of incorporating nested, higher resolution grids (*Debreu et al.*, 2008).

Recommendation: Further increase the model resolution to high-latitude eddy-resolving resolutions in important regions for decadal prediction (*e.g.* the NA SPG) using nested schemes and investigate the effects on predictability in NA SPG indices (and whether this arises through processes that are expected to be improved by resolution).

7.4.2 Implications for present and future observational networks

In Chapter 3 we noted that, even in the late twentieth century, sustained observations of key regions were not guaranteed. For example, there were winters in the 1990s in which there were no subsurface temperature or salinity observations in the Labrador Sea (Figure 3.1b). In addition, even in surface temperatures, the northern half of the NA SPG has had sustained observations for far less time than the rest of the North Atlantic Ocean (Figure 3.5). To this end, the global array of 'Argo' profiling floats, which began deployment in the year 2000, are a valuable tool, though these are limited to the upper 2000m of the ocean (*Gould et al.*, 2004). Although we have focused on this upper ocean decadal variability, we note that near surface variability can imprint on to deep ocean variability (*Mauritzen et al.*, 2012), observations of which (*e.g.* 'Deep Argo', *Johnson and Lyman*,

2014) would provide a further valuable tool to constrain the models.

Recommendation: Present observational network densities must be at least maintained if we are to be able to reliably characterise multi-annual/decadal variability.

The mode of variability we have diagnosed is manifest primarily in the NA SPG (Chapter 2, possible links to El Niño and tropical North Atlantic variability notwithstanding). Simulated heat transport variability on the southern edge of the NA SPG is found to be important, but even this is much farther north than the only cross-basin current array at 26.5°N (*Cunningham et al.*, 2007). Indeed, it is not clear (from model studies) to what extent variability in the NA SPG and tropical North Atlantic are coherent nor the precise phase relationships between the two (*Bingham et al.*, 2007; *Zhang*, 2010). We note that a promising first step in understanding the heat budget of the NA SPG is being taken with the recent 'Overturning in the Subpolar North Atlantic Program' (OSNAP[2]) that will constrain the heat/freshwater/volume fluxes into the NA SPG from the north. Nonetheless, fluxes from the south and, for example, the degree of recirculation in the NA SPG and 'leakage' through the Canary Current (*Swingedouw et al.*, 2013), are still not constrained.

Recommendation: The establishment of further cross-basin arrays, in particular at the southern boundary of the NA SPG (*e.g.* 45°N), to constrain the ocean heat and freshwater budget of the NA SPG.

Ocean temperatures represent the integration of both surface heat fluxes (SHFs) and the advective convergence/divergence of heat (see also Section 7.4.2). For this reason, given the lack of direct heat flux observations and estimates of subsurface circulation, it is not always clear whether well-modelled temperature evolution (or 'skill' if directly assessing against particular observations) is the result of a good representation of both surface and advective heat fluxes or due to compensating errors. Additionally, given that the variability in SHFs is an indication of the strength and location of ocean-atmosphere coupling (see also Chapter 4), they represent a good metric for discriminating between the mechanisms of variability suggested by climate models.

SHFs can be directly measured by moored buoys. However, there are a severely limited

[2]http://www.o-snap.org/ (October 2015)

number of buoys in the world oceans that are able to measure the relevant heat flux terms (in particular longwave heat fluxes, *Kato et al.*, 2013), the sum of which gives the net SHF between ocean and atmosphere. Even more starkly, there are no active moored buoys anywhere within the NA SPG (*Kato et al.*, 2013). The addition of moored buoys to measure SHFs would allow allow critical comparison of model mechanisms of variability against reality.

Recommendation: The placement of moored buoys in the NA SPG to be able to directly estimate SHFs between the ocean and atmosphere

7.4.3 Implications for investigations of simulated decadal variability, and further options for multi-model analysis

As repeatedly stated in this thesis, there is a severe paucity of direct observations within the NA SPG with which to constrain the decadal variability simulated in climate models, which likely does exist in the real world (*Mann et al.*, 1995; *Sicre et al.*, 2008; *Chylek et al.*, 2012). In the previous section (Section 7.4.2) we noted that our first priority should be to maintain, and then increase, the size of the observational network. In the meantime, as we have shown, a potentially powerful approach to future investigations of decadal variability is to diagnose the key processes within a range of models and examine where in this distribution the real world exists — and what implications this may have. Some specific examples are given next.

Recommendation: A combined approach to analysis of NA SPG decadal variability, combining multi-model analyses with critical comparison against observations, rather than describing variability in a single model and crudely contrasting with observations.

In general, in terms of the fundamental processes and drivers of simulated decadal variability in the NA SPG, there remain many competing theories (Chapter 1). Although conceptually appealing, it is far from clear that any potentially coupled ocean-atmosphere relationship can be diagnosed by separately analysing uncoupled ocean and atmosphere subcomponents (*Battisti et al.*, 1995; *Rodwell et al.*, 1999). In addition, even relatively

simple explanations of NA SPG periodicity require a coupled system (*Saravanan and McWilliams*, 1998). However, it is possible to test the relative importance of processes within a given system by artificially suppressing certain aspects, as has been shown for the relative importance of different eddy types in restratification of the Labrador Sea (*Gelderloos et al.*, 2011). We suggest a similar approach could be undertaken to assess the relative importance of, for example, Rossby Waves and ocean advection in transferring information around the NA SPG (both of which have similar, multi-annual timescales). In this example, this could be undertaken by either direct approaches, such as pinpointing Rossby Wave signals and damping them, or indirect approaches, such as artificially altering the thermocline depth to alter wave propagation timescales. This could be undertaken with either the climate prediction system or using the control simulation of HadGEM3.

Future work: Investigate methods to damp specific processes within a coupled model system in order to test the relative importance of particular processes in simulations of NA SPG decadal variability.

In Chapter 3 (and above in Section 7.3.1) we noted that in reality the driver of Labrador Sea density variability may not be stationary in time (Figure 3.10). Although the model control simulations we analysed in Chapter 6 were not drifting, it does not necessarily follow that the driver of density variability is stable throughout the simulation as this depends both on the bias and the actual variability in temperature or salinity. As such, it may be possible for climate models to capture these potential regime shifts. In addition, given that the control simulations were mostly run with external forcings appropriate for 1850, it would also be interesting to know whether the driver of density variability changes under historical forcings as well as under projected future climate change. Separately to analysis of hindcast simulations (see future work in Section 7.4.1), these analyses would shed light on the susceptibility of internal modes of variability to fundamental changes/breakdown and to what extent this may be important for the evolution of the NA SPG.

Future work: Test the stationarity of the driver of Labrador Sea density variability in the model control simulations, as well as in historical and future climate simulations.

The multi-model analysis presented in Chapter 6 focussed primarily on the ocean. However, as we have shown, the atmosphere plays an important role in reinforcing this ocean variability (Chapter 2, Figure 2.12 and Section 7.2.3) with clear feedbacks from ocean to atmosphere (Chapter 2, Figure 2.4c) and further interesting asymmetries between positive and negative phases of the variability (Chapter 4). *Gastineau et al.* (2013) recently highlighted similar lead/lag relationships between observed and simulated NA SPG SSTs and the NAO, although the strength of the coupling in their model was approximately half as strong as observed. As such, it would be interesting to investigate the characteristics of the NAO across the models and whether there are any links between the manifestation of NAO variability and the variability in the ocean. In addition, given our previous analysis that showed a link between the ocean model resolution in the North Atlantic and the locations of negative feedbacks with the Labrador Sea (Chapter 6, Figure 6.5 and Sections 7.3.3 and 7.4.1), it would be interesting to investigate any potential role for atmospheric resolution in the manifestation of variability in and around the North Atlantic.

Future work: Expand CMIP5 multi-model analysis to incorporate the NAO and characteristics of the atmosphere component of the models, such as resolution.

Finally, one caveat of our uncoupled atmosphere-only experiments (Chapter 4) was the use of globally-defined forcing fields, rather than just local (*e.g.* North Atlantic) fields (Section 4.3.1). This leaves open the possibility that non-local forcing, such as from the tropical Atlantic or tropical Pacific (via El Niño/Southern Oscillation, ENSO), may drive the NAO (and thus even the ocean) variability in the North Atlantic (*Ineson and Scaife*, 2009). Indeed, similar to our analysis of the link between Labrador Sea biases and variability, it has been shown that biases in the 'thermocline feedback' in the tropical Pacific in CMIP5 models result in incorrect characteristics of ENSO (*Michael et al.*, 2013). Given the size of our multi-model database (42 models), further analysis of the links between biases in the tropical Pacific and the manifestation of variability in the North Atlantic may reveal under what conditions/in which model setups variability in the ENSO regions more or less strongly affects the NA SPG.

Future work: Expand CMIP5 multi-model analysis to include ENSO diagnostics and investigate possible teleconnections to the NA SPG.

7.5 Concluding remarks

A paucity of observations means that understanding past and potential future variability in the NA SPG still relies heavily on climate models, whose multi-annual/decadal variability is poorly constrained. We have diagnosed the mechanism of decadal variability in the NA SPG in a new, high resolution climate model. We have highlighted, through a holistic analysis approach, how the specifics of this variability may be important for understanding — and having confidence in — decadal climate predictions. Continued critical comparison between climate models and expanding observational networks is crucial to improving the models and prediction systems. We have also shown how the various modes/mechanisms of NA SPG variability in climate models may be inextricably linked to something as simple as errors in their mean states. Understanding to what extent this link does affect the outcomes of decadal predictions should be a key subject of future study if we are to have confidence in the ability of climate models to faithfully represent the state, and fate, of our climate.

Appendix A

Tracer release experiments to examine simulated deep water formation

A.1 Introduction

In Chapter 2 we noted the interannual variability in wintertime mixed layer depth (MLD) in HadGEM3 (Figure 2.1d). The mean state mixed layers in HadGEM3 show a similar pattern but with the maximum depth stretching to around 500m in the eastern NA SPG, and to the full depth in the Labrador Sea (not shown). This is much greater than the real world, either through estimates applying the same *Kara et al.* (2000) methodology to the EN4 dataset (not shown), or published climatologies (*de Boyer Montégut et al.*, 2004). It seems possible that HadGEM3 overestimates the vigour of deep water formation in the Labrador Sea, and as such gives this region a more prominent role in NA SPG variability than it has in reality.

This can also be seen in the analysis of the depth coherence of near surface signals (Figure 3.7). In EN4 the region downstream of the Denmark Straits sill overflows is found to be particularly depth coherent, which is not the case in HadGEM3. Additionally, at the deepest of the chosen layers (Figure 3.7, bottom row) the positive correlation with near surface signals is much greater in HadGEM3 than in EN4 in the Labrador Sea. It is therefore possible that the prominence of the 17 year variability displayed in HadGEM3

compared to EN4 is related to the increased role for the Labrador Sea.

A.2 Simulated tracer release experiment

To investigate the importance of Labrador Sea deep water formation (which forms upper North Atlantic deep water, UNADW) in HadGEM3 and whether HadGEM3 is able to differentiate between UNADW and lower North Atlantic deep water (LNADW, formed in the GIN Seas) we design sensitivity experiments with the model. In these, the model is run for 30 years with a passive tracer released in the Labrador Sea region (Figure A.1a, red) and a separate passive tracer released in the GIN Seas region (Figure A.1a, blue) throughout the top 2000m of the water column. The tracers are set to a concentration of 100% within each of the three-dimensional volumes and allowed to evolve freely with the ocean advection-diffusion scheme outside of these regions. After 30 years the concentration profiles are stable, although the absolute values continue to increase as there is no sink for the tracers.

To compare the profiles of UNADW (Labrador Sea) and LNADW (GIN Seas) we show the zonal mean normalised tracer concentrations at a latitude of approximately 45°N (Figure A.1b, approximate latitude due to the curved nature of the model grid). We use normalised concentrations as the absolute concentrations are very different due to the differing proximities of the source regions. The core of southward flowing deep water exists at 1700m, which compares favourably to the observed depth of the deep western boundary current (*Meinen et al.*, 2013). However, the two profiles do not show any differentiation by source region, contrary to observations in which the core of Labrador Sea water (UNADW) sits more than 1000m above the core of LNADW of GIN Seas origin (*Toole et al.*, 2011). These simulated profiles are consistent with the view that LNADW, formed upstream in the GIN Seas, merely aids or preconditions deep convection in the Labrador Sea in HadGEM3, possibly contributing to the large vertical and horizontal extent (at depth) of Labrador Sea water.

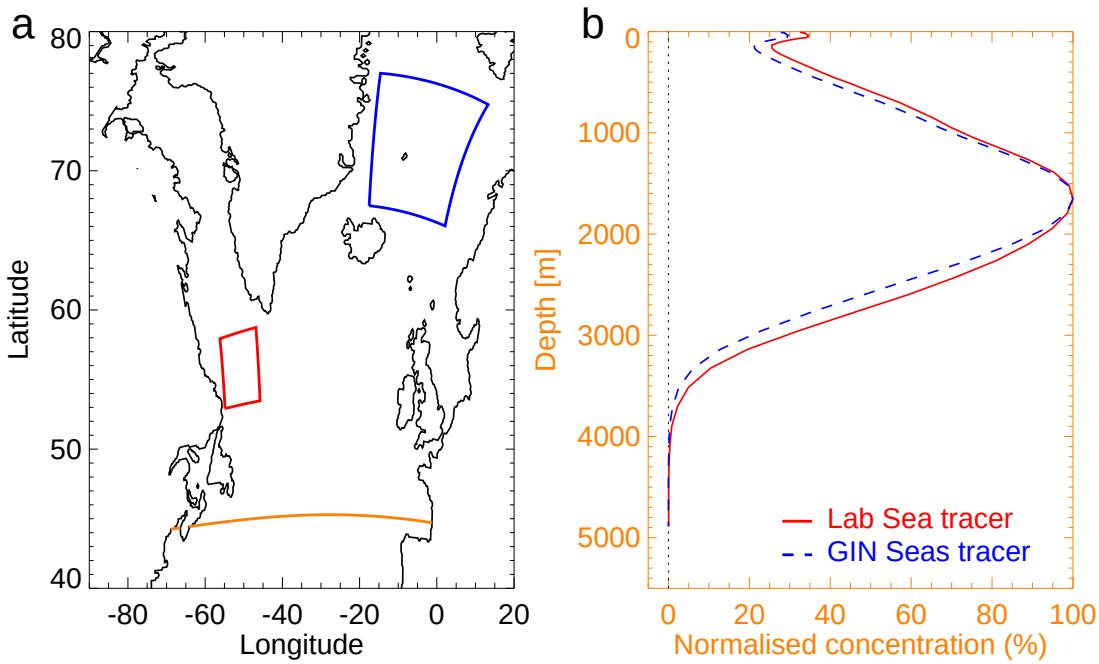

Figure A.1: a) The location of tracer release over the top 2000m in the Labrador Sea (red) and Greenland-Iceland-Norwegian (GIN) Seas (blue). b) The zonal mean profiles at 45°N (orange line on map) of the normalised concentration (each tracer normalised relative to its maximum absolute value) of the tracers released in the Labrador Sea (red) and GIN Seas (blue, dashed).

Bibliography

Ahmed, M., K. J. Anchukaitis, A. Asrat, H. P. Borgaonkar, M. Braida, B. M. Buckley, U. Büntgen, B. M. Chase, D. A. Christie, E. R. Cook, et al. (2013), Continental-scale temperature variability during the past two millennia, *Nature Geoscience*, *6*, 339–346.

Alexander, M. A., and C. Deser (1995), A mechanism for the recurrence of wintertime midlatitude SST anomalies, *Journal of Physical Oceanography*, *25*(1), 122–137.

Allan, R., and T. Ansell (2006), A new globally complete monthly historical gridded mean sea level pressure dataset (HadSLP2): 1850-2004, *Journal of Climate*, *19*(22), 5816–5842.

Allison, L., E. Hawkins, and T. Woollings (2015), An event-based approach to understanding decadal fluctuations in the Atlantic meridional overturning circulation, *Climate Dynamics*, *44*(1-2), 163–190.

Alvarez-Garcia, F., M. Latif, and A. Biastoch (2008), On multidecadal and quasi-decadal North Atlantic variability, *Journal of Climate*, *21*(14), 3433–3452.

Ba, J., N. S. Keenlyside, M. Latif, W. Park, H. Ding, K. Lohmann, J. Mignot, M. Menary, O. H. Otterå, B. Wouters, et al. (2014), A multi-model comparison of Atlantic multidecadal variability, *Climate Dynamics*, *43*(9-10), 2333–2348.

Baines, P. G., and C. K. Folland (2007), Evidence for a rapid global climate shift across the late 1960s, *Journal of Climate*, *20*(12), 2721–2744.

Balmaseda, M. A., K. Mogensen, and A. T. Weaver (2013), Evaluation of the ECMWF ocean reanalysis system ORAS4, *Quarterly Journal of the Royal Meteorological Society*, *139*(674), 1132–1161.

Barnston, A. G., and R. E. Livezey (1987), Classification, seasonality and persistence of low-frequency atmospheric circulation patterns, *Monthly Weather Review*, *115*(6), 1083–1126.

Battisti, D., U. Bhatt, and M. Alexander (1995), A modeling study of the interannual variability in the wintertime North Atlantic Ocean, *Journal of Climate*, *8*(12), 3067–3083.

Behrens, E., A. Biastoch, and C. W. Böning (2013), Spurious AMOC trends in global ocean sea-ice models related to subarctic freshwater forcing, *Ocean Modelling*, *69*, 39–49.

Bell, C. J., L. J. Gray, A. J. Charlton-Perez, M. M. Joshi, and A. A. Scaife (2009), Stratospheric communication of El Niño teleconnections to European winter, *Journal of Climate*, *22*(15), 4083–4096.

Biastoch, A., C. W. Boening, J. Getzlaff, J.-M. Molines, and G. Madec (2008a), Causes of interannual-decadal variability in the meridional overturning circulation of the midlatitude North Atlantic ocean, *Journal of Climate*, *21*(24), 6599–6615.

Biastoch, A., C. W. Boening, and J. R. E. Lutjeharms (2008b), Agulhas leakage dynamics affects decadal variability in Atlantic overturning circulation, *Nature*, *456*(7221), 489–492.

Bindoff, N., P. Stott, K. AchutaRao, M. Allen, N. Gillett, D. Gutzler, K. Hansingo, G. Hegerl, Y. Hu, S. Jain, I. Mokhov, J. Overland, J. Perlwitz, R. Sebbari, and X. Zhang (2013), Chapter 10: Detection and Attribution of Climate Change: from Global to Regional., *In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* .

Bingham, R. J., C. W. Hughes, V. Roussenov, and R. G. Williams (2007), Meridional coherence of the North Atlantic meridional overturning circulation, *Geophysical Research Letters*, *34*(23).

Boer, G. (2000), A study of atmosphere-ocean predictability on long time scales, *Climate Dynamics*, *16*(6), 469–477.

Booth, B., N. Dunstone, P. Halloran, N. Bellouin, and T. Andrews (2012), Aerosols indicated as prime driver of 20th century North Atlantic climate variability, *Nature*, *484*(7393), 228–232.

Born, A., and J. Mignot (2012), Dynamics of decadal variability in the Atlantic subpolar gyre: A stochastically forced oscillator, *Climate Dynamics*, *39*(1-2), 461–474.

Born, A., T. F. Stocker, C. C. Raible, and A. Levermann (2013), Is the Atlantic subpolar gyre bistable in comprehensive coupled climate models?, *Climate Dynamics*, *40*(11-12), 2993–3007.

Branstator, G., and H. Teng (2012), Potential impact of initialization on decadal predictions as assessed for CMIP5 models, *Geophysical Research Letters*, *39*(12).

Brayshaw, D. J., B. Hoskins, and M. Blackburn (2008), The storm-track response to idealized SST perturbations in an aquaplanet GCM, *Journal of the Atmospheric Sciences*, *65*(9), 2842–2860.

Bretherton, C. S., and D. S. Battisti (2000), Interpretation of the results from atmospheric general circulation models forced by the time history of the observed sea surface temperature distribution, *Geophysical Research Letters*, *27*(6), 767–770.

Brohan, P., J. J. Kennedy, I. Harris, S. F. Tett, and P. D. Jones (2006), Uncertainty estimates in regional and global observed temperature changes: A new data set from 1850, *Journal of Geophysical Research: Atmospheres (1984–2012)*, *111*(D12).

Burkholder, K. C., and M. S. Lozier (2011), Mid-depth lagrangian pathways in the North Atlantic and their impact on the salinity of the eastern subpolar gyre, *Deep Sea Research Part I: Oceanographic Research Papers*, *58*(12), 1196 – 1204.

Burkholder, K. C., and M. S. Lozier (2014), Tracing the pathways of the upper limb of the North Atlantic meridional overturning circulation, *Geophysical Research Letters*, *41*(12), 4254–4260.

Cabanes, C., T. Lee, and L.-L. Fu (2008), Mechanisms of interannual variations of the meridional overturning circulation of the North Atlantic ocean, *Journal of Physical Oceanography*, *38*(2), 467–480.

Carton, J. A., and B. S. Giese (2008), A reanalysis of ocean climate using Simple Ocean Data Assimilation (SODA), *Monthly Weather Review*, *136*(8), 2999–3017.

Chang, Y. S., Z. D. Garraffo, H. Peters, and T. M. Özgökmen (2009), Pathways of Nordic overflows from climate model scale and eddy resolving simulations, *Ocean Modelling*, *29*(1), 66–84.

Chelton, D. B., R. A. Deszoeke, M. G. Schlax, K. El Naggar, and N. Siwertz (1998), Geographical variability of the first baroclinic rossby radius of deformation, *Journal of Physical Oceanography*, *28*(3), 433–460.

Chylek, P., C. Folland, L. Frankcombe, H. Dijkstra, G. Lesins, and M. Dubey (2012), Greenland ice core evidence for spatial and temporal variability of the Atlantic multidecadal oscillation, *Geophysical Research Letters*, *39*(9).

Clement, A., K. Bellomo, L. N. Murphy, M. A. Cane, T. Mauritsen, G. Rdel, and B. Stevens (2015), The Atlantic multidecadal oscillation without a role for ocean circulation, *Science*, *350*(6258), 320–324.

Collins, M. (2002), Climate predictability on interannual to decadal time scales: The initial value problem, *Climate Dynamics*, *19*(8), 671–692.

Collins, M., M. Botzet, A. Carril, H. Drange, A. Jouzeau, M. Latif, S. Masina, O. Otteraa, H. Pohlmann, A. Sorteberg, et al. (2006), Interannual to decadal climate predictability in the North Atlantic: A multimodel-ensemble study, *Journal of Climate*, *19*(7), 1195–1203.

Collins, M., R. Knutti, J. Arblaster, J. Dufresne, T. Fichefet, P. Friedlingstein, X. Gao, W. Gutowski, T. Johns, G. Krinner, M. Shongwe, C. Tebaldi, A. Weaver, and M. Wehner (2013), Chapter 12: Long-term climate change: Projections, commitments and irreversibility., *In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*.

Corti, S., A. Weisheimer, T. Palmer, F. Doblas-Reyes, and L. Magnusson (2012), Reliability of decadal predictions, *Geophysical Research Letters*, *39*(21).

Cunningham, S. A., T. Kanzow, D. Rayner, M. O. Baringer, W. E. Johns, J. Marotzke, H. R. Longworth, E. M. Grant, J. J. M. Hirschi, L. M. Beal, C. S. Meinen, and H. L.

Bryden (2007), Temporal variability of the Atlantic meridional overturning circulation at 26.5°N, *Science*, *317*(5840), 935–938.

Curry, R., and M. McCartney (2001), Ocean gyre circulation changes associated with the North Atlantic Oscillation, *Journal of Physical Oceanography*, *31*(12), 3374–3400.

Dai, A., A. Hu, G. Meehl, W. Washington, and W. Strand (2005), Atlantic thermohaline circulation in a coupled general circulation model: Unforced variations versus forced changes, *Journal of Climate*, *18*(16), 3270–3293.

Danabasoglu, G. (2008), On multidecadal variability of the Atlantic meridional overturning circulation in the community climate system model version 3, *Journal of Climate*, *21*(21), 5524–5544.

de Boyer Montégut, C., G. Madec, A. S. Fischer, A. Lazar, and D. Iudicone (2004), Mixed layer depth over the global ocean: An examination of profile data and a profile-based climatology, *Journal of Geophysical Research: Oceans (1978–2012)*, *109*(C12).

Debreu, L., C. Vouland, and E. Blayo (2008), AGRIF: Adaptive grid refinement in Fortran, *Computers & Geosciences*, *34*(1), 8–13.

Delworth, T., S. Manabe, and R. Stouffer (1993), Interdecadal variations of the thermohaline circulation in a coupled ocean-atmosphere model, *Journal of Climate*, *6*(11), 1993–2011.

Delworth, T. L., and R. J. Greatbatch (2000), Multidecadal thermohaline circulation variability driven by atmospheric surface flux forcing, *Journal of Climate*, *13*(9), 1481–1495.

Delworth, T. L., R. Zhang, and M. E. Mann (2007), Decadal to centennial variability of the Atlantic from observations and models, *Geophysical Monograph-American Geophysical Union*, *173*, 131.

Deser, C., J. E. Walsh, and M. S. Timlin (2000), Arctic sea ice variability in the context of recent atmospheric circulation trends, *Journal of Climate*, *13*(3), 617–633.

Dickson, R. R., J. Meincke, S.-A. Malmberg, and A. J. Lee (1988), The “great salinity anomaly” in the northern North Atlantic 1968–1982, *Progress in Oceanography*, *20*(2), 103–151.

Doblas-Reyes, F., I. Andreu-Burillo, Y. Chikamoto, J. García-Serrano, V. Guemas, M. Kimoto, T. Mochizuki, L. Rodrigues, and G. Van Oldenborgh (2013), Initialized near-term regional climate change prediction, *Nature Communications*, *4*, 1715.

Dong, B., and R. Sutton (2001), The dominant mechanisms of variability in Atlantic ocean heat transport in a coupled ocean-atmosphere GCM, *Geophysical Research Letters*, *28*(12), 2445–2448.

Dong, B., and R. Sutton (2005), Mechanism of interdecadal thermohaline circulation variability in a coupled ocean-atmosphere GCM, *Journal of Climate*, *18*(8), 1117–1135.

Donlon, C. J., M. Martin, J. Stark, J. Roberts-Jones, E. Fiedler, and W. Wimmer (2012), The operational sea surface temperature and sea ice analysis (OSTIA) system, *Remote Sensing of Environment*, *116*, 140–158.

Drouard, M., G. Rivière, and P. Arbogast (2013), The North Atlantic Oscillation response to large-scale atmospheric anomalies in the northeastern Pacific, *Journal of the Atmospheric Sciences*, *70*(9), 2854–2874.

Duchez, A., J.-M. Hirschi, S. Cunningham, A. Blaker, H. Bryden, B. de Cuevas, C. Atkinson, G. McCarthy, E. Frajka-Williams, D. Rayner, et al. (2014), A new index for the Atlantic meridional overturning circulation at 26°N, *Journal of Climate*, *27*(17), 6439–6455.

Dufresne, J.-L., M.-A. Foujols, S. Denvil, A. Caubel, O. Marti, O. Aumont, Y. Balkanski, S. Bekki, H. Bellenger, R. Benshila, S. Bony, L. Bopp, P. Braconnot, P. Brockmann, P. Cadule, F. Cheruy, F. Codron, A. Cozic, D. Cugnet, N. de Noblet, J.-P. Duvel, C. Ethé, L. Fairhead, T. Fichefet, S. Flavoni, P. Friedlingstein, J.-Y. Grandpeix, L. Guez, E. Guilyardi, D. Hauglustaine, F. Hourdin, A. Idelkadi, J. Ghattas, S. Joussaume, M. Kageyama, G. Krinner, S. Labetoulle, A. Lahellec, M.-P. Lefebvre, F. Lefevre, C. Levy, Z. Li, J. Lloyd, F. Lott, G. Madec, M. Mancip,

M. Marchand, S. Masson, Y. Meurdesoif, J. Mignot, I. Musat, S. Parouty, J. Polcher, C. Rio, M. Schulz, D. Swingedouw, S. Szopa, C. Talandier, P. Terray, N. Viovy, and N. Vuichard (2013), Climate change projections using the IPSL-CM5 earth system model: From CMIP3 to CMIP5, *Climate Dynamics*, *40*(9-10), 2123–2165.

Dunstone, N., D. Smith, and R. Eade (2011), Multi-year predictability of the tropical Atlantic atmosphere driven by the high latitude North Atlantic ocean, *Geophysical Research Letters*, *38*(14).

Eade, R., D. Smith, A. Scaife, E. Wallace, N. Dunstone, L. Hermanson, and N. Robinson (2014), Do seasonal-to-decadal climate predictions underestimate the predictability of the real world?, *Geophysical Research Letters*, *41*(15), 5620–5628.

Eden, C., and R. J. Greatbatch (2003), A damped decadal oscillation in the North Atlantic climate system, *Journal of Climate*, *16*(24), 4043–4060.

Eden, C., and J. Willebrand (2001), Mechanism of interannual to decadal variability of the North Atlantic circulation, *Journal of Climate*, *14*(10), 2266–2280.

Escudier, R., J. Mignot, and D. Swingedouw (2013), A 20-year coupled ocean-sea ice-atmosphere variability mode in the North Atlantic in an AOGCM, *Climate Dynamics*, *40*(3-4), 619–636.

Fevrier, S., J. Sirven, and C. Herbaut (2007), Interaction of a coastal kelvin wave with the mean state in the Gulf Stream separation area, *Journal of Physical Oceanography*, *37*(6), 1429–1444.

Flatau, M. K., L. Talley, and P. P. Niiler (2003), The North Atlantic Oscillation, surface current velocities, and SST changes in the subpolar North Atlantic, *Journal of Climate*, *16*(14), 2355–2369.

Flato, G., J. Marotzke, B. Abidun, P. Braconnot, S. Chou, W. Collins, P. Cox, F. Driouech, S. Emori, V. Eyring, C. Forest, E. Gleckler, E. Guilyardi, C. Jakob, V. Kattsov, C. Reason, and M. Rummukainen (2013), Chapter 9: Evaluation of Climate Models., *In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* .

Folland, C., T. Palmer, and D. Parker (1986), Sahel rainfall and worldwide sea temperatures, 1901–85, *Nature*, *320*(6063), 602–607.

Frankcombe, L. M., M. H. England, M. E. Mann, and B. A. Steinman (2015), Separating internal variability from the externally forced climate response, *Journal of Climate*, (2015).

Frankcombe, L. M., A. von der Heydt, and H. A. Dijkstra (2010), North Atlantic multidecadal climate variability: An investigation of dominant time scales and processes, *Journal of Climate*, *23*(13), 3626–3638.

Frankignoul, C., P. Muller, and E. Zorita (1997), A simple model of the decadal response of the ocean to stochastic wind forcing, *Journal of Physical Oceanography*, *27*(8), 1533–1546.

Gaetani, M., and E. Mohino (2013), Decadal prediction of the Sahelian precipitation in CMIP5 simulations, *Journal of Climate*, *26*(19), 7708–7719.

Gastineau, G., F. D'Andrea, and C. Frankignoul (2013), Atmospheric response to the North Atlantic ocean variability on seasonal to decadal time scales, *Climate Dynamics*, *40*(9-10), 2311–2330.

Gelderloos, R., C. A. Katsman, and S. S. Drijfhout (2011), Assessing the roles of three eddy types in restratifying the Labrador Sea after deep convection, *Journal of Physical Oceanography*, *41*(11), 2102–2119.

Goldenberg, S. B., C. W. Landsea, A. M. Mestas-Nuez, and W. M. Gray (2001), The recent increase in Atlantic hurricane activity: Causes and implications, *Science*, *293*(5529), 474–479.

Good, S. A., M. J. Martin, and N. A. Rayner (2013), EN4: Quality controlled ocean temperature and salinity profiles and monthly objective analyses with uncertainty estimates, *Journal of Geophysical Research: Oceans*, *118*(12), 6704–6716.

Gould, J., D. Roemmich, S. Wijffels, H. Freeland, M. Ignaszewsky, X. Jianping, S. Pouliquen, Y. Desaubies, U. Send, K. Radhakrishnan, et al. (2004), Argo profil-

ing floats bring new era of in situ ocean observations, *EOS, Transactions American Geophysical Union*, *85*(19), 185–191.

Grotzner, A., M. Latif, and T. Barnett (1998), A decadal climate cycle in the North Atlantic ocean as simulated by the ECHO coupled GCM, *Journal of Climate*, *11*(5), 831–847.

Gulev, S. K., M. Latif, N. Keenlyside, W. Park, and K. P. Koltermann (2013), North Atlantic ocean control on surface heat flux on multidecadal timescales, *Nature*, *499*(7459), 464–467.

Hakkinen, S., and P. B. Rhines (2004), Decline of subpolar North Atlantic circulation during the 1990s, *Science*, *304*(5670), 555–559.

Häkkinen, S., P. B. Rhines, and D. L. Worthen (2015), Heat content variability in the North Atlantic ocean in ocean reanalyses, *Geophysical Research Letters*, *42*(8), 2901–2909.

Hallberg, R., and A. Gnanadesikan (2006), The role of eddies in determining the structure and response of the wind-driven southern hemisphere overturning: Results from the modeling eddies in the Southern Ocean (MESO) project, *Journal of Physical Oceanography*, *36*(12), 2232–2252.

Ham, Y.-G., M. M. Rienecker, M. J. Suarez, Y. Vikhliaev, B. Zhao, J. Marshak, G. Vernieres, and S. D. Schubert (2014), Decadal prediction skill in the GEOS-5 forecast system, *Climate Dynamics*, *42*(1-2), 1–20.

Hartmann, D., A. Klein Tank, M. Rusticucci, L. Alexander, S. Brönnimann, Y. Charabi, F. Dentener, E. Dlugokencky, D. Easterling, A. Kaplan, B. Soden, P. Thorne, M. Wild, and P. Zhai (2013), Chapter 2: Observations: Atmosphere and surface., *In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*.

Hasselmann, K. (1976), Stochastic climate models part I. Theory, *Tellus*, *28*(6), 473–485.

Hawkins, E., and R. Sutton (2007), Variability of the Atlantic thermohaline circulation

described by three-dimensional empirical orthogonal functions, *Climate Dynamics*, *29*(7-8), 745–762.

Hawkins, E., and R. Sutton (2009), The potential to narrow uncertainty in regional climate predictions, *Bulletin of the American Meteorological Society*, *90*(8), 1095–1107.

Hawkins, E., J. Robson, R. Sutton, D. Smith, and N. Keenlyside (2011), Evaluating the potential for statistical decadal predictions of sea surface temperatures with a perfect model approach, *Climate Dynamics*, *37*(11-12), 2495–2509.

Hermanson, L., and R. Sutton (2010), Case studies in interannual to decadal climate predictability, *Climate Dynamics*, *35*(7-8), 1169–1189.

Hodson, D. L., and R. T. Sutton (2012), The impact of resolution on the adjustment and decadal variability of the Atlantic meridional overturning circulation in a coupled climate model, *Climate Dynamics*, *39*(12), 3057–3073.

Hodson, D. L., J. I. Robson, and R. T. Sutton (2014), An anatomy of the cooling of the North Atlantic ocean in the 1960s and 1970s, *Journal of Climate*, *27*(21), 8229–8243.

Hoerling, M. P., J. W. Hurrell, and T. Xu (2001), Tropical origins for recent North Atlantic climate change, *Science*, *292*(5514), 90–92.

Holland, M., C. Bitz, M. Eby, and A. Weaver (2001), The role of ice-ocean interactions in the variability of the North Atlantic thermohaline circulation, *Journal of Climate*, *14*(5), 656–675.

Hoskins, B. J., and P. J. Valdes (1990), On the existence of storm-tracks, *Journal of the Atmospheric Sciences*, *47*(15), 1854–1864.

Huang, W., B. Wang, L. Li, L. Dong, P. Lin, Y. Yu, T. Zhou, L. Liu, S. Xu, K. Xia, et al. (2014), Variability of Atlantic meridional overturning circulation in FGOALS-g2, *Advances in Atmospheric Sciences*, *31*(1), 95–109.

Hughes, C. W., and M. P. Meredith (2006), Coherent sea-level fluctuations along the global continental slope, *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, *364*(1841), 885–901.

Hurrell, J., Y. Kushnir, G. Ottersen, and M. Visbeck (2003), *The North Atlantic Oscillation: Climatic Significance and Environmental Impact*, vol. 134, Amer Geophysical Union.

Ineson, S., and A. Scaife (2009), The role of the stratosphere in the European climate response to El Niño, *Nature Geoscience*, *2*(1), 32–36.

Ingleby, B., and M. Huddleston (2007), Quality control of ocean temperature and salinity profiles. historical and real-time data, *Journal of Marine Systems*, *65*(1), 158–175.

Johns, W. E., M. O. Baringer, L. Beal, S. Cunningham, T. Kanzow, H. L. Bryden, J. Hirschi, J. Marotzke, C. Meinen, B. Shaw, et al. (2011), Continuous, array-based estimates of Atlantic ocean heat transport at 26.5N, *Journal of Climate*, *24*(10), 2429–2449.

Johnson, G. C., and J. M. Lyman (2014), Oceanography: Where's the heat?, *Nature Climate Change*, *4*(11), 956–957.

Johnson, H., and D. Marshall (2002), Localization of abrupt change in the North Atlantic thermohaline circulation, *Geophysical Research Letters*, *29*(6).

Josey, S. (2001), A comparison of ECMWF, NCEP-NCAR, and SOC surface heat fluxes with moored buoy measurements in the subduction region of the northeast Atlantic, *Journal of Climate*, *14*(8), 1780–1789.

Jungclaus, J., H. Haak, M. Latif, and U. Mikolajewicz (2005), Arctic-North Atlantic interactions and multidecadal variability of the meridional overturning circulation, *Journal of Climate*, *18*(19), 4013–4031.

Kara, A., P. Rochford, and H. Hurlburt (2000), An optimal definition for ocean mixed layer depth, *Journal of Geophysical Research: Oceans*, *105*(C7), 16,803–16,821.

Karspeck, A. R., S. Yeager, G. Danabasoglu, T. Hoar, N. Collins, K. Raeder, J. Anderson, and J. Tribbia (2013), An ensemble adjustment kalman filter for the CCSM4 ocean component, *Journal of Climate*, *26*(19), 7392–7413.

Kato, S., N. G. Loeb, F. G. Rose, D. R. Doelling, D. A. Rutan, T. E. Caldwell, L. Yu, and R. A. Weller (2013), Surface irradiances consistent with CERES-derived top-of-atmosphere shortwave and longwave irradiances, *Journal of Climate*, *26*(9), 2719–2740.

Kelly, B. C. (2007), Some aspects of measurement error in linear regression of astronomical data, *The Astrophysical Journal*, *665*(2), 1489.

Kennedy, J., N. Rayner, R. Smith, D. Parker, and M. Saunby (2011), Reassessing biases and other uncertainties in sea surface temperature observations measured in situ since 1850: 2. Biases and homogenization, *Journal of Geophysical Research: Atmospheres (1984–2012)*, *116*(D14).

Kirtman, B., S. Power, J. Adedoyin, G. Boer, R. Bojariu, I. Camilloni, F. Doblas-Reyes, A. Fiore, M. Kimoto, G. Meehl, M. Prather, A. Sarr, C. Schar, R. Sutton, G. van Oldenborgh, V. G, and W. HJ (2013), Chapter 11: Near-term climate change: projections and predictability., *In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* .

Kirtman, B. P., C. Bitz, F. Bryan, W. Collins, J. Dennis, N. Hearn, J. L. Kinter III, R. Loft, C. Rousset, L. Siqueira, et al. (2012), Impact of ocean model resolution on CCSM climate simulations, *Climate Dynamics*, *39*(6), 1303–1328.

Knight, J., R. Allan, C. Folland, M. Vellinga, and M. Mann (2005), A signature of persistent natural thermohaline circulation cycles in observed climate, *Geophysical Research Letters*, *32*(20).

Knutti, R., D. Masson, and A. Gettelman (2013), Climate model genealogy: Generation CMIP5 and how we got there, *Geophysical Research Letters*, *40*(6), 1194–1199.

Kohl, A. (2005), Anomalies of meridional overturning: Mechanisms in the North Atlantic, *Journal of Physical Oceanography*, *35*(8), 1455–1472.

Kröger, J., W. A. Müller, and J.-S. von Storch (2012), Impact of different ocean reanalyses on decadal climate prediction, *Climate Dynamics*, *39*(3-4), 795–810.

Kwon, Y.-O., and C. Frankignoul (2012), Stochastically-driven multidecadal variability of the Atlantic meridional overturning circulation in CCSM3, *Climate Dynamics*, *38*(5-6), 859–876.

Kwon, Y.-O., and C. Frankignoul (2014), Mechanisms of multidecadal Atlantic meridional overturning circulation variability diagnosed in depth versus density space, *Journal of Climate*, *27*(24), 9359–9376.

Kwon, Y.-O., M. A. Alexander, N. A. Bond, C. Frankignoul, H. Nakamura, B. Qiu, and L. A. Thompson (2010), Role of the Gulf Stream and Kuroshio-Oyashio systems in large-scale atmosphere-ocean interaction: A review, *Journal of Climate*, *23*(12), 3249–3281.

Large, W. G., and S. G. Yeager (2004), *Diurnal to decadal global forcing for ocean and sea-ice models: The data sets and flux climatologies*, Citeseer.

Levitus, S., J. Antonov, T. Boyer, R. Locarnini, H. Garcia, and A. Mishonov (2009), Global ocean heat content 1955–2008 in light of recently revealed instrumentation problems, *Geophysical Research Letters*, *36*(7).

Lindzen, R. S., and S. Nigam (1987), On the role of sea surface temperature gradients in forcing low-level winds and convergence in the tropics, *Journal of the Atmospheric Sciences*, *44*(17), 2418–2436.

Liu, Z. (2012), Dynamics of interdecadal climate variability: A historical perspective, *Journal of Climate*, *25*(6), 1963–1995.

Lorenz, E. (1975), The physical bases of climate and climate modelling, *Climate predictability*, *16*, 132–136.

Lorenz, E. N. (1963), Deterministic nonperiodic flow, *Journal of the Atmospheric Sciences*, *20*(2), 130–141.

Lozier, M. S., S. Leadbetter, R. G. Williams, V. Roussenov, M. S. Reed, and N. J. Moore (2008), The spatial pattern and mechanisms of heat-content change in the North Atlantic, *Science*, *319*(5864), 800–803.

MacLachlan, C., A. Arribas, K. A. Peterson, A. Maidens, D. Fereday, A. A. Scaife, M. Gordon, M. Vellinga, A. Williams, R. E. Comer, J. Camp, P. Xavier, and G. Madec (2015), Global seasonal forecast system version 5 (GloSea5): A high-resolution seasonal forecast system, *Quarterly Journal of the Royal Meteorological Society*, *141*(689), 1072–1084.

MacMartin, D. G., E. Tziperman, and L. Zanna (2013), Frequency domain multimodel analysis of the response of Atlantic meridional overturning circulation to surface forcing, *Journal of Climate*, *26*(21), 8323–8340.

Madec, G. (2008), NEMO ocean engine: Note du pole de modélisation, Institut Pierre-Simon Laplace (IPSL), France, No 27 ISSN No 1288-1619, available at: http://www.nemo-ocean.eu, *Tech. rep.*, IPSL LSCE, UVSQ, CEA CNRS, Unite Mixte, Bat 712, F-91191 Gif Sur Yvette, France.

Magnusson, L., M. Alonso-Balmaseda, S. Corti, F. Molteni, and T. Stockdale (2013), Evaluation of forecast strategies for seasonal and decadal forecasts in presence of systematic model errors, *Climate Dynamics*, *41*(9-10), 2393–2409.

Mann, M., J. Park, and R. Bradley (1995), Global interdecadal and century-scale climate oscillations during the past 5 centuries, *Nature*, *378*(6554), 266–270.

Marzocchi, A., J. J.-M. Hirschi, N. P. Holliday, S. A. Cunningham, A. T. Blaker, and A. C. Coward (2015), The North Atlantic subpolar circulation in an eddy-resolving global ocean model, *Journal of Marine Systems*, *142*(0), 126 – 143.

Matei, D., J. Baehr, J. H. Jungclaus, H. Haak, W. A. Müller, and J. Marotzke (2012), Multiyear prediction of monthly mean Atlantic meridional overturning circulation at 26.5N, *Science*, *335*(6064), 76–79.

Matsueda, M., R. Mizuta, and S. Kusunoki (2009), Future change in wintertime atmospheric blocking simulated using a 20-km-mesh atmospheric global circulation model, *Journal of Geophysical Research: Atmospheres (1984–2012)*, *114*(D12).

Mauritzen, C., A. Melsom, and R. T. Sutton (2012), Importance of density-compensated temperature change for deep North Atlantic ocean heat uptake, *Nature Geoscience*, *5*(12), 905–910.

McCarthy, G. D., D. A. Smeed, W. E. Johns, E. Frajka-Williams, B. I. Moat, D. Rayner, M. O. Baringer, C. S. Meinen, J. Collins, and H. L. Bryden (2015), Measuring the Atlantic meridional overturning circulation at 26ºN, *Progress in Oceanography*, *130*, 91–111.

Mecking, J., N. S. Keenlyside, and R. J. Greatbatch (2014), Stochastically-forced multi-decadal variability in the North Atlantic: A model study, *Climate Dynamics*, *43*(1-2), 271–288.

Medhaug, I., H. Langehaug, T. Eldevik, T. Furevik, and M. Bentsen (2012), Mechanisms for decadal scale variability in a simulated Atlantic meridional overturning circulation, *Climate Dynamics*, *39*(1-2), 77–93.

Meehl, G. A., L. Goddard, J. Murphy, R. J. Stouffer, G. Boer, G. Danabasoglu, K. Dixon, M. A. Giorgetta, A. M. Greene, E. Hawkins, et al. (2009), Decadal prediction: Can it be skilful?, *Bulletin of the American Meteorological Society*, *90*(10), 1467–1485.

Meehl, G. A., H. Teng, and J. M. Arblaster (2014), Climate model simulations of the observed early-2000s hiatus of global warming, *Nature Climate Change*, *4*(10), 898–902.

Megann, A., D. Storkey, Y. Aksenov, S. Alderson, D. Calvert, T. Graham, P. Hyder, J. Siddorn, and B. Sinha (2014), GO5.0: The joint NERC–Met Office NEMO global ocean model for use in coupled and forced applications, *Geoscientific Model Development*, *7*(3), 1069–1092.

Meinen, C. S., W. E. Johns, S. L. Garzoli, E. van Sebille, D. Rayner, T. Kanzow, and M. O. Baringer (2013), Variability of the deep western boundary current at 26.5N during 2004–2009, *Deep Sea Research Part II: Topical Studies in Oceanography*, *85*, 154–168.

Menary, M. B., and A. A. Scaife (2014), Naturally forced multidecadal variability of the Atlantic meridional overturning circulation, *Climate Dynamics*, *42*(5-6), 1347–1362.

Menary, M. B., D. L. Hodson, J. I. Robson, R. T. Sutton, and R. A. Wood (2015a), A mechanism of internal decadal Atlantic ocean variability in a high-resolution coupled climate model, *Journal of Climate*, *28*(19), 7764–7785.

Menary, M. B., D. L. R. Hodson, J. I. Robson, R. T. Sutton, R. A. Wood, and J. A. Hunt (2015b), Exploring the impact of CMIP5 model biases on the simulation of North Atlantic decadal variability, *Geophysical Research Letters*, *42*(14), 5926–5934, 2015GL064360.

Menary, M. B., W. Park, K. Lohmann, M. Vellinga, M. D. Palmer, M. Latif, and J. H. Jungclaus (2012), A multimodel comparison of centennial Atlantic meridional overturning circulation variability, *Climate Dynamics*, *38*(11-12), 2377–2388.

Michael, J.-P., V. Misra, and E. P. Chassignet (2013), The El Niño and Southern Oscillation in the historical centennial integrations of the new generation of climate models, *Regional Environmental Change*, *13*(1), 121–130.

Minobe, S., A. Kuwano-Yoshida, N. Komori, S.-P. Xie, and R. J. Small (2008), Influence of the Gulf Stream on the troposphere, *Nature*, *452*(7184), 206–209.

Moss, R. H., J. A. Edmonds, K. A. Hibbard, M. R. Manning, S. K. Rose, D. P. Van Vuuren, T. R. Carter, S. Emori, M. Kainuma, T. Kram, et al. (2010), The next generation of scenarios for climate change research and assessment, *Nature*, *463*(7282), 747–756.

Msadek, R., K. Dixon, T. Delworth, and W. Hurlin (2010), Assessing the predictability of the Atlantic meridional overturning circulation and associated fingerprints, *Geophysical Research Letters*, *37*(19).

Murphy, J. (1990), Assessment of the practical utility of extended range ensemble forecasts, *Quarterly Journal of the Royal Meteorological Society*, *116*(491), 89–125.

Nonaka, M., and S.-P. Xie (2003), Covariations of sea surface temperature and wind over the Kuroshio and its extension: Evidence for ocean-to-atmosphere feedback, *Journal of Climate*, *16*(9), 1404–1413.

Omrani, N.-E., N. S. Keenlyside, J. Bader, and E. Manzini (2014), Stratosphere key for wintertime atmospheric response to warm Atlantic decadal conditions, *Climate Dynamics*, *42*(3-4), 649–663.

Ortega, P., E. Hawkins, and R. Sutton (2011), Processes governing the predictability of

the Atlantic meridional overturning circulation in a coupled GCM, *Climate Dynamics*, *37*(9-10), 1771–1782.

Oschlies, A. (2002), Improved representation of upper-ocean dynamics and mixed layer depths in a model of the North Atlantic on switching from eddy-permitting to eddy-resolving grid resolution, *Journal of Physical Oceanography*, *32*(8), 2277–2298.

Otterå, O. H., M. Bentsen, H. Drange, and L. Suo (2010), External forcing as a metronome for Atlantic multidecadal variability, *Nature Geoscience*, *3*(10), 688–694.

Park, W., and M. Latif (2008), Multidecadal and multicentennial variability of the meridional overturning circulation, *Geophysical Research Letters*, *35*(22).

Pascoe, C., L. Gray, and A. Scaife (2006), A GCM study of the influence of equatorial winds on the timing of sudden stratospheric warmings, *Geophysical Research Letters*, *33*(6).

Penduff, T., M. Juza, L. Brodeau, G. Smith, B. Barnier, J.-M. Molines, A.-M. Tréguier, and G. Madec (2010), Impact of global ocean model resolution on sea-level variability with emphasis on interannual time scales, *Ocean Science*, *6*(1), 269–284.

Pierce, D. W. (2001), Distinguishing coupled ocean–atmosphere interactions from background noise in the North Pacific, *Progress in Oceanography*, *49*(1), 331–352.

Pohlmann, H., M. Botzet, M. Latif, A. Roesch, M. Wild, and P. Tschuck (2004), Estimating the decadal predictability of a coupled AOGCM, *Journal of Climate*, *17*(22), 4463–4472.

Pohlmann, H., D. M. Smith, M. A. Balmaseda, N. S. Keenlyside, S. Masina, D. Matei, W. A. Muller, and P. Rogel (2013), Predictability of the mid-latitude Atlantic meridional overturning circulation in a multi-model system, *Climate Dynamics*, *41*(3-4), 775–785.

Polo, I., J. Robson, R. Sutton, and M. A. Balmaseda (2014), The importance of wind and buoyancy forcing for the boundary density variations and the geostrophic component of the AMOC at 26°N, *Journal of Physical Oceanography*, *44*(9), 2387–2408.

Pozo-Vázquez, D., S. Gámiz-Fortis, J. Tovar-Pescador, M. Esteban-Parra, and Y. Castro-Díez (2005), North Atlantic winter SLP anomalies based on the autumn ENSO state, *Journal of Climate*, *18*(1), 97–103.

Prahl, F. G., L. A. Muehlhausen, and D. L. Zahnle (1988), Further evaluation of long-chain alkenones as indicators of paleoceanographic conditions, *Geochimica et Cosmochimica Acta*, *52*(9), 2303–2310.

Qiu, B., and K. A. Kelly (1993), Upper-ocean heat balance in the Kuroshio extension region, *Journal of Physical Oceanography*, *23*(9), 2027–2041.

Quillmann, U., T. M. Marchitto, A. E. Jennings, J. T. Andrews, and B. F. Friestad (2012), Cooling and freshening at 8.2 ka on the NW Iceland Shelf recorded in paired δ^{18}O and Mg/Ca measurements of the benthic foraminifer Cibicides lobatulus, *Quaternary Research*, *78*(3), 528–539.

Rayner, N., D. E. Parker, E. Horton, C. Folland, L. Alexander, D. Rowell, E. Kent, and A. Kaplan (2003), Global analyses of sea surface temperature, sea ice, and night marine air temperature since the late nineteenth century, *Journal of Geophysical Research: Atmospheres (1984–2012)*, *108*(D14).

Rhein, M., D. Kieke, S. Hüttl-Kabus, A. Roessler, C. Mertens, R. Meissner, B. Klein, C. W. Böning, and I. Yashayaev (2011), Deep water formation, the subpolar gyre, and the meridional overturning circulation in the subpolar North Atlantic, *Deep Sea Research Part II: Topical Studies in Oceanography*, *58*(17), 1819–1832.

Richardson, P. L. (1980), *The Benjamin Franklin and Timothy Folger charts of the Gulf Stream*, Springer.

Roberts, C., J. Waters, K. Peterson, M. Palmer, G. McCarthy, E. Frajka-Williams, K. Haines, D. Lea, M. Martin, D. Storkey, et al. (2013), Atmosphere drives recent interannual variability of the Atlantic meridional overturning circulation at 26.5N, *Geophysical Research Letters*, *40*(19), 5164–5170.

Robson, J., R. Sutton, K. Lohmann, D. Smith, and M. D. Palmer (2012), Causes of the rapid warming of the North Atlantic ocean in the mid-1990s, *Journal of Climate*, *25*(12), 4116–4134.

Robson, J., D. Hodson, E. Hawkins, and R. Sutton (2014), Atlantic overturning in decline?, *Nature Geoscience*, *7*(1), 2–3.

Robson, J. I. (2010), Understanding the performance of a decadal prediction system, Ph.D. thesis, The University of Reading, Reading, UK.

Rodwell, M., D. Rowell, and C. Folland (1999), Oceanic forcing of the wintertime North Atlantic Oscillation and European climate, *Nature*, *398*(6725), 320–323.

Roessler, A., M. Rhein, D. Kieke, and C. Mertens (2015), Long-term observations of North Atlantic current transport at the gateway between western and eastern Atlantic, *Journal of Geophysical Research: Oceans*, *120*(6), 4003–4027.

Roussenov, V. M., R. G. Williams, C. W. Hughes, and R. J. Bingham (2008), Boundary wave communication of bottom pressure and overturning changes for the North Atlantic, *Journal of Geophysical Research: Oceans*, *113*(C8).

Sarafanov, A., A. Falina, A. Sokov, and A. Demidov (2008), Intense warming and salinification of intermediate waters of southern origin in the eastern subpolar North Atlantic in the 1990s to mid-2000s, *Journal of Geophysical Research: Oceans*, *113*(C12).

Saravanan, R., and J. C. McWilliams (1997), Stochasticity and spatial resonance in interdecadal climate fluctuations, *Journal of Climate*, *10*(9), 2299–2320.

Saravanan, R., and J. C. McWilliams (1998), Advective ocean-atmosphere interaction: An analytical stochastic model with implications for decadal variability, *Journal of Climate*, *11*(2), 165–188.

Scaife, A., A. Arribas, E. Blockley, A. Brookshaw, R. Clark, N. Dunstone, R. Eade, D. Fereday, C. Folland, M. Gordon, et al. (2014), Skillful long-range prediction of European and North American winters, *Geophysical Research Letters*, *41*(7), 2514–2519.

Scaife, A. A., D. Copsey, C. Gordon, C. Harris, T. Hinton, S. Keeley, A. O'Neill, M. Roberts, and K. Williams (2011), Improved Atlantic winter blocking in a climate model, *Geophysical Research Letters*, *38*(23).

Schauer, U., and A. Beszczynska-Möller (2009), Problems with estimation and interpretation of oceanic heat transport — conceptual remarks for the case of Fram Strait in the Arctic ocean, *Ocean Science*, *5*(4), 487–494.

Schlesinger, M. E., and N. Ramankutty (1994), An oscillation in the global climate system of period 65-70 years, *Nature*, *367*(6465), 723–726.

Screen, J. A. (2013), Influence of Arctic sea ice on European summer precipitation, *Environmental Research Letters*, *8*(4), 044,015.

Sévellec, F., and A. V. Fedorov (2013), The leading, interdecadal eigenmode of the Atlantic meridional overturning circulation in a realistic ocean model, *Journal of Climate*, *26*(7), 2160–2183.

Sicre, M.-A., P. Yiou, J. Eiriksson, U. Ezat, E. Guimbaut, I. Dahhaoui, K.-L. Knudsen, E. Jansen, and J.-L. Turon (2008), A 4500-year reconstruction of sea surface temperature variability at decadal time-scales off North Iceland, *Quaternary Science Reviews*, *27*(21-22), 2041–2047.

Smeed, D., G. McCarthy, S. Cunningham, E. Frajka-Williams, D. Rayner, W. Johns, C. Meinen, M. Baringer, B. Moat, A. Duchez, et al. (2014), Observed decline of the Atlantic meridional overturning circulation 2004–2012, *Ocean Science*, *10*(1), 29–38.

Smith, D. M., R. Eade, N. J. Dunstone, D. Fereday, J. M. Murphy, H. Pohlmann, and A. A. Scaife (2010), Skilful multi-year predictions of Atlantic hurricane frequency, *Nature Geoscience*, *3*(12), 846–849.

Smith, D. M., A. A. Scaife, and B. P. Kirtman (2012), What is the current state of scientific knowledge with regard to seasonal and decadal forecasting?, *Environmental Research Letters*, *7*(1), 15,602–15,612.

Smith, D. M., R. Eade, and H. Pohlmann (2013), A comparison of full-field and anomaly initialization for seasonal to decadal climate prediction, *Climate Dynamics*, *41*(11-12), 3325–3338.

Smith, D. M., S. Cusack, A. W. Colman, C. K. Folland, G. R. Harris, and J. M. Murphy

(2007), Improved surface temperature prediction for the coming decade from a global climate model, *Science*, *317*(5839), 796–799.

Snoussi, E. H. (1998), Necessary conditions for multistationarity and stable periodicity, *Journal of Biological Systems*, *6*(01), 3–9.

Spence, P., O. A. Saenko, W. Sijp, and M. England (2011), The role of bottom pressure torques on the interior pathways of North Atlantic deep water, *Journal of Physical Oceanography*, *42*(1), 110–125.

Strong, C., and G. Magnusdottir (2008), Tropospheric Rossby wave breaking and the NAO/NAM, *Journal of the Atmospheric Sciences*, *65*(9), 2861–2876.

Sutton, R., and M. Allen (1997), Decadal predictability of North Atlantic sea surface temperature and climate, *Nature*, *388*(6642), 563–567.

Sutton, R., and D. Hodson (2005), Atlantic Ocean forcing of North American and European summer climate, *Science*, *309*(5731), 115–118.

Sutton, R. T., W. A. Norton, and S. P. Jewson (2000), The North Atlantic Oscillation — what role for the ocean?, *Atmospheric Science Letters*, *1*(2), 89–100.

Swingedouw, D., C. B. Rodehacke, E. Behrens, M. Menary, S. M. Olsen, Y. Gao, U. Mikolajewicz, J. Mignot, and A. Biastoch (2013), Decadal fingerprints of freshwater discharge around Greenland in a multi-model ensemble, *Climate Dynamics*, *41*(3-4), 695–720.

Swingedouw, D., P. Ortega, J. Mignot, E. Guilyardi, V. Masson-Delmotte, P. G. Butler, M. Khodri, and R. Séférian (2015), Bidecadal North Atlantic ocean circulation variability controlled by timing of volcanic eruptions, *Nature Communications*, *6*.

Taylor, K. E., R. J. Stouffer, and G. A. Meehl (2012), An overview of CMIP5 and the experiment design, *Bulletin of the American Meteorological Society*, *93*(4), 485–498.

Tebaldi, C., R. L. Smith, D. Nychka, and L. O. Mearns (2005), Quantifying uncertainty in projections of regional climate change: A bayesian approach to the analysis of multimodel ensembles, *Journal of Climate*, *18*(10), 1524–1540.

Timmermann, A., M. Latif, R. Voss, and A. Grotzner (1998), Northern hemispheric interdecadal variability: A coupled air-sea mode, *Journal of Climate*, *11*(8), 1906–1931.

Toole, J., R. Curry, T. Joyce, M. McCartney, and B. Peña-Molino (2011), Transport of the North Atlantic deep western boundary current about 39N, 70W: 2004–2008, *Deep Sea Research Part II: Topical Studies in Oceanography*, *58*(17), 1768–1780.

Tréguier, A.-M., S. Theetten, E. P. Chassignet, T. Penduff, R. Smith, L. Talley, J. Beismann, and C. Böning (2005), The North Atlantic subpolar gyre in four high-resolution models, *Journal of Physical Oceanography*, *35*(5), 757–774.

Tréguier, A. M., J. Deshayes, J. Le Sommer, C. Lique, G. Madec, T. Penduff, J.-M. Molines, B. Barnier, R. Bourdalle-Badie, and C. Talandier (2014), Meridional transport of salt in the global ocean from an eddy-resolving model, *Ocean Science*, *10*(2), 243–255.

Trenberth, K. E. (1997), The definition of El Niño, *Bulletin of the American Meteorological Society*, *78*(12), 2771–2777.

Uppala, S. M., P. Kållberg, A. Simmons, U. Andrae, V. Bechtold, M. Fiorino, J. Gibson, J. Haseler, A. Hernandez, G. Kelly, et al. (2005), The ERA-40 re-analysis, *Quarterly Journal of the Royal Meteorological Society*, *131*(612), 2961–3012.

Vecchi, G. A., R. Msadek, T. L. Delworth, K. W. Dixon, E. Guilyardi, E. Hawkins, A. R. Karspeck, J. Mignot, J. Robson, A. Rosati, et al. (2012), Comment on "multiyear prediction of monthly mean Atlantic meridional overturning circulation at 26.5N", *Science*, *338*(6107), 604–604.

Visbeck, M., H. Cullen, G. Krahmann, and N. Naik (1998), An ocean model's response to North Atlantic Oscillation-like wind forcing, *Geophysical Research Letters*, *25*(24), 4521–4524.

Volkov, D. L., T. Lee, and L.-L. Fu (2008), Eddy-induced meridional heat transport in the ocean, *Geophysical Research Letters*, *35*(20).

Walters, D., M. Best, A. Bushell, D. Copsey, J. Edwards, P. Falloon, C. Harris, A. Lock, J. Manners, C. Morcrette, et al. (2011), The Met Office Unified Model global atmo-

sphere 3.0/3.1 and JULES global land 3.0/3.1 configurations, *Geoscientific Model Development Discussions*, *4*(2), 1213–1271.

Wang, C., L. Zhang, S.-K. Lee, L. Wu, and C. R. Mechoso (2014), A global perspective on CMIP5 climate model biases, *Nature Climate Change*, *4*(3), 201–205.

Watanabe, M., M. Kimoto, T. Nitta, and M. Kachi (1999), A comparison of decadal climate oscillations in the north Atlantic detected in observations and a coupled GCM, *Journal of Climate*, *12*(9), 2920–2940.

Weigel, A. P., M. A. Liniger, and C. Appenzeller (2007), The discrete brier and ranked probability skill scores, *Monthly Weather Review*, *135*(1), 118–124.

Wilks, D. (1997), Resampling hypothesis tests for autocorrelated fields, *Journal of Climate*, *10*(1), 65–82.

Williams, K., C. Harris, A. Bodas-Salcedo, J. Camp, R. Comer, D. Copsey, D. Fereday, T. Graham, R. Hill, T. Hinton, et al. (2015), The Met Office Global Coupled model 2.0 (GC2) configuration, *Geoscientific Model Development Discussions*, *8*(1), 521–565.

Williams, R. G., V. Roussenov, D. Smith, and M. S. Lozier (2014), Decadal evolution of ocean thermal anomalies in the North Atlantic: The effects of Ekman, overturning, and horizontal transport, *Journal of Climate*, *27*(2), 698–719.

Wittenberg, A. T., A. Rosati, T. L. Delworth, G. A. Vecchi, and F. Zeng (2014), ENSO modulation: Is it decadally predictable?, *Journal of Climate*, *27*(7), 2667–2681.

Wohlleben, T., and A. Weaver (1995), Interdecadal climate variability in the subpolar North-Atlantic, *Climate Dynamics*, *11*(8), 459–467.

Woollings, T., C. Franzke, D. Hodson, B. Dong, E. Barnes, C. Raible, and J. Pinto (2014), Contrasting interannual and multidecadal NAO variability, *Climate Dynamics*, *45*(1-2), 539–556.

Wouters, B., W. Hazeleger, S. Drijfhout, v. G. Oldenborgh, and V. Guemas (2013), Multi-year predictability of the North Atlantic subpolar gyre, *Geophysical Research Letters*, *40*(12), 3080–3084.

Xie, L., T. Yan, L. J. Pietrafesa, J. M. Morrison, and T. Karl (2005), Climatology and interannual variability of North Atlantic hurricane tracks, *Journal of Climate*, *18*(24), 5370–5381.

Yang, S., and J. H. Christensen (2012), Arctic sea ice reduction and european cold winters in CMIP5 climate change experiments, *Geophysical Research Letters*, *39*(20).

Yashayaev, I. (2007), Hydrographic changes in the Labrador Sea, 1960–2005, *Progress in Oceanography*, *73*(3), 242–276.

Yeager, S., and G. Danabasoglu (2012), Sensitivity of Atlantic meridional overturning circulation variability to parameterized Nordic Sea overflows in CCSM4, *Journal of Climate*, *25*(6), 2077–2103.

Zanna, L. (2012), Forecast skill and predictability of observed Atlantic sea surface temperatures, *Journal of Climate*, *25*(14), 5047–5056.

Zhang, L., and C. Zhao (2015), Processes and mechanisms for the model SST biases in the North Atlantic and North Pacific: A link with the Atlantic meridional overturning circulation, *Journal of Advances in Modeling Earth Systems*, *7*(2), 739–758.

Zhang, R. (2008), Coherent surface-subsurface fingerprint of the Atlantic meridional overturning circulation, *Geophysical Research Letters*, *35*(20).

Zhang, R. (2010), Latitudinal dependence of Atlantic meridional overturning circulation (AMOC) variations, *Geophysical Research Letters*, *37*.

Zhang, R., and T. L. Delworth (2006), Impact of Atlantic multidecadal oscillations on India/Sahel rainfall and Atlantic hurricanes, *Geophysical Research Letters*, *33*(17).

Zhang, R., and G. K. Vallis (2007), The role of bottom vortex stretching on the path of the north Atlantic western boundary current and on the northern recirculation gyre, *Journal of Physical Oceanography*, *37*(8), 2053–2080.

Zhang, R., T. L. Delworth, R. Sutton, D. L. Hodson, K. W. Dixon, I. M. Held, Y. Kushnir, J. Marshall, Y. Ming, R. Msadek, et al. (2013), Have aerosols caused the observed Atlantic multidecadal variability?, *Journal of the Atmospheric Sciences*, *70*(4), 1135–1144.