University of Reading

A window on the deep ocean: the special value of ocean bottom pressure for monitoring the large-scale, deep-ocean circulation

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open access

Hughes, C. W., Williams, J., Blaker, A., Coward, A. and Stepanov, V. (2018) A window on the deep ocean: the special value of ocean bottom pressure for monitoring the large-scale, deep-ocean circulation. Progress in Oceanography, 161. pp. 19-46. ISSN 0079-6611 doi: 10.1016/j.pocean.2018.01.011 Available at https://centaur.reading.ac.uk/87015/

It is advisable to refer to the publisher's version if you intend to cite from the work. See Guidance on citing.
Published version at: http://dx.doi.org/10.1016/j.pocean.2018.01.011
To link to this article DOI: http://dx.doi.org/10.1016/j.pocean.2018.01.011

Publisher: Elsevier

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the End User Agreement.

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Progress in Oceanography 161 (2018) 19–46

Contents lists available at ScienceDirect

Progress in Oceanography

journal homepage: www.elsevier.com/locate/pocean

A window on the deep ocean: The special value of ocean bottom pressure for monitoring the large-scale, deep-ocean circulation

Chris W. Hughes[a,b,*], Joanne Williams[b], Adam Blaker[c], Andrew Coward[c], Vladimir Stepanov[b,1]

[a] University of Liverpool, UK
[b] National Oceanography Centre, Liverpool, UK
[c] National Oceanography Centre, Southampton, UK

ABSTRACT

We show how, by focusing on bottom pressure measurements particularly on the global continental slope, it is possible to avoid the "fog" of mesoscale variability which dominates most observables in the deep ocean. This makes it possible to monitor those aspects of the ocean circulation which are most important for global scale ocean variability and climate. We therefore argue that such measurements should be considered an important future component of the Global Ocean Observing System, to complement the present open-ocean and coastal elements. Our conclusions are founded on both theoretical arguments, and diagnostics from a fine-resolution ocean model that has realistic amplitudes and spectra of mesoscale variability. These show that boundary pressure variations are coherent over along-slope distances of tens of thousands of kilometres, for several vertical modes. We illustrate the value of this in the model Atlantic, by determining the time for boundary and equatorial waves to complete a circuit of the northern basin (115 and 205 days for the first and second vertical modes), showing how the boundary features compare with basin-scale theoretical models, and demonstrating the ability to monitor the meridional overturning circulation using these boundary measurements. Finally, we discuss applicability to the real ocean and make recommendations on how to make such measurements without contamination from instrumental drift.

1. Introduction

In monitoring the global ocean circulation we are faced with a major challenge in the form of the wide disparity in length scales involved. A recent review (Wunsch, 2016) highlighted how this challenge limits what can be said about large-scale, integral properties of the ocean. In essence, the issue is that ocean currents are dominated by mesoscale variability (Ferrari and Wunsch, 2009), with natural length scales of order 10–100 km, so that any one in situ measurement is only representative of a very small region of the ocean. Quantification of mapping accuracy requires a knowledge of the frequency-wavenumber spectrum of ocean variability. To this end, Wortham and Wunsch (2014) have made an effort to characterise this spectrum as seen in the primary physical variables of pressure (and sea level), velocity and density (or temperature and salinity). Their spectrum varies regionally, and most of this variation is designed to reflect the varying characteristics of mesoscale eddies around the world.

One method of obtaining large-scale information is to use a variable which intrinsically integrates some property. Earth rotation measurements are one such variable, but can be difficult to interpret because the integral involves the entire Earth system, not just the ocean. Somewhat more focused is the Earth's gravity field as measured by the GRACE satellite mission. This has provided extremely valuable information about variations in total ocean mass and the sources of water responsible for these changes (Dieng et al., 2015) and is a crucial element of the ocean and Earth observation system, although it does suffer from some of the same ambiguities as Earth rotation, the influence on long time scales of long term plastic deformation of the earth, particularly with respect to the pole tide, remains contentious (Wahr et al., 2015), and it is limited to providing relatively coarse resolution information on ocean bottom pressure variations.

A second way to obtain large-scale information is to have good sampling over the entire ocean. In this respect, satellite altimetry is a particularly powerful system, with sufficient sampling to average out most of the mesoscale variability. Once the trend and seasonal cycle has been removed, the measured variability in global mean sea level has a standard deviation of only 2.5 mm, a level of noise which allows for detection of a trend of 1 mm yr^{-1} from only 2 years of data, compared

* Corresponding author at: University of Liverpool, UK.
E-mail address: cwh@liv.ac.uk (C.W. Hughes).
[1] Now at Hydrometeorological Center of Russia, and University of Reading, UK.

https://doi.org/10.1016/j.pocean.2018.01.011
Received 18 October 2017; Received in revised form 16 January 2018; Accepted 17 January 2018
Available online 30 January 2018
0079-6611/ © 2018 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/BY/4.0/).

to a typical requirement for local sea level which is measured in decades (Hughes and Williams, 2010).

The Argo float program sampling is now sufficiently dense that a similar noise reduction is apparent in estimates of upper ocean heat content (Wunsch, 2016), although the sampling is significantly poorer than altimetry, and even altimetry leaves significant room for improvement with the present nadir-sampling systems only measuring thin lines along the ocean surface. These systems are providing very important inventory information; how much water there is in the ocean and in different density classes. What they cannot generally do is provide useful transport estimates.

To the extent that the ocean is in geostrophic balance, pressure and sea level represent naturally integrating variables, pressure difference at a particular latitude and depth being proportional to the integrated horizontal current perpendicular to the section. Unfortunately, to obtain a useful integral it must be from boundary to boundary, otherwise the end points are likely to be in regions of strong mesoscale variability and the integral will still be dominated by the mesoscale (Wunsch, 2008). For sea level this is a problem because the boundaries are in shallow water where locally-driven dynamics can dominate, as the direct effect of wind stress on sea level is inversely proportional to the depth. Viscous processes also become important in shallow water, so geostrophic balance does not hold. Furthermore, the boundaries are the most troublesome region for satellite altimetry. Here, special measures must be taken to apply the standard path-length corrections to altimetry, tidal variability is typically larger and more complicated than in the open ocean, and temporal aliasing is more important (Vignudelli et al., 2011).

To give an idea of the size of the signals we are interested in, a good rule of thumb is that, at mid-latitudes where the Coriolis parameter f is about $10^{-4}\,s^{-1}$, a sea level difference of 1 cm (or a pressure difference of 1 mbar = 1 hPa) reflects a transport of 1 Sv (Sv stands for sverdrup, a unit of $10^6\,m^3\,s^{-1}$), on the assumption that the associated geostrophic flow penetrates to 1000 m depth. This is the transport associated with about a 5% change in the Atlantic meridional overturning circulation (AMOC), for example, and is the size of change we might aspire to monitor if changes in global ocean circulation are considered.

To put these numbers into context with the mesoscale variability, Fig. 1 (top) shows the standard deviation of sea level from 20 years of satellite altimetry (trend, annual and semiannual cycle removed). This is deliberately plotted using a saturated colour scale, in order to show how few regions approach variability of only a few centimetres.

It is not just the amount of variability that matters, but also its spectrum in both space and time. For the frequency spectrum, given a certain standard deviation, it is helpful for climate monitoring if the variability is dominated by the highest frequencies, since high frequencies can be averaged out more effectively if sampling frequency is high enough. Fig. 1 (bottom), updated from Hughes and Williams (2010), illustrates the variability in the shape of the spectrum in a relatively intuitive way: it simply shows the colours which would be perceived if the spectrum of sea level variability was translated to a light spectrum, with periods 2–24 weeks mapped on to the visible range, corresponding to wavelengths of 380–760 nm.

More detailed explanation of these colour plots and their scale bars is given in the appendix, but they should not be interpreted in a very quantitative way. For present purposes, the value of these colour spectrum plots is as a qualitative condensation of a combination of information about amplitude of variability (brightness) and spectral shape (colour), which we can also exploit when looking at model diagnostics. Blue colours tend to represent relatively higher variability at high frequencies, and similar colours are often representative of similar processes, but more detailed analysis is needed to confirm this. We will not attempt similar diagnostics for the spatial spectrum because, as we will find, bottom pressure is strongly influenced by topography, so the along-slope and across-slope variations can be very different, something which is difficult to account for with wavenumber spectra in the presence of complex topography.

Our purpose in this paper is to illustrate the value of ocean bottom pressure measurements, and to make the case that such measurements, in particular regions, should be a major part of a global ocean observing system. In the following sections, we will see that bottom pressure is quieter than sea level, and has a "whiter" characteristic spectrum (meaning that it will appear more blue in the spectral colour plots). We will also find that mesoscale variability is strongly damped by steep topography, and give a theoretical reason why that should be expected. Focusing on the steep topography of the continental slope, we will show how this allows us to see global scale ocean processes and to access diagnostics which test simple theoretical representations of the global ocean circulation, particularly the AMOC.

We will make these arguments based on diagnostics from a fine resolution global ocean model. While we will only illustrate these arguments with one model, we have investigated a number of different models with a range of resolutions and architectures, and the general findings we present are robust.

Section 2 describes the model runs, and general aspects of the data analysis, Section 3 discusses the variability and spectra of model sea level and bottom pressure, demonstrating how different bottom pressure is and describing some general features. Section 4 presents a theoretical argument explaining why the mesoscale signal is so strongly suppressed in bottom pressure, particularly over steep topography. Section 5 focuses on the Atlantic continental slope, illustrating the striking coherence of dynamical signals over large distances, and making some links to theoretical ideas and simple models, particularly in the context of the AMOC. Finally, in Section 6 we discuss how this can be applied in the real ocean, highlighting the capabilities and deficiencies of present technology and some possibilities for the future.

2. Model descriptions

The model diagnostics are mainly from the National Oceanography Centre run N006 of the 1/12° global NEMO model. This is a single integration of NEMO v3.6 encompassing years 1958–2012 (inclusive), though it has more recently been extended to 2015. The model is forced by the Drakkar Surface Forcing data set version 5.2, which supplies surface air temperature, winds, humidity, surface radiative heat fluxes and precipitation (Dussin et al., 2014; Brodeau et al., 2010). To prevent excessive drifts in global salinity due to deficiencies in the fresh water forcing, sea surface salinity is relaxed toward climatology with a piston velocity of 33.33 mm day^{-1} psu^{-1}. Sea ice is represented by the Louvain-la-Neuve Ice Model version 2 (LIM2) sea-ice model (Timmerman et al., 2005). Bottom topography is represented as partial steps and bathymetry is derived from ETOPO2 (U.S. Department of Commerce, 2006). Climatological initial conditions for temperature and salinity were taken in January from PHC2.1 (Steele et al., 2001) at high latitudes, MEDATLAS (Jourdan et al., 1998) in the Mediterranean, and Levitus et al. (1998) elsewhere. More details of the model and validation of its representation of the AMOC can be found in Moat et al. (2016).

There is no atmospheric pressure forcing, so the sea level can be considered to be equivalent to the inverse barometer-corrected dynamic topography provided in the satellite data. The output data are averaged over 5-day periods which start at the beginning of each year, giving 73 × 5-day means per year (the last day of leap years is thus not saved). The model is volume conserving (Boussinesq), so we calculate bottom pressure from sea level (multiplied by acceleration due to gravity and surface density) plus an integral of gravity times density using hydrostatic balance exactly as implemented in the model, then subtract off the global area-averaged pressure at each time to enforce mass conservation. The corresponding adjustment to global area-averaged sea level was also made, as described by Greatbatch (1994).

The nominal 1/12° resolution is on the tripolar ORCA12 grid, which is regular in longitude south of 20°N, with Mercator latitude spacing

Fig. 1. Sea level variability from satellite altimetry, after subtraction of a linear trend, annual and semiannual cycle, showing (top) standard deviation over 22 years, 1993–2014 inclusive, on a saturated scale, and (bottom) the spectrum of variability (from 1024 weeks from 1993 to mid-2012) as described in the text, the appendix, and in Hughes and Williams (2010).

chosen to make square grid cells. In the northern hemisphere, the grid distorts to avoid producing a pole in the Arctic Ocean, instead having two poles over land. Over ocean points, the linear resolution (square root of grid cell area) thus varies between 9.27 km at the equator, 2.09 km at 77°S (the southernmost point), and 1.75 km in the Canadian Arctic.

Preliminary analyses showed that annual and semiannual signals frequently show clearly above the background spectrum, but higher harmonics generally do not, though they can with sufficient averaging of spectra. Most of the diagnostics presented here are from the 54-year period 1959 to 2012 inclusive, after subtraction of a mean, annual cycle, semiannual cycle and linear trend, determined by simultaneous least squares fitting. The exception is the spectral colour plots, which use a period from 1980 to January 2008 (this gives a time series length of 2048 values, representing 28 years and 20 days).

In addition to the NEMO data, we also show some diagnostics from the Advanced Global Barotropic Ocean Model (AGBOM). This is a global ocean model with constant water density, at 0.25° resolution, based on that described in Stepanov and Hughes (2004). It uses a simple parameterisation of self-attraction and loading with a proportionality

factor $\beta = 0.12$ as described in that paper (this is of importance mainly for the variability at periods of a few days or less, including the tides). It is forced using 6-hourly atmospheric pressure and wind stress from the ERA-40 reanalysis (Uppala et al., 2005) as well as 17 components of diurnal and semidiurnal tidal forcing. For present purposes, tides are removed by filtering before forming 5-day means. We use a time series of 1080 5-day means, which covers approximately 14.8 years.

Ideally we would have a single-layer model with exactly the same configuration and forcing as the NEMO run, but this is not available. However, there are some advantages to using AGBOM, as it fills in some missing physics such as tides and the response to varying air pressure, which are absent from the NEMO run. As we will see, the AGBOM results do appear consistent with NEMO in the relevant frequency ranges, which testifies to the robustness of the barotropic modelling, since the model configurations and discretisations are quite different. We expect the main influence of atmospheric pressure variability on bottom pressure to be seen on time scales shorter than a month (Stepanov and Hughes, 2006), and this expectation is consistent with the spectra we present in the next section.

3. Spectral analyses

In this section we start by examining and discussing global maps related to the amplitude and spectra of sea level variability, showing how it is partitioned into signals related to steric and mass (bottom pressure) variability. We discuss the known forms of variability, and draw attention to the relative quietness of bottom pressure, especially along the continental slope. We then present more detailed spectral analyses, averaged over different regions chosen based on what we have learned from the maps. These highlight the fact that the part of the spectrum which represents mesoscale variability is much weaker in bottom pressure than in sea level, especially on the steep continental slopes.

Fig. 2 shows the same diagnostics, standard deviation and spectrum of sea level, from the NEMO model as were shown from altimetry in Fig. 1. It is immediately clear that the model is performing well in reproducing much of the variability of the real ocean, not just in amplitude but also in spectral characteristics. The two figures are not precisely comparable, being based on different lengths of time series, with different temporal sampling (5-day means for the model, nominal 20-day low-pass filter for altimetry) and different spatial filtering (determined by resolution and parameterised friction in the model, and by a compromise between satellite track spacing, observed scales of variability, instrumental noise and high frequency ocean processes for the altimetry). Nonetheless, many features of the observations are reproduced in the model at very similar amplitudes.

The spectrum plots in Figs. 1 and 2 focus on timescales associated with mesoscale variability, which is appropriate as that is the main source of "noise" with respect to our aim of measuring large-scale variability. As noted by Hughes and Williams (2010), we see that in many regions the coastal and shelf sea variability has a blue tinge in the spectrum, and is separated by a region of lower variability near the top of the continental slope from a differently-coloured spectrum of high variability offshore. This is particularly apparent at mid to high latitudes near the western boundary of the North and South Atlantic, but also occurs elsewhere (as many of the features described in this section are at quite small scales, we recommend that the reader uses the electronic version of this paper to zoom in on the regions being discussed). As a reminder, blue represents spectral slopes significantly shallower than f^{-2}, so "white noise" would appear intensely blue. We interpret this blue colour on the shelf as the result of the strong influence of wind stress forcing in this region. The minimum of variability near the top of the continental slope is also visible in the standard deviation. This drop-off in eddy amplitude near the western boundary is referred to by Zhai et al. (2010) as the "eddy graveyard".

There are some clear differences between model and observations, for example the plume of high variability stretching to the west and steadily north from the tip of South Africa in the model which suggests a too-regular path of Agulhas rings penetrating into the South Atlantic. However, overall, the similarity in both amplitude and spectral shape (colour) is sufficient to give us confidence that the model can be used to investigate the influence of mesoscale variability on observation of large-scale processes.

For comparison, Fig. 3 repeats these sea level diagnostics for ocean bottom pressure. We report bottom pressure in mbar, equivalent to hPa. Applying hydrostatic balance as a scaling factor, 1 mbar pressure is equivalent to approximately 1 cm of sea level (the scaling can vary by about 2–3% depending on the local water density and gravity used). Note that, for standard deviation, the colour scale amplitude has been halved compared to the sea level plots, but still the area occupied by saturated scale values has markedly reduced. Similarly, for the spectral colour, the brightness of the plot is exaggerated compared to Fig. 2; the effect is equivalent to multiplying the spectral power by a factor of 10 (i.e. the time series are multiplied by $\sqrt{10}$ before calculating the colours), which is why the Arctic and Southern Ocean, among other regions, are so much brighter than in the sea level spectrum shown in Fig. 2. The quieter nature of bottom pressure has been noted before, in both coarse resolution (Vinogradova et al., 2007) and eddying models (Bingham and Hughes, 2008a), and the remarkably quiet nature of much of the tropics has led to ocean bottom pressure measurements being used to determine the annual cycle of mass exchange between ocean and land (Hughes et al., 2012; Williams et al., 2014; Hsu and Velicogna, 2017).

Here, we focus particularly on the spatial distribution of the variability and its spectrum. In addition to the tropics, we see a very quiet region along the eastern boundary continental slope. Compared to sea level (Figs. 1 and 2) we see an even more marked and, especially in the Atlantic, a broader minimum in bottom pressure variability along the western boundary continental slope. The spectrum is generally less "red" than sea level (hence the overall blue colour), with the exception of regions of intense eddy activity where the eddy variability also dominates bottom pressure.

Several features stand out in the bottom pressure variability. There is often evidence of strong topographic influence, even in regions where the colour of the spectrum suggests a mesoscale influence, with entire sub-basins having a rather constant colour. In the case of the Arctic, the colour is the blue which we will usually come to associate with wind stress forcing, and the entire deep Arctic basin has extremely well correlated variability (not shown), consistent with the barotropic mode first noted by Hughes and Stepanov (2004) and nicely elucidated by Fukumori et al. (2015). The Mediterranean is similar, though a difference in character between the eastern and western basins is apparent. The well-known regions of high barotropic variability in the subpolar North Pacific, the south east Pacific and Indian Ocean sectors of the Southern Ocean (Chao and Fu, 1995; Fu and Davidson, 1995) show up clearly. In the Argentine Basin, the range of colours is indicative of the complex interplay of mesoscale eddies and barotropic basin modes which is known to occur here (Fu et al., 2001; Weijer et al., 2007; Hughes et al., 2007; Fu, 2007). Perhaps the most obvious feature, though, is in the Caribbean Sea, which shows up as bright red in the spectral colour plot. It was this feature of the spectrum, repeated across a wide range of ocean models, which led to the discovery of the Rossby whistle: a 120-day baroclinic basin mode in the region, excited by baroclinic instability of the Caribbean Current (Hughes et al., 2016). This clear, coherent mode dominates the regional bottom pressure, despite having a standard deviation of less than 2 mbar.

We have interpreted much of the bottom pressure variability in terms of barotropic modes, but it is not clear that bottom pressure has to be dominated by barotropic dynamics (by which we mean dynamics associated with depth-independent pressure variations). Baroclinic modes also have a bottom pressure signature, especially over weakly sloping topography and close to the equator. In addition, topography

Fig. 2. Sea level variability from 54 years of ocean model data (5-day means), after subtraction of a linear trend, annual and semiannual cycle, showing (top) standard deviation on a saturated scale, and (bottom) the spectrum of variability from a 28-year subset starting in 1980, as described in the text, the appendix, and in Hughes and Williams (2010).

causes mode coupling between barotropic and baroclinic modes, so that barotropic signals in bottom pressure can be the result of baroclinic processes, as in the Caribbean Sea mode. Such mode coupling can also result from nonlinear dynamics. The intense blue bands in the spectrum of sea level either side of the Pacific equator (Figs. 1 and 2) are associated with Tropical Instability Waves (Legeckis, 1977). Farrar (2011) has recently shown how these waves, which have periods of about 30 days, can radiate beyond the permitted latitude range for baroclinic Rossby waves of that period by partially converting to barotropic waves. The interesting colours in the bottom pressure spectrum in this region thus represent a combination of baroclinic signals and barotropic signals induced by the baroclinic variability.

A simple example of such mode coupling is a baroclinic equatorial Kelvin wave which, when it reaches the eastern boundary, will still cause a coastal sea level change despite the water depth becoming too shallow for this to be a baroclinic mode (i.e. shallower than the thermocline depth). This may be what we are seeing on the Atlantic coast of Africa, where the coastal bottom pressure variability has a pink hue in the spectral colour plot. Baroclinic processes, such as the Kelvin wave and associated coastal-trapped waves which may have a baroclinic

Fig. 3. Bottom pressure variability from 54 years of ocean model data (5-day means), after subtraction of a linear trend, annual and semiannual cycle, showing (top) standard deviation on a saturated scale, and (bottom) the spectrum of variability from a 28 year subset starting in 1980, shown brighter than the equivalent sea level spectrum (power increased by a factor of 10).

component, cause a coastal signal which is seen in bottom pressure but not in steric sea level variability. The latter point is confirmed in Fig. 4, which shows the variability and spectral colour of the steric component of sea level variability (i.e. the part that is attributable to water column density changes rather than bottom pressure changes). The low variability around coasts and most shelf sea regions confirms that the variability here is predominantly barotropic, with sea level and bottom pressure varying in step as they would in a homogeneous ocean. In the limit of zero depth this is inevitable, as the steric signal, being a depth integral, must tend to zero, meaning density changes do not significantly disrupt the relationship between sea level and bottom pressure.

In comparison with Fig. 2, Fig. 4 shows that most of the blue-purple "haze" visible in the sea level spectral colour originates from the bottom pressure variability, and steric variability produces a sharper-looking plot. We can also see how, for example over many small islands and seamounts in the Indian and Pacific ocean, the sea level spectrum is continuous from deep to shallow water, but the partitioning of

Fig. 4. Steric variability from 54 years of ocean model data (5-day means), after subtraction of a linear trend, annual and semiannual cycle, showing (top) standard deviation on a saturated scale, and (bottom) the spectrum of variability from a 28 year subset starting in 1980.

variability shifts sharply from being dominated by steric variations in the deep water to bottom pressure variations in shallow water. It seems that such small topographic features offer little dynamical obstacle to the propagation of open-ocean sea level signals to the coast, despite the switch between steric and bottom pressure-related sea level change. In fact, Williams and Hughes (2013) showed (in a different high-resolution model) that the relationship between small-island sea level and surrounding open ocean sea level is practically unaffected by the existence of the island topography, for 5-day mean values.

From a single model run, mode coupling and nonlinear processes mean that it is impossible to determine which variability is the result of purely barotropic processes. In Fig. 5, we show the variability of inverse barometer-corrected sea level (or bottom pressure, which is equivalent in this model) from the AGBOM model. Since this model has constant density, all processes in the model are due to barotropic dynamics. The spectral colour is almost everywhere blue to purple, supporting the idea that much of the bottom pressure variability, particularly at high latitudes and in broad shelf seas, is the result of purely barotropic processes.

In contrast, most of the more colourful features in Fig. 3 are absent

Fig. 5. Sea level or bottom pressure variability from 15 years of barotropic ocean model data (5-day means), after subtraction of a linear trend, annual and semiannual cycle, showing (top) standard deviation on a saturated scale, and (bottom) the spectrum of variability, shown brighter than the equivalent sea level spectrum (power increased by a factor of 10).

from Fig. 5, strongly suggesting that these features are the result, either directly or indirectly, of baroclinic processes. This includes the deep-water tropics, and most of the eastern boundary regions. It is striking how the low variability region in the barotropic model extends right round the continental slope of the North Pacific. This is not so dramatic in the North Atlantic where even barotropic variability is significant in the western basins. Though this variability is generally smaller than either the steric variability seen in Fig. 4, or the total bottom pressure variability in Fig. 3, it is comparable in magnitude to the total bottom pressure variability on much of the western continental slope, suggesting that purely barotropic processes are a significant fraction of the total in this region.

Putting these diagnostics together we see that pure barotropic variability accounts for much of the blue-white "wash" seen in the sea level variability in Figs. 1 and 2, and especially on broad shelf seas, though much of the shelf sea variability on narrower shelves and close to the shelf break is, though locally barotropic, induced by baroclinic processes. The blue-white "wash" is removed when focusing on steric sea level, as shown in Fig. 4. Remembering our rule of thumb suggesting we want to measure large-scale signals of order 1 cm, the large

amplitude of this mesoscale variability and barotropic shelf sea variability means that, without spatial averaging, sea level is a poor monitor of large-scale ocean circulation except perhaps in a few very quiet regions near eastern boundaries and near the poles.

Bottom pressure can be significantly quieter than sea level. Even in bottom pressure the mesoscale variability is too large for climatological monitoring in many regions, especially in the western basins. We have seen that steep topography alone is not sufficient to suppress such variability, as the sea level variability over many small tropical islands and shallow seamounts is almost the same as in nearby deep water. The continental slope, however, is both steep and long. These are the characteristics needed to suppress mesoscale variability, as we will discuss in the next section. The result is illustrated in Fig. 3, which shows strong suppression of variability over the continental slope. Combined with Fig. 5, which suggests that a significant fraction of even the small variability seen over these regions results from (presumably large-scale) barotropic processes, this suggests that bottom pressure on the continental slope is the most promising variable to monitor in order to measure large-scale, climatological changes in the ocean circulation.

While the spectral colour maps are a useful qualitative indicator of the dynamics, they do not provide good quantitative information, and are limited to representing a particular subset of the frequencies which can be resolved. To address this issue, we use the maps to guide us in choice of regions to average over, in order to investigate the spectral shapes and amplitudes in more detail. The most obvious geographical division, in the deep ocean at least, is latitude, so in Fig. 6 we show power spectra averaged over the deep ocean in latitude bands, with a secondary partitioning based on how energetic the steric sea level variability is (as an attempt to isolate the influence of nonlinear mesoscale eddy variability). We define "quiet" as regions where the standard deviation of steric sea level variability is less than 6 cm, and "energetic" as where it is greater than 9 cm.

Fig. 6 contains a lot of information. Focusing first on sea level (red) at latitudes equatorward of 30°, we see that the spectrum is divided into two regions with a very clear breakpoint. The higher frequency part of the spectrum is often a very straight line, indicating a power law. However, the gradient of the line varies, between almost −4 at the lowest latitudes to about −3 at 15–30°, and becomes shallower still at higher latitudes. As noted above, the annual cycle and its harmonics are very clear with such averaging, with harmonics as high as 5 cycles per year being clearly visible in some cases.

As we move to higher latitudes, a difference between the quiet and energetic regions becomes more apparent. In the quiet regions, the breakpoint in the spectrum is still visible out to latitudes of 45–60°, but in the energetic regions it becomes more of a smooth, rounded transition rather than a breakpoint. This suggests that the rounded spectrum is indicative of the more nonlinear processes in the higher latitude energetic regions, whereas the breakpoint is related to more linear processes.

Finally, at the highest latitudes, the breakpoint is lost and, particularly in the Arctic, the spectrum becomes much more noisy. This is informative in itself. The spectra are formed by area-weighted averaging of the modulus-squared fourier transforms of the time series from each grid point in each region. For a single gridpoint, this would produce a very noisy spectrum, but for multiple gridpoints the different realisations of time series with the same (or similar) underlying spectrum, but different phasing of the components, averages out that noise. However, in the Arctic, the time series are all strongly correlated, so there are no different realisations, and no averaging effect is found. The noisiness of the spectra is an indication of the number of spatial degrees of freedom in the variability, so large-scale processes (or small regions) will result in noisier spectra. We see this effect at the high frequency end of many of the curves, where large-scale barotropic processes become dominant.

Lin et al. (2008) and Hughes and Williams (2010) attributed the breakpoints to Rossby wave processes. In linear dynamics, Rossby waves are limited to frequencies below a strongly latitude dependent cutoff frequency. This frequency is given by $f = 2\pi/\beta R_1$, where R_1 is the baroclinic Rossby radius and β is the meridional gradient of the Coriolis parameter. The green boxes in Fig. 6 show the range of this cutoff frequency for each latitude range, based on zonally averaged values of the Rossby radius taken from Chelton et al. (1998). Where a clear breakpoint is visible, it does indeed lie in this band.

Turning now to the other curves in Fig. 6, blue shows NEMO bottom pressure, and black shows the AGBOM sea level or bottom pressure. We see that bottom pressure is much quieter than sea level at low frequencies, but approaches sea level at high frequencies where barotropic processes become dominant. In fact, the barotropic model spectra are often more energetic than the NEMO bottom pressure spectra at the very highest frequencies (periods shorter than about two weeks), a fact which may be attributable to the effect of atmospheric pressure forcing (compare Stepanov and Hughes (2006), who find that the inverse barometer correction ceases to be a good approximation at these timescales on a basin scale). Otherwise, the AGBOM and NEMO bottom pressure spectra tend to be very similar in quiet regions. In the more energetic regions, however, the difference between NEMO and AGBOM bottom pressure spectra tends to have a shape reflecting the sea level spectra, suggesting that bottom pressure is being influenced by the energetic mesoscale variability as we deduced from the spectral colour maps.

It is also worth remarking that, at periods longer than 10 years, the NEMO bottom pressure spectra start to slope up again. Given that this is not the case for wind stresses (not shown), and that barotropic processes have no memory on timescales longer than a few months, we would not expect this to be the case in AGBOM, though the time series is too short to confirm this. This decadal-to-multidecadal bottom pressure variability may be related to climate-scale ocean dynamics, or to baroclinic ocean model adjustment processes.

Since we have identified the continental slope as an important region, Fig. 7 shows the spectra averaged over different regions defined by topography. Here, shallow means all regions shallower than 200 m, and deep means regions deeper than 3200 m as before. For the range of depths in between, however, we only include points which are on the global continental slope, excluding seamounts and isolated islands. This region is defined by starting at a latitude on the Atlantic coast of the USA where the continental slope happens to be monotonic, and following depth contours from this section until they close after passing round Europe, Africa, Asia and the Americas. Some deep contours also pass round Australia and New Zealand, and some shallow contours enter and pass round the Arctic and Mediterranean (more detail in the Atlantic is given later in this paper, including a map of the Atlantic part of the continental slope, Fig. 8). The definition of the continental slope is thus limited to very long contours in the depth range 200–3200 m.

Although not especially energetic in terms of sea level, the shallow region is the most energetic in bottom pressure. Because the definition of "energetic" is based on steric sea level, and steric signals are small in shallow water, there are no points which are both energetic and shallow. As expected, the shallow signals are dominated by bottom pressure, but at the lower frequencies a large fraction of this represents a locally barotropic signal which is induced by baroclinic variability, as it is absent in the purely barotropic AGBOM.

In the deep ocean we see a pattern consistent with Fig. 6. The pure barotropic AGBOM spectrum is indifferent to whether or not there is energetic steric variability, but some fraction of the steric variability is seen in the NEMO bottom pressure, lifting the spectrum above the barotropic spectrum in energetic regions, but much less so in quiet regions.

Over the continental slope, a more interesting result arises. Here, the NEMO bottom pressure is comparable to that in quiet deep ocean regions, whatever the local steric variability. Particularly at frequencies above about 0.3 cycles per year (periods shorter than about 3 years) bottom pressure on the slope does not seem to be significantly

Fig. 6. Spectra of NEMO sea level (red), NEMO bottom pressure (blue) and barotropic model sea level (black) averaged over the deep ocean (depth > 3200 m) in latitude bands. The spectra are further divided according to the amplitude of steric sea level variability in NEMO, with "quiet' defined as less than 6 cm standard deviation, and "energetic" as more than 9 cm standard deviation. Green bars mark the range of linear baroclinic cut-off frequencies for each latitude range. Some representative power laws are shown in the bottom left plot. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Fig. 7. Spectra of NEMO sea level (red), NEMO bottom pressure (blue) and barotropic model sea level (black) averaged over the ocean in regions defined by topography. Shallow is all regions shallower than 200 m, deep is all regions deeper than 3200 m, and in between is the continental slope (see text for a more detailed definition). The top and bottom rows show spectra further divided according to the amplitude of steric sea level variability in NEMO, with definitions as in Fig. 6. The green box simply provides a constant reference level for comparison. Some representative power laws are shown in the top left plot. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

contaminated by effects associated with local steric variability, although there is significant variability above the purely barotropic dynamics in AGBOM. This suggests that, on the continental slope, we have a window into processes beyond the purely barotropic or purely mesoscale. It is worth noting that even the purely barotropic processes (black) are less energetic over the slope than either the deep or shallow ocean.

We will investigate this in more detail following a discussion of the theory, but we conclude this section by noting the clear message of these diagnostics. Although mesoscale variability may dominate pressures, currents and sea level in the ocean interior, a quite different spectrum is apparent in bottom pressure in many regions, and especially on the continental slope. Bottom pressure is a special variable which gives access to different processes in different frequency ranges than those accessible by other physical ocean variables. This opens up particular opportunities for ocean monitoring.

4. Theory of mesoscale suppression on the continental slope

The argument for suppression of mesoscale energy in bottom pressure on the continental slope centres on the relationship between bottom pressure p_b and vertical velocity w. The kinematic boundary condition on near-bottom velocity (neglecting a viscous boundary layer) is $\boldsymbol{v}_b \cdot \nabla H = -w_b$, where $\boldsymbol{v}_b$ is the horizontal velocity at the bottom, w_b is vertical velocity at the bottom, and H is ocean depth, with the sea floor at $z = -H$. If f is the Coriolis parameter ($f = 2\Omega \sin\phi$ where Ω is the Earth's angular rotation rate and ϕ is latitude), p is pressure, and $\hat{\boldsymbol{k}}$ is the upward unit vector, we can use geostrophic balance $\rho f \boldsymbol{v} = \hat{\boldsymbol{k}} \times \nabla p$ to substitute for $\boldsymbol{v}_b$ in terms of pressure. Writing this in a coordinate system in which x is measured along the direction of the depth gradient (positive towards deep water) and y is along the depth contour, we obtain

$$w_b = -u_b \frac{\partial H}{\partial x} = \frac{1}{\rho f} \frac{\partial p_b}{\partial y} \frac{\partial H}{\partial x}, \tag{1}$$

which leads to the scaling

$$w_b \sim u_b S, \tag{2}$$

where S is the bottom slope, and u_b is the bottom horizontal flow toward deeper water, related to the along-slope gradient of bottom pressure.

A second scaling for w_b can be obtained from the vorticity equation. Consider the inviscid equation of motion, with Boussinesq and "conventional" approximations:

$$\boldsymbol{v}_t + (f\hat{\boldsymbol{k}} + \boldsymbol{\omega}) \times \boldsymbol{v} = -\nabla \left(\frac{p}{\rho_0} + \frac{|\boldsymbol{v}|^2}{2} + \Phi \right), \tag{3}$$

where $\boldsymbol{v}$ is the horizontal velocity, $\boldsymbol{\omega} = \nabla \times \boldsymbol{u}$ is the relative vorticity

Fig. 8. The Atlantic continental slope defined as described in the text for the NEMO model. Dots indicate reference nodes every 500 km along the 2000 m contour, with circled numbers representing along-slope distance in units of 10,000 km. Colours indicate how points at other depths in the range 100–3200 m are mapped to the distance variable. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

(curl of the 3D velocity), and Φ is the gravity potential. If we take $\hat{\boldsymbol{k}}\cdot\nabla\times$ of (3) we obtain the vertical vorticity equation:

$$\zeta_t + \boldsymbol{v}\cdot\nabla(f+\zeta) = \nabla\cdot[w(f\hat{\boldsymbol{k}} + \boldsymbol{\omega})], \tag{4}$$

where ζ is the vertical component of vorticity $\boldsymbol{\omega}$. Taking representative horizontal velocity to be U, horizontal eddy length scale L, vertical length scale H and time scale T, this allows us to derive a scaling for w_b. First, we note that $\boldsymbol{\omega}/f$ scales as the Rossby number $\mathrm{Ro} = U/fL$, so that, dimensionally, $f\hat{\boldsymbol{k}} + \boldsymbol{\omega} \sim f\hat{\boldsymbol{k}}(1 \pm \mathrm{Ro})$. After an integral from top to bottom, the right hand side of (4) becomes $f(w_a - w_b)(1 \pm \mathrm{Ro})$, where w_a is vertical velocity at the surface, which can be taken as zero for the mesoscale (the Ekman pumping velocity is much smaller than other vertical velocities in this scaling). Dimensionally, introducing a factor H on the left hand side to account for the vertical integral, (4) can then be written

$$H[\zeta_t + \boldsymbol{v}\cdot\nabla(f+\zeta)] \sim fw_b(1 \pm \mathrm{Ro}). \tag{5}$$

Scaling the remaining terms and rearranging, this becomes

$$w_b \sim \left(\frac{HU}{L(1 \pm \mathrm{Ro})}\right)\left(\frac{1}{Tf},\ \frac{L}{R},\ \mathrm{Ro}\right), \tag{6}$$

where the terms on the right hand side derive from the time dependence, $\boldsymbol{v}\cdot\nabla f$, and $\boldsymbol{v}\cdot\nabla\zeta$ terms respectively in (4). Here, R is the Earth's radius (arising from $f/\beta = R\tan\phi$, approximated as R for mid-latitude regions).

Setting these two scalings for bottom velocity, (2) and (6), to be equal, gives

$$\frac{u_b}{U} \sim \left(\frac{H}{LS(1 \pm \mathrm{Ro})}\right)\left(\frac{1}{Tf},\ \frac{L}{R},\ \mathrm{Ro}\right). \tag{7}$$

The three bracketed terms are precisely the terms which are assumed to be small in the quasigeostrophic approximation, so they are small for the (large) proportion of the mesoscale eddy field which can be described by quasigeostrophic scaling. In more detail, for the terms to be small, timescales must be long compared to the inertial period and length scales short compared to the Earth's radius, both of which are clearly true of mesoscale eddies. The third requirement, that the Rossby

number be small, is effectively the definition of mesoscale (or larger) rather than submesoscale.

If the terms in the final brackets are small, then u_b can only become comparable to U if the first term H/LS is large. This term can be interpreted as the aspect ratio of the eddies divided by the slope, and clearly shows how steeper slopes result in smaller u_b/U ratios. Another interpretation is that H/LS is the width of the "extended slope" divided by the eddy length scale, where "extended slope" means an imaginary slope with constant gradient S, and its width is the horizontal distance over which it extends from top to bottom of the water column. Note that the eddies can be smaller than the extended slope, and the total scaling can still be small if the Rossby number is small.

Thus, we see that, for mesoscale eddies, the relative suppression of bottom velocities in comparison with mid-water velocities is determined mainly by the Rossby number, the bottom slope, and the eddy scale. For example, a slope of 0.1 in water 2 km deep leads to $H/LS = 1$ at an eddy length scale of 20 km, meaning that bottom velocity suppression occurs for all eddies larger than 20 km at $Ro = 1$, or 2 km at $Ro = 0.1$. The continental slope can be about five times steeper or gentler than a slope of 0.1 in extreme cases, showing that a large fraction of the mesoscale energy cannot penetrate to the bottom over typical continental slopes. There are clearly processes with small length scales or large Rossby numbers which can penetrate to a steeply sloping bottom without attenuation (and even processes which are bottom-trapped), but the bulk of the mesoscale variability cannot do so. Near the foot of the continental slope, the topographic constraint weakens, and bottom-trapped Rossby waves are often observed (e.g. Hogg (2000)), particularly at periods shorter than about 10 days such that $1/Tf$ is only moderately small. Over the abyssal plain, slopes may not be strong enough to constrain the bottom velocity so strongly. In these regions, the small bottom pressure variability indicated by NEMO within most subtropical gyres (Fig. 3) must be attributed to the fact that the energy input is at the surface, and the surface-intensified stratification (especially at lower latitudes) results in weakened flow below the thermocline. When the Rossby number is the main constraint we can write the mesoscale suppression factor as

$$\frac{H\mathrm{Ro}}{LS(1 \pm \mathrm{Ro})} = \frac{HU}{fL^2S} = \frac{gH\eta}{f^2L^3S}, \tag{8}$$

where the last equality invokes geostrophic balance to express the velocity scale in terms of a sea level scale η (hence $\mathrm{Ro} = g\eta/f^2L^2$). The factor $(1 \pm \mathrm{Ro})$ has been dropped from all except the first form because, as this makes clear, there is no constraint on the scaling when the Rossby number approaches 1. From the final scaling we see that, for a given amplitude of sea level variability, the suppression is sharply dependent on length scale. For example, for $\eta = 0.1$ m, we obtain $\mathrm{Ro} = 1/4$ at mid-latitudes, and hence a mesoscale suppression factor of 1/4 for length scale 20 km. However, for a 40 km length scale, this suppression factor becomes 1/32.

Furthermore, the influence of this mesoscale suppression is cumulative. Since most of the mesoscale variability is generated in the open ocean and propagates toward the western boundary (e.g. Zhai et al., 2010), once it encounters the continental slope and the mesoscale suppression scaling becomes of order 1 or less, this interaction will influence the eddy propagation so that it does not enter the shallower slope region at all. The raw scaling only applies to eddies generated over the slope.

The argument so far has been in terms of velocities and, in particular, the horizontal velocity component that is constrained at the bottom, u_b, is that perpendicular to depth contours; there is no constraint on the velocity along depth contours as these have no associated vertical velocity. Translating this into pressures, the constraint on u_b/U should be interpreted as a constraint tending to reduce the bottom pressure gradient along depth contours in relation to a typical mid-depth pressure gradient. Locally, this means that bottom pressure will be close to being a function of H, with that function varying slowly with distance along the continental slope. The ability of the mesoscale to excite only flows with much longer scales along than across depth contours therefore means that the mesoscale-induced along-topography flows tend to be an integral of mesoscale influences over length scales typically much larger than the mesoscale itself. As this integral will tend to include forcing of both signs, the integration will usually have the effect of reducing even this component in comparison with typical mesoscale pressure gradients. The exception to this is when the depth contours are closed over distances which are not large compared to the mesoscale length scale, so no averaging occurs. This explains why small islands do not act as a significant barrier to bottom pressure signals, and why the suppression is limited in the case of the Caribbean Sea Rossby whistle mode (Hughes et al., 2016). In fact, closed depth contours (more strictly, closed contours of H/f) lead to a coupling of stratification and bottom pressure which excites a barotropic circulation around the closed contours, enabling baroclinic disturbances to rapidly skip across such closed contour regions (Marshall, 2011). The averaging effect of long contours is clearly an important consideration.

To summarise this scaling analysis in simple terms, the vorticity balance places a constraint on the vertical stretching of water columns, which limits the possible size of the vertical velocity at the bottom. The vertical velocity at the bottom is coupled to the horizontal velocity via the fact that flow cannot pass through the seafloor, so horizontal bottom velocities lead to larger vertical velocities where the slope is steep. For steep continental slopes and typical mesoscale conditions, these two scalings turn out to be incompatible if we use the same horizontal velocity scaling in each case. As a result, the horizontal bottom velocity (and hence the along-slope pressure gradient) at the bottom must be much smaller than the typical near-surface velocities (pressure gradients), and mesoscale variability is suppressed in bottom pressures on the continental slope. We have neglected the viscous boundary layer in this scaling on the grounds that it is not generally found to produce a large local perturbation in pressure. That is not to say that it is unimportant in indirectly shaping the larger-scale pressure field. On the contrary, we suspect that it is important for models to represent the slopes and the frictional processes on those slopes as well as possible, and that this may be one of the main limitations of the present generation of ocean models.

So far we have presented diagnostics concerning the size and the spectrum of variability, and made a scaling analysis. These suggest that we should see large-scale, coherent variability in bottom pressure on the continental slope, and hence that this would be a good place to monitor large-scale ocean circulation variability while minimising the mesoscale noise. In the next section, we will test this suggestion in more detail.

5. Atlantic variability and continental slope bottom pressure

For the sake of providing a concrete example, we will focus on the Atlantic Ocean, for which there is an established interest in the long-term climate variability particularly related to the AMOC. Studies of the Pacific and Indian oceans would be of interest in their own right, but only one basin is necessary in order to establish the general principles.

A coordinate system for the Atlantic continental slope is devised as shown in Fig. 8. This allows us to plot bottom pressure values in two dimensions, distance along the slope (colours), and depth (time is the third dimension). As this coordinate system is used in many subsequent plots, it is worth describing in some detail.

Initially, the 2000 m depth contour is followed around the basin as a reference, starting deliberately on the Pacific side of South America so any link between Pacific and Atlantic values can be seen. The contour consists of a set of straight line (actually very short great-circle) segments between positions where the linear interpolation of depths between neighbouring grid points is 2000 m, thus avoiding any rectangular zigzagging around grid boxes. The along-slope distance is then calculated as the sum of the lengths of these short line segments. The

black dots in Fig. 8 are every 500 km distance along this contour, and are referred to as "nodes". The ringed numbers are every 10,000 km along, and thus represent distance in units of 10,000 km. This is used as the distance axis in later plots. Colours are an additional indicator of distance, with a colour change every 1000 km. Note that the distance is defined following the 2000 m contour which, though not as convoluted as a coastline, can be quite convoluted in places. Thus, although the nodes are separated by 500 km along the 2000 m contour, the great-circle distance between nodes is typically (median value) about 390 km.

Contours are then followed at other depths, to a maximum of 3200 m, at an interval of 1 m, to identify all the gridpoints associated with each continuous contour (excluding seamounts and other closed contour regions on the slope). Below 3200 m, contours do not pass through Drake Passage and also start to spread down the mid-Atlantic ridge. For each contour, the point nearest to each node of the reference contour is labelled as having the same distance as that node, with distance linearly interpolated between points matched to nodes. If the distance to a node is greater than 1000 km, the point is ignored, thus preventing the use of shallow contours which pass round the Arctic, Mediterranean, or parts of the Gulf of Mexico and Caribbean Sea. This procedure provides a usable distance value for points off the reference contour.

For each point on the continental slope, we now have a depth and a distance. Bottom pressures were extracted from these points, and a mean, trend, annual and semiannual cycle were fitted and removed. The two dimensional fields at each time were then regridded onto a regular grid every 50 km in distance and 10 m in depth, from 100 m to 3200 m, using a Delauney triangulation, which treats depths in metres and distances in kilometres as equivalent.

This interpolation gives values at all depths and distances, which is unrealistic given the gaps in shallow contours. To account for this, we calculate a mapping error estimate which is the geometric mean of the horizontal distance of each point on the regular grid from the points used in the triangulation. Points for which this distance is greater than 200 km are left blank.

5.1. Variability and vertical (cross-slope) structure

Fig. 9 (top) shows the standard deviation of bottom pressure as a function of distance and depth. Note that depth here is the depth of the bottom of the ocean. In the diagnostics presented here, horizontal position is a function of depth, and any signal that is described as 'independent of depth' could equally well be described as 'independent of cross-slope distance', for a given along-slope distance as defined in Fig. 8. There may or may not be variation in the vertical above a given point. Thus, the usual ideas of barotropic and baroclinic modes are not appropriate interpretations when looking at bottom pressure on the continental slope, unless it is vertical. To emphasize this, the depth-averaged bottom pressure at a given distance will be represented by $\langle p \rangle_H$, where the subscript H emphasizes that the averaging variable is depth of the ocean floor rather than the vertical coordinate z.

Fig. 9 can thus be imagined as looking into the Atlantic from the south, and unwrapping the continental slope into a straight strip. Vertical lines mark the equator, and grey masks regions of missing data. Reading from the left, the first grey patch at distance about 1.9–2.2 represents the Caribbean Sea and the Gulf of Mexico, with the tip of Florida at about 2.2, and Cape Hatteras at 2.4. Around 3.0–3.1 is the top of the Labrador Sea, where some shallow contours are lost into Baffin Bay or to circumnavigate the Arctic. Between about 3.7 and 4.0, a deeper range connects to the Arctic between Iceland and Scotland. The Strait of Gibraltar is at about 4.57.

The position of the actual gridpoints used to map this standard deviation is shown in the top panel of Fig. 10. In addition to the gaps discussed above, a sparse region in the top kilometre at distances of 1.1–1.2 reflects the Vitoria-Trindade ridge, a feature off the Brazilian coast which diverts deep contours (including the 2000 m reference) far away from the coast, but has no equivalent in shallow regions. In addition to these gaps, the density of points is strongly influenced by the steepness of the slope, with greater density along the northern boundary where the slope is gentle, and much more sparse coverage in regions of steep slopes. From this we can see that, even at 1/12° resolution, parts of the continental slope are barely resolved (not resolved, if we take the criterion of no more than one vertical grid step per horizontal step as the definition of resolved). Typical climate models will have much lower resolution than this, which calls into question how realistic their representation of boundary processes can be.

Returning to the bottom pressure standard deviation in Fig. 9, we see that a large fraction of the slope has a standard deviation below 2 mbar, as we suspected based on the map (Fig. 3). Variability is lower near the equator and on the eastern boundary (distances above about 4). There are also deep maxima in variability in the northern Labrador Sea (3.0–3.2) and south of Iceland (3.6–3.8), locations where deep water formation and mixing processes are likely to be important (these are large Rossby number processes, so not subject to damping according to the scaling argument above). Another deep maximum occurs near Cape Hatteras (2.4) where the deep western boundary current passes beneath the Gulf Stream.

We suspected that a significant portion of this variability would be due to large-scale barotropic processes. To remove these, we calculated a depth-averaged bottom pressure $\langle p \rangle_H$ at each distance and time, and subtracted this off. The residual standard deviation is shown in the middle panel of Fig. 9. It is clear that this mode accounts for a significant part of the variability (typically about half the variance).

After subtracting the depth average, $\langle p \rangle_H$ to look at higher vertical mode structures, we performed an Empirical Orthogonal Function (EOF) analysis of the time series at each distance, returning a set of basis functions in depth and time, for each distance. The middle and lower panels of Fig. 10 shows composites of the first two depth basis functions. These functions are calculated independently at each distance. The only processing to improve the appearance of the plot is the choice of arbitrary sign at each point to maintain continuity in distance.

It is striking how uniform the structure is in these EOFs. EOF1 has a single node at almost all distances, that node being between about 800 and 1500 m depth except where it is forced deeper by the absence of shallow data (together with the subtraction of the average over the available depth range). The increased deep variability in some locations has an influence on the EOF structure, but this does not seem to dominate.

In the case of EOF2, there are generally two nodes, though there is more variability in the depths of these except along the eastern boundary, where the structure is very uniform. The deep variability has more influence on EOF2 than on EOF1.

The depth average, $\langle p \rangle_H$, plus the first two EOFs typically explain about 90% of the total variance, though this can be as little as 50% in small regions. This is illustrated in the bottom panel of Fig. 9, which shows the standard deviation of the residual after subtracting the depth average and the first 2 EOFs. This is below 0.5 mbar in most of the basin, though larger in northern regions, near Cape Hatteras, and along parts of the South American coast, particularly those influenced by the Antarctic Circumpolar Current. In terms of the residual after subtracting the depth average, the first two EOFs typically explain over 80% of the variance. The depth average, $\langle p \rangle_H$, and first two EOFs together are therefore sufficient to describe most of what is seen along the continental slope.

5.2. Coherence along the slope

Each of the three modes ($\langle p \rangle_H$, EOF1, EOF2) has an associated time series at each distance. In Fig. 11, we show the cross-correlations between those time series at each distance with each other distance.

For the depth-independent mode, $\langle p \rangle_H$, this confirms that the

Fig. 9. Standard deviation of bottom pressure on the Atlantic continental slope as a function of distance and depth of the ocean floor (compare with Fig. 8). The top panel shows the total (after removing mean, trend, annual and semiannual cycles). The middle panel shows the residual after removing the depth average (cross-slope average, $\langle p \rangle_H$) at each distance and time. The bottom panel shows the residual after further removing the signal explained by the first two EOFs at each distance.

variability is highly coherent over the whole basin, with positive correlations almost everywhere. The correlations are for Fourier-filtered time series, showing periods shorter than 1.5 years below the diagonal, and periods between 1.5 and 10 years above. The correlations are stronger at the lower frequencies, except for those involving the strip at distances 0.0–0.29 which is on the Pacific side of South America, where long-period correlations drop off sharply, even becoming slightly negative, showing that there is a clear distinction between basins for this large-scale mode.

Perhaps more striking is the result for EOF1. At low frequencies (above the diagonal), this shows three large blocks of strong correlation, and one smaller one. The large blocks are at distances 0–1.4 (the east Pacific and southern hemisphere western boundary of the Atlantic to about the easternmost tip of Brazil), 2.4–3.8 (Cape Hatteras to eastern Iceland), and 4–6 (Scotland to South Africa, the entire eastern boundary). The scales of these regions imply strong correlations over distances of measured in tens of thousands of kilometres. In contrast to the depth-independent mode, $\langle p \rangle_H$, there is clear communication between the Pacific and Atlantic (distances 0–0.29 in the Pacific are correlated with the western South Atlantic distances 0.29–1.4), although the tip of South America provokes a drop in correlation.

The smaller block is from about 1.4 to 1.93, with a weaker extension to 2.35 (the north coast of South America and the eastern Caribbean, with extension to north Florida). In addition, there is an off-diagonal block showing correlation between the eastern and western boundary equator positions (marked with black lines), which shows strong correlations between the equator-spanning small block and the entire eastern boundary.

Fig. 10. Position of the points used in this analysis (top), and the structures of the first two EOFs of bottom pressure at each distance, after subtracting the depth average (cross-slope average, $\langle p \rangle_H$) (middle and bottom). EOFs are dimensionless and have a variance of 1.

The northern block (Cape Hatteras to Iceland) is split into two subregions separated at around the northern limit of the Labrador Sea (3.1). Similarly, the eastern block (Scotland to South Africa) is split into subregions at the Strait of Gibraltar (4.57), although in this case it is more like the start of a reduction in correlation to the north rather than the separation of two clear regions. Within the subregions, the correlations are especially strong, usually above 0.8 or 0.9. It is noticeable that the boundaries between blocks and sub-blocks often lie in the regions of large deep variability identified before.

At higher frequencies (below the diagonal), the northern block is similar if a bit weaker, but the other blocks are significantly weaker, especially the eastern boundary block, and the off-diagonal correlation between eastern and western equators is missing. However, the curious high-correlation feature extending down and to the right of the diagonal from the eastern equator gives a clue to the reason for this. This indicates that points on the eastern boundary correlate with one another if they are at equal distances from the equator. This suggests a signal propagating away from the equator sufficiently slowly that lags in signal propagation become important enough to reduce the correlation in the higher frequency band. Points at equal distances from the equator can then still correlate because they are at equal lags.

Assessment of statistical significance depends on the spectral content at each point, but the correlations within blocks are so strong as to be clearly significant with any reasonable estimate of degrees of freedom (only 14 degrees of freedom are needed for a correlation of 0.5 to be significant at the 95% level; we have 54 years of data and, as we shall see, the data are not dominated by the longest periods). We will, however, consider this in more detail when looking at lagged correlations later on.

For EOF2, we see a picture broadly similar to that for EOF1, but

Fig. 11. Cross correlations between the time series for each vertical (cross-slope) mode at each distance along the slope, with that at each other distance. Values above the diagonal are for periods between 1.5 and 10 years, and values below the diagonal are for periods shorter than 1.5 years. Linear trend, annual, and semiannual cycles were removed at an earlier stage. Black lines indicate the positions of the equator on the western and eastern boundaries.

Fig. 12. Hovmoeller plots showing the temporal variations in the three vertical (cross-slope) modes as a function of distance along the continental slope and time, after filtering to pass periods shorter than 1.5 years. Linear trend, annual, and semiannual cycles were removed at an earlier stage.

Fig. 13. Hovmoeller plots showing the temporal variations in the three vertical (cross-slope) modes as a function of distance along the continental slope and time, after filtering to pass periods between 1.5 and 10 years. Linear trend, annual, and semiannual cycles were removed at an earlier stage.

with more interruption by local noise. The same broad blocks are visible, though more weakly, and the same off-diagonal equatorial correlation and eastern boundary lagged correlation structure are also apparent.

Figs. 12–14 show Hovmoeller plots of the different modes in different frequency bands. The first, Fig. 12 shows only a representative 5.5-year period so that the structure can be clearly seen. Again, this highlights the basin-scale nature of the depth-independent mode, $\langle p \rangle_H$, with perhaps additional variability in the northern region (note that the boundary values are coherent over the entire basin, but this does not preclude anticorrelation with values in the interior). No lags can be discerned, which is to be expected as the largest-scale signals should propagate with a barotropic Kelvin wave speed of around 200 m s^{-1}, as was found in the western North Atlantic from bottom pressure measurements on the continental slope (Elipot et al., 2013).

More interesting are the plots for EOF1 and EOF2. Here, we can clearly see the effect of finite propagation speeds. Propagation is toward the western equator (equators are again marked as vertical lines), and away from the eastern equator. Wave speeds are slowest near the equator, and faster further away, but it must be remembered that the concept of along-slope distance depends on length scales and on how convoluted the 2000 m contour is. Various techniques have been applied in an attempt to estimate wave speeds, but we do not generally find robust numbers. We can see that speeds near the equator are between about 2 and 3 m s^{-1} for EOF1, and between about 1 and 1.5 m s^{-1} for EOF2. Speeds clearly get faster away from the equator, but meaningful quantification of these speeds is difficult. These speeds highlight the importance of the continental slope still further. They are roughly consistent with baroclinic Kelvin wave speeds at the equator: Brandt et al. (2016) calculate Kelvin wave speeds of 2.47 and 1.32 m s^{-1} for the first and second modes respectively. However, the Kelvin wave speed for a vertical boundary will become steadily slower as the waves propagate polewards, because the stratification weakens. The faster propagation speeds seen here are consistent with the fact that the true boundary waves on a slope are not Kelvin waves, but a hybrid mixture between baroclinic Kelvin waves and topographic Rossby waves, which become more barotropic and faster as the slope scale comes to exceed the Kelvin wave trapping scale, which is the relevant Rossby radius (Huthnance, 1978). For example, Cartwright et al. (1980) calculate a barotropic shelf wave speed of about 7.9 m s^{-1} off the Scottish coast (translating from the quoted wavenumber of 53 degrees per 100 km at the K_1 tidal period).

At longer periods (now showing the entire range of time), Fig. 13 shows that these wave propagation lags become insignificant, as lags along the eastern boundary can no longer be seen. For EOF1 there are hints of a slower equatorward propagation along the western boundary (particularly 2.5–3.1), but this is irregular in nature and too slow to be related to the boundary waves, it may be an advective phenomenon, if it is significant.

Another feature which becomes apparent at longer periods is the contrast in amplitudes between eastern and western boundaries. Signals are much weaker, and rather constant amplitude on the east, with a hint of amplification north of about the Strait of Gibraltar (4.57). This is consistent with idealised theories, e.g. Kawase (1987) and Johnson and

Fig. 14. Hovmoeller plots showing the temporal variations in the three vertical (cross-slope) modes as a function of distance along the continental slope and time, after filtering to pass periods longer than 10 years. Linear trend, annual, and semiannual cycles were removed at an earlier stage.

Marshall (2002a,b) which assume that low frequency signals propagate effectively instantly along eastern boundaries without attenuation or amplification. Amplification is expected poleward of the highest latitude at which Rossby waves can propagate for a given frequency. Marshall and Johnson (2013) give a nice overview of how theoretical models treat boundary waves (usually with a vertical sidewall), and of the importance of this process for a wide range of issues. On western boundaries, they are found (in theory) to decay in amplitude toward the equator, unlike our model results (except perhaps close to the equator), where we see an amplification in many places. There is clearly rapid transmission of information along the western boundary, but a more complex response than the simple, linear, vertical sidewall theory would predict. It should be noted though, that the theoretical results are for buoyancy forcing only. With wind stress variability too, angular momentum balance at each latitude requires that there be changing sidewall pressures as seen in the AMOC response to winds (Elipot et al., 2017). Here we have a rare case where diagnostics from a realistic model can address a problem of basin scale dynamics rather than being overwhelmed by mesoscale variability.

For completeness, Fig. 14 shows periods longer than 10 years. The contrast between east and west is now even more evident for EOF1, partial barriers to communication of western boundary signals appear at about Cape Hatteras (2.4) and the Gulf of Mexico (2.1). Some signal does appear to eventually propagate to the equator, from where it rapidly appears all along the eastern boundary. However, we caution against interpreting this plot in too much detail. Although periods as short as 10 years should be present, the variability is dominated by a roughly 50-year period, which is the longest possible within this dataset after detrending. While the model physics should still be consistent, the probability of being dominated by long-term thermohaline adjustment is high, and the effective number of degrees of freedom is so low that it is dangerous to conclude anything from apparent correlations. The large, quadratic-in-time variability in the region 2.1–3.0 seems to explain the unusual oscillation in the low frequency spectrum of bottom pressure in energetic continental slope regions, which can be seen in Fig. 7.

Returning to the high frequency data, we look in more detail at the long distance coherence of the signal when accounting for lags. We calculated lags relative to four different regions distinguished by colours in Fig. 15. Roughly, these regions are the eastern boundary (red), a northern region (orange), a north tropical western boundary region south of the Gulf of Mexico and north of the equator (green), and a southern hemisphere western boundary region (blue). The results of this analysis are shown in Fig. 15, which shows both correlations (Fig. 15a and c) and calculated lags (Fig. 15b and d. The precise extent of the reference regions is marked by the paired dots of matching colour at the top of Fig. 15a and c.

Absolute values of the lags are chosen to be consistent between the different curves as described below, so they can be interpreted as the time taken for a signal to propagate to each point from the eastern boundary equator (or, for negative values, minus the time taken to propagate from that point to the eastern boundary equator). The correlations are plotted as thin lines, with large dots on top only for points that are statistically significant. Lags are plotted only for these

Fig. 15. Best lagged correlations, and the lag at which the best correlation is found, for EOF1 and EOF2 time series compared with reference time series based on data from between the matching-coloured dots. Lags are shown only where the correlations are significant at the 99% level, and correlations are plotted with a thin line where they are not significant at this level. The time series are filtered to pass periods between 20 days and 1.5 years. Linear trend, annual, and semiannual cycles were removed at an earlier stage. See main text for more details. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

significant points.

The calculation was performed as follows. We started by selecting the expected coherent regions by eye, from the Hovmoeller plot, Fig. 12. The procedure worked on time series filtered to pass periods between 20 days and 1.5 years and normalised to a standard deviation of 1 (periods shorter than 20 days were removed in order to permit shifts of the time series of less than 5 days, using a Fourier method; normalisation avoids skewing the results to focus on regions of highest variability). For each region, an iterative method was then applied to determine the best lagged correlations. Initially, all the time series in the central 2000 km of that region were averaged together to provide a reference time series T_0. Then, all time series in the region were correlated with T_0, with a range of different lags, to identify the lag at which the correlation was best. Each time series was then shifted by this best lag, and the regional average (now over the entire region) was recalculated to obtain a new reference time series T_1. This process was iterated to convergence (four iterations was sufficient). The best lag, and the associated best correlation, were then calculated relative to the converged regional average time series T_4, for all points, not just those in the selected region.

Finally, a constant was added to the lags derived for each region. For the eastern region (red), the constant was chosen so that the lag is

Fig. 16. Cross correlations for sea level variability in the NEMO model, along a meridional section of the Atlantic at 23°W from 55°S to 64°N. Values above the diagonal use time series filtered to pass periods between 1.5 and 10 years, those below use 0–1.5 years. Linear trend, annual, and semiannual cycles were removed at an earlier stage. Black lines mark the equator.

zero at the eastern equator. As signals propagate along the eastern boundary away from the equator, the lag grows. This same growth is seen in lags with respect to the northern region (orange), so the constant for that region is chosen to align it with the eastern region lags. Similarly, where both are significant, the lags for the north tropical western boundary region (green) have very similar variations to the eastern region (red), so the constant lag for this region is found by aligning the two (tracing the signals back in time to the western boundary at the equator). For the northern boundary region (orange) we have two options, as the lags show it preceding the north tropical western boundary region (green), and lagging behind the eastern boundary region (red), so two different constant lags can be calculated, with a difference which is the time taken for a signal to perform a complete circuit. This region is therefore plotted with both lags (the eastern region is also repeated using the same offset). Finally, the constant lag for the south western boundary region (blue) is chosen to align it with the other lags at the western boundary equator. The whole graph can be considered to be periodic in y, with a period of 115 days for EOF1 and 205 days for EOF2.

It is worth noting that, for EOF1, significant correlations for the northern region (orange) overlap the significant lags from the north tropical western region (green) on both the eastern and western boundaries. On the east, the lags show signals propagating into the northern region. On the west, they show signals propagating away from the northern region. This means an estimate for the basin circuit time can be made purely from these two region, without needing the eastern region. The fact that all three lags agree in the eastern region is therefore an independent test of this circuit time. The circuit time estimate for EOF2 is more fragile as it relies purely on the small region of overlap between significant lags for the northern (orange) and north tropical western (green) regions, which also involve rather low correlations.

Significance of the correlations was assessed by a Monte Carlo method: 1000 time series with the same spectrum as each reference time series were generated. These were each correlated with the actual reference time series at all possible lags, and the correlation at the best lag chosen. These 1000 correlations were then sorted to determine the 99% confidence level (this ranged between 0.10 and 0.25 for the different time series).

What we find is that, when accounting for lags, the correlations are indeed coherent over very large distances, as they are at lower frequencies without accounting for the lags. Furthermore, the lagged correlations remain significant to some distance either side of the reference regions, with consistently varying lags where the curves overlap. For example, in EOF1, the eastern boundary signal (red) remains detectable as far as Greenland (3.3) and, though the correlation drops off rapidly, the north tropical western mode (green) also obtains a small part of its variance from the Greenland coast, as well as communicating with much of the eastern boundary. The south western boundary signal (blue) involves signals propagating from the Pacific, and extends to the Atlantic equator. For EOF2 we find a similar pattern, though correlations tend to drop off more rapidly, reducing the overlap of regions with significant correlations, and the corresponding lags are larger, corresponding to the slower propagation speeds. These lags are all consistent with expectations for the propagation of continental shelf waves and equatorial Kelvin waves.

The consistency of the lags allows us to estimate various propagation times. We find times to cross from the western boundary to the east at the equator of 28 and 40 days respectively for EOF1 and EOF2 (using slightly different regions or weightings varies these numbers by ± 1 day). Given a distance of approximately 5900 km, these times correspond to speeds of 2.44 and 1.71 m s^{-1}, comparable to the first and second baroclinic mode Kelvin wave speeds calculated by Brandt et al. (2016) which are 2.47 and 1.32 m s^{-1}. We can also determine the time to complete a full circuit of the North Atlantic, which is 115 days for EOF1 and 205 days for EOF2 (these estimates have larger uncertainties of about 5 days for EOF1 and 20 days for EOF2). The travel time along the continental slope, from eastern equator to western equator, is therefore 87 days for EOF1 and 165 days for EOF2.

Translating these delays to mean propagation speeds is not straightforward. The along-slope distance from eastern equator to western equator is 39,000 km, but the sum of great-circle distances between nodes over the same stretch of continental slope is only 28,133 km. Using these two extremes for the distance, we obtain a mean along-slope propagation speed for EOF1 of 3.7–5.2 m s^{-1}, and for EOF2 we obtain 2.0–2.7 m s^{-1}. Although the range is wide, the mean values are clearly faster than the corresponding equatorial Kelvin wave speeds.

These are important parameters for understanding timescales of basin adjustment, which could not realistically be computed otherwise. Again, the clean nature of the continental slope bottom pressure signal has made it possible to distinguish basin scale dynamics which could not be seen in other, mesoscale-contaminated signals. Note, though, that these numbers may depend on the model representation of topography and friction, so they may be different in the real world, and different again in coarser resolution models.

To conclude this subsection, it is worth emphasizing the exceptional nature of these long correlation scales. For comparison with Fig. 11, Fig. 16 shows cross-correlations for NEMO model sea level along a mid-Atlantic meridional section at 23°W. The distance is again measured in units of 10,000 km, so the scale is comparable to Fig. 11 (the whole distance scale is less than a quarter of the total scale for the continental slope). These correlations are typical of the open ocean, with some large-scale correlation apparent near the equator (black lines), but nothing approaching the clarity and large-scale nature of the boundary signals.

5.3. The meridional overturning circulation

The most obvious quantity to attempt to recover in the Atlantic is the meridional overturning circulation (AMOC), which has received a

great deal of attention as a significant climatic mode of the Earth system in recent decades. It is not our purpose here to revisit all this work. Bingham and Hughes (2008b) showed that much of the AMOC at 42°N could be recovered simply from boundary pressure measurements in a 1/4° resolution ocean model (and almost all of that from the western boundary only, at interannual periods), and Bingham and Hughes (2009a) looked in more detail at 50°N in the same model, and in 100 years of a coarser-resolution climate model run. These ideas have been used to measure the AMOC using the WAVE array at 43°N (Hughes et al., 2013; Elipot et al., 2013, 2014), and it has been shown how the measurements of the RAPID array at 26°N (McCarthy et al., 2012) can be interpreted in the same way (Elipot et al., 2017). The argument that eddies do not dominate the AMOC signal has also been made based on observations, and theory in a vertical sidewall context (Kanzow et al., 2009) (though note that this argument relates to the integrated transports, not to the pathways of flow and tracer transport, which are very strongly influenced by eddy variability).

Here, we have focused on the boundary signals themselves rather than their relevance to the AMOC, but it is still worth demonstrating their link to the AMOC explicitly. The theoretical argument is straightforward. Integration of geostrophic balance from west to east across the ocean basin at constant latitude and depth, leads to the following balance:

$$fT = p_E - p_W, \tag{9}$$

where T is the zonally integrated northward mass transport (in kilogrammes per metre of depth per second), and p_E and p_W are bottom pressures at the eastern and western end of the section respectively. Assuming a midlatitude value $f = 10^{-4}\,\mathrm{s}^{-1}$, and a density of about $1000\,\mathrm{kg\,m}^{-3}$, this leads to a 1 mbar pressure difference producing a net northward volume transport of $1\,\mathrm{Sv\,km}^{-1}$. It is therefore straightforward to see from Figs. 12–14 how the eastern and western boundary pressures contribute to net meridional transport.

The east-west symmetry of the depth-independent mode, $\langle p \rangle_H$, is testament to the fact that there can be rather little net northward transport across each latitude section (a net transport must be balanced by either an accumulation of mass, a flow through the Bering Strait, or a net gain or loss of mass from evaporation and precipitation, so it is tightly constrained by these integral properties). The AMOC represents a flow which is to the north at some depths and south at others, and is therefore reflected in the other EOF modes. From the vertical structures of the EOFs (Fig. 10), a positive value of EOF1 on the western boundary would correspond to a southward flow above about 1300 m, and a northward return flow below that depth; a negative AMOC anomaly. The same on the east would produce a positive AMOC anomaly.

Similarly, a positive value of EOF2 on the west would lead to a southward transport above about 500 m and at depth (typically below about 1.8 km but somewhat variable), and a northward transport at intermediate depths. Again, the same on the east would produce the opposite AMOC change.

We test that these relationships between boundary pressure and AMOC hold in the model by diagnosing the AMOC (in the sense of zonally-integrated meridional transport per unit depth) and comparing with that predicted based on the boundary pressures. This involves some complications because the pressure values at a given distance along the slope are not all at the same latitude. Equally, the model grid is not perfectly aligned with latitude lines north of about 20°N, though it remains within 1 degree of a constant in the Atlantic to about 55°N. Accordingly, we use the latitude associated with the reference 2000 m depth contour to define the latitudes of the pressure measurements, and average the AMOC in 1-degree bins for comparison, relying on the spatial coherence of the signals in order for the two datasets to match. The sign associated with the boundary pressure's contribution to the AMOC is determined by the direction of the contour at that point; it is negative where increasing distance moves north along the contour, and positive where it moves south. In this case we also use monthly means rather than 5-day means as an extra filter on ageostrophic high frequency variability.

Fig. 17 shows the resulting AMOC variability (left) in three different frequency bands, and the percentage of AMOC variance explained (right) by using only boundary pressure measurements. It should be recalled that these do not represent all the contributions to the zonal integral in this depth range. There are also contributions from the Mediterranean, Gulf of Mexico and Caribbean Sea which we are missing out, as well as any portions of the Mid-Atlantic Ridge which rise above 3200 m depth.

Nonetheless, with the exception of a region around the equator where the geostrophic argument breaks down, a large part of the AMOC variability is captured by these boundary measurements, and this proportion generally increases at longer timescales. The exception, at mid-depths for the longest periods, is because there is little variability to capture in this depth range. At these longest periods, the AMOC anomalies truly do seem to consist of a northward shallow flow and a southward return flow, with a node at intermediate depths, across the entire Atlantic basin.

In contrast, at the shortest periods, there is variability at all depths, and the amplitude rises toward the equator. This suggests a scaling proportional to $1/f$, and that the boundary signals may be best thought of as pressure signals at the boundary, with the AMOC as an incidental result of their presence. That is consistent with the independent propagation we see along eastern and western boundaries, with similar amplitudes.

At intermediate periods, the deep AMOC variability is small, but there is no systematic latitude dependence outside the tropics. This is the regime in which the eastern boundary signals are much smaller than those on the west but also extremely coherent in space (also true at longer periods).

The reconstruction has difficulties in the shallow North Atlantic, and particularly the Gulf of Mexico latitudes (about 20–30°N), but this is to be expected given the complicated geometry here, and even these regions improve at the longest timescales. Overall, at intermediate and long timescales, the unexplained residual standard deviation (not shown) is less than $0.3\,\mathrm{Sv\,km}^{-1}$ over most of this depth range, for latitudes more than 10 degrees from the equator.

In summary, the highly-correlated, basin-scale boundary pressure signals are indeed a good diagnostic of the AMOC.

6. Application to the real ocean

We have shown above that, in the NEMO model, continental slope bottom pressure measurements provide a means of extracting large-scale information about the ocean circulation, and the AMOC in particular. If this carries over to the real ocean, and if such measurements can be made with the necessary accuracy, then this would make a strong case for such measurements to be considered an important component of the Global Ocean Observing System.

The continental slope represents a very small fraction of the global ocean area and, as a result, there are rather few measurements made in this region. It is poorly monitored using Argo floats which rarely enter these regions and tend to spend very little time there when they do enter. The difficulty seems to be, in some ways, one of perception. Because the slope region is so small, it seems disproportionate to spend resources here which could be used to improve open-ocean sampling.

However, if we view the ocean in the latitude-depth plane, rather than latitude-longitude, the continental slope spans most of the active ocean depth range. This is the appropriate viewpoint for the meridional overturning circulation. Just by looking at the ocean "sideways", we obtain a very different viewpoint on where to expend resources.

6.1. Real ocean illustrations

First we need to show that this mesoscale suppression occurs in the

Fig. 17. Diagnostics of zonally-integrated northward water volume transport across the Atlantic (MOC) in the NEMO model as a function of depth and latitude, using monthly values after subtracting mean, linear trend, and annual and semiannual cycles. Left hand panels show the standard deviation in Sv km^{-1} (or, equivalently, $10^3\,m^2\,s^{-1}$). Right hand panels show the percentage of variance in the MOC which is explained by the boundary pressures using Eq. (9). Negative values are plotted as grey. The time series are divided into long periods (longer than 10 years, top), medium (1.5–10 years, middle), and short periods (2 months to 1.5 years, bottom).

real ocean and is not a model artefact. We showed in Section 3 that the model represents the mesoscale energy and its spectrum well. In fact, a number of experiments have demonstrated this at the Atlantic western boundary. In Fig. 18 we give an example from the continental slope near Halifax, Nova Scotia (distance 2.8 in the model slope diagnostics), from the RAPID-Scotian array (Hughes et al., 2013).

We see that, after 5-day averaging (blue), the variability over almost 4 years drops below 2 mbar except at the deepest point. Furthermore, subtracting an estimate of the depth average based on the four shallowest instruments, this drops below 1 mbar (below 0.6 mbar for the two central depths). Comparing with Fig. 9 at distance 2.8, this matches expectations both for amplitude and structure, with the amplitudes rising both in deeper and in shallower water. Our model continental slope diagnostics do not extend below 3200 m depth, but it can be seen from Fig. 3 that this amplitude increase continues at greater depths as a result of weakly-damped mesoscale variability at the foot of the continental slope.

Other sites show similar amplitudes, and a similar reduction on subtracting the depth average, making it possible to discern large-scale correlations and basin-scale responses to forcing (Elipot et al., 2013, 2014, 2017).

We have focused here on the continental slope, for monitoring ocean dynamics, but an even quieter region is the deep tropical ocean. Here, the dynamical signals are so small that bottom pressure can, in principle, be used to monitor changes in global ocean mass. In practice, this has so far been limited to the annual cycle (Hughes et al., 2012; Williams et al., 2014), which has been determined to an accuracy of ± 0.3 mbar, or about ± 1100 Gt of water, but the principle works well at long timescales too (Hughes et al., 2012). Thus, tropical ocean bottom pressure monitoring would provide a valuable contribution to the global sea level budget.

To some extent, the continental slope signals can be discerned in sea level measurements, again confirming the relevance of these arguments to the real ocean. Hughes and Meredith (2006) used satellite altimetry to show that signals with periods shorter than 1 year were correlated over very long distances around the global continental slope. One block of correlation they found stretched from Cape Hatteras right round the north of the Atlantic to tropical Africa (their Figs. 3 and 5), which suggests that the dominant signal seen was the depth-independent mode, consistent with the fact that no propagation lag could be seen over this distance.

The altimetry result highlights further the need for bottom pressure measurements: without them it is impossible to distinguish the different vertical modes, and the depth-independent mode is liable to dominate. Fig. 18 shows the importance of the depth-independent mode. Once this is subtracted, the vertical structures of the remaining observed variability are quite robust, with the first EOF being almost linear with depth (Elipot et al., 2013) apart from instruments on the gently-sloping tail of the continental slope, and having a quite different spectral structure from the depth-independent mode, with less variability at the shortest periods.

To the extent that it has been possible to test it, the real ocean does indeed seem to behave in a manner similar to the model.

6.2. Achieving the accuracy required

We are left with the question of how to make these measurements to the required accuracy. While bottom pressure recorders can be excellent for high-frequency measurements, they are prone to a long-period drift which, although it is well characterised as an exponential plus linear drift, can be very different from instrument to instrument and even between deployments of the same instrument (Watts and

Fig. 18. Standard deviations of bottom pressures as a function of depth, from 4 years of measurements at the RAPID-Scotian array near Halifax, Nova Scotia. Black shows the result from the raw 20-min sampled data after removal of tides and instrumental drift. Blue shows the variability based on 5-day means of the same data. Red is after subtracting the depth average based on the 4 shallowest points (those above 3200 m depth, as used in the model analysis), with pale red for 20-min sampling and dark for 5-day means. The dashed blue line shows the expected reduction resulting from this subtraction if the time series at each depth were uncorrelated. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Kontoyiannis, 1990; Polster et al., 2009). We must find an independent way to remove this drift. Several possibilities are available with present technology.

The most obvious technique is to use bottom pressures determined from satellite gravimetry. Indeed, some success has been claimed for this method by Landerer et al. (2015), who construct a time series of the AMOC from GRACE data. This is an interesting development, but should be treated with some caution. The patterns of bottom pressure seen by GRACE are limited to large horizontal scales, and the illustrations in Landerer et al. (2015) show a mode in which contrasting pressure anomalies appear over the entire shelf versus a wide region of the western abyssal plain. This is in contrast to what we see in ocean models where the relevant region is just the thin strip of the continental slope. Furthermore, at 26°N the inference is made based on broad pressure anomalies to east and west being representative of the deep branch of the AMOC, where they could equally well have been interpreted as representing the shallow branch, producing a time series with the opposite sign. It is possible that the broad-scale patterns seen by GRACE are indirectly reflecting the AMOC, especially if the variations are dominated by a large-scale response to wind stress variations over the time period considered, as suggested by Elipot et al. (2017). Or perhaps GRACE is seeing transient changes in broad water mass properties associated with AMOC transport anomalies. What is clear is that satellite gravimetry does not have the spatial resolution to distinguish between the shallow and deeper parts of the continental slope, especially in the presence of neighbouring larger signals both on the continental shelf and at the foot of the slope, so any such measurement relies on larger-scale correlations. While the GRACE results are very interesting, they cannot be used as a means to monitor the AMOC on longer timescales without detailed validation and understanding of how the measured signal relates to the AMOC and further understanding of the influence of plastic deformation of the earth on the observed signals. Similar arguments limit the use of upper-ocean and sea level measurements to derive the deep, large scale circulation. Although the influence of the AMOC has been noted (Bingham and Hughes, 2009b; Duchez et al., 2014; Goddard et al., 2015; McCarthy et al., 2015), it is an indirect link partially masked by other effects such as the influence of coastal winds (Woodworth et al., 2014, 2017).

Following satellite gravimetry, the most straightforward method to obtain bottom pressure at a point is to combine satellite altimetry with hydrographic measurements which allow the calculation of the steric sea level. The difference between the two is then a measure of ocean bottom pressure. Williams et al. (2015) investigated this approach using collocated tall moorings and bottom pressure recorders as part of the RAPID array at 26°N. They found that it was crucial to have information right to the surface, that very careful calibration of the moored instruments was required at the start and end of each deployment and that, even with ideal sampling and calibration, accuracy is limited to about 1–2 mbar in 5000 m depths (somewhat less in shallower water, but errors tend to be dominated by surface waters). This approach is therefore marginally feasible, but requires great care in its implementation. It may be the best approach for tropical measurements related to global ocean mass. Such a system should incorporate a good quality bottom pressure recorder as a check on system integrity, and to dramatically reduce errors at periods much shorter than the instrument deployment length.

The method pursued by the WAVE group (Hughes et al., 2013) focuses not on the pressures themselves, but on differences between pressures at different depths on the slope. Since the most interesting information is not in the depth-independent mode, $\langle p \rangle_H$, but in the depth-varying modes, this is sufficient to capture these modes. For a vertical sidewall, this would simply be a matter of measuring density at the boundary and using hydrostatic balance to compute the pressure differences. For a sloping wall, the horizontal pressure differences are also important (in fact dominant, on all but the steepest slopes, at intra-annual periods). The horizontal differences can be measured by measuring near-bottom currents and invoking geostrophic balance. The detailed theory, known as the Stepping Method, is given by Hughes et al. (2013), who applied this method to the first deployment of the RAPID-Scotia array, which was designed with this methodology in mind. Elipot et al. (2013) also applied the method at the predecessors to this array, for which the relevant data was available, but less well sampled. This method requires only bottom and near-bottom measurements, thus making the array more robust and simple to maintain than one which also includes moorings reaching to the surface.

We reproduce here (Fig. 19), Fig. 11 from Hughes et al. (2013), which shows how accurately the pressure difference budget can be closed with this measurement system. The residual errors are 0.16 mbar for the depth distance of 600 m between RS1 and RS2, 0.37 mbar for the 1675 m distance between RS4 and RS1, and 0.48 mbar for the 2300 m distance between RS5 and RS1. This method therefore meets the sub-millibar requirement for resolving the signals of interest. Note, though, that RS3 was excluded from this analysis because it showed multiple adjustment periods, and there may be other kinds of deviation from exponential plus linear drift in other cases.

Fig. 19 also demonstrates the effectiveness of removal of the high frequency tidal signals when data are sampled in such a way as to resolve the tides. This can become more of an issue with satellite data, either altimetry or gravimetry, for which a tidal model must be used and aliasing to long periods remains an issue. Long period tides remain in the data but their departure from equilibrium (the equilibrium tide can be calculated accurately from first principles) is expected to be limited to periods shorter than a few months, and to be predominantly geostrophic. In this sense, the non-equilibrium tides are simply part of the signal to be monitored. The largest long period tides have amplitudes of up to about 3 mbar (Egbert and Ray, 2003). It is, therefore, important at the 1 mbar level to calculate and subtract off the long period equilibrium tides, including the pole tide.

Fig. 19. Differences between directly-measured bottom pressure differences and those inferred from density and current measurement. Differences are between (RS2, RS4 RS5) and RS1 (see Fig. 18). The top panel shows the raw values, together with the exponential-plus-linear instrument drift that is fitted. The bottom panel shows the remainder after removing drift, together with its standard deviation in Pa (100 Pa = 1 hPa = 1 mbar). Reproduced from Hughes et al. (2013).

Despite the success of this method, it would be nice to have a more straightforward means of monitoring bottom pressures without the need for drift corrections. As Fig. 19 shows, the present technology works well apart from the exponential-plus-linear drift, so all that is needed is a small number of measurements without drift to correct this.

One possibility has been put forward by Sasagawa and Zumberge (2013), who use a known weight acting over a known area to produce a reference pressure at intermittent times. They claim a precision of 1.16 mbar per calibration point, which can reach our required precision if sufficient (order 10 or more) calibrations can be made per deployment. However, although a recent deployment of two instruments demonstrated a clear improvement of the drift, almost eliminating the exponential part and significantly reducing the linear trend, a remaining, mostly linear trend difference between the instruments remained (Sasagawa et al., 2016), amounting to about 5 mbar after 1.5 years. This is very good progress, but not quite at the required accuracy yet.

Another proposed system is the use of a reference pressure in a closed chamber, together with a differential pressure sensor (Gennerich and Villinger, 2015). The targeted trend accuracy is 0.1 mbar yr^{-1}, which would be a major step forward, but the instrument currently exists in concept only.

We would like to encourage these developments, and other initiatives to measure ocean bottom pressures. The required precision and stability is of order 10^{-7}–10^{-8} over a year or longer, if we are aiming at 0.1 mbar (1 mm) in depths of 1000–5000 m. Perhaps a method using fluids rather than crystals would avoid the unknown source of the drift in the quartz crystal sensors presently used? Perhaps some optical technique could provide a better solution? Although the technology is improving, we are still using essentially the system that was developed in the 1970s, it may be time to explore new options.

Finally, if we did have a very stable pressure sensor, this would then bring into focus the fact that the seafloor itself is not perfectly stable. Secular vertical land motions of order 1 mm yr^{-1} are a very familiar

feature to those working with tide gauges, and much larger values can be found in tectonically-active areas (indeed, the aim of several bottom pressure measurement experiments has been to measure motion of the seafloor, not ocean dynamics). There is no equivalent of Global Navigation Satellite Systems like GPS to monitor the seafloor, so for the present we would have to rely on placement of instruments in stable positions, and/or accurate models of vertical land movement. In fact, it is not the geometrical motion that is relevant for ocean dynamics, but the motion relative to geopotential surfaces: what we need to know is what level the instrument is at, which involves knowing the Earth's gravity field. Satellite data can help with this, but an intriguing possibility for the future may be the use of highly accurate optical clocks. These are now approaching an accuracy of 10^{-18} (Margolis, 2014). General relativity predicts that a change in geopotential height causes a clock to run faster if it is raised through a height δh by a factor of approximately $g\delta h/c^2$, where c is the speed of light. The factor g/c^2 is about 10^{-16}, meaning that a clock accuracy of 10^{-18} would make it possible to measure level differences of 1 cm. The technological challenges may be formidable, but there is nothing which cannot be overcome in principle.

The vertical land movement problem is only an issue for direct measurement of bottom pressure (or satellite gravity measurements, where it becomes a rather different issue). For the indirect methods, either altimetry plus a tall mooring or the Stepping Method, precise positioning of the instruments in the ocean is not a major issue and it suffices to know depths to a few metres.

6.3. Final thoughts

To summarise, continental slope bottom pressure measurements can make an extremely valuable contribution to a Global Ocean Observing System, monitoring basin-scale dynamics uncontaminated by the mesoscale variability which dominates most other measurements, and allowing for measurement of the AMOC. This is possible with current technology if we combine bottom pressure recorders with near-bottom density and current measurements, and are willing to sacrifice the less dynamically-interesting depth-independent mode (though a combination of altimetry and a full-depth mooring, plus bottom pressure recorders, could also provide this mode to about 1–2 mbar accuracy, permitting monitoring of variations in the Arctic throughflow plus precipitation minus evaporation, though only with an accuracy of a few sverdrups). Future developments may make direct measurements of bottom pressure more accessible, and these should be encouraged, but an important start can be made with the present technology. Though the accuracies required may seem daunting, these are simply a reflection of assuming a 1 sverdrup target accuracy, together with an equivalence of 1 mbar to 1 sverdrup per kilometre at mid latitudes (the pressure signals reduce, and hence the required accuracy increases, at low latitudes). Any system which aims for 1 Sv transport accuracy is implicitly determining bottom pressure at the 1 mbar level. It would seem sensible to make that accuracy an explicit aim, to be sure we are not fooling ourselves when we consider how well the ocean state is monitored.

Acknowledgments

We would like to thank Rory Bingham, Shane Elipot, Miguel Ángel Morales Maqueda, Mark Tamisiea, Mike Meredith, John Loder, Blair Greenan, John Toole, and the members of the NOC, BIO and WHOI seagoing and engineering teams who played a large role in the RAPID-WAVE and Weighing the Oceans projects and related discussions which form the background to this paper. CWH and JW were supported by NERC National Capability funding and NERC grant NE/I023384/1. ACC would like to acknowledge the support of NERC National Capability funding for the development and production of the N006 simulation, which was conducted on the ARCHER UK National Supercomputing Service (http://www.archer.ac.uk). ATB and JW would like to acknowledge the support of the European Union's Horizon 2020 research and innovation program under grant agreement No. 633211 (AtlantOS). We would also like to thank the two reviewers for their careful reading and constructive comments. The research materials supporting this publication can be accessed by contacting the corresponding author.

Appendix A. Interpreting the colour spectrum plots

The colour spectrum plots in Figs. 1–5 are designed to give a qualitative representation of the different spectra at each point in the ocean, in a manner that can be naturally interpreted. We do this by exploiting the way the eye interprets spectra of light.

In a person with normal vision, three numbers are sufficient to describe any colour including its brightness, because the eye has three different colour-sensitive receptors with different ranges of sensitivity. Thus, the eye is only sensitive to these three different weighted averages of the visible spectrum. A particular perceived colour can be produced from many different spectra. For example, the "sodium orange" of many old street lamps is produced by a spectrum dominated by a very narrow band at about 590 nm wavelength, but the same perceived colour can be produced by a combination of red and yellow light.

The details of how we calculate colours from a spectrum are given in the appendix to Hughes and Williams (2010), and we use exactly the same parameters here except that we divide by a ten times smaller normalizing factor of 1.6×10^5 in the bottom pressure plots, in order to make them brighter (note however that there is an error in that appendix: the matrices given as A4 and A5 are both transposed). In this appendix we describe in more detail how to interpret the colours and the scales attached to the figures in this paper.

The main thing to note is that, just as with light, there is no way to infer the spectrum from the colour, since many different spectra are consistent with the three weighted averages which determine a given colour. Only the forward calculation can be done, to find the colour given a particular spectrum. This is how we produce the different colour bars on the plots in Figs. 1–5. A particular shape of spectrum is chosen, and various parameters changed, with the resulting colour plotted as a function of those varying parameters.

In the first three colour bars, the chosen spectrum is a Gaussian as a function of period (i.e. a Gaussian Spectral Power Distribution, as described in the appendix of Hughes and Williams (2010)). The two varied parameters are the period of the peak (x-axis), and the overall brightness (y-axis). Thus, the first colour bar shows the intense colours which result from a sharply-peaked spectrum peaking at different periods, the second shows the same for a broader Gaussian, and the third for a still broader Gaussian (the white lines show the actual Gaussian for a particular period).

The fourth colour bar shows the colours for particular power laws, i.e. spectra in which power is proportional to frequency raised to powers in the range -4 to 0. Again, the y-axis represents varying overall brightness. In this case the colours range smoothly from orange-red, through grey-white near a power of -2, to blue at power zero. Note that we do not generally expect steep power laws (-2 or below) to persist to the lowest measured frequencies (there are great difficulties with calculating spectra if they do). However, the mapping to colour is only sensitive to our chosen range of periods from 2 to 24 weeks, so the power law only need extend over this range to dominate the colour. In an unfortunate clash of terminology, it turns out that broadband white light has what is usually referred to as a "red" spectrum, with power spectral density proportional to frequency raised to the power -2. As a result, such "red" spectra appear in these spectral colour plots as a grey scale (the Wortham and Wunsch (2014) spectrum is "red" in the 2–24 week period range, and hence appears as grey). Similarly, a "white" spectrum (power spectral density independent of

frequency, index 0) appears as an intense blue.

A third spectral shape was used in the generation of Fig. 1b of Hughes and Williams (2010). In that case the spectra were modelled as two power laws: zero for low frequencies, and -4 for high frequencies, with the varying parameter being the period at which the switch between power laws occurs. This was designed to represent an extreme version of sea level spectra like those in Fig. 6, many of which show very steep power laws at high frequency and much gentler power laws at low frequency, with a sharp transition at a period related to the shortest baroclinic Rossby wave period. The range of colours produced with that spectral shape was very similar to that in our second colour bar, for periods shorter than 100 days (i.e. the more intense reds were absent, as they are in the power law colour bar).

Looking at the relationship between spectrum and colour, we can draw some general conclusions. Pale–mid-blue and orange-red are quite common in smooth spectra, and generally represent more or less power in the high frequencies than would be found with a -2 power law (again, over the 2–24 week period range). Other colours, and more intense reds and blues, require more structure in the spectrum and tend to reflect sharp changes in gradient or peaks near a particular period, as shown in the first colour bar. Furthermore, there are colours seen in Figs. 1–5 which are not seen in any of our colour bars, especially pinks and light purples. These may be suggestive of multiple peaks or more complex spectral shapes.

As the above shows, there is no simple, intuitive way to uniquely translate a colour to a spectrum, but there are general guiding principles which the colour bars help with. In many ways, a better way to use these plots is to treat them like false colour astronomical photographs spanning a wide range of the electromagnetic spectrum: different processes result in different colours, and those colours add in a way we are used to interpreting with visible light, but the colours are only qualitative pointers to processes. Further detailed investigation into the actual spectra and time series is needed once regions with a particular colour have been identified.

After calculating these diagnostics for a number of variables and different models, we have found them to be a good way of getting a quick overview of model performance in the mesoscale band. Similar features are often found in different models, and different features stand out strongly and invite investigation.

References

Bingham, R.J., Hughes, C.W., 2008a. The relationship between sea-level and bottom pressure variability in an eddy permitting ocean model. Geophys. Res. Lett. 35, L03602. http://dx.doi.org/10.1029/2007GL032662.

Bingham, R.J., Hughes, C.W., 2008b. Determining North Atlantic meridional transport variability from pressure on the western boundary: a model investigation. J. Geophys. Res. (Oceans) 113, C09008. http://dx.doi.org/10.1029/2007JC004679.

Bingham, R.J., Hughes, C.W., 2009a. Geostrophic dynamics of meridional transport variability in the subpolar North Atlantic. J. Geophys. Res. (Oceans) 114, C12029. http://dx.doi.org/10.1029/2009JC005492.

Bingham, R.J., Hughes, C.W., 2009b. Signature of the Atlantic meridional overturning circulation in sea level along the east coast of North America. Geophys. Res. Lett. 36, L02603. http://dx.doi.org/10.1029/2008GL036215.

Brandt, P., Claus, M., Greatbatch, R.J., Kopte, R., Toole, J.M., Johns, W.E., Boning, C.W., 2016. Annual and semiannual cycle of equatorial Atlantic circulation associated with basin-mode resonance. J. Phys. Oceanogr. 46, 3011–3029. http://dx.doi.org/10.1175/JPO-D-15-0248.1.

Brodeau, L., Barnier, B., Treguier, A.M., Penduff, T., Gulev, S., 2010. An ERA40-based atmospheric forcing for global ocean circulation models. Ocean Model. 31, 88–104. http://dx.doi.org/10.1016/j.ocemod.2009.10.005.

Cartwright, D.E., Huthnancem, J.M., Spencer, R., Vassie, J.M., 1980. On the St Kilda shelf tidal regime. Deep-Sea Res. 27A, 61–70. http://dx.doi.org/10.1016/0198-01409(80)90072-2.

Chao, Y., Fu, L.-L., 1995. A comparison between the TOPEX/POSEIDON data and a global ocean general circulation model during 1992–1993. J. Geophys. Res. 100 (C12), 24,965–24,976. http://dx.doi.org/10.1029/95JC02260.

Chelton, D.B., de Szoeke, R.A., Schlax, M.G., 1998. Geographical variability of the first baroclinic Rossby radius of deformation. J. Phys. Oceanogr. 28, 433–460. http://dx.doi.org/10.1175/1520-0485(1998)028<0433:GVOTFB>2.0.CO;2.

Duchez, A., Hirschi, J.J.-M., Cunningham, S.A., Blaker, A.T., Bryden, H.L., de Cuevas, B.A., Atkinson, C.P., McCarthy, G.D., Frajka-Williams, E., Rayner, D., Smeed, D., 2014. A new index for the Atlantic meridional overturning circulation at 26°N. J. Clim. 27, 6439–6455. http://dx.doi.org/10.1175/JCLI-D-13-00052.1.

Dieng, H.B., Palanisamy, H., Cazenave, A., Meyssignac, B., von Schuckmann, K., 2015. The sea level budget since 2003: inference on the deep ocean heat content. Surv. Geophys. 36, 209–229. http://dx.doi.org/10.1007/s10712-015-9314-6.

Egbert, G.D., Ray, R.D., 2003. Deviation of long-period tides from equilibrium: kinematics and geostrophy. J. Phys. Oceanogr. 33, 822–839. http://dx.doi.org/10.1175/1520-0485(2003)33<822:DOLTFE>2.0.CO;2.

Dussin, R., Barnier, B., Brodeau, L., 2014. The Making of Drakkar Forcing Set DFS5. DRAKKAR/MyOcean Rep. 05-10-14, LGGE, Grenoble, France.

Elipot, S., Hughes, C., Olhede, S., Toole, J., 2013. Coherence of western boundary pressure at the RAPID WAVE array: boundary wave adjustments or deep western boundary current advection? J. Phys. Oceanogr. 43, 744–765. http://dx.doi.org/10.1175/JPO-D-12-067.1.

Elipot, S., Frajka-Williams, E., Hughes, C.W., Willis, J.K., 2014. The observed North Atlantic meridional overturning circulation: its meridional coherence and ocean bottom pressure. J. Phys. Oceanogr. 44, 517–537. http://dx.doi.org/10.1175/JPO-D-13-026.1.

Elipot, S., Frajka-Williams, E., Hughes, C.W., Olhede, S., Lankhorst, M., 2017. Observed basin-scale response of the North Atlantic meridional overturning circulation to wind stress forcing. J. Clim. 47, 2029–2054. http://dx.doi.org/10.1175/JCLI-D-16-0664.1.

Farrar, J.T., 2011. Barotropic Rossby waves radiating from tropical instability waves in the Pacific Ocean. J. Phys. Oceanogr. 41, 1160–1181. http://dx.doi.org/10.1175/2011JPO4547.1.

Ferrari, R., Wunsch, C., 2009. Ocean circulation kinetic energy: reservoirs, sources, and sinks. Annu. Rev. Fluid Mech. 41, 253–282. http://dx.doi.org/10.1146/annurev.fluid.40.111406.102139.

Fu, L.-L., 2007. Interaction of mesoscale variability with large-scale waves in the Argentine Basin. J. Phys. Oceanogr. 37, 787–793. http://dx.doi.org/10.1175/JPO2991.1.

Fu, L.-L., Davidson, R.A., 1995. A note on the barotropic response of sea level to time-dependent wind forcing. J. Geophys. Res. 100 (C12), 24,955–24,963. http://dx.doi.org/10.1029/95JC02259.

Fu, L.-L., Cheng, B., Qiu, B., 2001. 25-day period large-scale oscillations in the Argentine Basin revealed by the TOPEX/Poseidon altimeter. J. Phys. Oceanogr. 31, 506–517. http://dx.doi.org/10.1175/1520-0485(2001)031<0506:DPLSOI>2.0.CO;2.

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. Prog. Oceanogr. 134, 152–172. http://dx.doi.org/10.1016/j.pocean.2015.01.013.

Gennerich, H.-H., Villinger, H., 2015. A new concept for an ocean bottom pressure meter capable of precision long-term monitoring in marine geodesy and oceanography. Earth Space Sci. 2, 181–186. http://dx.doi.org/10.1002/2014EA000053.

Goddard, P.B., Yin, J., Griffies, S.M., Zhang, S., 2015. An extreme event of sea-level rise along the Northeast coast of North America in 2009–2010. Nat. Commun. 6, 6346. http://dx.doi.org/10.1038/ncomms7346.

Greatbatch, R.J., 1994. A note on the representation of steric sea level in models that conserve volume rather than mass. J. Geophys. Res. (Oceans) 99 (C6), 12,767–12,771. http://dx.doi.org/10.1029/94JC00847.

Hogg, N.G., 2000. Low-frequency variability on the western flanks of the Grand Banks. J. Mar. Res. 58, 523–545. http://dx.doi.org/10.1357/002224000321511007.

Hsu, C.-W., Velicogna, I., 2017. Detection of sea level fingerprints derived from GRACE gravity data. Geophys. Res. Lett. 44, 8953–8961. http://dx.doi.org/10.1002/2017GL074070.

Hughes, C.W., Meredith, M.P., 2006. Coherent sea-level fluctuations along the global continental slope. Philos. Trans. Roy. Soc. A 364, 885–901. http://dx.doi.org/10.1098/rsta.2006.1744.

Hughes, C.W., Stepanov, V.N., 2004. Ocean dynamics associated with rapid J_2 fluctuations: importance of circumpolar modes and identification of a coherent Arctic mode. J. Geophys. Res. (Oceans) 109, C06002. http://dx.doi.org/10.1029/2003JC002176.

Hughes, C.W., Williams, S.D.P., 2010. The color of sea level: importance of spatial variations in spectral shape for assessing the significance of trends. J. Geophys. Res. (Oceans) 115, C10048. http://dx.doi.org/10.1029/2010JC006102.

Hughes, C.W., Stepanov, V.N., Fu, L.-L., Barnier, B., Hargreaves, G.W., 2007. Three forms of variability in Argentine Basin ocean bottom pressure. J. Geophys. Res. (Oceans) 112 (C1), C01011. http://dx.doi.org/10.1029/2006JC003679.

Hughes, C.W., Tamisiea, M.E., Bingham, R.J., Williams, J., 2012. Weighing the ocean: using a single mooring to measure changes in the mass of the ocean. Geophys. Res. Lett. 39, L17602. http://dx.doi.org/10.1029/2012GL052935.

Hughes, C.W., Elipot, S., Morales Maqueda, M.Á., 2013. Test of a method for monitoring the geostrophic meridional overturning circulation using only boundary measurements. J. Atmos. Ocean. Technol. 30, 789–809. http://dx.doi.org/10.1175/JTECH-D-12-00149.1.

Hughes, C.W., Williams, J., Hibbert, A., Boening, C., Oram, J., 2016. A Rossby whistle: a resonant basin mode observed in the Caribbean Sea. Geophys. Res. Lett. 43, 7036–7043. http://dx.doi.org/10.1002/2016GL069573.

Huthnance, J.M., 1978. On coastal trapped waves: analysis and numerical calculation by inverse iteration. J. Phys. Oceanogr. 8, 74–92. http://dx.doi.org/10.1175/1520-0485(1978)008<0074:OCTWAA>2.0.CO;2.

Johnson, H.L., Marshall, D.P., 2002a. Localization of abrupt change in the North Atlantic thermohaline circulation. Geophys. Res. Lett. 29 (6). http://dx.doi.org/10.1029/2001GL014140.

Johnson, H.L., Marshall, D.P., 2002b. A theory for the surface Atlantic response to thermohaline variability. J. Phys. Oceanogr. 32, 1121–1132. http://dx.doi.org/10.1175/1520-0485(2002)032<1121:ATFTSA>2.0.CO;2.

Jourdan, D., Balopoulos, E., Garcia-Fernandez, M., Maillard, C., 1998. Objective analysis of temperature and salinity historical data set over the mediterranean basin. In:

OCEANS '98 Conference Proceedings, vol. 1. IEEE, Nice, France, pp. 82–87. http://dx.doi.org/10.1109/OCEANS.1998.725649.
Kanzow, T., Johnson, H.L., Marshall, D.P., Cunningham, S.A., Hirschi, J.J.M., Mujahid, A., Bryden, H.L., Johns, W.E., 2009. Basinwide integrated volume transports in an eddy-filled ocean. J. Phys. Oceanogr. 39, 3091–3109. http://dx.doi.org/10.1175/2009JPO4185.1.
Kawase, M., 1987. Establishment of deep ocean circulation driven by deep-water production. J. Phys. Oceanogr. 17, 2294–2317. http://dx.doi.org/10.1175/1520-0485(1987)017<2294:EODOCD>2.0.CO;2.
Landerer, F.W., Weise, D.N., Bentel, K., Boening, C., Watkins, M.M., 2015. North Atlantic meridional overturning circulation variations from GRACE ocean bottom pressure anomalies. Geophys. Res. Lett. 42, 8114–8121. http://dx.doi.org/10.1002/2015GL065730.
Legeckis, R., 1977. Long waves in the eastern equatorial Pacific Ocean: a view from a geostationary satellite. Science 197, 1179–1181. http://dx.doi.org/10.1126/science.197.4309.1179.
Levitus, S., Conkright, M., Boyer, T.P., OBrian, T., Antonov, J., Stephens, C., Johnson, L.S. D., Gelfeld, R., 1998. World Ocean Database 1998. Tech. Rep. NOAA Atlas NESDIS 18, U.S. Government Printing Office, Washington, D.C.
Lin, X., Yang, J., Wu, D., Zhai, P., 2008. Explaining the global distribution of peak-spectrum variability of sea surface height. Geophys. Res. Lett. 35, L14602. http://dx.doi.org/10.1029/2008GL034312.
Margolis, H., 2014. Timekeepers of the future. Nat. Phys. 10, 82–83. http://dx.doi.org/10.1038/nphys2834.
Marshall, D.P., 2011. Rossby wormholes. J. Mar. Res. 69, 309–330.
Marshall, D.P., Johnson, H.L., 2013. Propagation of meridional circulation anomalies along western and eastern boundaries. J. Phys. Oceanogr. 43, 2699–2717. http://dx.doi.org/10.1175/JPO-D-13-0134.1.
McCarthy, G., Frajka-Williams, E., Johns, W.E., Barringer, M.O., Meinen, C.S., Bryden, H.L., Rayner, D., Duchez, A., Roberts, C., Cunningham, S.A., 2012. Observed interannual variability of the Atlantic meridional overturning circulation at 26.5°N. Geophys. Res. Lett. 39, L19609. http://dx.doi.org/10.1029/2012GL052933.
McCarthy, G., Haigh, I.D., Hirschi, J.J.-M., Grist, J.P., Smeed, D.A., 2015. Ocean impact on decadal Atlantic climate variability revealed by sea-level observations. Nature 521, 508–510. http://dx.doi.org/10.1038/nature14491.
Moat, B.I., Josey, S.A., Sinha, B., Blaker, A.T., Smeed, D.A., McCarthy, G.D., Johns, W.E., Hirschi, J.J.-M., Frajka-Williams, E., Rayner, D., Duchez, A., Coward, A.C., 2016. Major variations in subtropical North Atlantic heat transport at short (5 day) timescales and their causes. J. Geophys. Res. (Oceans) 121, 3237–3249. http://dx.doi.org/10.1002/2016JC011660.
Polster, A., Fabian, M., Villinger, H., 2009. Effective resolution and drift of Paroscientific pressure sensors derived from long-term seafloor measurements. Geochem. Geophys. Geosyst. 10, Q08008. http://dx.doi.org/10.1029/2009GC002532.
Sasagawa, G., Zumberge, M.A., 2013. A self-calibrating pressure recorder for detecting seafloor height change. IEEE J. Ocean. Eng. 38, 447–454. http://dx.doi.org/10.1109/JOE.2012.2233312.
Sasagawa, G., Cool, M.J., Zumberge, M.A., 2016. Drift-corrected seafloor pressure observations of vertical deformation at Axial Seamount. Earth Space Sci. 3, 381–385. http://dx.doi.org/10.1002/2016EA000190.
Steele, M., Morley, R., Ermold, W., 2001. PHC: a global ocean hydrography with a high quality Arctic Ocean. J. Clim. 14, 2079–2087. http://dx.doi.org/10.1175/1520-0442(2001)014<2079:PAGOHW>2.0.CO;2.
Stepanov, V., Hughes, C.W., 2004. Parameterization of ocean self-attraction and loading in numerical models of the ocean circulation. J. Geophys. Res. (Oceans) 109, C03037. http://dx.doi.org/10.1029/2003JC002034.
Stepanov, V., Hughes, C.W., 2006. Propagation of signals in basin-scale ocean bottom pressure from a barotropic model. J. Geophys. Res. (Oceans) 111, C12002. http://dx.doi.org/10.1029/2005JC003450.
Timmerman, A., Goosse, H., Madec, G., Fichefet, T., Ethe, C., Dulire, V., 2005. On the representation of high latitude processes in the ORCA-LIM global coupled sea-ice ocean model. Ocean Model. 8, 175–201. http://dx.doi.org/10.1016/j.ocemod.2003.12.009.
Uppala, S.M., Kållberg, P.W., Simmons, A.J., et al., 2005. The ERA-40 re-analysis. Quart. J. Roy. Meteorol. Soc. 131, 2961–3012. http://dx.doi.org/10.1256/qj.04.176.
U.S. Department of Commerce, 2006. U.S. Department of Commerce. National Oceanic and Atmospheric Administration, National Geophysical Data Center: 2-minute Gridded Global Relief Data (ETOPO2v2). Available at < http://www.ngdc.noaa.gov/mgg/global/etopo2.html > .
Vignudelli, S., Kostianoy, A., Cipollini, P., Benveniste, J. (Eds.), 2011. Coastal Altimetry. Springer-Verlag, Berlin Heidelberg, pp. 565. http://dx.doi.org/10.1007/978-3-642-12796-0.
Vinogradova, N.T., Ponte, R., Stammer, D., 2007. Relation between sea level and bottom pressure and the vertical dependence of oceanic variability. Geophys. Res. Lett. 34, L03608. http://dx.doi.org/10.1029/2006GL028588.
Wahr, J., Nerem, R.S., Bettadpur, S.V., 2015. The pole tide and its effect on GRACE time-variable gravity measurements: implications for estimates of surface mass variations. J. Geophys. Res. Solid Earth 120, 4597–4615. http://dx.doi.org/10.1002/2015JB011986.
Watts, D.R., Kontoyiannis, H., 1990. Deep-ocean bottom pressure measurement: drift removal and performance. J. Atmos. Ocean. Technol. 7, 296–306. http://dx.doi.org/10.1175/1520-0426(1990)007<0296:DOBPMD>2.0.CO;2.
Weijer, W., Vivier, F., Gille, S.T., Dijkstra, H.A., 2007. Multiple oscillatory modes of the Argentine Basin. Part I: Statistical analysis. J. Phys. Oceanogr. 37, 2855–2868. http://dx.doi.org/10.1175/2007JPO3527.1.
Williams, Joanne, Hughes, C.W., 2013. The coherence of small island sea level with the wider ocean: a model study. Ocean Sci. 9, 111–119. http://dx.doi.org/10.5194/os-9-111-2013.
Williams, J., Hughes, C.W., Tamisiea, M.E., Williams, S.D.P., 2014. Weighing the ocean with bottom-pressure sensors: robustness of the ocean mass annual cycle estimate. Ocean Sci. 10, 701–718. http://dx.doi.org/10.5194/os-10-701-2014.
Williams, J., Hughes, C.W., Tamisiea, M.E., 2015. Detecting trends in bottom pressure measured using a tall mooring and altimetry. J. Geophys. Res. (Oceans) 120, 5216–5232. http://dx.doi.org/10.1002/2015JC010955.
Woodworth, P.L., Morales Maqueda, M.Á., Roussenov, V.M., Williams, R.G., Hughes, C.W., 2014. Mean sea-level variability along the northeast American Atlantic coast and the roles of the wind and the overturning circulation. J. Geophys. Res. (Oceans) 119, 8916–8935. http://dx.doi.org/10.1002/2014JC010520.
Woodworth, P.L., Morales Maqueda, M.Á., Gehrels, W.R., Roussenov, V.M., Williams, R.G., Hughes, C.W., 2017. Variations in the difference between mean sea level measured either side of Cape Hatteras and their relation to the North Atlantic Oscillation. Clim. Dyn. 49, 2451–2469. http://dx.doi.org/10.1007/s00382-016-3464-1.
Wortham, C., Wunsch, C., 2014. A multidimensional spectral description of ocean variability. J. Phys. Oceanogr. 44, 944–966. http://dx.doi.org/10.1175/JPO-D-13-0113.1.
Wunsch, C., 2008. Mass and volume transport variability in an eddy-filled ocean. Nat. Geosci. 1 (3), 165–168. http://dx.doi.org/10.1038/ngeo126.
Wunsch, C., 2016. Global ocean integrals and means, with trend implications. Ann. Rev. Mar. Sci. 8, 1–33. http://dx.doi.org/10.1146/annurev-marine-122414-034040.
Zhai, X., Johnson, H.L., Marshall, D.P., 2010. Significant sink of ocean-eddy energy near western boundaries. Nat. Geosci. 3, 608–612. http://dx.doi.org/10.1038/NGEO943.