

RESEARCH ARTICLE

10.1002/2018JC013774

Local and Remote Influences on the Heat Content of the Labrador Sea: An Adjoint Sensitivity Study
Daniel C. Jones¹ , **Gael Forget²** , **Bablu Sinha³**, **Simon A. Josey³**, **Emma J. D. Boland¹** ,
Andrew J. S. Meijers¹ , and **Emily Shuckburgh¹** 
¹British Antarctic Survey, Natural Environment Research Council, Cambridge, UK, ²Massachusetts Institute of Technology, Cambridge, MA, USA, ³National Oceanography Centre, Southampton, UK

Key Points:

- We use adjoint sensitivity fields to quantify possible influences on Labrador Sea heat content
- We identify a basin-scale adjustment mechanism involving the West African and European shelves
- Nonlocal heat fluxes can have a considerable impact on Labrador Sea heat content

Supporting Information:

- Supporting Information S1
- Video S1
- Video S2
- Video S3
- Video S4
- Video S5
- Video S6
- Video S7

Correspondence to:

D. C. Jones,
dannes@bas.ac.uk

Citation:

Jones, D. C., Forget, G., Sinha, B., Josey, S. A., Boland, E. J. D., Meijers, A. J. S., & Shuckburgh, E. (2018). Local and remote influences on the heat content of the Labrador Sea: An adjoint sensitivity study. *Journal of Geophysical Research: Oceans*, 123, 2646–2667. <https://doi.org/10.1002/2018JC013774>

Received 10 JAN 2018

Accepted 21 MAR 2018

Accepted article online 6 APR 2018

Published online 19 APR 2018

Abstract The Labrador Sea is one of the few regions on the planet where the interior ocean can exchange heat directly with the atmosphere via strong, localized, wintertime convection, with possible implications for the state of North Atlantic climate and global surface warming. Using an observationally constrained ocean adjoint model, we find that annual-mean Labrador Sea heat content is sensitive to temperature/salinity changes (1) along potential source water pathways (e.g., the subpolar gyre, the North Atlantic Current, the Gulf Stream) and (2) along the West African and European shelves, which are not significant source water regions for the Labrador Sea. The West African coastal/shelf adjustment mechanism, which may be excited by changes in along-shelf wind stress, involves pressure anomalies that propagate along a coastal waveguide toward Greenland, changing the across-shelf pressure gradient in the North Atlantic and altering heat convergence in the Labrador Sea. We also find that nonlocal (in space and time) heat fluxes (e.g., in the Irminger Sea, the seas south of Iceland) can have a strong impact on Labrador Sea heat content. Understanding and predicting the state of the Labrador Sea and its potential impacts on North Atlantic climate and global surface warming will require monitoring of oceanic and atmospheric properties at remote sites in the Irminger Sea, the subpolar gyre, and along the West African and European shelf/coast system, among others.

Plain Language Summary There are only a handful of locations on Earth where natural processes can rapidly inject heat and carbon into the interior ocean, where it can remain for decades to centuries, potentially slowing global surface warming. One of these locations is the Labrador Sea, which features strong exchanges of heat with the atmosphere and exceptionally deep mixing between the surface ocean and interior ocean. In this paper, we examine the factors that influence the heat content of the Labrador Sea. Using a numerical model, we find that although the heat content is most sensitive to local exchanges with the atmosphere, there is an unexpected connection between the heat content of the Labrador Sea and wind strength along the coast of West Africa and Europe. Sustained changes in wind strength in those regions can change the large-scale circulation of the entire North Atlantic, ultimately changing the amount of heat that gets transported into the Labrador Sea and potentially impacting North Atlantic climate and global surface warming.

1. Introduction

The Labrador Sea (LS) is a semienclosed marginal sea of the North Atlantic Ocean flanked by the continental shelves of North America and Greenland (Figure 1a). Because of its partially enclosed geometry and significant seasonal buoyancy loss, the Labrador Sea features some of the deepest mixed layers in the world ocean, reaching over 2,000 m in some years (Lazier et al., 2002; Piron et al., 2017; Spall, 2004). Temperature anomalies can enter the deep interior ocean via the Labrador Sea, potentially impacting oceanic uptake and storage of heat and carbon, with implications for global and regional climate (Lozier et al., 2017; Pérez et al., 2013, and references therein). For instance, an increase in heat uptake and intermediate-depth heat storage in the subpolar North Atlantic (among other regions) during the first decade of the 21st century has been connected to a hiatus in global surface warming (Chen & Tung, 2014; Drijfhout et al., 2014). Record low densities in the Labrador Sea have been connected to reduced northward ocean heat transport and

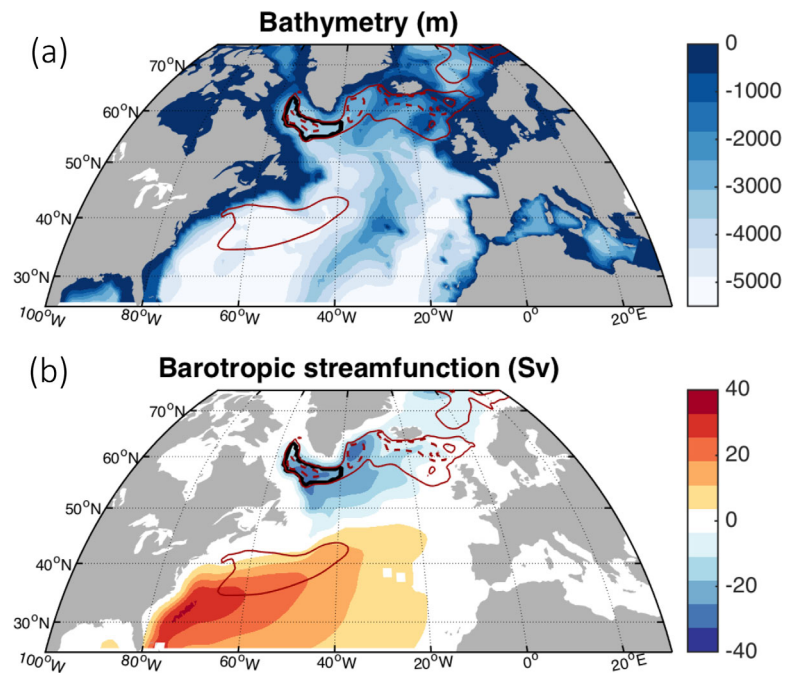


Figure 1. ECCOv4-r2 (a) bathymetry and (b) multiyear mean barotropic stream function for 1992–2011, constructed from annual-mean stream functions. The thick, solid black line indicates the Labrador Sea region wherein the March–April–May (MAM) mean mixed layer depths exceed 300 m. Also shown are the 250 m (red, solid) and 500 m (red, dashed) MAM mean mixed layer depth contours for 1992–2011.

significant cooling of the upper North Atlantic (Robson et al., 2014, 2016). A recent high-resolution climate model study found that such negative Labrador Sea density trends appear to be followed by positive winter states of the North Atlantic Oscillation, which can ultimately reverse the sign of the density trend through multidecadal atmosphere–ocean interactions (Ortega, 2017; Sutton et al., 2017). Understanding the factors that can alter Labrador Sea heat content is thus especially important for predicting the state of the North Atlantic sector and more broadly for predicting global surface warming.

Like most of the global ocean, the Labrador Sea has a long memory in that it may be affected by processes and properties in remote regions across a wide range of time scales (Robson et al., 2012). For example, changes in the nearby Irminger Sea and the remote Nordic Seas can influence stratification in the Labrador Sea (Pickart et al., 2003). Understanding how both local and remote oceanic and atmospheric properties affect the Labrador Sea is important for understanding the climate system and may help guide the design of future observational/monitoring networks (Heimbach et al., 2011; Liu & Alexander, 2007). In this study, we aim to understand how local and remote ocean properties (e.g., potential temperature) and surface forcing can affect the heat content of the Labrador Sea. We will address the following three questions:

- What are the potential source waters of the Labrador Sea?
- What are the possible influences of local and remote ocean properties on the heat content of the Labrador Sea?
- What are the possible influences of local and remote net heat fluxes and wind stresses on the heat content of the Labrador Sea?

In order to address these questions, we will use an adjoint method to calculate the linear sensitivities of the annual-mean Labrador Sea heat content to the time-evolving ocean state and surface forcing. In section 2, we describe the model setup and the particular adjoint sensitivity experiment we performed. Because adjoint methods are well described in many places, we refer the reader to these works for a more thorough and general description to adjoint modeling (e.g., Fukumori et al., 2007; Griewank & Walther, 2012; Heimbach, 2008; Marotzke et al., 1999; Mazloff et al., 2010; Thacker & Long, 1988; Verdy et al., 2014). In section 3, we discuss the results of our adjoint sensitivity experiments. In particular, we identify and examine an

adjustment mechanism that involves a teleconnection between the West African shelf and the Labrador Sea. Our conclusions are summarized in section 4.

2. Model Description and Experimental Design

We use the modeling setup associated with ECCOv4 (release 2, hereafter ECCOv4-r2 or just ECCOv4), an observationally constrained ocean state estimate, to calculate sensitivity fields. The model setup is available for download on Github (https://github.com/gaelforget/ECCO_v4_r2) as an instance of the MIT general circulation model (MITgcm, <http://mitgcm.org/>). ECCOv4-r2 is a product of the Estimating the Circulation and Climate of the Ocean (ECCO) consortium, which has produced a large variety of state estimation products that are freely available for download via <http://www.ecco-group.org/>. The adjoint model used in this work was generated using the algorithmic differentiation tool TAF (Giering & Kaminski, 1998, <http://www.fastopt.com/>). We briefly describe the model setup and state estimation process below; readers interested in a more detailed description are referred to Forget et al. (2015a) and references therein.

ECCOv4 uses a Lat-Lon-Cap (LLC) grid referred to as LLC90 that covers the global ocean, including the Arctic Ocean. The horizontal grid size ranges from around 40 to 50 km in the Arctic up to 110 km at the equator. Parameterized diffusion includes diapycnal and isopycnal components, simple convective adjustment, and the GGL mixed layer turbulence closure scheme (Gaspar et al., 1990). The along-isopycnal effect of unresolved eddies is parameterized as a bolus transport (Gent & McWilliams, 1990, hereafter GM). In this work, we use diffusivity and GM intensity parameters that have been optimized by the ECCOv4-r2 state estimation process, all of which are time-invariant, three-dimensional fields (Forget et al., 2015b). ECCOv4 features fully interactive, dynamic sea ice, so buoyancy, and mass fluxes are recalculated based on the thermodynamic balance of Losch et al. (2010). Open ocean rain, evaporation, and runoff simply carry (advect through the free surface) the local SST and a salinity value of zero, and runoff is provided by a monthly climatology (Fekete et al., 2002). Surface salinity restoring is *not* used here. Buoyancy, radiative, and mass fluxes are calculated using the bulk formulae of (Large & Yeager, 2009) using six-hourly ERA-Interim reanalysis fields (Dee et al., 2011) as a “first guess” for the forcing fields. Specifically, we use wind stress, 2 m air temperature, 2 m specific humidity, wind speed, downward longwave radiation, and downward shortwave radiation as model inputs. These fields have been iteratively adjusted by the state estimation process in order to minimize model-data misfits.

2.1. Validation of the ECCOv4 Global Ocean State Estimate

ECCOv4 is constrained by a global set of observations and represents the Labrador Sea and more generally the North Atlantic at sufficient accuracy for our purposes. ECCOv4-r2 captures the annual cycle and interannual variability of Labrador Sea deep convection, as seen by comparison with the gridded Argo product of (Roemmich & Gilson, 2009, RG09, Figure 2). Even though RG09 shows more high-frequency variability than ECCOv4, the two products are in good agreement at seasonal and interannual time scales.

We compare individual, nongridded Argo profiles with ECCOv4 “profiles” taken at the locations and times of the Argo profiles in the Labrador Sea (Figures 3a–3d). Using this approach offers a direct comparison with observations at specific locations and times, thus it is a particularly stringent test of the validity of the ECCOv4 solution. At 100 and 750 m during the Argo period, the mean ECCOv4 temperature and salinity lie within roughly 5% of the mean Argo values, although individual profiles may feature much larger differences (for temperatures, the 95% misfit interval is typically around 20% and up to roughly 50% of the Argo mean value in extreme cases). The influence of deep convection from 1992 to 1996 can be seen in the ECCOv4 temperature and salinity profiles. ECCOv4-r2 also captures the seasonal cycle, interannual variability, and long-term trend in sea level height as measured by altimetry (Figure 3e, for more details see Forget & Ponte, 2015). Labrador Sea bottom pressure is somewhat noisier, with a correlation of approximately 0.45 between ECCOv4-r2 and the GRACE-mascons product of Watkins et al. (2015). ECCOv4-r2 is also in good agreement with sea surface temperatures from the HadISST 1.1 product (Rayner et al., 2003), with a correlation of 0.95 (supporting information Figure S1), although ECCOv4-r2 sea surface temperatures are consistently colder than HadISST 1.1 in the winter. Because some of the same data (e.g., Argo, altimetry) have been used to constrain ECCOv4 and the other products, good agreement between them is perhaps not surprising. The presented comparisons, however, provide confirmation that the Labrador Sea and the broader

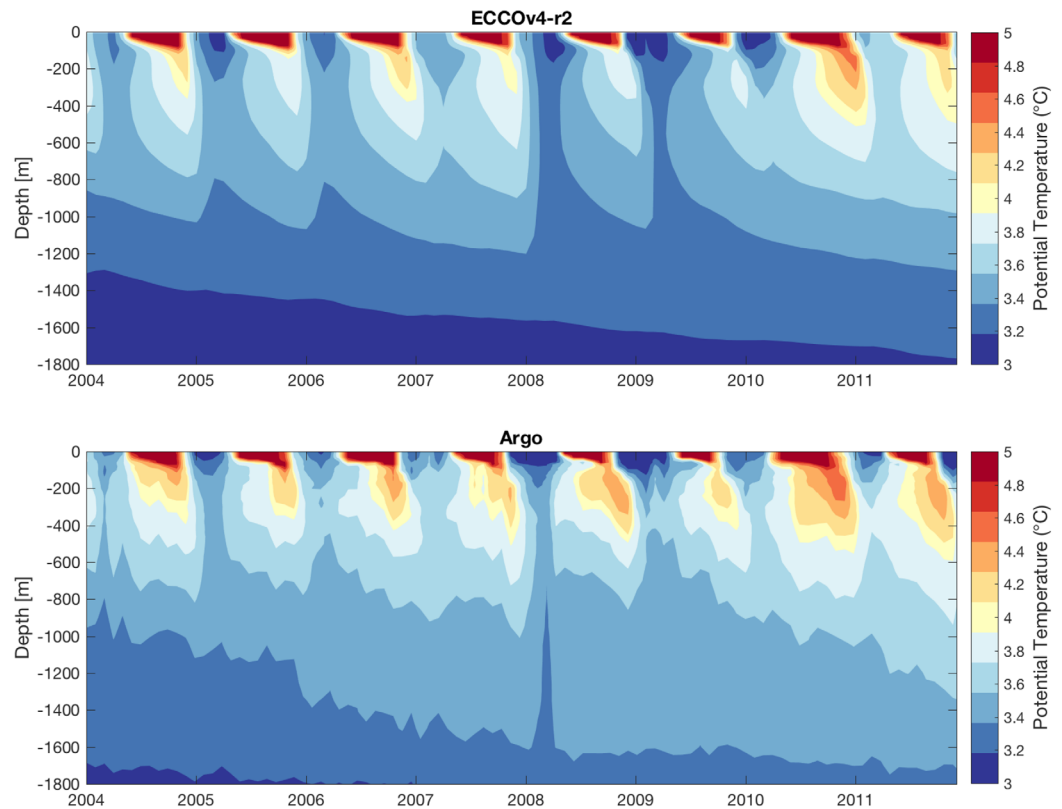


Figure 2. Comparison of ECCOv4-r2 and Argo potential temperatures ($^{\circ}\text{C}$), averaged between $55\text{--}50^{\circ}\text{W}$ and $55\text{--}60^{\circ}\text{N}$. Argo data taken from Scripps gridded product (Roemmich & Gilson, 2009, http://www.argo.ucsd.edu/Gridded_fields.html).

North Atlantic are both well represented in ECCOv4-r2, giving us realistic circulation and hydrography well-suited for adjoint sensitivity experiments.

2.2. Design of the Adjoint Sensitivity Experiment

Adjoint methods allow for sensitivity calculations that would be extremely impractical by more conventional means. In a typical “forward” perturbation experiment, the input of a numerical model (e.g., net heat flux) is perturbed by a chosen finite amount at a particular set of locations and times, and the effects are observed in various output fields (e.g., sea surface temperature). The effects propagate away from the perturbation site at a range of speeds, expressing the time scales of various adjustment processes. By contrast, in an adjoint sensitivity experiment, one defines a single quantity of interest (which may be an integral over some chosen region and time period), and the adjoint method simultaneously calculates the sensitivities to *every* selected input at *all* locations and times that are included in the numerical model (Figure 4). Thus, a single adjoint sensitivity run calculates sensitivities that would otherwise require an unfeasibly large number of forward perturbation experiments.

It is worth noting that adjoint sensitivity fields are not simply correlations between variables. Adjoint sensitivity fields indicate causal relationships contained in the model equations, whereas correlations describe how two variables change together, irrespective of whether or not they are causally related. Of course, the causal relationships highlighted by adjoint methods are those of the model, which are only approximations of real processes.

One possible caveat is that adjoint sensitivities are linearized about a time-varying reference state, which is a sufficiently accurate approach for some applications but not necessarily for others. The linear approximation is generally expected to hold for sufficiently small perturbations and short time scales. In this work, we use objective functions that are averaged over 1 year and over the entire water column in part to ensure

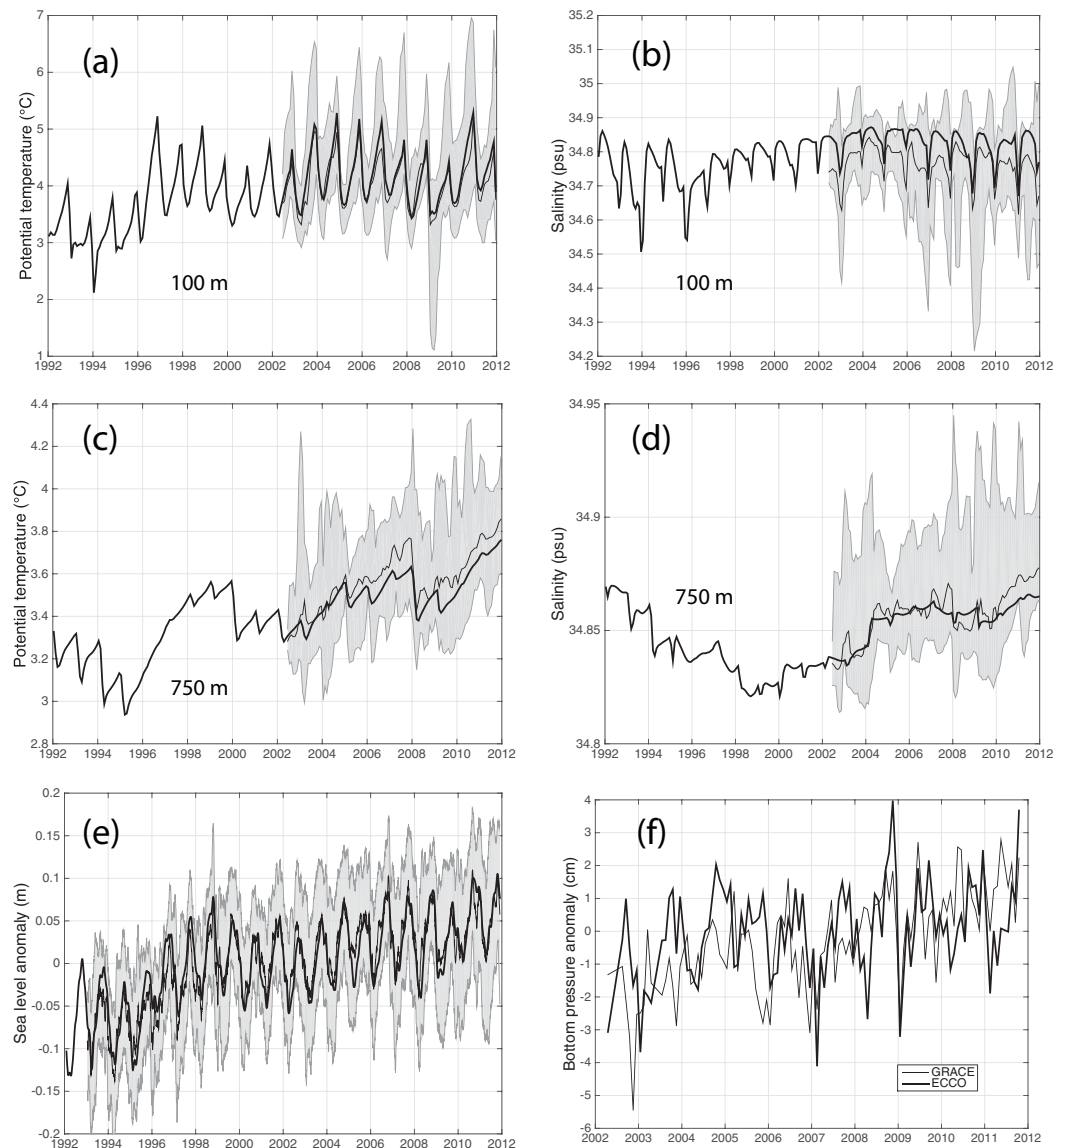


Figure 3. Validation of ECCOV4-r2 with observational data in the Labrador Sea (Figure 1). Comparison of ECCOV4-r2 and Argo (a and c) temperature profiles and (b and d) salinity profiles. Mean ECCOV4-r2 values are shown as thick black lines. Argo-ECCOV4 misfits are calculated as $m_i = (a_i - e_i) + \bar{e}_i$, where a_i is the Argo value, e_i is the corresponding ECCOV4-r2 value, and \bar{e}_i is the mean ECCOV4 value. Median values of m_i are shown as thin black lines, and the shading indicates the 95% interval for m_i , i.e., between the 2.5th and 97.5th quantiles. Comparisons are shown at 100 and 750 m. (e) Comparison of ECCOV4 (thick black line) sea level anomaly with Topex-Poseidon-Jason family of altimeters (mean is thin black line, shading shows 95% misfit interval, Forget & Ponte, 2015). (f) Comparison of ECCOV4 (thick black line) bottom pressure with GRACE/mascons data (thin black line), downloaded from http://grace.jpl.nasa.gov/data/get-data/jpl_global_mascons/.

that the linear approximation is suitable—the response of spatially and temporally averaged objective functions tends to be more linear than that of more localized and/or instantaneous quantities. The suitability of the linear approximation is confirmed in Appendix A.

For our adjoint sensitivity study, we use the average heat content over a control volume V and time interval Δt as our objective function:

$$J = \frac{1}{V\Delta t} \int_V \int_{\Delta t} H(\mathbf{r}, t) dt dV, \quad (1)$$

where $H = \rho_0 c_p \theta(\mathbf{r}, t)$ is the heat content, $\theta(\mathbf{r}, t)$ is the potential temperature, \mathbf{r} is the position vector, and t is time. The reference density is set as $\rho_0 = 1,027 \text{ kg/m}^3$ and the heat capacity is $c_p = 3,850 \text{ J/(kg K)}$. The

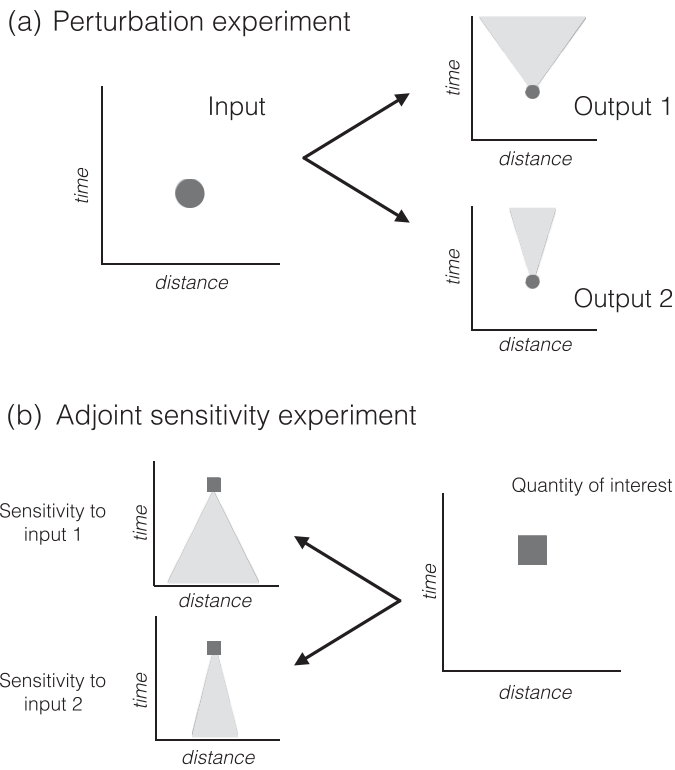


Figure 4. Schematic of (a) a traditional forward perturbation experiment and (b) an adjoint sensitivity experiment. The output of the forward perturbation experiment is a set of perturbed fields (Δy), whereas the output of an adjoint sensitivity experiment is a collection of gradients (i.e., sensitivities of the form $\partial J/\partial x$, where J is the objective function and x is an input variable).

pretation, we include maps of the standard deviation fields for surface forcing in the online supplemental information.

For the purpose of plotting three-dimensional sensitivity fields (e.g., $\partial J/\partial T$), it is sensible to scale the sensitivity fields by the thickness of the depth level Δz (Heimbach et al., 2011). The scaled sensitivity takes the form:

$$\frac{1}{\Delta z} \frac{\partial J}{\partial T}, \quad (4)$$

which has units of $1/m$. This scaling prevents the relatively large grid boxes in the deep ocean interior from dominating the sensitivity. In each section, we explicitly describe the type of scaling applied for each type of analysis. Since ECCOV4 does not feature an adjoint representation of the sea ice model, sensitivities to air-sea fluxes are corrected by a factor of $1-f$, where f is the fractional coverage of sea ice area in each model grid cell. For instance, the sensitivity of Labrador Sea heat content to air-sea heat fluxes in a completely ice-covered grid cell ($f = 1$) is set to zero ($1-f$).

2.3. Kinematic and Dynamic Sensitivities

In this work, we decompose some of the adjoint sensitivity fields into kinematic and dynamic components following Marotzke et al. (1999). This allows us to distinguish between sensitivities to changes that propagate along isopycnals (i.e., kinematic) with sensitivities to changing density structures (i.e., dynamic). Formulating the annual-mean and volume-mean heat content as a function of density and temperature $J = J[\rho(T, S), T]$ allows us to write the sensitivity of the heat content to temperature variations at constant salinity as follows:

$$\left(\frac{\partial J}{\partial T}\right)_S = \left(\frac{\partial J}{\partial \rho}\right)_T \left(\frac{\partial \rho}{\partial T}\right)_S + \left(\frac{\partial J}{\partial T}\right)_\rho. \quad (5)$$

averaging volume V covers the entire Labrador Sea water column, delineated by the 300 m March–April–May mixed layer depth contour (averaged over 1992–2011) in the Labrador Sea. The time integral covers a 1 year period from 1 January to 31 December. We analyze a 10 member ensemble of 11 years adjoint sensitivity runs, with the objective function covering the last year of the run, specifically from 2002 to 2011. The ensemble approach allows us to describe the sensitivity fields in terms of ensemble means and standard deviations about the mean that reflect interannual variability over 2002–2011.

Our adjoint model calculates the linear sensitivities of the objective function J to a set of independent variables x . For a selected independent variable x , an adjoint model calculates a set of time-evolving sensitivity fields:

$$\frac{\partial J}{\partial x}(\mathbf{r}, t). \quad (2)$$

The objective function J is a scalar, but the sensitivity field $\partial_x J$ may have rich spatial and temporal structure. Throughout this work, we use 14 day averaged sensitivity fields for analysis. Adjoint sensitivity fields can be scaled in various ways depending on the question at hand (Heimbach et al., 2011; Verdy et al., 2014). One choice is to scale by a value of the standard deviation. For an independent variable x , we compute

$$dJ_x(\mathbf{r}, t) = \left[\frac{\partial J}{\partial x}(\mathbf{r}, t) \right] \sigma_x(\mathbf{r}), \quad (3)$$

where $\sigma_x(\mathbf{r})$ is the spatially varying standard deviation in time (relative to 14 day averages) after the seasonal cycle has been removed. This choice means that we are using an interannual standard deviation together with an annual-mean objective function. To help with inter-

The first term on the right-hand side of equation (5) is the “dynamic” component of the sensitivity (i.e., sensitivity to changes in density), and the second term on the right-hand side is the “kinematic” component (i.e., dynamically inactive sensitivities to temperature anomalies). Using the coefficient of thermal expansion α and coefficient of haline contraction β , defined as:

$$\alpha \equiv -\frac{1}{\rho} \left(\frac{\partial \rho}{\partial T} \right)_S \text{ and } \beta \equiv \frac{1}{\rho} \left(\frac{\partial \rho}{\partial S} \right)_T, \quad (6)$$

we can write

$$\left(\frac{\partial J}{\partial S} \right)_T = \left(\frac{\partial J}{\partial \rho} \right)_T \left(\frac{\partial \rho}{\partial S} \right)_T = \beta \rho \left(\frac{\partial J}{\partial \rho} \right)_T, \quad (7)$$

and the dynamic sensitivity becomes:

$$F_{dyn} = \left(\frac{\partial J}{\partial \rho} \right)_T \left(\frac{\partial \rho}{\partial T} \right)_S = \frac{1}{\beta \rho} \left(\frac{\partial J}{\partial S} \right)_T \left(\frac{\partial \rho}{\partial T} \right)_S = -\frac{\alpha}{\beta} \left(\frac{\partial J}{\partial S} \right)_T. \quad (8)$$

The kinematic sensitivity can also be written as a function of sensitivities to temperatures and salinities,

$$F_{kin} = \left(\frac{\partial J}{\partial T} \right)_S + \frac{\alpha}{\beta} \left(\frac{\partial J}{\partial S} \right)_T. \quad (9)$$

In these forms, the dynamic and kinematic sensitivities can be calculated directly from standard MITgcm adjoint model output, which includes sensitivities to potential temperature and salinity throughout the entire model run. We use monthly 1992–2011 averaged, three-dimensional α/β fields derived from ECCOV4-r2 potential temperatures and salinities using the TEOS-10 toolbox (McDougall & Barker, 2011).

Sensitivity fields (e.g., F_{kin} , F_{dyn} , more generally written $\partial J/\partial x$) can be converted into impacts ΔJ by multiplying by perturbations $\Delta x = \sigma_x(\mathbf{r})$, i.e., $\Delta J = (\partial J/\partial x)\Delta x$. Physically, applying a unit increase of $\Delta T = 1^\circ\text{C}$ to a dynamic sensitivity field F_{dyn} can be interpreted as instead imposing a density-equivalent decrease in salinity ($\Delta S = -\Delta T\alpha/\beta$) due to the presence of the factor $-\alpha/\beta$ in equation (8). Here the phrase “density-equivalent” refers to the fact that if the condition $\alpha\Delta T = -\beta\Delta S$ is satisfied, then the small perturbations ΔT and ΔS have the same impact on the density via the linear equation of state for seawater, i.e., $\rho = \rho_0(1 - \alpha\Delta T + \beta\Delta S)$. In contrast, applying a perturbation of $\Delta T = 1^\circ\text{C}$ to a kinematic sensitivity field F_{kin} can be interpreted as simultaneously imposing both a $\Delta T = 1^\circ\text{C}$ change in potential temperature and a change in salinity given by $\Delta S = \Delta T\alpha/\beta$ (see equation (9)). The combination of these changes ensures that the density remains constant, i.e., the perturbation is carried out along a density surface in T/S space.

3. Results: Adjoint Pathways and Processes

We examine sensitivity fields from a 10 member ensemble of 11 years adjoint sensitivity experiments, with one experiment for each objective function year in the 2002–2011 range, in order to quantify the sensitivity of the Labrador Sea heat content to local and remote influences. We decompose the sensitivity fields into kinematic and dynamic components as described in section 2.3.

3.1. Sensitivity to Changes at Constant Density

Positive kinematic sensitivities indicate potential “source regions” for a given control volume of interest (e.g., the Labrador Sea) by quantifying the extent to which potential temperature anomalies may directly get transported into the region of interest at constant density. Any selected region of the global ocean integrates influences from increasingly remote regions as we consider increasingly distant times in the past. Thus, the volume covered by nonzero values of kinematic sensitivity tends to increase with longer lags, reflecting the action of adjoint advection, diffusion, and mixing at constant density (Figure 5).

For short lags (–0.8 year in Figure 5, right-hand column), the sensitivities are concentrated in the Labrador Sea and the wider subpolar gyre, with varying lateral influences at different depths. In the upper 500 m, sensitivity signals propagate along the eastern coast of Greenland via the East Greenland Current and the Denmark Strait Overflow, the cold and fresh currents underneath that connect the Irminger Sea and the Nordic Seas via the Denmark Strait. Below 500 m, the sensitivities are confined to the Irminger Sea and the Iceland

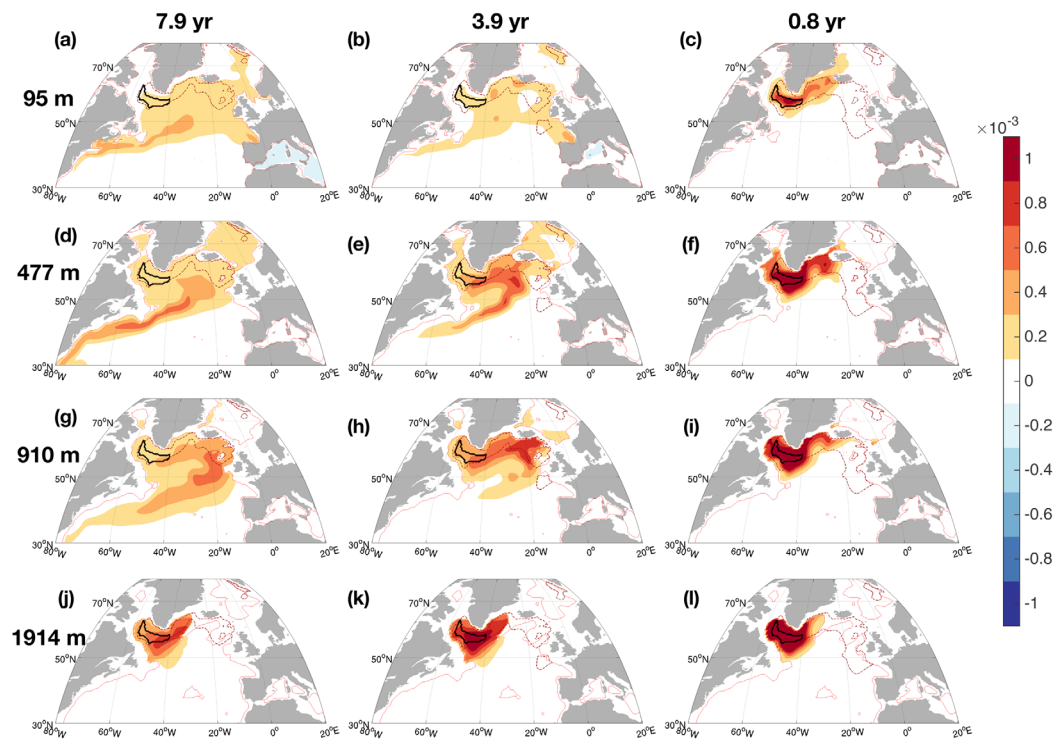


Figure 5. Ensemble mean kinematic sensitivities for the annual-mean and column-mean heat content in the Labrador Sea, shown at three different lags (-7.9 years, -3.9 years, and -0.8 year) and at four different depth levels. The objective function is defined as an average over the entire LS water column, and these plots show cuts of the sensitivity fields at 95, 477, 910, and 1,914 m. The fields are scaled as $\partial_{\theta_n} J / (J_0 \Delta z)$, where J is the annual-mean Labrador Sea heat content, θ_n is the potential temperature, J_0 is the ensemble mean annual heat content $7.9 \times 10^6 \text{ J/m}^3$, and Δz is the thickness of the vertical level. The fields are scaled such that in a region with sensitivity $1 \times 10^{-3} [\text{m}^\circ\text{C}]^{-1}$, a unit perturbation of $\Delta T = 1^\circ\text{C}$, together with the simultaneous salinity perturbation $\Delta S = \Delta T \alpha / \beta$ required to keep the density constant, applied over a 14 day period in a single grid cell with 1 m thickness will induce a linear perturbation in the annual-mean heat content of roughly $1.3 \times 10^{-10} J_0$.

Basin, as the shallow bathymetry of the Denmark Strait and the ridge to the east of Iceland prevent exchange with the Nordic Seas. At lag -3.9 years (Figure 5, middle column), the influence of the subpolar gyre is apparent from the surface down to roughly 1,000 m. The imprint of the North Atlantic Current is especially visible at 477 m. At lag -7.9 years, we find sensitivities in the Gulf Stream concentrated in the upper 900 m. At 477 m, we see the broadest sensitivity pattern, with nonzero values stretching from the Gulf of Mexico to the Nordic Seas and continuing into the Arctic. In contrast, for all lags considered, the deep ocean sensitivities remain largely confined to the Labrador Sea and Irminger Sea, highlighting the vastly different circulation time scales and pathways found in the upper, intermediate, and deep zones of the North Atlantic.

Although the kinematic sensitivity field is positive nearly everywhere, we find small negative sensitivities in the near-surface Mediterranean Sea (see Figures 5a and 5b), which is a region of anomalously high salinity relative to the North Atlantic. The predicted linear sensitivity of LS heat content to an increase in Mediterranean Sea temperature, together with the simultaneous decrease in salinity required to keep the density constant, is a decrease in LS heat content. Although this potential adjustment pathway is interesting, we do not investigate it further here.

3.2. Sensitivity to Changes in Density

A change in buoyancy in a region of the global ocean can potentially influence Labrador Sea heat content via re-arrangements in hydrostatic pressure fields and geostrophy, even if that region is not a “source” of water for the Labrador Sea. Changes in temperature can thus influence the dynamics of the ocean in various ways, for example, by changing the tilt of density surfaces and associated geostrophic transports, and/or by exciting barotropic and baroclinic motions with characteristics similar to Kelvin waves and Rossby

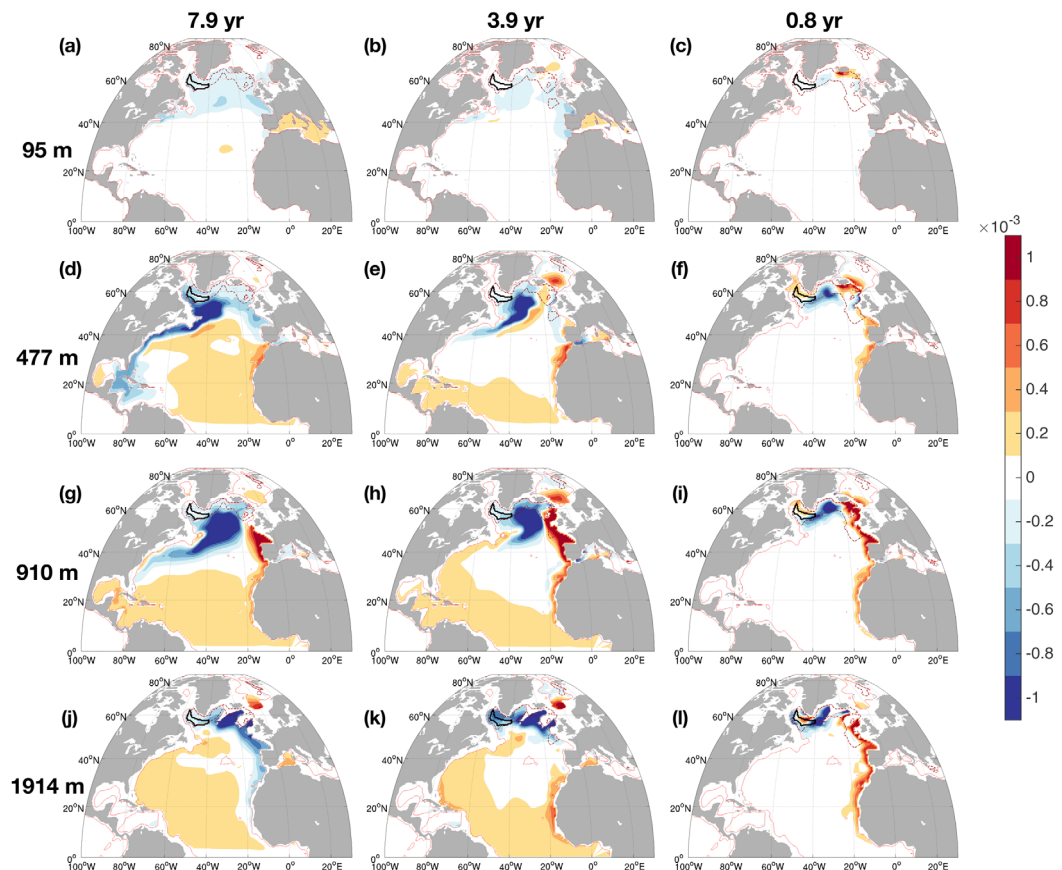


Figure 6. Ensemble mean dynamic sensitivities for the annual-mean and column-mean heat content in the Labrador Sea, shown at three different lags (-7.9 years, -3.9 years, and -0.8 year) and at four different depth levels. The fields are scaled in the same fashion as Figure 5.

waves modified by the presence of bottom topography. These mechanisms can potentially affect heat convergence and thereby heat content in the LS. Like the kinematic fields, the dynamic sensitivity fields are four-dimensional (three spatial dimensions, one time dimension) and thus contain a tremendous amount of information.

At short lags (-0.8 year, Figure 6, right column) below 100 m, we see a positive sensitivity anomaly along the entire eastern boundary of the subtropical and subpolar North Atlantic, extending from the coast/shelf of West Africa to the coast/shelf of Iceland. For lags longer than about 5 years, the sensitivity field becomes increasingly baroclinic, with variations between positive and negative values with depth (Figure 6, left column). This coastal/shelf sensitivity field reflects a complex superposition of mechanisms that can potentially act to change the basin-wide meridional pressure gradient, thereby altering the associated circulation and ultimately heat convergence in the LS. Even though the eastern subtropical Atlantic is not a source water region for the Labrador Sea, it can influence the Labrador Sea dynamically. The kinematic sensitivities all along the African and most of the European shelf are negligibly small compared with the sensitivities in the subpolar gyre, NAC, and Gulf Stream, but the dynamic sensitivities are relatively large.

Positive-negative dipoles in the dynamic sensitivity fields can indicate locations where changes in temperature can alter stratification, the tilt of density surfaces, and the associated transport (Figure 6). The dipoles seen at 477 m and 910 m across all lags tend to straddle the regions of maximum kinematic sensitivities, both of which are broadly oriented along large-scale circulation features (e.g., the eastern edge of the subpolar gyre, the NAC, the Gulf Stream). Increasing potential temperature in the region of positive sensitivity and/or decreasing potential temperature in the region of negative sensitivity leads to an *increase* in Labrador Sea heat content by changing the transport and convergence of heat. The linear response of the heat content is the product of the sensitivity and an anomaly, i.e., $\Delta J = (\partial_x J) \Delta x$, so to understand the sign of the

response we must consider both the sign of the sensitivity and the sign of the anomaly in the independent variable x . This is broadly consistent with a transport-driven mechanism identified by Williams et al. (2015) in which an increase in Labrador Sea density *enhances* overturning and produces stronger heat convergence in the subpolar gyre. Animations of the dynamic sensitivity and ensemble standard deviations are available as supplemental information.

Kinematic and dynamic changes may partially offset one another. For instance, if a perturbation in potential temperature is applied just north of the core of the NAC kinematic sensitivities, at a depth of 477 m and a lag of roughly 8 years (see Figures 5 and 6), then the kinematic and dynamic effects would partially cancel each other. In a region of positive kinematic sensitivity, an increase in potential temperature will ultimately get advected into the LS and increase its heat content. However, in a region of negative dynamic sensitivity, an increase in potential temperature will induce a change in density that ultimately decreases LS heat content. The ratio of the kinematic and dynamic responses would depend on how the perturbation is applied, but nevertheless, the presence of opposing kinematic and dynamic sensitivities highlight the presence of potentially complex adjustment mechanisms.

3.3. Sensitivity to Changes in Different Regions

Labrador Sea heat content is influenced by increasingly remote regions as we consider more negative lags (i.e., as we look further back in time). By dividing the North Atlantic and Arctic Oceans into different regions based on geographic and dynamic considerations, we can quantify the time scales over which these regions can contribute to variability in the Labrador Sea. We use nine analysis regions covering the North Atlantic and Arctic (Figure 7a). Regions 1, 2, and 3 are the Labrador Sea, Irminger Sea, and broader subpolar gyre (which does not include regions 1 and 2), respectively. Region 4 includes Hudson Bay, Baffin Bay, and part of the Northwest Passage. Region 5 consists of the Nordic Seas, with a southern boundary delineated by relatively shallow bathymetric features, and Region 6 is the Arctic Ocean, which is only partially shown in the chosen map projection. The subtropical gyre is divided into three regions based approximately on the structure of the barotropic stream function (Figure 1). Region 7 contains the Gulf Stream, the Caribbean Sea, and the Gulf of Mexico, with an eastern boundary that coincides with the maximum eastward extent of the 30 Sv contour of the subpolar gyre. Region 8 is the central subtropics, with a maximum eastward boundary that coincides with the 5 Sv contour of the barotropic stream function. Region 9 includes the Eastern Subtropics and the Mediterranean Sea, so it will be affected by the along-shelf propagating wave signals discussed in previous sections. The boundary between region 3 and the subtropical regions (7, 8, and 9) is the 0 Sv contour of the barotropic stream function.

In each region, the sensitivity of the objective function J to a perturbation Δx at time t is generally written:

$$R_{pos}(t) = \sum_{i,j,k} \left(\frac{\partial J}{\partial x} \right)_{i,j,k,n} \sigma_x(\mathbf{r}), \quad (10)$$

where the sum is over grid cell indices within the selected region, and $\sigma_x(\mathbf{r})$ is the three-dimensional, time-independent, deseasonalized standard deviation in potential temperature. For convenience, we will refer to this as the “pos” response, which is the sensitivity to a positive perturbation with spatially varying magnitude $\sigma_x(\mathbf{r})$. Here we use 2 week averaged sensitivities, so the scaled sensitivity indicates the change in annual-mean Labrador Sea heat content brought about by the linear sensitivity to a change in the 14 day averaged ocean state, which in this instance is taken to be a perturbation in potential temperature ΔT (and/or its density-equivalent perturbation in salinity $\Delta S = \Delta T(\alpha/\beta)$) (Figure 7).

We start by analyzing the regional kinematic sensitivities. Local kinematic sensitivities (i.e., sensitivities to region 1) can be described by two-term exponential decay with a fast decay rate of 1.0 ± 0.1 years and a slow decay rate of 12 ± 0.6 years (ensemble mean and ensemble standard deviation). Peak sensitivity to the Irminger Sea has some spread across the model ensemble, with the maximum occurring at lag -2.1 ± 0.4 years. Maximum sensitivity to the subpolar gyre occurs at lag -1.0 ± 0.8 years; for longer lags it decreases roughly linearly at a rate of $1.2 \pm 0.2\%/yr$. The Nordic Seas sensitivity peaks at -4.7 ± 0.7 years. The contribution of the Hudson remains negligible, probably due to its small size and its relatively inaccessible geography. On the short, 10 year time scale of these experiments, the Arctic Ocean makes a small contribution to the response, but for lags longer than -1.8 years, the kinematic sensitivity increases roughly linearly at a rate of $0.2 \pm 0.02\%/yr$.

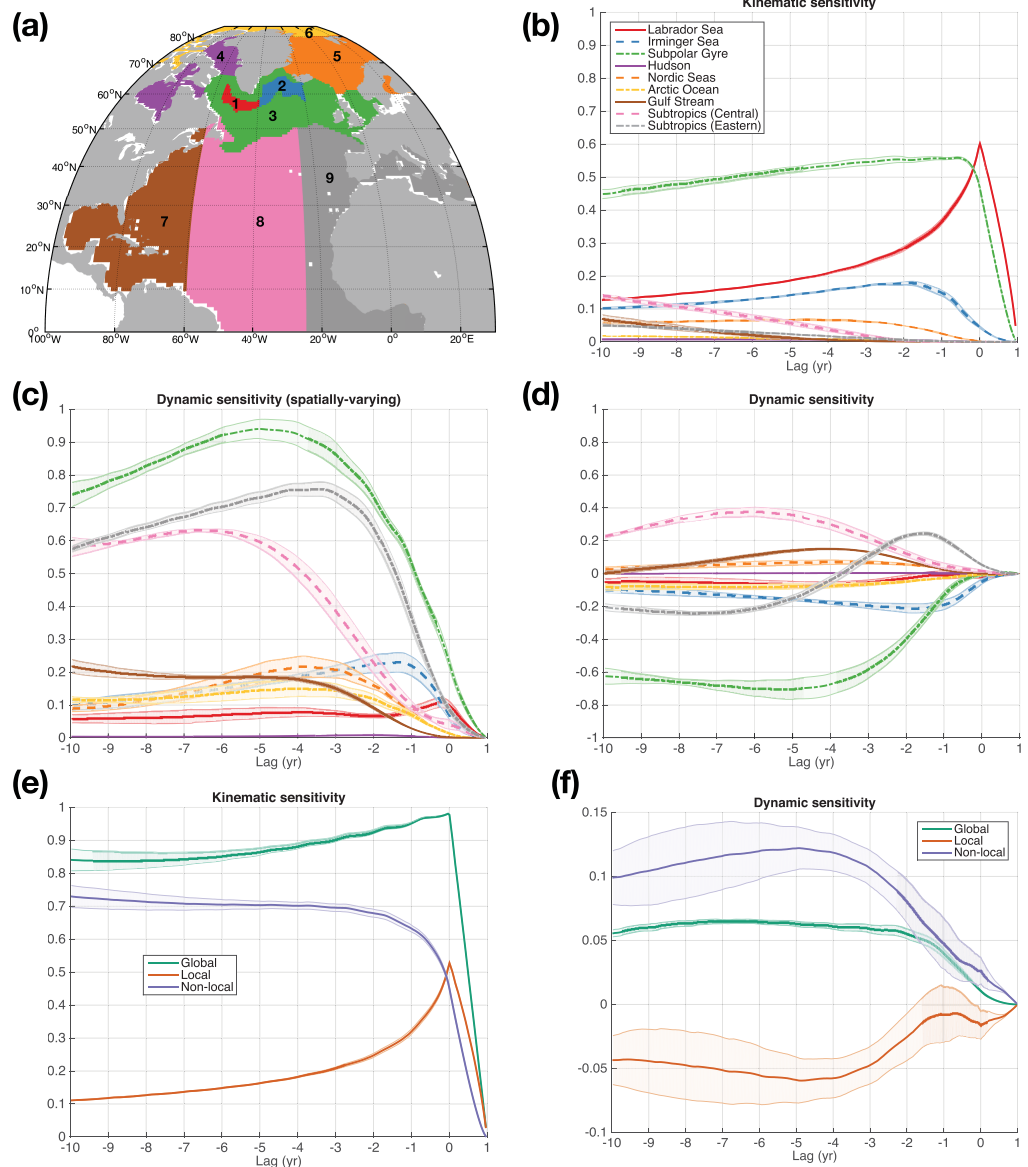


Figure 7. Regional, scaled sensitivity time series for the areas indicated in Figure 7a. The lines indicate ensemble means, and the shading indicates the standard deviation across ensemble members. Shown are the scaled sensitivities for (b) the kinematic sensitivity (R_{pos}), (c) the dynamic sensitivity (R_{var}), and (d) the dynamic sensitivity (R_{pos}). All the time series have been scaled by the same constant and so are directly comparable. Scaled sensitivities are total impacts of a given region, i.e., they are not scaled by the size of each region. The objective function is a year-long integral starting at lag 0. Also shown are simplified time series plots with local (i.e., in the Labrador Sea), nonlocal, and global (local plus nonlocal) scaled sensitivities R_{pos} for the (e) kinematic sensitivity and (f) dynamic sensitivity. Note that plot (f) features a different vertical scale than the other plots.

The Gulf Stream region (region 7, which also includes the Gulf of Mexico) is not a major source region for the LS on 10 year time scales. In terms of the kinematic sensitivity, it reaches a relative value of 1% at lag -5.3 ± 0.4 years. The small value of the relative contribution may be an artifact of the choice of region decomposition, but the time at which this maximum is reached is not sensitive to the value of the maximum. The kinematic sensitivity fields do show the imprint of the Gulf Stream at lag -7.9 years at a depth of 477 m, although by this time the core of the sensitivity has not yet reached the Gulf Stream itself, as it is still located in the range of the NAC (Figure 5). The central subtropics shows zero sensitivity for lags shorter than -2 years, and going further back in time it increases at a rate of $1.6 \pm 0.1\%/yr$. The Eastern Subtropics is not a source region for the Labrador Sea, with sensitivities well

below 5% for the entire 10 year experiment. For lags longer than -2 years, the sensitivity increases at a rate of $0.4 \pm 0.03\%/yr$.

The dynamic sensitivity time series show a very different picture. We display the sensitivity in two different fashions. We preserve the sign of the dynamic sensitivity, which can be either positive or negative; the scaled sensitivity is calculated as shown in equation (10), so the sign of the scaled sensitivity associated with each grid cell comes from the sign of the sensitivity (Figure 7d). As the scaled sensitivities are summed up in space, this approach may lead to cancelations within a region that contains positive and negative responses. Alternatively, we can sum up the absolute value of the dynamic sensitivity, i.e.:

$$R_{var}(t) = \sum_{i,j,k} \left| \frac{\partial J}{\partial X} \sigma_x(\mathbf{r}) \right|_{i,j,k}. \quad (11)$$

Conceptually, this is equivalent to performing a convolution between the sensitivity field and an anomaly field wherein the anomalies have the same sign as the sensitivities. As this is extremely unlikely to be realized in any particular evolution of the ocean state, one should consider the spatially varying sensitivity an upper bound (the largest possible impact, in terms of the positive/negative structure of the response).

The Irminger Sea displays a negative scaled sensitivity for all lags considered, i.e., an increase in temperature here would dynamically decrease the Labrador Sea heat content (Figure 7). The minimum R_{pos} occurs at lag -1.7 ± 0.3 years, and the maximum R_{var} occurs at lag -1.6 ± 0.5 years. For the broader subpolar gyre, the extremum R_{pos} occurs at lag -6.0 ± 2.0 years. The scaled sensitivity indicates a relatively strong dynamic sensitivity to the state of the subpolar gyre in 1992 and 1993, which were years of exceptionally strong mixed layer depth and subpolar gyre circulation within the 20 year ECCOv4 period. The Central Subtropics scaled sensitivity peaks at lag -6.1 ± 0.3 years (R_{pos}) and lag -6.8 ± 1.2 years (R_{var}).

The Nordic Seas maximum dynamic scaled sensitivity occurs at lag -3.6 ± 0.5 years (R_{pos}) and -3.7 ± 0.3 years (R_{var}). The Arctic Ocean has only a weak scaled sensitivity, peaking at lag -7.1 ± 1.3 years (R_{pos}) and -5.2 ± 2.5 years (R_{var}). As discussed above, the Eastern Subtropics impact the sensitivity via dynamics, although it is not a strong source region for the Labrador Sea. The R_{pos} peak occurs at lag -1.5 ± 0.1 years, whereas the R_{var} peak occurs at lag -3.6 ± 0.4 years. This contrast indicates that there are cancelations that may occur when the dynamic sensitivity is forced uniformly. The time scales should be interpreted with care, as they represent the combined effect of many different processes.

Here we present simplified time series showing the scaled sensitivity R_{pos} to local (in the Labrador Sea), non-local (everywhere except the Labrador Sea), and global perturbations (the sum of local and nonlocal) (Figures 7e and 7f). The nonlocal kinematic sensitivity exceeds the local kinematic sensitivity for lags longer than about 1 month, but as discussed above, the local sensitivity decays somewhat slowly with lag, remaining above 10% of the maximum global value for all lags considered. The global kinematic sensitivity also decays with lag, described empirically on the interval $[-10,0]$ by two-term exponential decay with time scales $\tau_1 = 8.4$ years and $\tau_2 = 22$ years. For dynamic sensitivities (Figure 7), the scaled sensitivity to nonlocal density anomalies is always positive and larger in magnitude than the negative scaled sensitivity to local density anomalies, thus the global scaled sensitivity is always positive.

3.4. Sensitivity to Surface Forcing

The heat content of the Labrador Sea can be affected by local and remote surface fluxes, such as zonal and meridional wind stress and net heat flux. Here we examine the 14 day mean sensitivity fields associated with these processes at the sea surface (Figure 8). Since our numerical model is an ocean-only model with imposed atmospheric forcing, sensitivities are relative to the imposed surface forcing, as opposed to a dynamic air-sea coupling. Ensemble standard deviations and animations of the sensitivity fields are available as supporting information.

3.4.1. Net Heat Flux

By convention, a positive heat flux *decreases* ocean temperature, i.e., ocean heat *loss* is positive. Large, negative sensitivities in the Labrador Sea at short lags thus indicate, as one would expect, that local heat gain increases heat content at short lags. At 3.9 year lag, the largest negative sensitivities are found south of Iceland and in the Nordic Seas (Figure 8). Anomalies in this region can get advected via the subpolar gyre into the LS. There is also a region of positive sensitivity along the European continental shelf. At 7.9 year lag, the

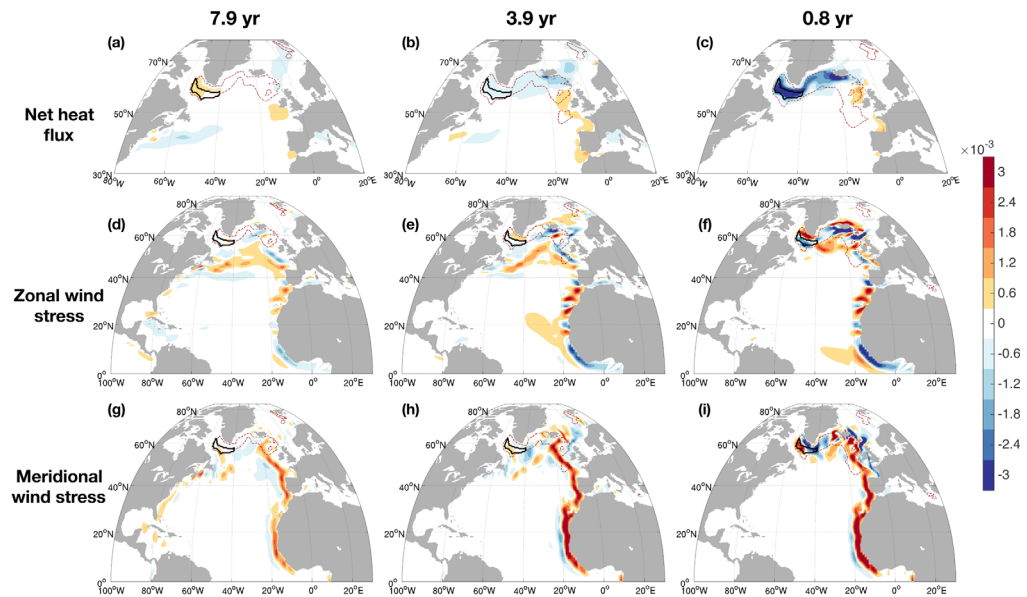


Figure 8. Ensemble mean sensitivities of the annual-mean and column-mean Labrador Sea heat content for objective function years 2002–2011, shown at three different lags (–7.9 years, –3.9 years, and –0.8 year). The fields have been scaled as $J_0^{-1}(\partial_x J)\Delta x$, where x is the independent variable and J_0 is the scaling constant $\rho_0 c_p (2.0^\circ\text{C})$. The result is a dimensionless measure of the sensitivity of LS heat content to a positive perturbation Δx applied at one grid point for 2 weeks, with $\Delta Q_{net} = 60 \text{ W/m}^2$ and $\Delta\tau_E = \Delta\tau_N = 0.06 \text{ N/m}^2$.

sensitivity of Labrador Sea heat content to local heat flux has changed sign to positive values, indicating that the linear, time-delayed scaled sensitivity to a local increase in heat loss is in fact an *increase* in Labrador Sea heat content. This counterintuitive result is broadly consistent with a mechanism identified by Williams et al. (2015), in which increasing the density of the Labrador Sea (e.g., through increased heat loss) accelerates the overturning and increases heat convergence in the subpolar gyre. However, these positive sensitivity values are much smaller than negative sensitivity at lag 0.

For a two-dimensional surface forcing field like net heat flux, the scaled sensitivity metric R_{pos} takes the form:

$$R_{pos}(t) = \sum_{ij} \left(\frac{\partial J}{\partial x} \right)_{ij,n} \sigma_x(\mathbf{r}), \quad (12)$$

where in this case $\sigma_x(\mathbf{r})$ is the two-dimensional, time-independent, deseasonalized standard deviation in net heat flux. The scaled sensitivity varies with region and time scale, and an annual cycle is present in each time series, with extrema in late winter to early spring (see Figure 9).

To compare the timing of the scaled sensitivity extrema and mixed layer depth, we construct a mean seasonal cycle for the monthly mean scaled sensitivity R_{pos} and mixed layer depth and calculate various lag correlations between the seasonal cycles. In each region considered, the monthly mean scaled sensitivity leads the monthly mean mixed layer depth by about 1 month, so forcing anomalies that occur roughly 1 month before maximum mixed layer depth tend to produce the largest linear scaled sensitivities in annual-mean LS heat content. At this lag (–1 month), correlations between scaled sensitivities and mixed layer depth are very high, explaining over 80% of the variance independently of the region.

Considering the full time series again, we see that Labrador Sea heat content is most sensitive to heat fluxes during the year over which the objective function is defined (Figure 9a). The maximum magnitude scaled sensitivity occurs at 2.2 ± 0.8 months (positive lag), which is between February and April in the year over which the objective function is calculated. Strong vertical mixing over this period enables heat flux anomalies to mix over the largest possible fraction of the water column, thereby increasing the storage of heat in the relatively quiescent deep interior Labrador Sea. The heat content is still sensitive to heat fluxes from the previous 3–4 years, highlighting the importance of preconditioning from previous years in encouraging

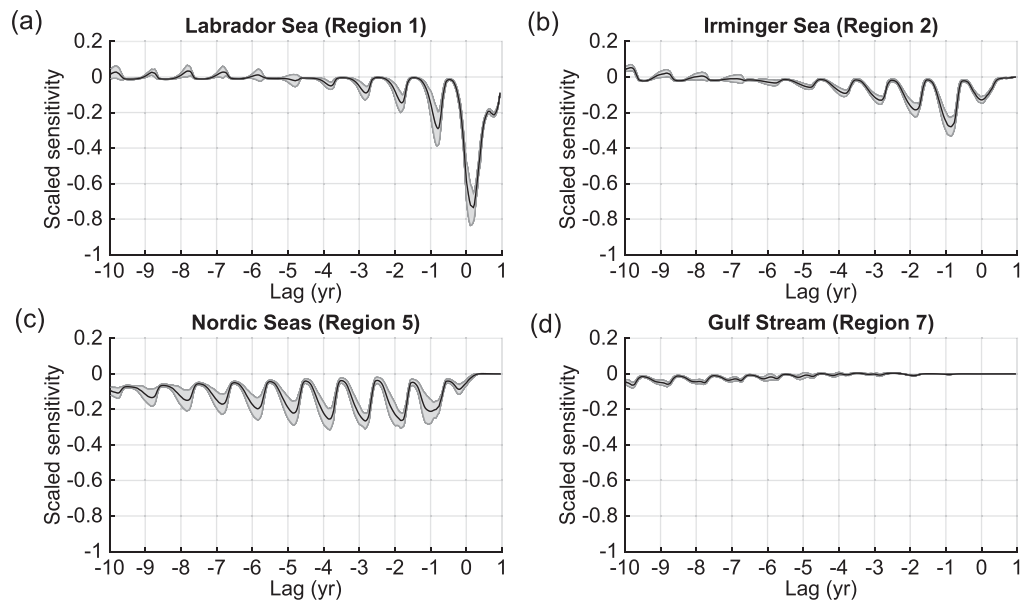


Figure 9. Scaled sensitivity R_{pos} of the Labrador Sea heat content to a uniformly signed net heat flux perturbation in the past for the (a) Labrador Sea, (b) Irminger Sea, (c) Nordic Seas, and (d) the Gulf Stream. The objective function is the annual-mean and column-mean Labrador Sea heat content for objective function years in the 2002–2011 range. The thick lines indicate ensemble means, and the shading indicates one standard deviation across the ensemble members. The fields are shown on the same relative scale. We use the convention that positive net heat fluxes *decrease* ocean surface potential temperatures, so a negative scaled sensitivity indicates LS heat loss/gain due to increased/decreased (more positive/negative) air-sea heat flux.

deep convection. After roughly 5–7 years, the local sensitivity switches sign, but it has a much smaller magnitude than the sensitivity to the target year (i.e., the year on the lag interval [0,1]).

The most negative sensitivities to the Irminger sea heat flux occur around lag -10 ± 0.8 months, which is roughly the previous February–March (Figure 9b). R_{pos} for the Irminger Sea is larger than R_{pos} for the Labrador Sea for lags longer than about 10 months. Sensitivities to fluxes in the broader subpolar gyre (not shown) are nonzero for nearly the entire 10 year integration period, with decreasing effect each previous year. Sensitivities to heat fluxes in the Nordic Seas have their greatest magnitudes between lags -4 and -2 years, although there is considerable spread in the ensemble in winter (Figure 9c). Sensitivities to fluxes in the Gulf Stream region display a complex, double-peaked annual cycle, although for more negative lags a clearer seasonal signal emerges. By lag -10 years, the Gulf Stream region scaled sensitivities reach roughly 10% of the local, target year (i.e., between lag 0 and 1) scaled sensitivity to heat fluxes in the Labrador Sea (Figure 9d). These results are broadly consistent with oscillating adjoint sensitivity patterns of the AMOC to changes in the Labrador Sea region, as reported by Czeschel et al. (2010). Sensitivities to other regions are small (not shown).

3.4.2. Wind Stress

Both the zonal and meridional wind stress sensitivity patterns feature numerous positive/negative sensitivity dipoles (Figure 8). For wind stress, these dipoles indicate regions where a change in wind position and/or wind stress curl can induce changes in transport via Ekman pumping/suction. The meridional sensitivity fields feature strong, coastally guided, somewhat stationary signals along the eastern edge of the Atlantic basin. Considering the meridional and zonal sensitivity fields together, we see that the sensitivity pattern roughly aligns with the local coastline and shelf bathymetry, suggesting that alongshore winds are important for the scaled sensitivity. Although the eastern Atlantic basin is not a strong source region for the Labrador Sea, changes in these locations can alter dynamics and heat/salt convergence. This region of positive sensitivity extends from the west coast of North Africa all the way up to the seas south of Iceland. The adjoint sensitivity fields suggest that if this region is forced by an increase in northward wind stress, the associated enhanced coastal downwelling will ultimately induce an *increase* in LS heat and salt content (a positive anomaly acting on a positive sensitivity region will increase the objective function) and vice versa.

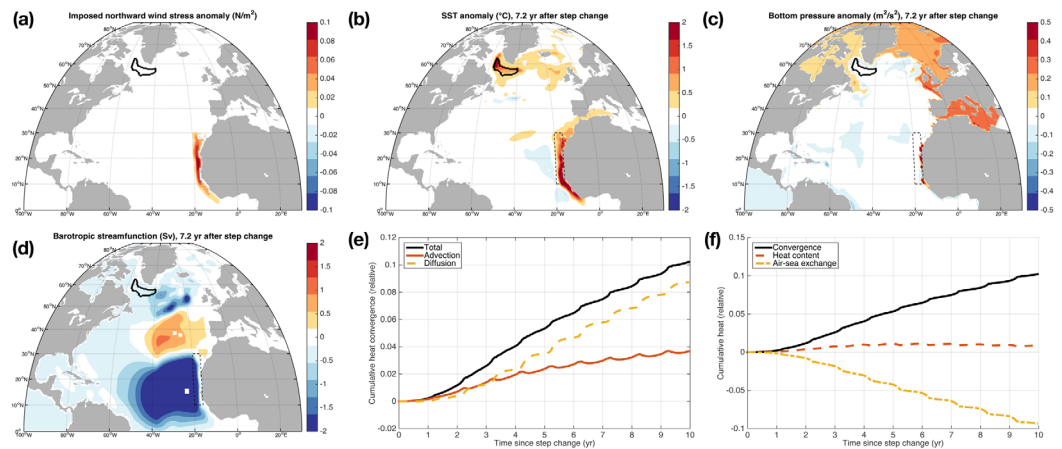


Figure 10. Results of a northward wind stress step response experiment. (a) Spatial pattern of the imposed change in northward wind stress. Anomalies relative to a control run, 7.2 years after the step change is imposed, are shown for (b) sea surface temperature, (c) bottom pressure, and (d) barotropic stream function (negative values indicate counterclockwise rotation). The Labrador Sea region is indicated by a thick black line, and the approximate region of the wind stress perturbation is shown by a thin dashed line. (e) Time series of cumulative heat convergence relative to the control run, split into advective and diffusive flux convergence components. (f) Time series of cumulative heat convergence, cumulative heat exchange with the atmosphere, and heat storage relative to the control run. Time series (e) and (f) are scaled by the decadal-mean Labrador Sea heat content, 7.9×10^6 J.

In order to test the hypothesis that an increase in wind stress along the West African shelf will eventually increase the LS heat content, as suggested by the adjoint model, we perform a 10 year step response experiment using the ECCOv4 setup. After a 10 year spin up under control conditions, we impose a permanent step change in wind stress along the coast with a sign structure that matches the sensitivity field and a maximum magnitude of 0.1 N/m^2 (Figure 10a). The change in wind stress along the West African shelf induces a change in Ekman transport across the bathymetry that enhances downwelling of warm surface waters along the coast, creating an across-bathymetry pressure anomaly.

The direct effect of the change in wind stress is largely local, i.e., the warming signal detected in the vicinity of the wind stress perturbation is not connected to the Labrador Sea, as West Africa is not a source region for the Labrador Sea on the time scales considered here (Figure 10b). It is instead the across-bathymetry pressure anomaly, which excites a combination of barotropic and baroclinic motions, that ultimately induces a change in Labrador Sea heat content. An initial, rapid bottom pressure anomaly roughly follows f/H contours along the Atlantic side of the Greenland-Iceland-Scotland ridge, reaching the Labrador sea in less than two weeks. In the following 2–3 months, the pressure anomaly makes its way over the Greenland-Iceland-Scotland Ridge, spreading rapidly across the Nordic Seas and the broader Arctic Ocean. The pressure change propagates southward through the Denmark Strait, setting up an across-bathymetry pressure gradient anomaly along the entire northern boundary of the Atlantic Ocean (Figure 10c). This basin-wide, across-bathymetry pressure anomaly adjusts for 2–3 years after the step change. The change in basin-scale pressure gradient across the North Atlantic speeds up the circulation of the subpolar gyre (Figure 10d). Diffusive heat convergence and advective heat convergence into the LS both increase as the gyre spins up (Figure 10e). The net increase in LS heat convergence is strongly offset by an increase in ocean heat loss to the atmosphere, which is likely to encourage convection into the deep ocean and the resulting increase in LS heat storage (Figure 10f).

The response of the Labrador Sea heat content to this imposed change in wind stress is well approximated by the linear approach used in the adjoint model. We verified this by examining results from four different step response experiments, with maximum values of ± 0.1 and $\pm 0.4 \text{ N/m}^2$. For the 0.1 N/m^2 step response, the nonlinear component of the response remains small (less than 5% of the maximum linear response for the duration of the model run). For the 0.4 N/m^2 step response experiment, the nonlinear component is larger (less than 20% of the maximum linear response). Thus, it appears the linear approximation works well for modest wind stress perturbations, but it starts to break down for large values of wind stress, as one may expect.

3.4.3. Relative Importance of Heat Flux and Wind Stress

In order to summarize the complex spatiotemporal information contained in the adjoint sensitivity fields, we use two different formulations of the scaled sensitivity following Verdy et al. (2014). The scaled sensitivity of the Labrador Sea heat content with respect to each of the three surface forcing fields is:

$$R_{mean,F}(t) = \left\langle \left| \frac{\partial J}{\partial Q_{net}} \sigma_{Q_{net}} \right| \right\rangle + \left\langle \left| \frac{\partial J}{\partial \tau_E} \sigma_{\tau_E} \right| \right\rangle + \left\langle \left| \frac{\partial J}{\partial \tau_N} \sigma_{\tau_N} \right| \right\rangle, \quad (13)$$

where $\sigma_x = \sigma_x(\mathbf{r})$, the sensitivity fields are functions of space and time, and the angular brackets represent sums of the impacts $(\partial_x J) \sigma_x(\mathbf{r})$ over chosen areas. Each term in equation (13) represents the impact of one particular surface forcing variable, either the net heat flux Q_{net} , eastward wind stress τ_E , or northward wind stress τ_N . In this metric, positive and negative impacts may offset each other in the spatial sum. For example, suppose that net heat flux Q_{net} becomes more positive everywhere in the selected ocean basin. Locations and times with positive sensitivities $\partial_Q J > 0$ contribute to an *increase* in J , whereas locations and times with negative sensitivities $\partial_Q J < 0$ contribute to a *decrease* in J . The metric $R_{mean,F}$ may also be interpreted as the impact of basin-scale changes in forcing (Verdy et al., 2014).

The metric $R_{mean,F}$ represents one extreme along a spectrum of possible responses. The other extreme is the very unlikely case in which the sign of the perturbations exactly match the signs of the sensitivity field, such that the impacts are always positive:

$$R_{var,F}(t) = \left\langle \left| \frac{\partial J}{\partial Q_{net}} \sigma_{Q_{net}} \right| \right\rangle + \left\langle \left| \frac{\partial J}{\partial \tau_E} \sigma_{\tau_E} \right| \right\rangle + \left\langle \left| \frac{\partial J}{\partial \tau_N} \sigma_{\tau_N} \right| \right\rangle. \quad (14)$$

In this metric, there are no cancelations of differently signed impacts. Locations and times with positive sensitivities contribute to an *increase* in J , and locations and times with negative sensitivities *also* contribute to a *increase* in J . In order for J to respond in this way, the imposed perturbation must have some spatial structure on scales smaller than basin-scale. Note that equation (14) is a variant of equation (11), in that equation (14) uses standard deviations for the anomalies and includes multiple terms.

Considered together, the two components of wind stress make a much larger relative contribution to cumulative $R_{var,F}$ (89%) than to $R_{mean,F}$ (49%), highlighting the importance of spatial structure in the wind-driven sensitivity of LS heat content (Table 1). Spatially varying wind forcing that matches the sign structure of the sensitivity fields drives a much larger heat content response than a basin-wide change in wind forcing. This is consistent with the large number of dipoles present in the adjoint sensitivity fields. Under a change in basin-scale forcing (measured by $R_{mean,F}$), the scaled sensitivity to a dipole (with equal magnitudes) is zero, whereas under a change in forcing that matches the sign structure (measured by $R_{var,F}$), the response from a dipole is additive. Although the exact partitioning of the scaled sensitivity between zonal wind stress and meridional wind stress is a result of the decomposition of the wind stress vector into zonal and meridional components, the total scaled sensitivity from the wind stress is independent of the rotation of the wind stress vector. The ensemble standard deviations for cumulative $R_{mean,F}$ and $R_{var,F}$ are all less than 10%, so by this measure the sensitivity fields are fairly stationary for years in the range 2002–2011.

3.4.4. Local Versus Nonlocal Sensitivity to Surface Forcing

Here we quantify the local and nonlocal contributions of three surface forcing fields, as well as kinematic and dynamic sensitivities, to the scaled sensitivity R_{pos} (Figure 11). The largest sensitivity is to net heat flux, particularly to local perturbations in the target year (i.e., on the lag interval [0,1] year). In the target year, the scaled sensitivity to local forcing is larger than the scaled sensitivity to nonlocal forcing, but this situation quickly reverses for negative lags. The responses to zonal and meridional wind stress display a complex range of responses and time scales, including a strong seasonal cycle and a slow, multiyear adjustment that reflects the sensitivity of the circulation field to aspects of the wind stress (e.g., gyres responding to wind stress curl).

The cumulative sensitivity (summing responses from lag +1 year to more negative lags) can be used to quantify when cumulative nonlocal effects exceed cumulative local effects. This local-to-remote transition time scale t_{LN} offers a simple measure of the relative responses of

Table 1
Cumulative $R_{mean,F}$ and $R_{var,F}$ for Each Variable, Summed Over the Entire 11 Year Duration of the Adjoint Sensitivity Experiments

Variable	Cumulative $R_{mean,F}$	Cumulative $R_{var,F}$
Net heat flux	51% ± 3%	12% ± 0.3%
Zonal wind stress	25% ± 6%	47% ± 1%
Meridional wind stress	24% ± 2%	41% ± 1%

Note. Values are displayed as ensemble means and ensemble standard deviations for each variable, scaled by the total $R_{mean,F}$ and $R_{var,F}$ including all three variables.

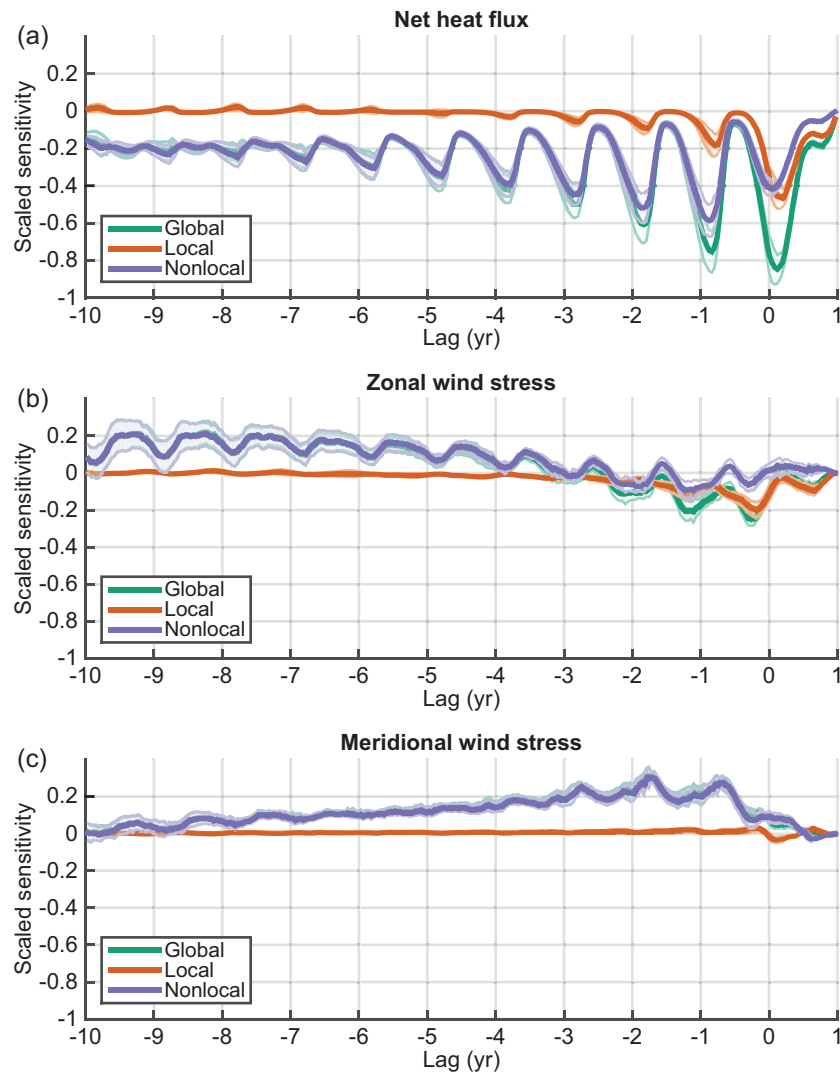


Figure 11. Scaled sensitivities R_{pos} of the Labrador Sea heat content to local and nonlocal surface forcing, scaled by the maximum magnitude sensitivity to surface forcing. “Local” is defined as within the Labrador Sea region (Figure 1), and “nonlocal” is the rest of the global ocean. The sum of the two is denoted by the sensitivity to the “global” forcing. To calculate the scaled sensitivity, the sensitivities are multiplied by a spatially varying standard deviation $\sigma_x(r)$ as described in the text and plotted in the supporting information. The lines indicate ensemble means across 2002–2012, and the shading indicates one standard deviation across ensemble members. Results are shown relative to the maximum value of the ensemble mean sensitivity to net heat flux.

the LS heat content to local and nonlocal forcing. We estimate the transition time scale by using cumulative sums of both R_{mean} and R_{var} and report the result as an ordered pair (R_{mean}, R_{var}) . For net heat flux, the transition time scale is $(-0.69, -0.60)$ year, which is 7.2–8.3 months before the start of the objective function integral at lag 0. For lags that are more negative than $(-0.69, -0.60)$ year, the cumulative sensitivity to nonlocal changes in net heat flux exceeds the cumulative sensitivity to local changes in net heat flux. There is a sharp difference between the cumulative sensitivity to zonal wind stress and the cumulative sensitivity to meridional wind stress. The transition time scale for meridional (northward) wind stress is short and positive, (0.9 year, 1.0 year), whereas the transition time scale for the zonal (eastward) wind stress spans a much larger range $(-4.8 \text{ years}, 1.0 \text{ year})$. The nonlocal effect of the meridional wind stress is rapid, dominated by the across-shelf pressure gradient adjustment mechanism discussed in section 3.4.2. But the zonal wind stress sensitivity contains many positive-negative dipoles that partially cancel each other when measured by R_{mean} .

4. Conclusions

Using a realistic, observationally constrained ocean model in adjoint mode (Forget et al., 2015a), we examined the sensitivity of the column-averaged, annual-mean heat content of the Labrador Sea to (1) changes in potential temperature at constant density, (2) changes in density, and (3) changes in net heat fluxes and wind stresses on 10 year time scales. We presented key aspects of these complex, temporally varying and spatially varying sensitivity fields and examined some of the adjustment mechanisms highlighted by the sensitivity fields. By decomposing the sensitivity fields into kinematic and dynamic components, we tracked potential source waters for the Labrador Sea and identified both local and remote regions in which density changes can alter circulation and ultimately change Labrador Sea heat convergence.

Positive kinematic sensitivity fields indicate pathways along which potential temperature changes can affect LS heat content for a fixed circulation pattern. In this way, calculating positive kinematic sensitivities is conceptually similar to performing “reverse passive tracer experiments” in which a tracer is allowed to propagate backward in time following a fixed pattern of circulation and mixing. In this sense, the kinematic sensitivity fields can also be thought of as highlighting the “source waters” of the Labrador Sea (Marotzke et al., 1999; Song et al., 2016). Our source water calculations indicate that potential Labrador Sea source regions include the broader subpolar gyre, the Nordic Seas, the North Atlantic Current, and the Gulf Stream, although the structure of the sensitivity patterns changes considerably with depth and time scale. The difference in the areal extent of the sensitivity fields reflects differences in circulation, e.g., the influence of perturbations spreads more rapidly in the upper 500 m of the North Atlantic than at 2,000 m.

By contrast, dynamic sensitivities indicate the linear perturbations that will result in the largest possible changes in LS heat content via changes in density, the associated wave field, and circulation. In the upper 100 m, we find mostly negative sensitivities in the subpolar gyre, indicating that an increase in upper ocean temperature can *reduce* the depth-averaged heat content by decreasing surface density. In the interior ocean, we find negative-positive dipoles in dynamic sensitivity that are coincident with regions of high kinematic sensitivity, indicating an underlying sensitivity to changes in the across-streamline tilt of density surfaces and the associated geostrophic transport. For example, cooling the Labrador Sea will ultimately *increase* LS heat content via a change in Gulf Stream heat transport and LS heat convergence. This is consistent with a heat convergence adjustment mechanism identified in historical temperature and salinity data as well as in idealized numerical experiments (Klöwer et al., 2014; Williams et al., 2015).

We also find relatively large dynamic sensitivities along the coast/shelf system of West Africa and Western Europe. This region of dynamic sensitivity is *not* a source region for the LS, i.e., kinematic sensitivities in this region are negligibly small. A similar pattern is also found in the sensitivity to meridional wind stress, indicating an adjustment mode related to changes in pressure. Perturbations in near-coastal, along-bathymetry wind stress induce cross-shelf pressure gradients by Ekman transport, and the resulting pressure anomalies propagate northward along the shelf. This mechanism eventually alters the pressure on the shelf all along the North Atlantic and into the Arctic Ocean, resulting in a change in subpolar gyre circulation and an associated increase/decrease in Labrador Sea heat convergence (Bell, 2011, and references therein). A similar adjustment pathway has been documented for Arctic Ocean bottom pressure, albeit for much faster barotropic Kelvin waves (Fukumori et al., 2015). Also, Pillar et al. (2016) found that the Atlantic Meridional Overturning Circulation is sensitive to meridional wind stress over the West African shelf, in consistency with our suggestion that wind stress perturbations in this region can ultimately affect the basin-scale pressure gradient over the entire North Atlantic, thereby altering large-scale gyre circulation and transport.

In terms of surface forcing, LS heat content is most sensitive to local (in space and time) heat fluxes, though other nonlocal locations/lags make significant contributions to the scaled sensitivity, highlighting the importance of preconditioning and advection of upstream temperature anomalies. Wind stress sensitivity patterns largely reinforce the pressure wave adjustment mechanism discussed above, as they feature significant positive alongshore sensitivities.

Here we summarize some of the dominant adjoint adjustment pathways revealed by the sensitivity fields (Figure 12). On short (less than 1 year) time scales, Labrador Sea heat content is most sensitive to perturbations in the Labrador Sea, Irminger Sea, the Greenland coast/shelf, and the eastern boundary of the Atlantic Ocean via pressure gradient adjustments (pathways A and B, Figure 12). On longer time scales, the LS becomes most sensitive to perturbations in the NAC and the Nordic Seas (pathways C and D). On the

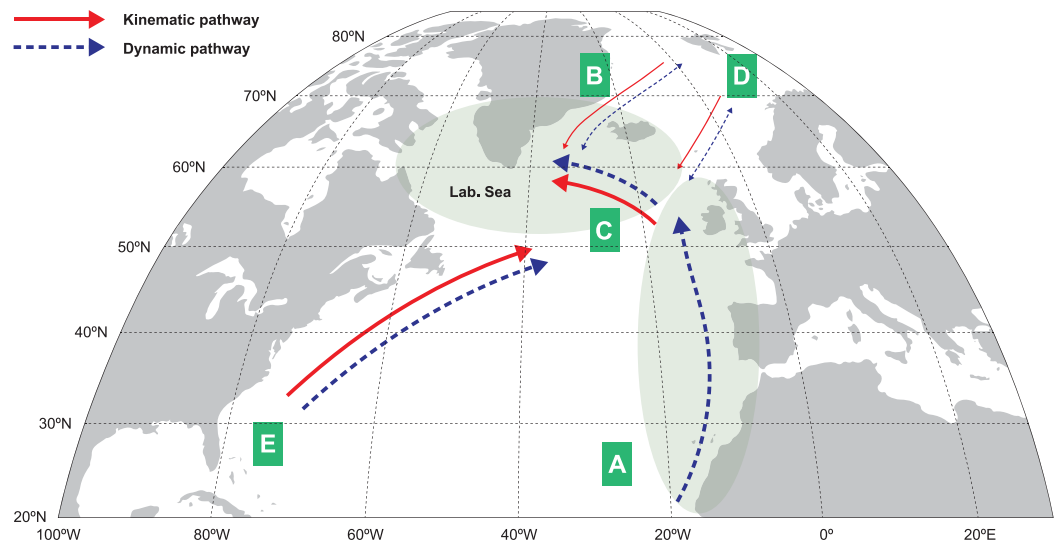


Figure 12. Schematic of major kinematic (solid red arrows) and dynamic (dashed green arrows) adjustment pathways for annual-mean Labrador Sea heat content. The pathways include (A) adjustment along the eastern boundary of the North Atlantic Ocean (dynamic only), which can affect the basin-wide pressure gradient and the associated circulation, (B) the coastal circulation of the East Greenland Current, (C) the circulation of the subpolar gyre, (D) exchanges with the Nordic Seas over the Greenland-Iceland-Scotland Ridge, and (E) circulation of the Gulf Stream and North Atlantic Current. The light green shading indicates regions that can affect LS heat content on time scales shorter than roughly 1 year. Changes in the unshaded regions will take longer than 1 year to affect LS heat content.

longest time scales considered in this study (5–10 years), we find increasingly large sensitivities in the Gulf Stream region, mainly in the top 500 m (pathway E). Although Figure 12 is a simplified representation, it provides a clear conceptual framework for understanding the adjustment pathways of LS heat content.

We have shown that adjoint sensitivity fields can be used to highlight and quantify potential adjustment pathways for heat content in a region of deep convection. We also examined the relative impact of net heat flux and wind stress on LS heat content. These sensitivity estimates can be used to inform future nonlinear forward perturbation experiments in both ocean-only and coupled models, which allow for a more thorough investigation of the mechanisms involved in each sensitivity pathway. In addition, the adjoint sensitivity fields presented here may also be used to inform the design of future observational networks (Heimbach et al., 2011). For instance, LS heat content is sensitive to net wintertime air-sea heat fluxes in the Irminger Sea and Nordic Seas over 10 year time scales, so long-term monitoring of fluxes and hydrography in these regions is needed to understand and predict the behavior of the Labrador Sea. Monitoring of wind stress along the West African and European shelf may also be important for projecting LS behavior, as it has an impact on the basin-scale pressure gradient of the entire North Atlantic Ocean. Our results highlight the numerous processes that control the climatically important heat content and the associated heat uptake in a critical region of the North Atlantic Ocean.

Appendix A: Examining the Validity of the Linear Approximation

Here we test the accuracy of the linear approximation for “typical” perturbation sizes using the forward, nonlinear ECCOV4-r2 model setup. To this end, we separate the linear and nonlinear responses of a given quantity by imposing positive and negative perturbations of the same magnitude in two different model runs (Verdy et al., 2014). Given a perturbation $\Delta Q = Q - Q_0$, in a quantity Q , then the response of a variable H can be approximated by Taylor series expansions as:

$$\Delta H = H - H_0 = \frac{\partial H}{\partial Q} (Q - Q_0) + \frac{1}{2} \frac{\partial^2 H}{\partial Q^2} (Q - Q_0)^2 + \dots, \quad (\text{A1})$$

where H_0 and Q_0 are reference values about which the partial derivatives are evaluated. We denote the response to a positive perturbation $Q > Q_0$ as ΔH_+ and the response to a negative perturbation $Q < Q_0$ as ΔH_- . We then estimate the linear response by the difference $(\Delta H_+ - \Delta H_-)/2 \approx (\partial_Q H)(Q - Q_0)$ and the nonlinear response by the sum $(\Delta H_+ + \Delta H_-)/2 \approx 0.5(\partial_{QQ} H)(Q - Q_0)^2$. This approach is expected to work well if

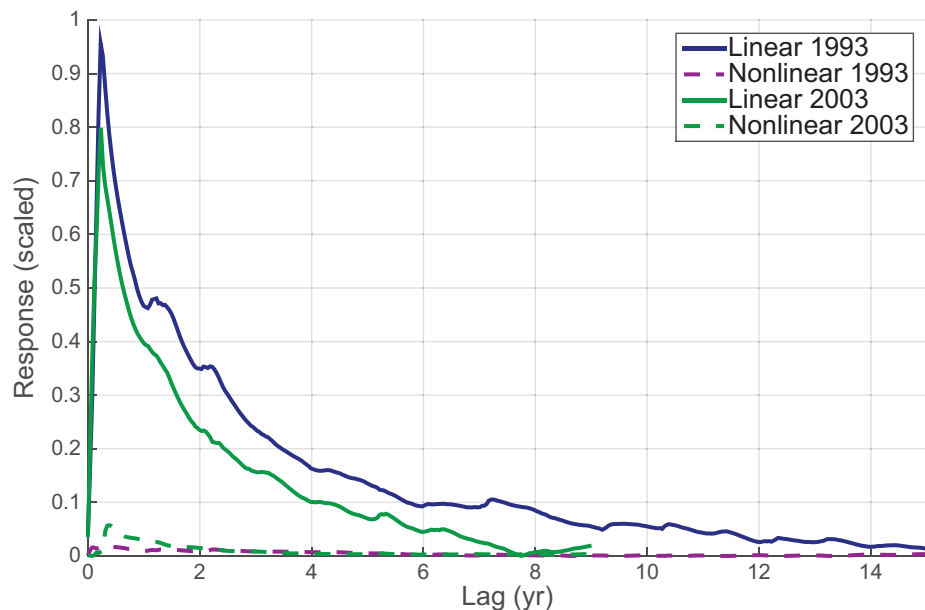


Figure A1. Normalized linear (solid lines) and nonlinear (dashed lines) responses of the depth-integrated potential temperature of the Labrador Sea to perturbations in net heat flux. The perturbations are applied over the entire Labrador Sea with magnitude 10 W/m^2 (not shown) and 40 W/m^2 (shown) in both 1993 (blue) and 2003 (green). When scaled by the magnitude of the heat flux perturbations, responses to the 10 and 40 W/m^2 are very nearly identical.

the response function in question can be well represented by a Taylor series expansion and if the first two nonconstant terms capture the majority of the variability of that response function.

We impose positive and negative perturbations of magnitude 10 and 40 W/m^2 over the Labrador Sea for the first 3 months (JFM) of both 1993 and 2003, a total of four different perturbation experiments. We chose these years because they represent end members for the stratification of the background state and deep convection, as 1993 features exceptionally weak stratification and deep mixing, while 2003 features relatively strong stratification and a relatively shallow winter mixed layer. Using 1993 and 2003 also allowed us to run the perturbation experiments for at least 9 years and exploit almost all of the ECCOv4 period (1993–2011). When scaled by the magnitude of the perturbations (i.e., 10 and 40 W/m^2), we find that the linear component of the response behaves nearly identically for both perturbation magnitudes, suggesting a high degree of linearity with respect to the magnitude of the perturbation, at least in the 10 – 40 W/m^2 range (Figure A1). In all cases, the nonlinear component of the response is small, becoming significant only when the total response itself becomes negligible. For the perturbations applied in 1993, the nonlinear response is small for at least 10–11 years, and for the perturbations applied in 2003 the linear response dominates for about 7 years, after which time the total response is small. Based on these results, we conclude that the linear approximation is suitably accurate on time scales of roughly 7 years for the problem of response to local air-sea heat flux. Notably, the responses show significant differences when comparing 1993 and 2003, suggesting that the ocean/climate background state does affect the sensitivity of the column-integrated heat content to net heat flux at the surface. In 1993, the potential temperature anomaly created by the change in heat flux penetrates much further into the interior ocean (down to roughly $2,000 \text{ m}$) due to deep convection. In 2003, the potential temperature anomaly induced by the perturbation stays confined to a relatively narrow depth range (roughly between 0 and 800 m). This contrast in mixed layer depth is consistent with the behavior of the observed ocean, in that heat loss of similar magnitudes can still lead to dramatically different mixed layers, highlighting the importance of preconditioning and stratification for deep mixing (Piron et al., 2017).

Acronyms

AMOC	Atlantic Meridional Overturning Circulation
DWBC	Deep Western Boundary Current

ECCOV4-r2	Estimating the Circulation and Climate of the Ocean (version 4, release 2)
LS	Labrador Sea
LSW	Labrador Sea Water
MITgcm	Massachusetts Institute of Technology general circulation model
NAC	North Atlantic Current
NADW	North Atlantic Deep Water
MAM	March–April–May time period
TAF	Transformation of Algorithms in Fortran (by FastOpt GmbH)

Acknowledgments

This study is supported by grants from the Natural Environment Research Council (NERC), including (1) The North Atlantic Climate System Integrated Study (ACSIS) (grant NE/N018028/1 (authors DJ, ES) and NE/N018044/1 (authors BS, SJ)), (2) Securing Multidisciplinary Understanding and Prediction of Hiatus and Surge events (SMURPHS) (grant NE/N006038/1 (author EB) and NE/N005686/1 (authors BS, SJ)), and (3) Ocean Regulation of Climate by Heat and Carbon Sequestration and Transports (ORCHESTRA) (grant NE/N018095/1 (authors EB, AM)). GF is supported by NASA award 1553749 and Simons Foundation award 549931. The ECCOV4-r2 model setup used in this work is available for download on Github (https://github.com/gaelforget/ECCO_v4_r2) as an instance of the MIT general circulation model (MITgcm, <http://mitgcm.org/>). Numerical model runs were carried out on ARCHER, the UK national HPC facility (<http://archer.ac.uk/>). Adjoint code was generated using the TAF software tool, created and maintained by FastOpt GmbH (<http://www.fastopt.com/>). Argo float data are available via <http://www.argo.ucsd.edu/>. The authors wish to thank Ric Williams, Chris Wilson, Yavor Kostov, Maïke Sonnewald, Mike Bell, Florian Sevellec, Nora Loose, and two anonymous reviewers for conversations that greatly improved the quality of this paper.

References

- Bell, M. J. (2011). Ocean circulations driven by meridional density gradients. *Geophysical and Astrophysical Fluid Dynamics*, *105*(2–3), 182–212. <https://doi.org/10.1080/03091929.2010.534468>
- Chen, X., & Tung, K.-K. (2014). Varying planetary heat sink led to global-warming slowdown and acceleration. *Science*, *345*(6199), 897–903. <https://doi.org/10.1126/science.1254937>
- Czeschel, L., Marshall, D. P., & Johnson, H. (2010). Oscillatory sensitivity of Atlantic overturning to high-latitude forcing. *Geophysical Research Letters*, *37*, L10601. <https://doi.org/10.1029/2010GL043177>
- Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., et al. (2011). The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society*, *137*(656), 553–597. <https://doi.org/10.1002/qj.828>
- Drijfhout, S. S., Blaker, A. T., Josey, S. A., Nurser, A. J. G., Sinha, B., & Balsaseda, M. A. (2014). Surface warming hiatus caused by increased heat uptake across multiple ocean basins. *Geophysical Research Letters*, *41*, 7868–7874. <https://doi.org/10.1002/2014GL061456>
- Fekete, B. M., Vörösmarty, C. J., & Grabs, W. (2002). High-resolution fields of global runoff combining observed river discharge and simulated water balances. *Global Biogeochemical Cycles*, *16*(3), 1042. <https://doi.org/10.1029/1999GB001254>
- Forget, G., Campin, J. M., Heimbach, P., Hill, C. N., Ponte, R. M., & Wunsch, C. (2015a). ECCO version 4: An integrated framework for non-linear inverse modeling and global ocean state estimation. *Geoscientific Model Development*, *8*(10), 3071–3104. <https://doi.org/10.5194/gmd-8-3071-2015>
- Forget, G., Ferreira, D., & Liang, X. (2015b). On the observability of turbulent transport rates by Argo: Supporting evidence from an inversion experiment. *Ocean Science*, *11*(5), 839–853. <https://doi.org/10.5194/os-11-839-2015>
- Forget, G., & Ponte, R. M. (2015). The partition of regional sea level variability. *Progress in Oceanography*, *137*, 173–195. <https://doi.org/10.1016/j.pocean.2015.06.002>
- Fukumori, I., Menemenlis, D., & Lee, T. (2007). A near-uniform basin-wide sea level fluctuation of the Mediterranean Sea. *Journal of Physical Oceanography*, *37*(2), 338–358. <https://doi.org/10.1175/JPO3016.1>
- Fukumori, I., Wang, O., Llovel, W., Fenty, I., & Forget, G. (2015). A near-uniform fluctuation of ocean bottom pressure and sea level across the deep ocean basins of the Arctic Ocean and the Nordic Seas. *Progress in Oceanography*, *134*, 152–172. <https://doi.org/10.1016/j.pocean.2015.01.013>
- Gaspar, P., Grégoris, Y., & Lefevre, J. M. (1990). A simple eddy kinetic energy model for simulations of the oceanic vertical mixing: Tests at station Papa and long-term upper ocean study site. *Journal of Geophysical Research*, *95*(C9), 16179–16193. <https://doi.org/10.1029/JC095iC09p16179>
- Gent, P. R., & McWilliams, J. C. (1990). Isopycnal mixing in ocean circulation models. *Journal of Physical Oceanography*, *20*(1), 150–155. [https://doi.org/10.1175/1520-0485\(1990\)020<0150:MIOCM>2.0.CO;2](https://doi.org/10.1175/1520-0485(1990)020<0150:MIOCM>2.0.CO;2)
- Giering, R., & Kaminski, T. (1998). Recipes for adjoint code construction. *ACM Transactions on Mathematical Software*, *24*(4), 437–474. <https://doi.org/10.1145/293686.293695>
- Griewank, A., & Walther, A. (2012). *Evaluating derivatives: Principles and techniques of algorithmic differentiation* (2nd ed.). Cambridge, UK: Society for Industrial and Applied Mathematics. <https://doi.org/10.1137/1.9780898717761>
- Heimbach, P. (2008). The MITgcm/ECCO adjoint modelling infrastructure. *CLIVAR Exchanges*, *13*(1), 13–17.
- Heimbach, P., Wunsch, C., Ponte, R. M., Forget, G., Hill, C., & Utke, J. (2011). Timescales and regions of the sensitivity of Atlantic meridional volume and heat transport toward observing system design. *Deep Sea Research, Part II*, *58*(17–18), 1858–1879. <https://doi.org/10.1016/j.dsr2.2010.10.065>
- Klöwer, M., Latif, M., Ding, H., Greatbatch, R. J., & Park, W. (2014). Atlantic meridional overturning circulation and the prediction of North Atlantic Sea surface temperature. *Earth and Planetary Science Letters*, *406*, 1–6. <https://doi.org/10.1016/j.epsl.2014.09.001>
- Large, W., & Yeager, S. (2009). The global climatology of an interannually varying air-sea flux data set. *Climate Dynamics*, *33*, 341–364. <https://doi.org/10.1007/s00382-008-0441-3>
- Lazier, J., Hendry, R., Clarke, A., Yashayaev, I., & Rhines, P. (2002). Convection and restratification in the Labrador Sea, 1990–2000. *Deep Sea Research, Part I*, *49*(10), 1819–1835. [https://doi.org/10.1016/S0967-0637\(02\)00064-X](https://doi.org/10.1016/S0967-0637(02)00064-X)
- Liu, Z., & Alexander, M. (2007). Atmospheric bridge, oceanic tunnel, and global climatic teleconnections. *Reviews of Geophysics*, *45*, RG2005. <https://doi.org/10.1029/2005RG000172>
- Losch, M., Menemenlis, D., Campin, J.-M., Heimbach, P., & Hill, C. (2010). On the formulation of sea-ice models. Part 1: Effects of different solver implementations and parameterizations. *Ocean Modelling*, *33*(1–2), 129–144. <https://doi.org/10.1016/j.ocemod.2009.12.008>
- Lozier, M. S., Bacon, S., Bower, A. S., Cunningham, S. A., de Jong, M. F., de Steur, L., et al. (2017). Overturning in the Subpolar North Atlantic Program: A new international ocean observing system. *Bulletin of the American Meteorological Society*, *98*(4), 737–752. <https://doi.org/10.1175/bams-d-16-0057.1>
- Marotzke, J., Giering, R., Zhang, K. Q., Stammer, D., Hill, C., & Lee, T. (1999). Construction of the adjoint MIT ocean general circulation model and application to Atlantic heat transport sensitivity. *Journal of Geophysical Research: Atmospheres*, *104*(C12), 29529–29547. <https://doi.org/10.1029/1999JC900236>
- Mazloff, M., Heimbach, P., & Wunsch, C. (2010). An eddy-permitting Southern Ocean state estimate. *Journal of Physical Oceanography*, *40*, 880–899.

- McDougall, T. J., & Barker, P. M. (2011). *Getting started with TEOS-10 and the Gibbs Seawater (GSW) Oceanographic Toolbox* (SCOR/IAPSO WG127, 28 pp.). SCOR/IAPSO. Retrieved from http://www.teos-10.org/pubs/Getting_Started.pdf
- Ortega, P. (2017). Mechanisms of decadal variability in the Labrador Sea and the wider North Atlantic in a high-resolution climate model. *Climate Dynamics*, *49*(7), 2625–2647. <https://doi.org/10.1007/s00382-016-3467-y>
- Pérez, F. F., Mercier, H., Vázquez-Rodríguez, M., Lherminier, P., Velo, A., Pardo, P. C., et al. (2013). Atlantic Ocean CO₂ uptake reduced by weakening of the meridional overturning circulation. *Nature Geoscience*, *6*(2), 146–152. <https://doi.org/10.1038/ngeo1680>
- Pickart, R. S., Straneo, F., & Moore, G. W. K. (2003). Is Labrador Sea Water formed in the Irminger basin? *Deep Sea Research, Part I*, *50*(1), 23–52. [https://doi.org/10.1016/S0967-0637\(02\)00134-6](https://doi.org/10.1016/S0967-0637(02)00134-6)
- Pillar, H. R., Heimbach, P., Johnson, H. L., & Marshall, D. P. (2016). Dynamical attribution of recent variability in Atlantic overturning. *Journal of Climate*, *29*(9), 3339–3352. <https://doi.org/10.1175/JCLI-D-15-0727.1>
- Piron, A., Thierry, V., Mercier, H., & Caniaux, G. (2017). Gyre-scale deep convection in the subpolar north Atlantic ocean during winter 2014–2015. *Geophysical Research Letters*, *44*, 1439–1447. <https://doi.org/10.1002/2016GL071895>
- Rayner, N. A., Parker, D. E., Horton, E. B., Folland, C. K., Alexander, L. V., Rowell, D. P., et al. (2003). Global analyses of sea surface temperature, sea ice, and night marine air temperature since the late nineteenth century. *Journal of Geophysical Research*, *108*(D14), 4407. <https://doi.org/10.1029/2002JD002670>
- Robson, J., Hodson, D., Hawkins, E., & Sutton, R. (2014). Atlantic overturning in decline? *Nature Geoscience*, *7*(1), 2–3. <https://doi.org/10.1038/ngeo2050>
- Robson, J., Ortega, P., & Sutton, R. (2016). A reversal of climatic trends in the North Atlantic since 2005. *Nature Geoscience*, *9*(7), 513–517. <https://doi.org/10.1038/ngeo2727>
- Robson, J., Sutton, R., Lohmann, K., Smith, D., & Palmer, M. D. (2012). Causes of the rapid warming of the North Atlantic Ocean in the Mid-1990s. *Journal of Climate*, *25*(12), 4116–4134. <https://doi.org/10.1175/jcli-d-11-00443.1>
- Roemmich, D., & Gilson, J. (2009). The 2004–2008 mean and annual cycle of temperature, salinity, and steric height in the global ocean from the Argo Program. *Progress in Oceanography*, *82*, 81–100. <https://doi.org/10.1016/j.pocean.2009.03.004>
- Song, H., Marshall, J., Follows, M. J., Dutkiewicz, S., & Forget, G. (2016). Source waters for the highly productive Patagonian shelf in the southwestern Atlantic. *Journal of Marine Systems*, *158*, 120–128. <https://doi.org/10.1016/j.jmarsys.2016.02.009>
- Spall, M. A. (2004). Boundary currents and watermass transformation in marginal seas. *Journal of Physical Oceanography*, *34*(5), 1197–1213. [https://doi.org/10.1175/1520-0485\(2004\)034<1197:BCAWTI>2.0.CO;2](https://doi.org/10.1175/1520-0485(2004)034<1197:BCAWTI>2.0.CO;2)
- Sutton, R. T., McCarthy, G. D., Robson, J., Sinha, B., Archibald, A., & Gray, L. J. (2017). Atlantic multi-decadal variability and the UK ACSIS programme. *Bulletin of the American Meteorological Society*, *99*, 415–425. <https://doi.org/10.1175/BAMS-D-16-0266.1>
- Thacker, W. C., & Long, R. B. (1988). Fitting dynamics to data. *Journal of Geophysical Research*, *93*(C2), 1227–1240. <https://doi.org/10.1029/JC093iC02p01227>
- Verdy, A., Mazloff, M. R., Cornuelle, B. D., & Kim, S. Y. (2014). Wind-driven sea level variability on the California coast: An adjoint sensitivity analysis. *Journal of Physical Oceanography*, *44*(1), 297–318. <https://doi.org/10.1175/JPO-D-13-018.1>
- Watkins, M. M., Wiese, D. N., Yuan, D.-N., Boening, C., & Landerer, F. W. (2015). Improved methods for observing Earth's time variable mass distribution with GRACE using spherical cap mascons. *Journal of Geophysical Research: Solid Earth*, *120*, 2648–2671. <https://doi.org/10.1002/2014JB011547>
- Williams, R. G., Roussenov, V., Lozier, M. S., & Smith, D. (2015). Mechanisms of heat content and thermocline change in the subtropical and subpolar North Atlantic. *Journal of Climate*, *28*(24), 9803–9815. <https://doi.org/10.1175/JCLI-D-15-0097.1>