Improved Climate Simulations through a Stochastic Parameterization of Ocean Eddies

Paul D. Williams

Department of Meteorology, University of Reading, Reading, United Kingdom

Nicola J. Howe

Risk Management Solutions, London, United Kingdom

Jonathan M. Gregory

National Centre for Atmospheric Science, University of Reading, Reading, and Met Office Hadley Centre, Exeter, United Kingdom

Robin S. Smith

National Centre for Atmospheric Science, University of Reading, Reading, United Kingdom

Manoj M. Joshi

Centre for Ocean and Atmospheric Sciences, University of East Anglia, Norwich, United Kingdom

(Manuscript received 23 October 2015, in final form 12 August 2016)

ABSTRACT

In climate simulations, the impacts of the subgrid scales on the resolved scales are conventionally represented using deterministic closure schemes, which assume that the impacts are uniquely determined by the resolved scales. Stochastic parameterization relaxes this assumption, by sampling the subgrid variability in a computationally inexpensive manner. This study shows that the simulated climatological state of the ocean is improved in many respects by implementing a simple stochastic parameterization of ocean eddies into a coupled atmosphere–ocean general circulation model. Simulations from a high-resolution, eddy-permitting ocean model are used to calculate the eddy statistics needed to inject realistic stochastic noise into a low-resolution, non-eddy-permitting version of the same model. A suite of four stochastic experiments is then run to test the sensitivity of the simulated climate to the noise definition by varying the noise amplitude and decorrelation time within reasonable limits. The addition of zero-mean noise to the ocean temperature tendency is found to have a nonzero effect on the mean climate. Specifically, in terms of the ocean temperature and salinity fields both at the surface and at depth, the noise reduces many of the biases in the low-resolution model and causes it to more closely resemble the high-resolution model. The variability of the strength of the global ocean thermohaline circulation is also improved. It is concluded that stochastic ocean perturbations can yield reductions in climate model error that are comparable to those obtained by refining the resolution, but without the increased computational cost. Therefore, stochastic parameterizations of ocean eddies have the potential to significantly improve climate simulations.

Corresponding author address: Paul D. Williams, Department of Meteorology, University of Reading, Earley Gate, Reading RG6 6BB, United Kingdom.
E-mail: p.d.williams@reading.ac.uk

1. Introduction

Numerical models of the atmosphere, ocean, and cryosphere represent physical processes in one of two broad ways. The first is through a complete representation,

in which the relevant processes are explicitly handled by solving fundamental physical equations that are well known and essentially exact. The second is through parameterization, which is an approximate closure scheme that handles the gross features of the relevant processes in a simplified manner with the help of constant parameters, the values of which are usually determined empirically.

Parameterization is implemented because a model has either insufficient complexity or insufficient grid resolution to capture the relevant processes. In the case of insufficient complexity, fundamental physical, chemical, or biological processes are simply missing from the model's basic equations and would not be captured even if the grid resolution were made infinitely fine. Atmospheric radiative transfer is an example in this category. In the case of insufficient grid resolution, the processes are captured by the underlying continuous partial differential equations, but the equations and boundary conditions are solved on a numerical grid that is too coarse (or with a time step that is too long) to explicitly resolve them. Ocean eddies and internal waves are examples in this category, because these small-scale features are captured by the classical fluid dynamics equations that are at the core of ocean general circulation models.

Faster computers are constantly permitting the development of climate models of greater complexity and higher resolution. Therefore, it might be argued that the need for parameterization is being gradually reduced over time. However, it is difficult to envisage any model ever being capable of explicitly simulating all of the climatically important components on all of the relevant time scales. Furthermore, it is known that the impact of the subgrid processes cannot necessarily be made vanishingly small simply by increasing the grid resolution, because information from arbitrarily small scales within the inertial subrange (down to the viscous dissipation scale) will always be able to contaminate the resolved scales in finite time (Palmer 2001). This feature of the subgrid dynamics perhaps explains why certain systematic errors are common to many different models (D'Andrea et al. 1998) and why numerical simulations are apparently not asymptoting as the resolution increases (e.g., Williamson 1999, 2008; Guemas and Codron 2011). Indeed, the Intergovernmental Panel on Climate Change (IPCC) has noted that the ultimate source of most large-scale errors is that "many important small-scale processes cannot be represented explicitly in models" (Randall et al. 2007, p. 601).

The major problem with conventional, deterministic parameterization schemes is their assumption that the impact of the subgrid scales on the resolved scales is uniquely determined by the resolved scales. This assumption can be made to sound plausible by invoking an analogy with the law of large numbers in statistical mechanics. According to this analogy, the subgrid processes are essentially random and of sufficiently large number per grid box that their integrated effect on the resolved scales is predictable (Williams 2005). In reality, however, the assumption is violated because the most energetic subgrid processes are only just below the grid scale, placing them far from the limit in which the law of large numbers applies. The implication is that the parameter values that would make deterministic parameterization schemes exactly correct are not simply uncertain; they are in fact indeterminate.

A possible solution would be to replace the traditional, deterministic parameterization schemes with stochastic versions that sample the subgrid variability in a computationally inexpensive manner. The suggestion that the climate system may be modeled using stochastic techniques was first made by Hasselmann (1976) and has been the subject of several recent review articles (Palmer 2001; Palmer and Williams 2008; Franzke et al. 2015; Berner et al. 2017). Stochastic parameterizations have demonstrated considerable success in modeling atmospheric convection (Lin and Neelin 2002), enhancing sea surface temperature predictability (Scott 2003), modeling El Niño–Southern Oscillation (Zavala-Garay et al. 2003), capturing regime transitions in rotating annulus laboratory experiments (Williams et al. 2003, 2004), improving the simulated atmospheric blocking frequency (Jung et al. 2005), and modeling sudden stratospheric warmings (Birner and Williams 2008).

Stochastic techniques have been applied widely to atmosphere models of the type used for short-term and medium-term weather prediction. For example, stochastic parameterization has been used operationally in the Ensemble Prediction Scheme (EPS) of the European Centre for Medium-Range Weather Forecasts (ECMWF) since 1998. It gives clear improvements in the skill of probabilistic predictions of precipitation (Buizza et al. 1999, 2005). Recently, there has been a growing interest in applying stochastic techniques to ocean models of the type used for longer-term seasonal forecasts and climate predictions (e.g., Sura and Penland 2002; Berloff 2005a; Berloff et al. 2007; Li and von Storch 2013; Porta Mana and Zanna 2014; Jansen and Held 2014; Kitsios et al. 2014; Andrejczuk et al. 2016). Dawson and Palmer (2015) have shown that it may not be necessary to represent the small scales accurately, or even explicitly, in order to improve the simulation of the large-scale climate. The question of whether stochastic closure schemes outperform their deterministic counterparts was listed by Williams et al. (2013) as a key

outstanding challenge in the field of mathematics applied to the climate system.

To examine the climatic impacts of stochastic perturbations in a coupled atmosphere–ocean general circulation model, Williams (2012) stochastically perturbed the air–sea buoyancy fluxes at the ocean surface. The stochastic perturbations were implemented to represent subgrid variability in the heat and freshwater fluxes associated with clouds, precipitation, and turbulent surface wind stresses. Significant changes were found in the century-mean ocean mixed layer depth, sea surface temperature, atmospheric Hadley circulation, and net upward water flux at the sea surface. These findings suggest that unresolved stochastic variability in air–sea fluxes may contribute to some of the biases exhibited by contemporary coupled climate models.

As much as 99% of the kinetic energy of the ocean is contained in the eddies (Open University 2001), which play a crucial role in transport and mixing (Treguier et al. 2014). Eddies are represented in ocean models using deterministic parameterizations such as the widely used Gent and McWilliams (1990) scheme. Ocean model resolutions are gradually increasing over time, to the extent that some ocean simulations are now described as eddy-permitting or eddy-resolving. However, Hallberg (2013) has shown that partially resolved eddies are not always superior to parameterized ones. Furthermore, turbulent eddies are energized from the mesoscale down to the Kolmogorov scale, implying that some of the eddy variability is missing even from simulations that are described as eddy-resolving. It follows that climate simulations may benefit from a stochastic representation of ocean eddies. Cooper and Zanna (2015) have developed a promising stochastic eddy parameterization for an idealized ocean model. However, stochastic eddy parameterizations have not previously been tested in long-term climate simulations using ocean general circulation models.

The present paper aims to investigate whether climate simulations can be improved by implementing a simple stochastic parameterization of ocean eddies in a coupled atmosphere–ocean general circulation model. The paper therefore complements the work of Williams (2012), by inserting noise throughout the three-dimensional ocean volume rather than at the two-dimensional ocean surface. The paper also complements the work of Cooper and Zanna (2015), by using a general circulation model rather than an idealized model. Our approach is to use simulations from a high-resolution, eddy-permitting ocean model to calculate the eddy statistics that are needed to inject realistic stochastic noise into a low-resolution, non-eddy-permitting version of the same model.

The layout of the paper is as follows. Section 2 describes how suitable noise statistics have been derived from a deterministic simulation using an eddy-permitting model. Section 2 also describes the suite of four stochastic experiments that have been run with the non-eddy-permitting model to test the sensitivity of the simulated climate to the noise definition by varying the noise amplitude and decorrelation time within reasonable limits. Section 3 presents the impacts of the noise on the simulated climate, considering ocean temperature and salinity both at the surface and at depth, as well as the global thermohaline circulation. The physical mechanisms that account for the results are stated and discussed in section 4. The paper concludes with a summary and discussion in section 5.

2. Methodology

a. Climate models

The two climate models that are used in this study, FAMOUS and HiGEM, differ mainly in their grid resolutions. The Fast Met Office/U.K. Universities Simulator (FAMOUS) is a computationally cheap ocean–atmosphere general circulation model that was developed collaboratively by the Met Office Hadley Centre and several U.K. universities (Smith et al. 2008). It has been used widely for studies of past, present, and future climate. Here we use version XFXWB (Smith 2012). FAMOUS is essentially a low-resolution version of the HadCM3 climate model (Gordon et al. 2000) with largely the same physics schemes, except for several simplifications such as the use of spatially constant coefficients in the Gent and McWilliams (1990) scheme and the removal of Iceland. FAMOUS has been tuned to reproduce the equilibrium climate and climate sensitivity of HadCM3 (Jones et al. 2005). The ocean model has a grid spacing of 2.5° in latitude and 3.75° in longitude, with 20 levels that increase in vertical resolution toward the surface. The ocean time step is 12 h.

The High-Resolution Global Environmental Model (HiGEM) is also a coupled ocean–atmosphere general circulation model from the Met Office Hadley Centre and U.K. universities. Here we use version 1.2 (Shaffrey et al. 2009). Compared to FAMOUS, the HiGEM ocean model has a finer grid spacing of ⅓° in latitude and longitude, with 40 levels that increase in vertical resolution toward the surface. The horizontal ocean resolution is therefore an order of magnitude finer than in FAMOUS, allowing eddies to be permitted. Consequently, the Gent and McWilliams (1990) eddy-transport

parameterization is switched off in HiGEM. The ocean time step is 20 min. The climatological errors relative to observations are reduced in HiGEM because of the high resolution (Shaffrey et al. 2009).

b. Stochastic noise

To design a stochastic parameterization of subgrid ocean eddies for FAMOUS, suitable characteristics must be chosen for the noise. First, we must choose which variables will be directly modified by the noise in the computer code. Then, we must choose the noise amplitude and probability distribution. Finally, we must choose whether to implement uncorrelated or autocorrelated noise, and in the latter case we must choose the decorrelation scales in time and space. Sura and Penland (2002) have emphasized that the details of stochastic perturbations introduced into numerical climate models must be physically justified to the fullest extent possible. To this end, high-resolution models may be regarded as truth and used to inform the choice of noise characteristics (Shutts and Palmer 2007).

In the present study, we objectively derive suitable noise statistics for FAMOUS, by regarding HiGEM as truth. We analyze 10 yr of the HiGEM control integration from Kuhlbrodt et al. (2015). The initial conditions for this control run were spun up for 110 yr, suggesting that it is reasonably well equilibrated. We use coarse graining to translate the statistics from HiGEM to FAMOUS. We coarse grain by spatially averaging from the HiGEM ocean grid onto the FAMOUS ocean grid in all three dimensions, and also by temporally averaging from the HiGEM ocean time step to the FAMOUS ocean time step.

To choose which variables will be directly modified by the noise in the FAMOUS model code, we use the approach of Brankart (2013) to investigate the relative importance of subgrid temperature and salinity structures in the FAMOUS ocean. Specifically, we analyze the impact on density of averaging the fine HiGEM data onto the coarse FAMOUS grid. The spatially averaged ocean density may be calculated by taking the temperature T and salinity S at each HiGEM grid point, first using them to calculate the density $\rho(T, S)$ at each HiGEM grid point, and then averaging these densities within each FAMOUS grid box to produce $\overline{\rho(T, S)}$. Alternatively, the average density may be calculated by taking the temperature and salinity at each HiGEM grid point, first averaging them within each FAMOUS grid box, and then using the average temperature $\overline{T}$ and average salinity $\overline{S}$ to calculate the average density $\rho(\overline{T}, \overline{S})$. These two calculations produce different average densities, because of nonlinearities in the equation of state of seawater related to cabelling (e.g., Williams et al. 2010). The difference $\overline{\rho(T, S)} - \rho(\overline{T}, \overline{S})$ between the two average densities,

which is always negative because of the sense of the curvature of isopycnal contours in the (T, S) plane, may be interpreted as a diagnostic of the eddy variability that is resolved in HiGEM but unresolved in FAMOUS.

The results of the above calculations are shown in Fig. 1. Figure 1a shows a global map of the density difference $\overline{\rho(T, S)} - \rho(\overline{T}, \overline{S})$ at a depth of 400 m. The density difference is greatest in the strongly eddying regions of the ocean, and it resembles global maps of observed eddy activity (Chelton et al. 2011). For example, the magnitude of the density difference reaches values of around 3×10^{-2} kg m^{-3} within the Gulf Stream in the western Atlantic Ocean and Kuroshio in the western Pacific Ocean. Figures 1b and 1c separate out the contributions of salinity and temperature to this density difference, by showing $\rho(\overline{T}, S) - \rho(\overline{T}, \overline{S})$ and $\rho(T, \overline{S}) - \rho(\overline{T}, \overline{S})$ at a depth of 400 m. These panels indicate that the contribution of salinity is at least an order of magnitude smaller than the contribution of temperature. For example, the root-mean-square density difference is 4×10^{-3} kg m^{-3} in Fig. 1c, but only 2×10^{-4} kg m^{-3} in Fig. 1b. Furthermore, the vertical profiles within the Gulf Stream shown in Fig. 1d indicate that subgrid fluctuations in temperature cause the majority of the density difference not only at 400 m but at all depths. Therefore, we consider it reasonable at this stage to stochastically perturb the temperature field only.

To explore whether the ocean temperature field is undervariable in FAMOUS compared to HiGEM, Fig. 2 shows global maps of the standard deviations of the temperature tendencies calculated from control integrations of the two models. We analyze the temperature tendencies rather than the temperatures themselves, because the tendencies are a cleaner diagnostic for isolating the effects of the eddies. To ensure a fair comparison, the temperature tendencies in HiGEM are first averaged onto the FAMOUS grid before computing the standard deviation. It is clear that FAMOUS generally exhibits less temperature variability than HiGEM at all depths and across the globe. For example, in the upper 100 m, temperatures in HiGEM show the largest variance along the equator in the Atlantic and Pacific basins and near the Gulf Stream and Kuroshio, but the corresponding variability in FAMOUS is substantially weaker. Similarly, at depths between 100 and 1000 m, most of the variability in HiGEM lies along the equator in the three ocean basins and along the southern polar front, but the corresponding variability in FAMOUS is again far too weak.

Having chosen to stochastically perturb the ocean temperature field in FAMOUS, we must next choose the noise amplitude and probability distribution. Figure 3 shows a
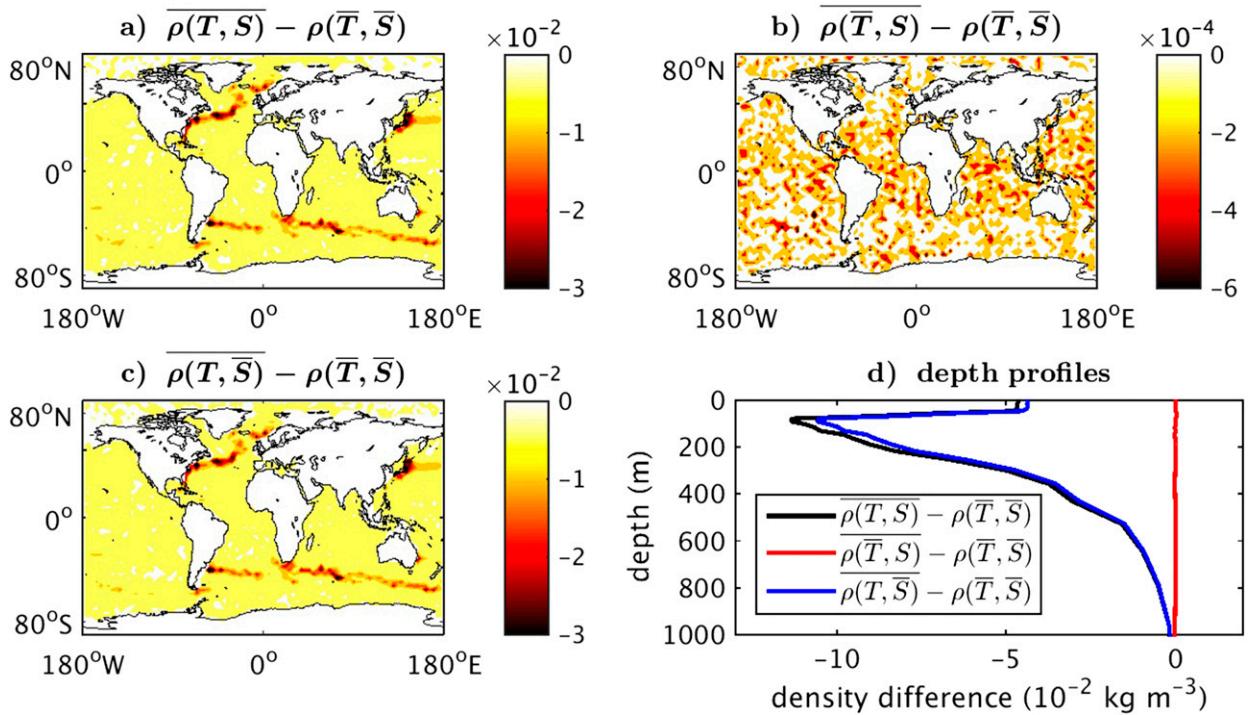
FIG. 1. (a) Ocean density difference (kg m^{-3}) at a depth of 400 m resulting from whether the spatial averaging from the fine HiGEM grid to the coarse FAMOUS grid is performed before or after calculating ocean density from temperature and salinity. (b) The contribution to the density difference (kg m^{-3}) at a depth of 400 m from averaging of salinity. (c) The contribution to the density difference (kg m^{-3}) at a depth of 400 m from averaging of temperature. (d) Depth profiles of the contributions to the density difference in the upper 1000 m at the point (43°N, 50°W) in the Gulf Stream. The temperature and salinity are averaged over an arbitrarily chosen 5-day period from the control integration of HiGEM.

scatterplot of the depth distribution of the temperature tendencies at all grid points in the global ocean from the control integration of HiGEM. Note that the temperature tendencies will include a contribution from the seasonal cycle, but this component is far smaller than the contribution from unforced variability, which dominates Fig. 3. The temperature variability is generally large in the upper ocean but declines with depth. At each depth, the temperature tendencies are found to be approximately normally distributed, which is consistent with the Gaussian eddy variability that was identified in observations by Biri et al. (2015). The mean, skewness, and kurtosis of the temperature tendencies are all found to be close to zero, but the standard deviation is nonzero and decreases with depth. Apart from a few outliers, the tendencies lie mainly within a logarithmic envelope function that has been fitted to the data by trial and error and is also shown in Fig. 3. Given these statistics from HiGEM, the noise that we apply to the temperature tendency field in FAMOUS is drawn from a Gaussian distribution with a mean value of zero and a standard deviation that decreases logarithmically with depth according to the fitted envelope function. At each depth, the applied noise amplitude is taken to be the same at all latitudes and longitudes.

Finally, for both space and time, we must choose whether to implement uncorrelated or autocorrelated noise in FAMOUS. Typical eddy horizontal length scales are much shorter than the horizontal grid spacing of 2.5° × 3.75° in the ocean model, implying that eddy correlations between horizontally neighboring grid boxes will be negligible. In contrast, typical eddy life times are much longer than the ocean time step of 12 h, implying that eddy correlations between successive time steps will be large. Diagnostics of temperature tendencies from the control integration of HiGEM confirm these hypotheses, giving horizontal correlations of essentially zero on the scale of the FAMOUS grid, but giving temporal correlations in the approximate range 5–30 days. In terms of vertical correlations, high-resolution global ocean simulations suggest that many eddies penetrate the full depth of the water column (Petersen et al. 2013). Therefore, we choose to implement noise that is autocorrelated (red) in time, uncorrelated (white) in latitude and longitude, and fully coherent in depth. The use of temporally correlated noise is further justified because temporally white noise has been found not to produce the same benefits as temporally red noise (Christensen et al. 2015).
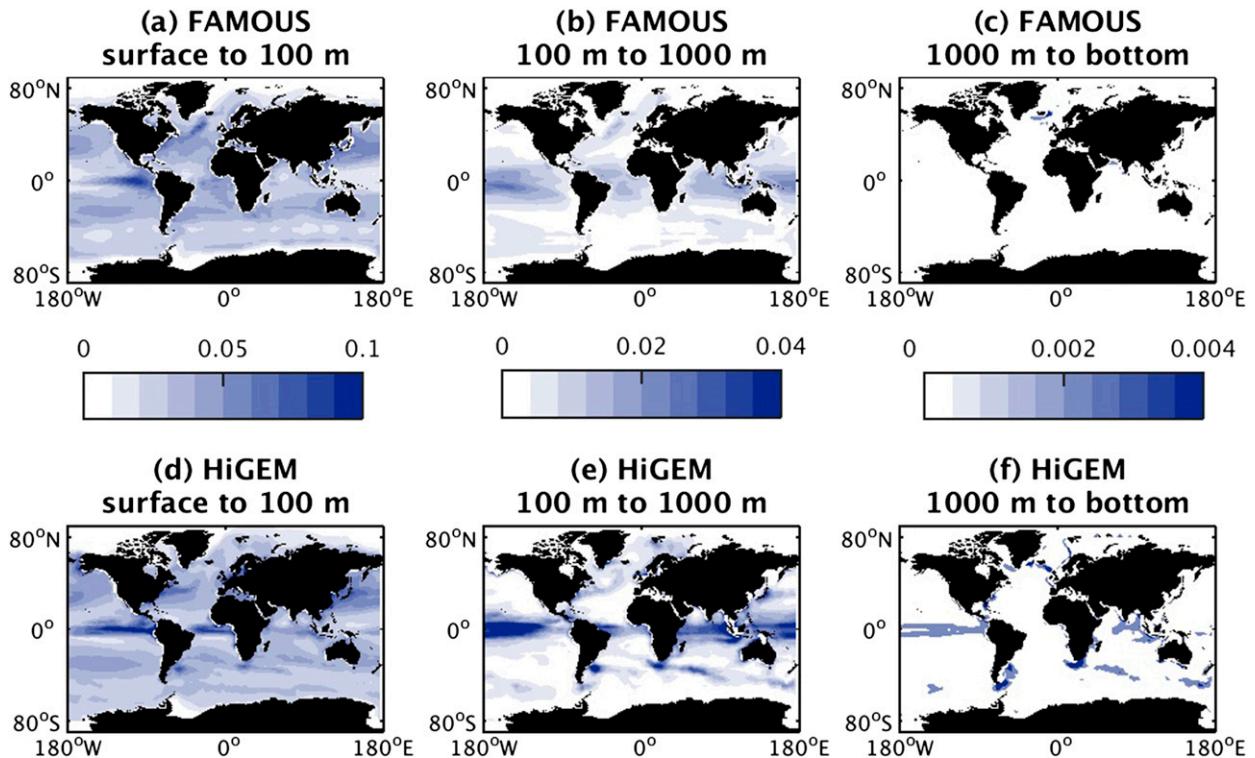
FIG. 2. Global maps of the standard deviation of the temperature tendency $[°C\,(12\,h)^{-1}]$, as calculated from the control runs of (top) FAMOUS and (bottom) HiGEM within the following three depth intervals: (a),(d) from the surface to 100 m (left color bar), (b),(e) from 100 to 1000 m (middle color bar), and (c),(f) from 1000 m to the bottom (right color bar).

c. Suite of experiments

We have run four stochastic experiments with FAMOUS, together with a deterministic control simulation for comparison. The details of the experiments are listed in Table 1. The experiments are designed to test the sensitivity of the simulated climate to the noise definition, by varying the noise amplitude and decorrelation time within reasonable limits. In each case, the noise is drawn from a Gaussian distribution with the depth profile derived in section 2b, and it is applied additively to perturb the ocean temperature tendency at each grid point and time step. The control simulation and the stochastic experiments are all initialized identically from the end of a long control run of duration 5000 yr, to ensure that the initial conditions are in equilibrium and lie on the attractor of the control model. Each experiment is then run for 50 yr.

The control simulation, hereafter referred to as CONT, is a deterministic run of the standard FAMOUS model. The first stochastic experiment, STOC_LOW_UNCOR, uses temporally uncorrelated (white) noise with a relatively low amplitude of half the value diagnosed from HiGEM. The next experiment, STOC_HIGH_UNCOR, also uses white noise but with double

the previous amplitude, to match the value diagnosed from HiGEM. The third experiment, STOC_HIGH_5d, uses the higher noise amplitude, but instead of white noise uses temporally autocorrelated (red) noise with a
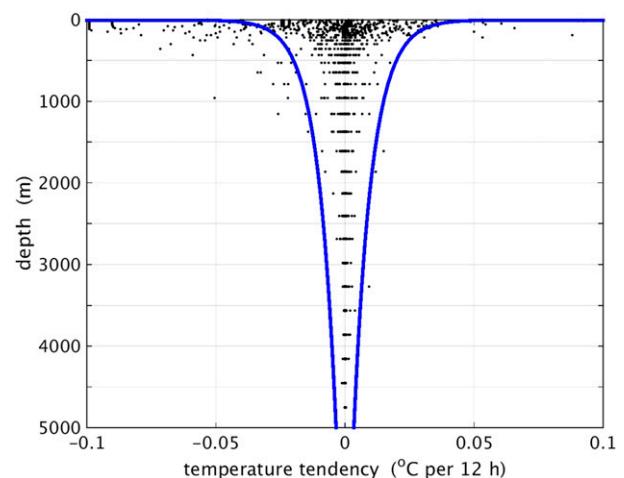

FIG. 3. Scatterplot of the depth dependence of the temperature tendencies $[°C\,(12\,h)^{-1}]$ diagnosed from a control integration of HiGEM (black dots). A logarithmic envelope function (blue curves) has been fitted to the distribution in such a manner that it contains the majority of the points in the scatterplot.

TABLE 1. Details of the numerical simulations that are performed in this study using the FAMOUS model. Four stochastic experiments, one of which has three ensemble members each with different noise realizations, are compared to a deterministic control run. In the stochastic experiments, the noise amplitude is the standard deviation of the zero-mean Gaussian distribution that is used to perturb the ocean temperature tendency at each grid point and time step. The noise amplitude listed here is the amplitude at a depth of 2.5 m, which is in the middle of the upper model layer. For reference, the amplitude of the envelope function in Fig. 3 at the same depth is $0.1°C (12 h)^{-1}$. In the deeper model layers, the noise amplitude decreases with depth according to the logarithmic envelope function shown in Fig. 3. Each experiment is run for 50 yr.

Experiment name	Noise amplitude at 2.5 m depth $[°C (12 h)^{-1}]$	Decorrelation time (days)	Number of ensemble members
CONT	0	—	1
STOC_LOW_UNCOR	0.05	0	1
STOC_HIGH_UNCOR	0.1	0	3
STOC_HIGH_5d	0.1	5	1
STOC_HIGH_30d	0.1	30	1

decorrelation time of 5 days. The final experiment, STOC_HIGH_30d, is similar but uses a decorrelation time of 30 days. To test the sensitivity of the simulated climate to the noise realization (i.e., the particular sequence of random numbers that is generated from a given initial seed), an ensemble of three members is run for the STOC_HIGH_UNCOR experiment. The ensemble members are identical except for the use of different seeds to initiate the random number generator.

The depth tapering of our generated noise makes it unlikely that the deep ocean temperatures will fall below the freezing point of $-1.8°C$. Nevertheless, there is a temperature limiter in the model, which takes corrective action if the freezing point is crossed. After each update to the ocean temperature field, any temperatures that are below the freezing point are reset to freezing by the temperature limiter. The pressure dependence of the freezing point is neglected, so a freezing point of $-1.8°C$ is applied at all depths.

3. Results

This section analyzes the climatological impacts of inserting the ocean eddy noise into FAMOUS. We examine the temperature and salinity fields averaged over the final decade of the simulations, both at the ocean surface and from a zonal-mean perspective. We then examine the trends and variability, focusing on time series and power spectra of the ocean heat content and meridional overturning circulation over 50 yr. Finally, we examine the approach to equilibrium. Where possible, we compare the impacts of the noise with the impacts of refining the resolution, by using the HiGEM control integration that was described in section 2b. We recall that the initial conditions for the HiGEM control integration were spun up for 110 yr, suggesting that it is reasonably well equilibrated. In particular, any subsequent adjustment toward equilibrium after 110 yr is

likely to be much smaller than the large differences between HiGEM and FAMOUS.

a. Sea surface temperature and salinity

Global maps of the sea surface temperature averaged over the final decade of the simulations are shown in Fig. 4. The spatial structure in CONT (Fig. 4a) captures the essential features seen in observations, including the bulk temperature gradient from the cold polar regions at less than 0°C to the warm tropical regions at over 30°C. The remaining panels in Fig. 4 show anomalies with respect to CONT. The anomaly pattern in STOCH_LOW_UNCOR (Fig. 4b) shows that weak, uncorrelated noise produces large-scale warming and cooling regions in each of the ocean basins. The magnitudes are generally small, with the root-mean-square anomaly being 0.40°C. The anomaly in STOCH_HIGH_UNCOR (Fig. 4c) shows that the response to strong, uncorrelated noise has roughly the same spatial pattern as the response to weak, uncorrelated noise. The amplitude is greater, however, with the root-mean-square anomaly increasing to 0.61°C. The introduction of temporal correlations to the noise, both on short time scales in STOCH_HIGH_5d (Fig. 4d) and on longer time scales in STOCH_HIGH_30d (Fig. 4e), maintains a similar response pattern but increases the amplitude further, with the root-mean-square anomalies increasing to 0.90° and 1.66°C, respectively. The root-mean-square sea surface temperature difference between any two decades of CONT is typically at least an order of magnitude smaller than the impacts of the noise in each of the stochastic experiments, demonstrating that the changes are robust.

For comparison, the sea surface temperature in the control integration of HiGEM relative to the control integration of FAMOUS is shown in Fig. 4f. The root-mean-square temperature difference is 1.92°C. There is a striking agreement between the sea surface

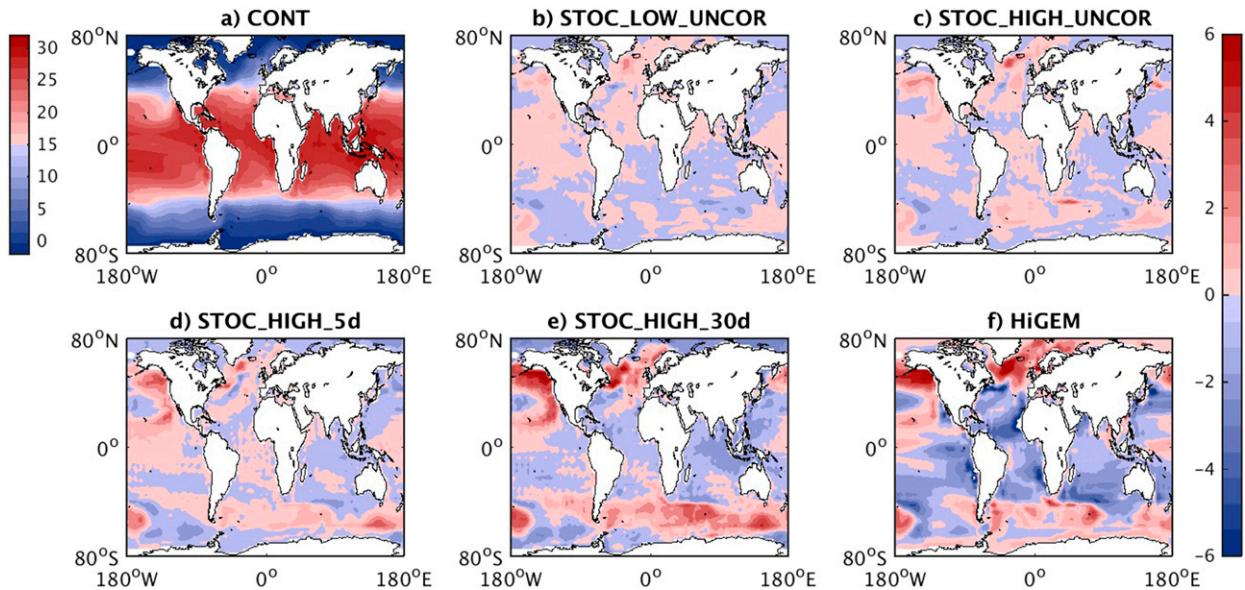FIG. 4. Global maps of sea surface temperature (SST) averaged over the final decade of the FAMOUS simulations. (a) SST (°C; left color bar) in the control simulation of FAMOUS. (b)–(e) SST anomaly (°C; right color bar) in the four stochastic simulations of FAMOUS, where the anomaly is calculated with respect to the control simulation of FAMOUS. (f) SST anomaly (°C; right color bar) in the control simulation of HiGEM, where the anomaly is calculated with respect to the control simulation of FAMOUS.

temperature anomalies in STOCH_HIGH_30d and HiGEM relative to CONT, especially considering that these responses are emergent phenomena that were not imposed a priori by the noise or the resolution increase. Notice for example the warming of the Atlantic Ocean north of 45°N and the cooling south of this latitude; the U-shaped warming of the northern, eastern, and southern boundaries of the North Pacific Ocean with cooling in the middle; the cooling of the Indian Ocean; and the warming of the Southern Ocean between 40° and 60°S at most longitudes. To quantify the similarities and how they improve as the noise definition varies, the pattern correlation coefficients are 0.39 between Figs. 4b and 4f, 0.41 between Figs. 4c and 4f, 0.47 between Figs. 4d and 4f, and 0.55 between Figs. 4e and 4f. The pattern correlation coefficient between Fig. 4f and the difference between any two decades of CONT is typically at least an order of magnitude smaller than these values, demonstrating that the changes are robust.

Global maps of the sea surface salinity averaged over the final decade of the simulations are shown in Fig. 5. Again, the spatial structure in CONT (Fig. 5a) captures the essential features seen in observations, including high salinity in the subtropical gyres where evaporation dominates over precipitation, low salinity in the subpolar gyres where precipitation dominates over evaporation, and higher salinity in the Atlantic than the Pacific. The remaining panels in Fig. 5 show anomalies with respect to CONT. The anomaly pattern in

STOCH_LOW_UNCOR (Fig. 5b) shows that weak, uncorrelated noise produces large-scale salinification and freshening regions in each of the ocean basins. The magnitudes are generally small, with the root-mean-square anomaly being 0.28 psu. The anomaly in STOCH_HIGH_UNCOR (Fig. 5c) shows that the response to strong, uncorrelated noise has roughly the same spatial pattern as the response to weak, uncorrelated noise. The amplitude is greater, however, with the root-mean-square anomaly increasing to 0.48 psu. The introduction of temporal correlations to the noise, both on short time scales in STOCH_HIGH_5d (Fig. 5d) and longer time scales in STOCH_HIGH_30d (Fig. 5e), maintains a similar response pattern but increases the amplitude further, with the root-mean-square anomalies increasing to 0.61 and 0.93 psu, respectively. The root-mean-square sea surface salinity difference between any two decades of CONT is typically at least an order of magnitude smaller than the impacts of the noise in each of the stochastic experiments.

For comparison, the sea surface salinity in the control integration of HiGEM relative to the control integration of FAMOUS is shown in Fig. 5f. The root-mean-square salinity difference is 2.82 psu. The agreement between the sea surface salinity anomalies in STOCH_HIGH_30d and HiGEM relative to CONT is generally poorer than the agreement between the sea surface temperature responses. Areas of agreement include the
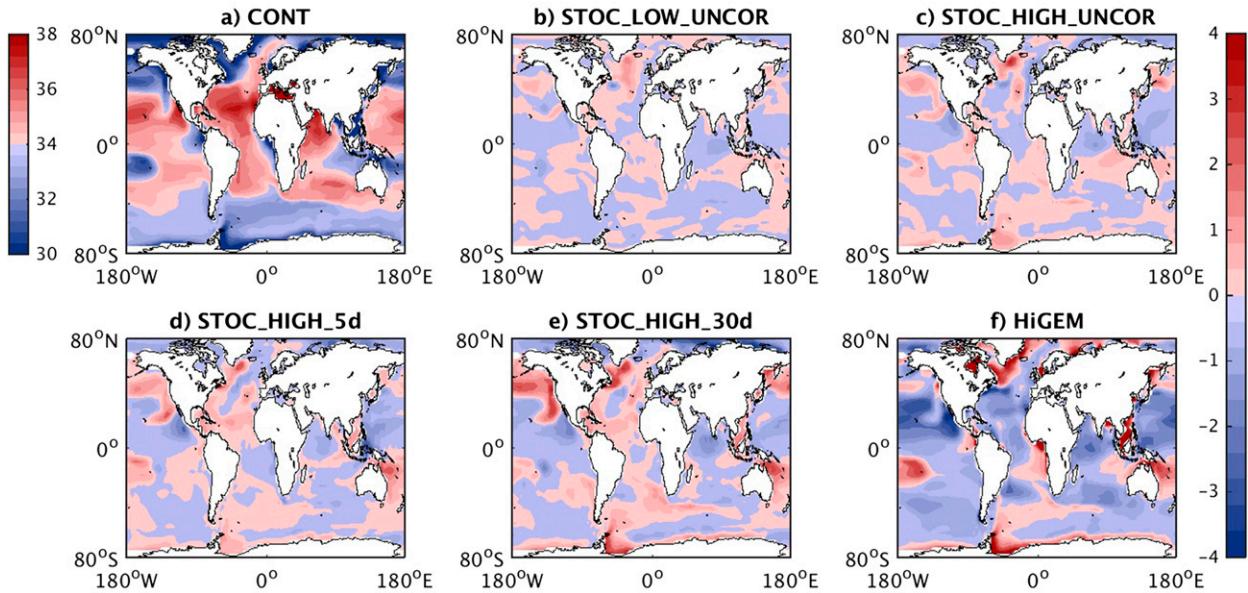
FIG. 5. As in Fig. 4, but for sea surface salinity (psu).

large-scale salinification east of Australia and east of Argentina and the large-scale freshening of the Indian Ocean and most of the North Pacific Ocean. However, the Atlantic Ocean is generally saltier in STOCH_HIGH_30d but fresher in HiGEM. To quantify the similarities, the pattern correlation coefficients are 0.14 between Figs. 5b and 5f, 0.23 between Figs. 5c and 5f, 0.20 between Figs. 5d and 5f, and 0.25 between Figs. 5e and 5f. The pattern correlation coefficient between Fig. 5f and the difference between any two decades of CONT is typically at least an order of magnitude smaller than these values.

b. Zonal-mean temperature and salinity

Latitude–depth plots of the global zonal-mean temperature averaged over the final decade of the simulations are shown in Fig. 6. The spatial structure in CONT (Fig. 6a) captures the essential features seen in observations, including the strong thermal stratification at low latitudes and the weak thermal stratification at high latitudes, with W-shaped isotherms caused by equatorial upwelling. The remaining panels in Fig. 6 show anomalies with respect to CONT. The anomaly pattern in STOCH_LOW_UNCOR (Fig. 6b) shows that weak, uncorrelated noise produces a general warming in the upper 2 km of the ocean and a cooling below. The magnitudes are generally small, with the root-mean-square anomaly being 0.28°C. The anomaly in STOCH_HIGH_UNCOR (Fig. 6c) shows that the response to strong, uncorrelated noise has roughly the same spatial pattern as the response to weak, uncorrelated noise. The

amplitude is greater, however, with the root-mean-square anomaly increasing to 0.65°C. The introduction of temporal correlations to the noise, both on short time scales in STOCH_HIGH_5d (Fig. 6d) and longer time scales in STOCH_HIGH_30d (Fig. 6e), maintains a similar response pattern but increases the amplitude further, with the root-mean-square anomalies increasing to 1.00° and 2.10°C, respectively. The root-mean-square zonal-mean temperature difference between any two decades of CONT is typically at least an order of magnitude smaller than the impacts of the noise in each of the stochastic experiments.

For comparison, the global zonal-mean temperature in the control integration of HiGEM relative to the control integration of FAMOUS is shown in Fig. 6f. The root-mean-square temperature difference is 1.25°C. There is again good agreement between the zonal-mean temperature anomalies in STOCH_HIGH_30d and HiGEM relative to CONT, except at the highest latitudes. To quantify the similarities, the pattern correlation coefficients are 0.63 between Figs. 6b and 6f, 0.82 between Figs. 6c and 6f, 0.89 between Figs. 6d and 6f, and 0.92 between Figs. 6e and 6f. The pattern correlation coefficient between Fig. 6f and the difference between any two decades of CONT is typically at least an order of magnitude smaller than these values.

The global zonal-mean temperature response is broken down into contributions from the Atlantic and Pacific Oceans in Figs. 7 and 8 . In the Atlantic Ocean, there is still a general cooling at depths below about 2 km in STOCH_HIGH_30d, but the warming above
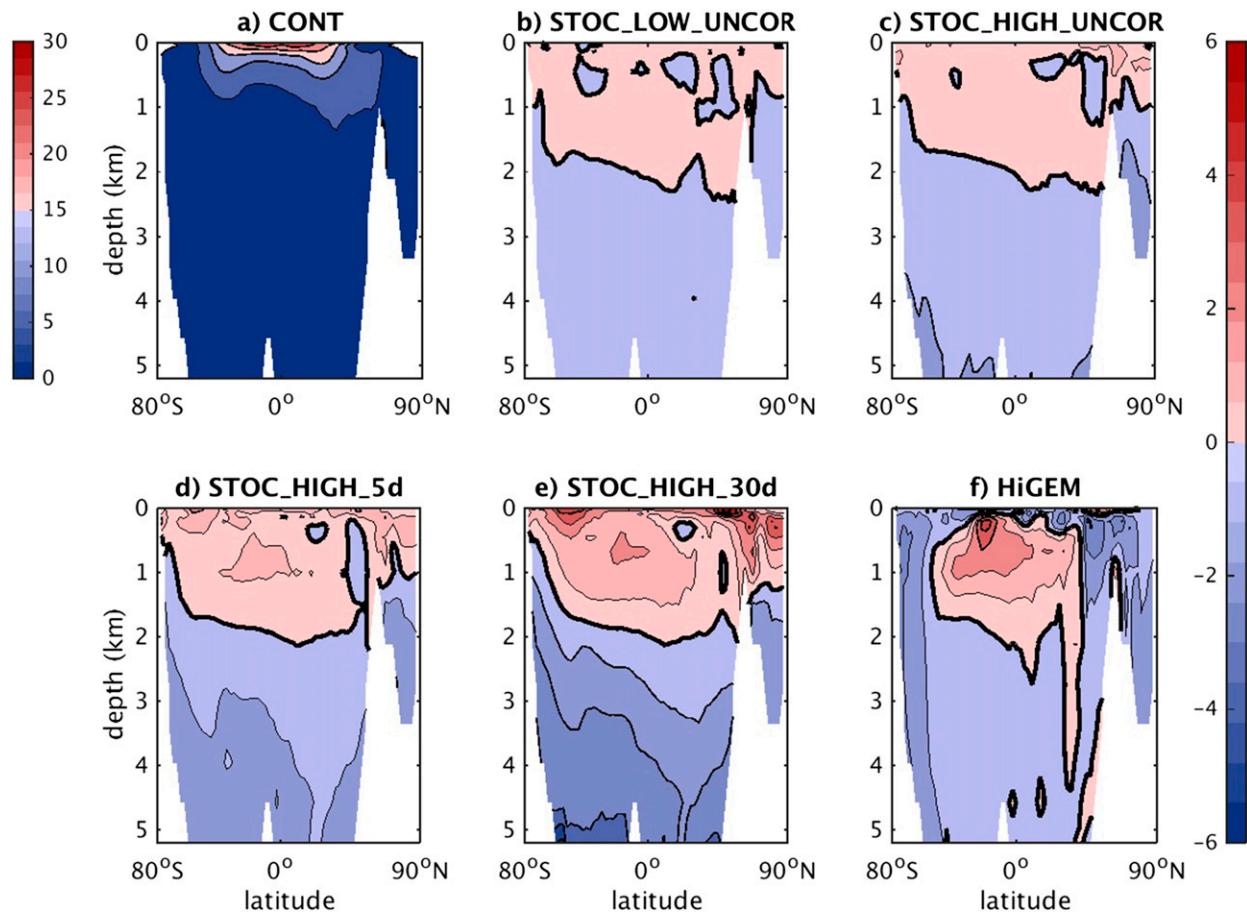
FIG. 6. Latitude–depth plots of the global zonal-mean temperature averaged over the final decade of the FAMOUS simulations. (a) Temperature (°C; left color bar) in the control simulation of FAMOUS. (b)–(e) Temperature anomaly (°C; right color bar) in the four stochastic simulations of FAMOUS, where the anomaly is calculated with respect to the control simulation of FAMOUS. (f) Temperature anomaly (°C; right color bar) in the control simulation of HiGEM, where the anomaly is calculated with respect to the control simulation of FAMOUS.

this level occurs only south of about 40°N, with cooling at higher latitudes. A similar pattern is seen in HiGEM, except for a strong but latitudinally localized warming that penetrates down to depths of about 4.5 km and is not captured by STOCH_HIGH_30d. In the Pacific Ocean, the temperature anomaly in HiGEM with respect to CONT is captured reasonably well by STOCH_HIGH_30d, except that the warming is perhaps too deep and that a cooling of the upper ocean at the northern-most latitudes is not captured.

Latitude–depth plots of the zonal-mean salinity in the Pacific Ocean averaged over the final decade of the simulations are shown in Fig. 9. The anomaly pattern in STOCH_LOW_UNCOR (Fig. 9b) shows that weak, uncorrelated noise produces a freshening everywhere except below about 3–4 km at all latitudes and above about 1 km at low latitudes, where there is a salinification. The magnitudes are generally small, with the root-mean-square

anomaly being 0.06 psu. The anomaly in STOCH_HIGH_UNCOR (Fig. 9c) shows that the response to strong, uncorrelated noise has roughly the same spatial pattern as the response to weak, uncorrelated noise. The amplitude is greater, however, with the root-mean-square anomaly increasing to 0.15 psu. The introduction of temporal correlations to the noise, both on short time scales in STOCH_HIGH_5d (Fig. 9d) and longer time scales in STOCH_HIGH_30d (Fig. 9e), maintains a similar response pattern but increases the amplitude further, with the root-mean-square anomalies increasing to 0.19 and 0.42 psu, respectively. The root-mean-square zonal-mean salinity difference between any two decades of CONT is typically at least an order of magnitude smaller than the impacts of the noise in each of the stochastic experiments.

For comparison, the Pacific zonal-mean temperature in the control integration of HiGEM relative to
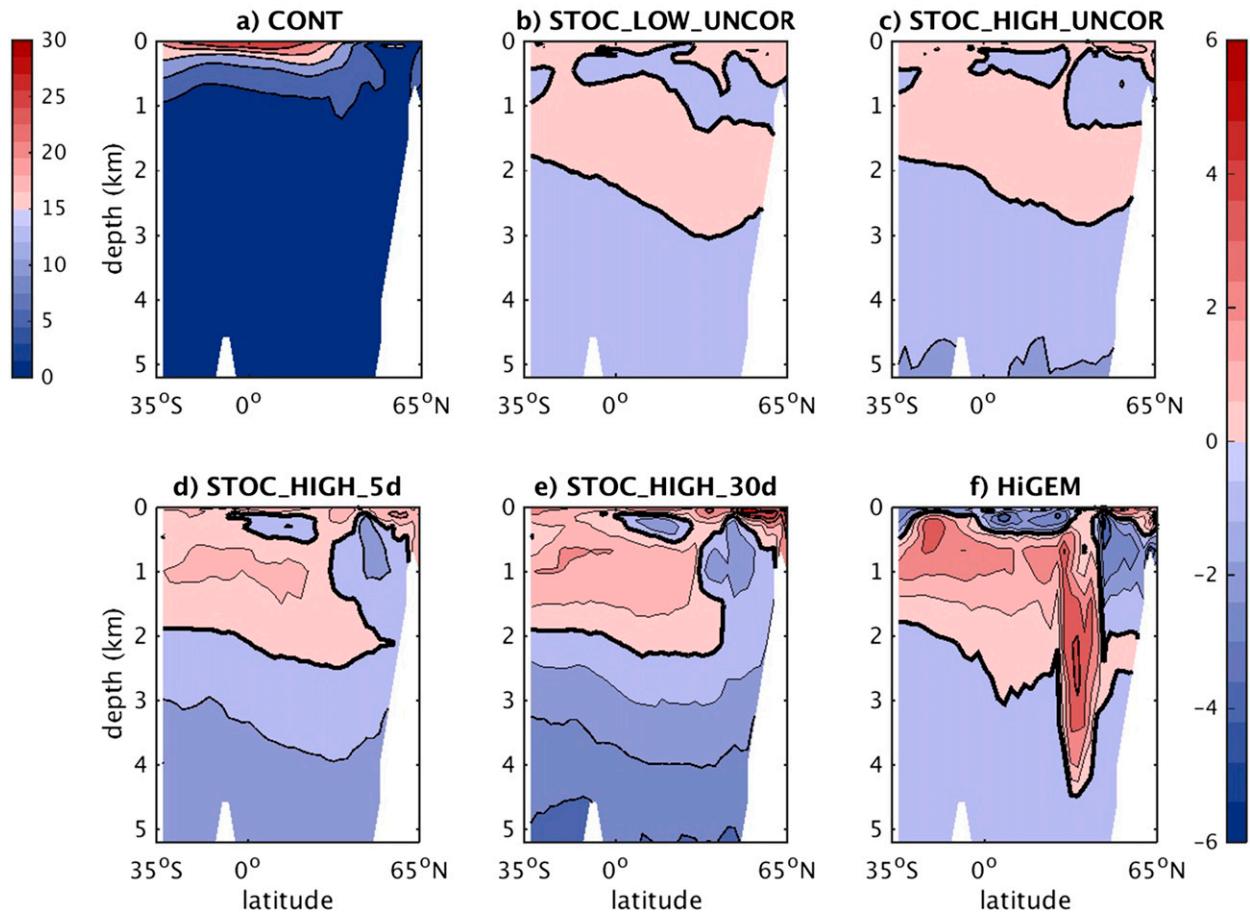
FIG. 7. As in Fig. 6, but for zonal-mean temperature in the Atlantic Ocean. The zonal average is taken across the entire Atlantic basin.

the control integration of FAMOUS is shown in Fig. 9f. The root-mean-square salinity difference is 0.98 psu. There is again excellent agreement between the Pacific zonal-mean salinity anomalies in STOCH_HIGH_30d and HiGEM relative to CONT, with similar results in the Atlantic (not shown). To quantify the similarities, the pattern correlation coefficients are 0.40 between Figs. 9b and 9f, 0.56 between Figs. 9c and 9f, 0.64 between Figs. 9d and 9f, and 0.69 between Figs. 9e and 9f. The pattern correlation coefficient between Fig. 9f and the difference between any two decades of CONT is typically at least an order of magnitude smaller than these values.

c. Ocean heat content

A consequence of the warming of the top 1–2 km and cooling of the bottom 3–4 km in each of the stochastic experiments is that there is a net cooling of the global ocean. To quantify this cooling, time series of the global ocean heat content are shown in Fig. 10. The heat content in CONT is stable and shows no significant trend over the period of 50 yr. In contrast, the heat contents in each of the stochastic experiments decrease from their initial values, as expected. In STOCH_LOW_UNCOR, the heat content appears to reach a new equilibrium after around 15 yr at 13×10^{22} J less than its initial value. In STOCH_HIGH_UNCOR, the heat content appears not to stabilize, reaching about 45×10^{22} J less than its initial value after 50 yr. The three ensemble members of STOC_HIGH_UNCOR follow each other reasonably well. In particular, their ensemble spread after 50 yr is much smaller than the differences between the experiments, showing that the long-term rate of change of ocean heat content is insensitive to the noise realization. In STOCH_HIGH_5d and STOCH_HIGH_30d, the loss of ocean heat content after 50 yr reaches 70×10^{22} and 140×10^{22} J, respectively.

To illustrate the order of magnitude of these heat losses, we note for comparison that the heat content of the global ocean is observed to have increased by around 20×10^{22} J between the mid-1950s and mid-1990s
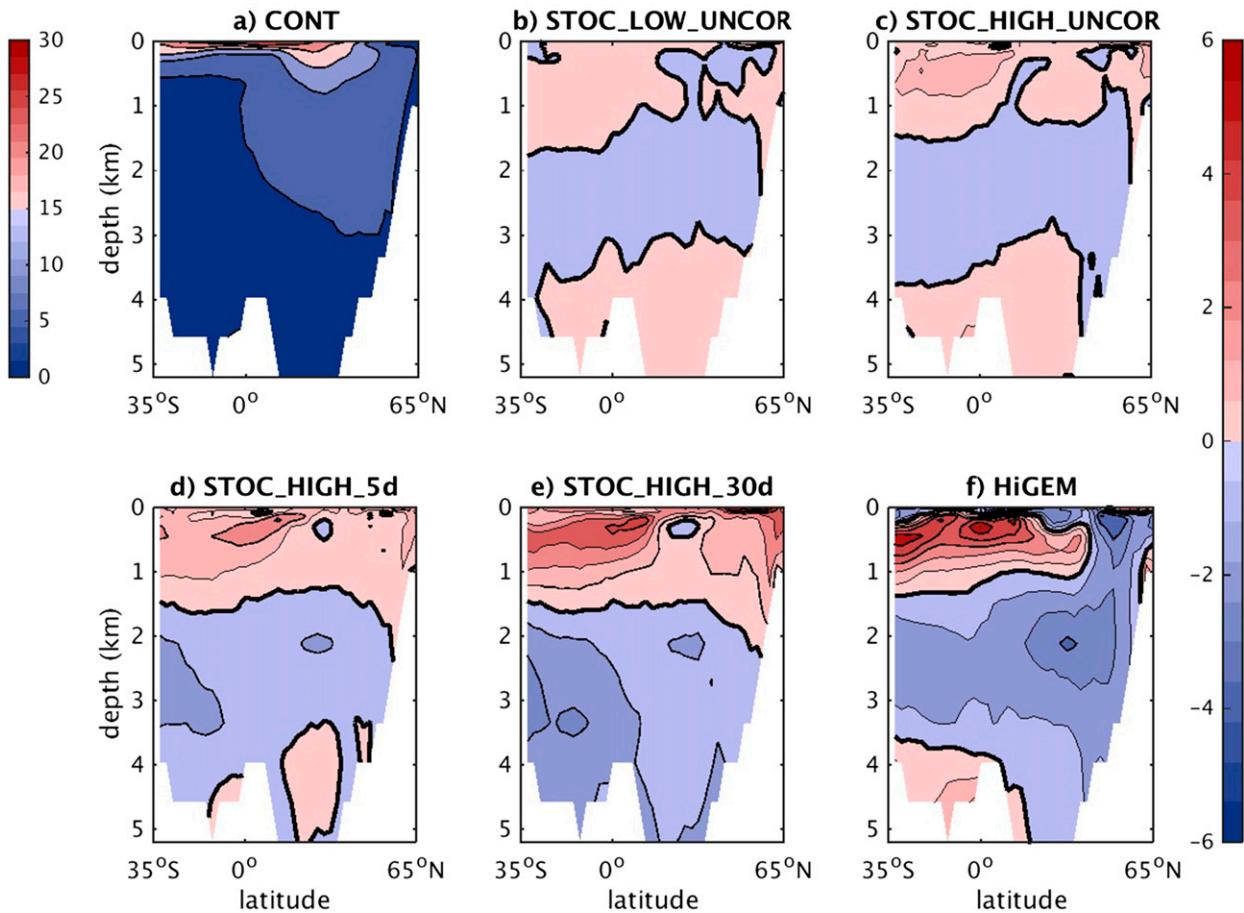
FIG. 8. As in Fig. 6, but for zonal-mean temperature in the Pacific Ocean. The zonal average is taken across the entire Pacific basin.

(Levitus et al. 2000). Therefore, the change in ocean heat content caused by weak, uncorrelated eddy noise is the same order of magnitude as the historic global warming signal, and the change caused by strong, correlated eddy noise is an order of magnitude greater than it. Figure 10 shows that eddy noise not only produces long-term trends in ocean heat content, but also increases the variability on interannual and decadal time scales. In particular, strong, uncorrelated noise produces more variability than weak, uncorrelated noise, and long temporal correlations produce more variability than short temporal correlations. Finally, it should be noted that in the limit of short decorrelation time scales, the effect of increasing the decorrelation time scale by a factor of 6 (from 5 to 30 days) is expected from basic theory and dimensional analysis to be similar to the effect of increasing the amplitude of the forcing by a factor of $\sqrt{6} \approx 2.4$, which is roughly in agreement with the results in Fig. 10.

d. Atlantic meridional overturning circulation

The Atlantic meridional overturning circulation in the FAMOUS simulations is shown in Fig. 11. The latitude–depth structure in CONT averaged over the final decade (Fig. 11a) captures the essential features seen in observations. These features include northward volume transport in approximately the upper 1.5 km of the Atlantic Ocean, which reaches a maximum value of about 18 Sv ($1 \text{ Sv} \equiv 10^6 \text{ m}^3 \text{ s}^{-1}$) at 26°N. This northward transport lies above the southward return transport of North Atlantic Deep Water (NADW), which in turn lies above a weak countercirculating cell of Antarctic Bottom Water (AABW).

The remaining panels in Fig. 11 show latitude–time Hovmöller (1949) plots of the annual-mean Atlantic meridional overturning circulation at a depth of 1000 m in the control simulation and four stochastic experiments. The relatively low interannual variability in CONT (Fig. 11b) is strengthened by the introduction of weak, uncorrelated eddy noise in STOCH_LOW_UNCOR (Fig. 11c) and more so by the strong, uncorrelated noise in STOCH_HIGH_UNCOR (Fig. 11d). The variability is enhanced even further by the introduction of temporally correlated noise in STOCH_HIGH_5d (Fig. 11e) and STOCH_HIGH_30d (Fig. 11f). The
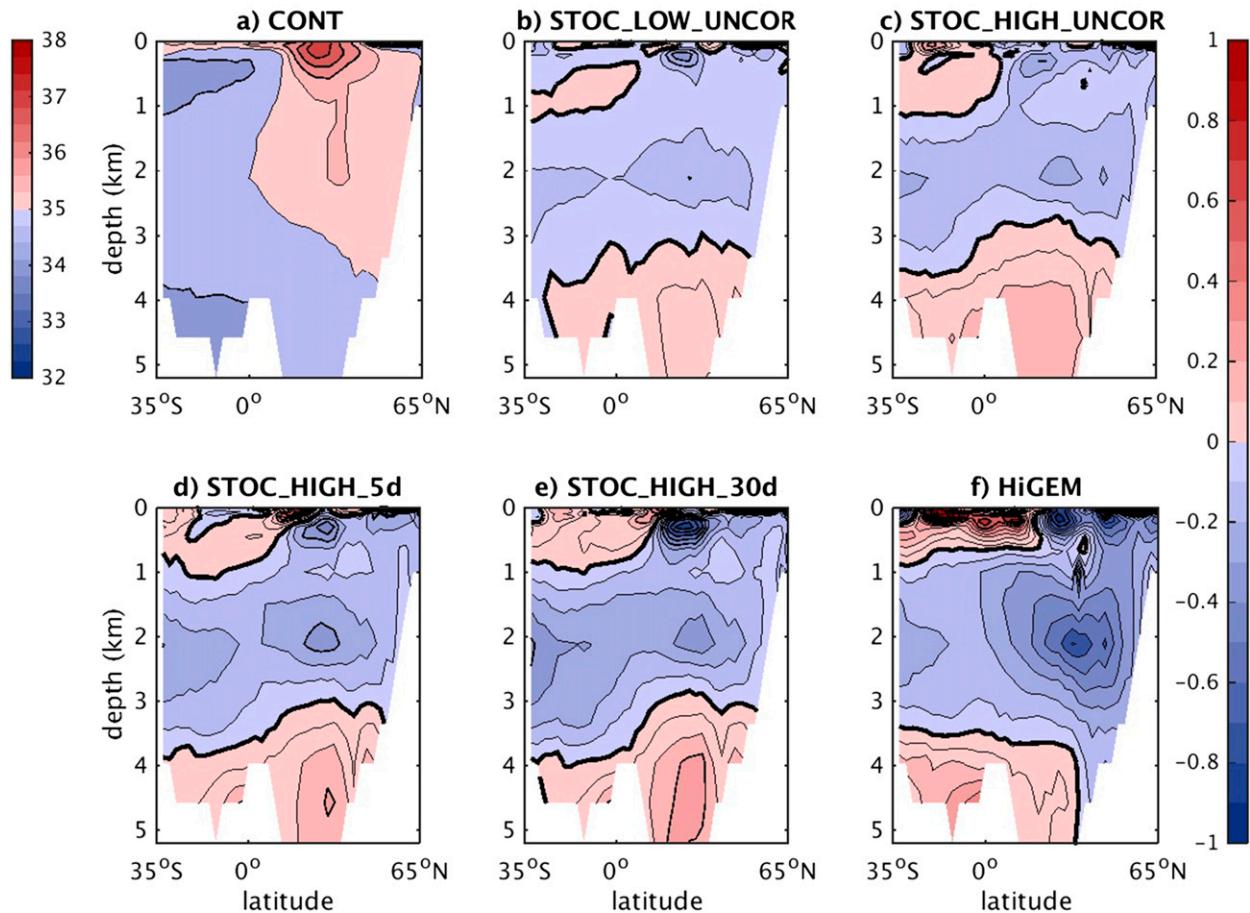
FIG. 9. As in Fig. 6, but for zonal-mean salinity in the Pacific Ocean (psu). The zonal average is taken across the entire Pacific basin.

time-averaged overturning circulation is also increased by the noise. Averaged over 50 yr at 26°N and at a depth of 1000 m, the strength of the circulation increases from 18.0 Sv in CONT to 19.1 Sv in STOCH_LOW_UNCOR, STOCH_HIGH_UNCOR, and STOCH_HIGH_5d and to 19.8 Sv in STOCH_HIGH_30d.

To examine the noise-driven variability increase in more detail, power spectra of the maximum Atlantic meridional overturning circulation at 26°N are shown in Fig. 12. Power spectra from the HiGEM control simulation and the RAPID observations (Srokosz and Bryden 2015) are included for comparison. All the power spectra display a prominent annual cycle, but the variabilities on time scales of 1–5 yr differ markedly. On these interannual time scales, CONT displays the least variability. The variability is increased by the addition of weak and strong uncorrelated noise in STOCH_LOW_UNCOR and STOCH_HIGH_UNCOR, and is increased further by the inclusion of fast and slow autocorrelations in STOCH_HIGH_5d and STOCH_HIGH_30d. HiGEM and RAPID are in close agreement with each other on these interannual time scales, demonstrating

the benefit of higher resolution. Three of the stochastic FAMOUS simulations are also in reasonable agreement with HiGEM and RAPID. The exceptions are CONT, in which the interannual variability is too weak, and STOCH_HIGH_30d, in which it is too strong. Note that on subannual time scales, HiGEM displays too much variability compared to RAPID, to the extent that FAMOUS (even without noise) is a closer match to the observations.

e. Approach to equilibrium

To check the equilibration of the stochastic integrations, we have continued the STOCH_HIGH_30d simulation for a further 250 yr, taking the total duration to 300 yr. According to the time evolution of the global horizontally averaged temperature, the upper 2 km of the ocean equilibrates within the first 50 yr, whereas the deep ocean takes around 100 yr. At equilibrium, the deep ocean temperatures are around −1.4°C, but they remain above both the freezing point (−1.8°C) and the coldest surface temperatures (also −1.8°C). After reaching equilibrium, the
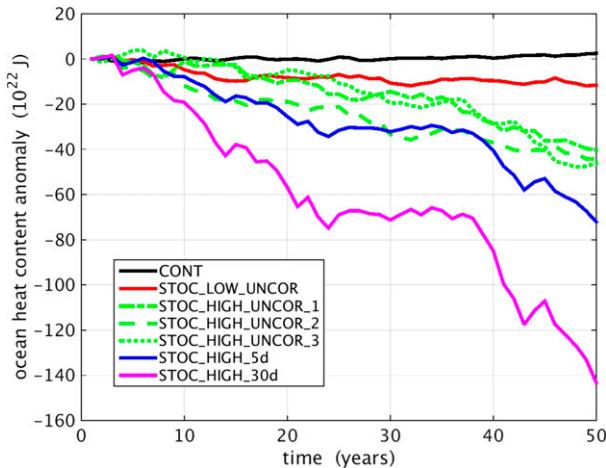
FIG. 10. Time series of the annual-mean global ocean heat content anomaly (with respect to its starting value) in each of the FAMOUS experiments.

simulation continues to run stably (without becoming unphysical) for at least a further 200 yr.

4. Physical interpretation

The changes to the mean climatological state that were identified in section 3 are a manifestation of what, in the field of stochastic dynamical systems, is called noise-induced drift or noise-induced rectification. This effect arises from interactions between the noise and nonlinearities in the model equations. It permits zero-mean noise to have non-zero-mean effects, as seen in our stochastic simulations. The noise-enhanced variability that was identified in section 3 has been documented previously when noise is added to coupled general circulation models used for climate simulation (Williams 2012).

There are two possible physical mechanisms that could account for the net upward heat transport and generation of stratification that occur in the ocean of our stochastic simulations. In the first mechanism, the temperature and density perturbations that are created by the stochastic noise trigger convective instabilities in the deep ocean. These instabilities are then removed by the convective adjustment scheme in FAMOUS (Smith et al. 2008), which mixes water columns vertically in such a manner that convective stability is reestablished. The net effect is an upward heat transport. In the second mechanism, the perturbations that are created by the stochastic noise generate horizontal variations in temperature and density, which are removed by the Gent and McWilliams (1990) parameterization. Again, the net effect is an upward heat transport, because the parameterization acts via an adiabatic rearrangement.

To investigate the physical mechanisms further, the depth profile of the contribution to the temperature tendency from the Gent and McWilliams (1990) parameterization in STOCH_HIGH_UNCOR relative to CONT is shown in Fig. 13. This temperature tendency anomaly has the correct shape to account for the temperature changes in Fig. 6c, indicating relative warming in the upper 1500 m of the ocean and cooling below. Furthermore, the magnitude of the warming reaches $0.13 \times 10^{-9}\,°\mathrm{C\,s}^{-1}$ at a depth of around 1000 m, which equates to a temperature increase of 0.2°C in 50 yr and is also consistent with Fig. 6c. Therefore, we conclude that the second of the two possible mechanisms discussed above is sufficient to account for the net upward heat transport that occurs in the ocean of our stochastic simulations. The same mechanism is presumed to operate in HiGEM, except that it occurs via the explicitly simulated eddies rather than the Gent and McWilliams (1990) parameterization. We speculate that if our stochastic perturbations were vertically uncorrelated, then the first physical mechanism would increase in importance. However, our use of vertically coherent noise inhibits the mechanism in the present simulations.

5. Summary and discussion

This study has found that the simulated climatological state of the ocean is improved in many respects by implementing a simple stochastic parameterization of ocean eddies into a coupled atmosphere–ocean general circulation model. Simulations from a high-resolution, eddy-permitting model (HiGEM) were used to calculate the eddy statistics needed to inject realistic stochastic noise into a low-resolution, non-eddy-permitting version of the same model (FAMOUS). A suite of four stochastic experiments was then run to test the sensitivity of the simulated climate to the noise definition, by varying the noise amplitude and decorrelation time within reasonable limits.

The addition of zero-mean noise to the ocean temperature tendency was found to have a nonzero effect on the mean climate. In terms of the ocean temperature and salinity fields both at the surface and at depth, the noise reduces many of the biases in the low-resolution model and causes it to more closely resemble the high-resolution model, as summarized in Table 2. The change in global ocean heat content caused by the noise is at least as large as the anthropogenic global warming signal. The variability of the strength of the global ocean thermohaline circulation is also improved. We conclude that stochastic ocean perturbations can yield reductions in climate model error that are comparable to those obtained by
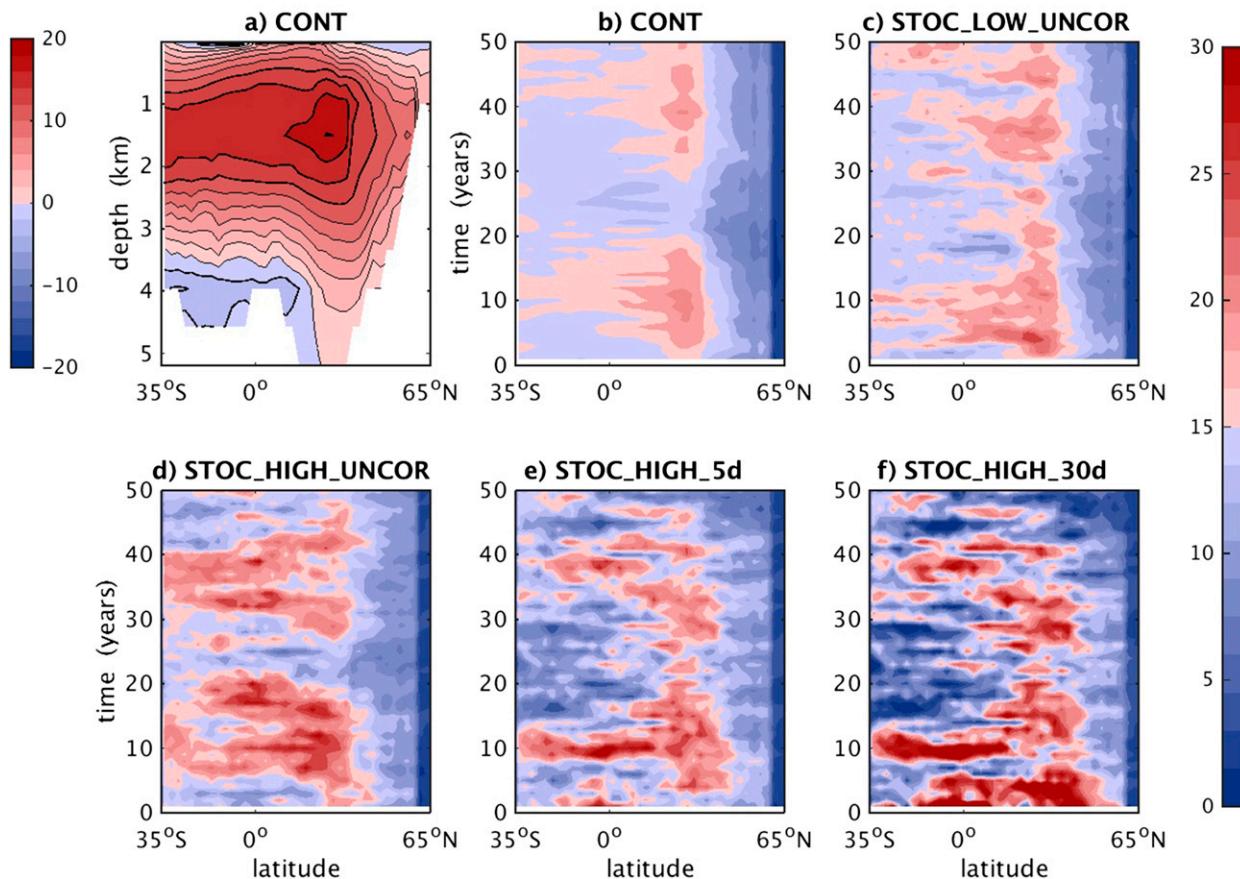
FIG. 11. Atlantic meridional overturning circulation in the FAMOUS experiments. (a) Latitude–depth plot of the streamfunction for the zonally integrated flow (Sv, where $1\,\mathrm{Sv} = 10^{6}\,\mathrm{m}^{3}\,\mathrm{s}^{-1}$; left color bar) in the control simulation averaged over the final decade. (b)–(f) Latitude–time plots of the annual-mean streamfunction for the zonally integrated flow (Sv; right color bar) at a depth of 1000 m in the control simulation and four stochastic experiments. When calculating the streamfunctions, the zonal integration is performed across the entire Atlantic basin.

refining the resolution, but without the increased computational cost.

In this latter respect, our findings are consistent with those of Berner et al. (2012), who studied the model error in an atmospheric general circulation model. They reported that, although the impact of adding stochastic noise is not universally beneficial in terms of model bias reduction, it is nevertheless beneficial across a range of variables and diagnostics. They also reported that, in terms of improving the magnitudes and spatial patterns of model biases, the impact of adding stochastic noise can be similar to the impact of increasing the resolution. Our results are consistent with these findings. We conclude that oceanic stochastic parameterizations join atmospheric stochastic parameterizations in having the potential to significantly improve climate simulations.

It is possible that some of the improvements to the time-averaged climate in FAMOUS that have been documented in this study could alternatively be

obtained by tuning the deterministic parameterizations. For example, the North Atlantic deep water formation rate affects the climatological temperature and salinity patterns and is known to be a function of the vertical diffusivity in ocean models (Schmittner and Weaver 2001). However, even if such improvements to the time-averaged climate were possible by this approach, it would be unlikely to improve the temporal variability. It would also be unlikely to capture any change in the probabilities of transitions between different metastable ocean circulation regimes, which are known to be exhibited by FAMOUS (Hawkins et al. 2011) and to be sensitive to noise (Monahan 2002).

Our results may have implications for centennial-scale climate and Earth system modeling. While many state-of-the-art models planned for phase 6 of the Coupled Model Intercomparison Project (CMIP6) will possess the resolution necessary to permit the largest ocean eddies, if not to fully resolve them, such resolution is
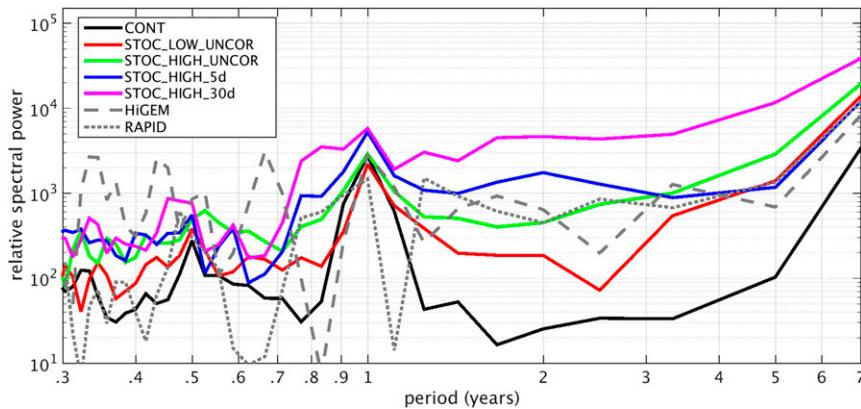
FIG. 12. Power spectra of the maximum value of the volume transport in the Atlantic meridional overturning circulation at 26°N in the FAMOUS simulations. The spectra from the three STOCH_HIGH_UNCOR ensemble members agree well with each other, but to avoid cluttering the figure, only one ensemble member is shown. Power spectra from the HiGEM control simulation and the RAPID observations are included for comparison.

still impractical for some key uses. Implementation of a stochastic parameterization could be of benefit to paleoclimate modeling studies, in which general circulation models must be integrated for centuries or even millennia in order to spin up components such as the deep ocean and coupled ice sheets. Stochastic parameterization could also be of benefit to Earth system models in which biogeochemical complexity necessitates many more tracers and prognostic variables than a standard ocean circulation model, making the model far more computationally demanding. A stochastic parameterization of ocean eddies may enable such models to approach the circulatory behavior of eddy-permitting ocean models without the extra computational cost.

Our results suggest several possible avenues for future work. First, models with ocean components that are structurally different from FAMOUS, such as those in which the vertical discretization is based on isopycnal coordinates rather than depth, might respond differently to our stochastic parameterization. Second, the implementation of our parameterization into Earth system models, such as FAMOUS-C, might improve their behavior and provide an opportunity for more evaluation of processes such as nutrient transport. Finally, our stochastic parameterization could be improved in several respects. For example, we have taken the noise amplitude to decrease with depth, but we have neglected its variations with latitude and longitude. We have taken the eddy decorrelation time scale to be a constant in each stochastic experiment, but this time scale is known to display geographic variations (e.g., Chelton et al. 2011). We have neglected horizontal correlations between neighboring

grid boxes, but it is possible that there will be some geographic locations and depths at which the horizontal eddy correlations are longer than the horizontal grid spacing. We have used vertically coherent noise, but not all eddies are perfectly correlated in depth (Petersen et al. 2013). We have extracted fluctuations from the eddy-permitting ocean model in a purely statistical way, but there are alternative dynamical approaches that can automatically extract fluctuations and project them onto the coarse grid (e.g., Berloff 2005b). Future work should explore refinements to our stochastic parameterization to address all of these limitations.
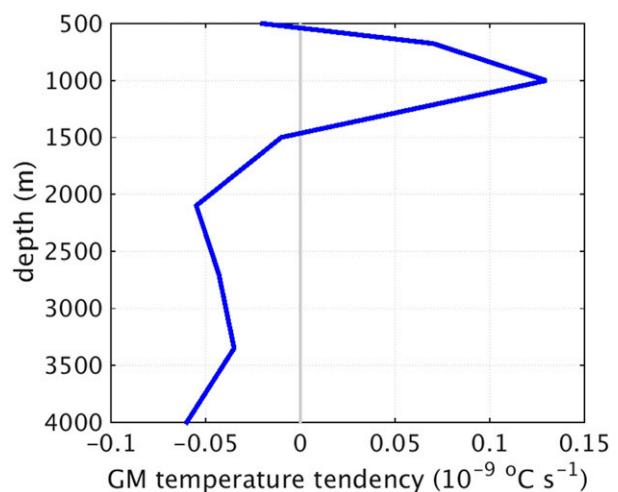

FIG. 13. Depth profile of the anomalous temperature tendency associated with the Gent and McWilliams (1990) scheme in STOCH_HIGH_UNCOR. The anomaly is relative to CONT and is averaged over the global ocean and final decade.

TABLE 2. Summary of FAMOUS model biases in decadal-mean temperature and salinity obtained in this study. The biases are calculated at the surface and in zonal means, and are expressed as pattern correlations and root-mean-square (RMS) differences with respect to CONT and HiGEM. In FAMOUS, the decadal average is taken over the final decade of the simulations.

		RUN					
		CONT	STOC_LOW_UNCOR	STOC_HIGH_UNCOR	STOC_HIGH_5d	STOC_HIGH_30d	HiGEM
Sea surface temperature	RMS difference between RUN and CONT (°C)	0	0.40	0.61	0.90	1.66	1.92
	RMS difference between RUN and HiGEM (°C)	1.92	1.83	1.76	1.67	1.55	0
	Pattern correlation between RUN–CONT and HiGEM–CONT	—	0.39	0.41	0.47	0.55	1
Sea surface salinity	RMS difference between RUN and CONT (psu)	0	0.28	0.48	0.61	0.93	2.82
	RMS difference between RUN and HiGEM (psu)	2.82	2.75	2.64	2.63	2.54	0
	Pattern correlation between RUN–CONT and HiGEM–CONT	—	0.14	0.23	0.20	0.25	1
Global zonal-mean temperature	RMS difference between RUN and CONT (°C)	0	0.28	0.65	1.00	2.10	1.25
	RMS difference between RUN and HiGEM (°C)	1.25	1.00	0.80	0.59	0.46	0
	Pattern correlation between RUN–CONT and HiGEM–CONT	—	0.63	0.82	0.89	0.92	1
Pacific zonal-mean salinity	RMS difference between RUN and CONT (psu)	0	0.06	0.15	0.19	0.42	0.98
	RMS difference between RUN and HiGEM (psu)	0.98	0.88	0.75	0.66	0.54	0
	Pattern correlation between RUN–CONT and HiGEM–CONT	—	0.40	0.56	0.64	0.69	1

REFERENCES

Andrejczuk, M., F. C. Cooper, S. Juricke, T. N. Palmer, A. Weisheimer, and L. Zanna, 2016: Oceanic stochastic parameterizations in a seasonal forecast system. *Mon. Wea. Rev.*, **144**, 1867–1875, doi:10.1175/MWR-D-15-0245.1.

Berloff, P. S., 2005a: Random-forcing model of the mesoscale oceanic eddies. *J. Fluid Mech.*, **529**, 71–95, doi:10.1017/S0022112005003393.

——, 2005b: On dynamically consistent eddy fluxes. *Dyn. Atmos. Oceans*, **38**, 123–146, doi:10.1016/j.dynatmoce.2004.11.003.

——, W. Dewar, S. Kravtsov, and J. McWilliams, 2007: Ocean eddy dynamics in a coupled ocean–atmosphere model. *J. Phys. Oceanogr.*, **37**, 1103–1121, doi:10.1175/JPO3041.1.

Berner, J., T. Jung, and T. N. Palmer, 2012: Systematic model error: The impact of increased horizontal resolution versus improved stochastic and deterministic parameterizations. *J. Climate*, **25**, 4946–4962, doi:10.1175/JCLI-D-11-00297.1.

——, and Coauthors, 2017: Stochastic parameterization: Towards a new view of weather and climate models. *Bull. Amer. Meteor. Soc.*, doi:10.1175/BAMS-D-15-00268.1, in press.

Biri, S., M. G. Scharffenberg, and D. Stammer, 2015: A probabilistic description of the mesoscale eddy field in the ocean. *J. Geophys. Res. Oceans*, **120**, 4778–4802, doi:10.1002/2014JC010681.

Birner, T., and P. D. Williams, 2008: Sudden stratospheric warmings as noise-induced transitions. *J. Atmos. Sci.*, **65**, 3337–3343, doi:10.1175/2008JAS2770.1.

Brankart, J.-M., 2013: Impact of uncertainties in the horizontal density gradient upon low resolution global ocean modelling. *Ocean Modell.*, **66**, 64–76, doi:10.1016/j.ocemod.2013.02.004.

Buizza, R., M. Miller, and T. N. Palmer, 1999: Stochastic representation of model uncertainties in the ECMWF ensemble prediction scheme. *Quart. J. Roy. Meteor. Soc.*, **125**, 2887–2908, doi:10.1002/qj.49712556006.

——, P. L. Houtekamer, Z. Toth, G. Pellerin, M. Z. Wei, and Y. J. Zhu, 2005: A comparison of the ECMWF, MSC, and NCEP global ensemble prediction systems. *Mon. Wea. Rev.*, **133**, 1076–1097, doi:10.1175/MWR2905.1.

Chelton, D. B., M. G. Schlax, and R. M. Samelson, 2011: Global observations of nonlinear mesoscale eddies. *Prog. Oceanogr.*, **91**, 167–216, doi:10.1016/j.pocean.2011.01.002.

Christensen, H. M., I. M. Moroz, and T. N. Palmer, 2015: Simulating weather regimes: Impact of stochastic and perturbed parameter schemes in a simple atmospheric model. *Climate Dyn.*, **44**, 2195–2214, doi:10.1007/s00382-014-2239-9.

Cooper, F. C., and L. Zanna, 2015: Optimisation of an idealised ocean model, stochastic parameterisation of sub-grid eddies. *Ocean Modell.*, **88**, 38–53, doi:10.1016/j.ocemod.2014.12.014.

D'Andrea, F., and Coauthors, 1998: Northern Hemisphere atmospheric blocking as simulated by 15 atmospheric general circulation models in the period 1979–1988. *Climate Dyn.*, **14**, 385–407, doi:10.1007/s003820050230.

Dawson, A., and T. N. Palmer, 2015: Simulating weather regimes: Impact of model resolution and stochastic parameterization. *Climate Dyn.*, **44**, 2177–2193, doi:10.1007/s00382-014-2238-x.

Franzke, C. L. E., T. J. O'Kane, J. Berner, P. D. Williams, and V. Lucarini, 2015: Stochastic climate theory and modeling. *Wiley Interdiscip. Rev.: Climate Change*, **6**, 63–78, doi:10.1002/wcc.318.

Gent, P. R., and J. C. McWilliams, 1990: Isopycnal mixing in ocean circulation models. *J. Phys. Oceanogr.*, **20**, 150–155, doi:10.1175/1520-0485(1990)020<0150:IMIOCM>2.0.CO;2.

Gordon, C., C. Cooper, C. A. Senior, H. Banks, J. M. Gregory, T. C. Johns, J. F. B. Mitchell, and R. A. Wood, 2000: The simulation of SST, sea ice extents and ocean heat transports in a version of the Hadley Centre Coupled Model without flux adjustments. *Climate Dyn.*, **16**, 147–168, doi:10.1007/s003820050010.

Guemas, V., and F. Codron, 2011: Differing impacts of resolution changes in latitude and longitude on the midlatitudes in the LMDZ atmospheric GCM. *J. Climate*, **24**, 5831–5849, doi:10.1175/2011JCLI4093.1.

Hallberg, R., 2013: Using a resolution function to regulate parameterizations of oceanic mesoscale eddy effects. *Ocean Modell.*, **72**, 92–103, doi:10.1016/j.ocemod.2013.08.007.

Hasselmann, K., 1976: Stochastic climate models. Part I. Theory. *Tellus*, **28A**, 473–485, doi:10.1111/j.2153-3490.1976.tb00696.x.

Hawkins, E., R. S. Smith, L. C. Allison, J. M. Gregory, T. J. Woollings, H. Pohlmann, and B. de Cuevas, 2011: Bistability of the Atlantic overturning circulation in a global climate model and links to ocean freshwater transport. *Geophys. Res. Lett.*, **38**, L10605, doi:10.1029/2011GL047208.

Hovmöller, E., 1949: The trough-and-ridge diagram. *Tellus*, **1A**, 62–66, doi:10.1111/j.2153-3490.1949.tb01260.x.

Jansen, M. F., and I. M. Held, 2014: Parametrizing subgrid-scale eddy effects using energetically consistent backscatter. *Ocean Modell.*, **80**, 36–48, doi:10.1016/j.ocemod.2014.06.002.

Jones, C., J. Gregory, R. Thorpe, P. Cox, J. Murphy, D. Sexton, and P. Valdes, 2005: Systematic optimisation and climate simulation of FAMOUS, a fast version of HadCM3. *Climate Dyn.*, **25**, 189–204, doi:10.1007/s00382-005-0027-2.

Jung, T., T. N. Palmer, and G. J. Shutts, 2005: Influence of a stochastic parametrization on the frequency of occurrence of North Pacific weather regimes in the ECMWF model. *Geophys. Res. Lett.*, **32**, L23811, doi:10.1029/2005GL024248.

Kitsios, V., J. Frederiksen, and M. Zidikheri, 2014: Scaling laws for parametrizations of subgrid interactions in simulations of oceanic circulations. *Philos. Trans. Royal Soc. London*, **372A**, 20130285, doi:10.1098/rsta.2013.0285.

Kuhlbrodt, T., J. M. Gregory, and L. C. Shaffrey, 2015: A process-based analysis of ocean heat uptake in an AOGCM with an eddy-permitting ocean component. *Climate Dyn.*, **45**, 3205–3226, doi:10.1007/s00382-015-2534-0.

Levitus, S., J. I. Antonov, T. P. Boyer, and C. Stephens, 2000: Warming of the World Ocean. *Science*, **287**, 2225–2229, doi:10.1126/science.287.5461.2225.

Li, H., and J.-S. von Storch, 2013: On the fluctuating buoyancy fluxes simulated in a 1/10° OGCM. *J. Phys. Oceanogr.*, **43**, 1270–1287, doi:10.1175/JPO-D-12-080.1.

Lin, J. W.-B., and J. D. Neelin, 2002: Considerations for stochastic convective parameterization. *J. Atmos. Sci.*, **59**, 959–975, doi:10.1175/1520-0469(2002)059<0959:CFSCP>2.0.CO;2.

Monahan, A. H., 2002: Stabilization of climate regimes by noise in a simple model of the thermohaline circulation. *J. Phys. Oceanogr.*, **32**, 2072–2085, doi:10.1175/1520-0485(2002)032<2072:SOCRBN>2.0.CO;2.

Open University, 2001: *Ocean Circulation*. 2nd ed. Oceanography Series, Vol. 3, Butterworth–Heinemann, 286 pp.

Palmer, T. N., 2001: A nonlinear dynamical perspective on model error: A proposal for non-local stochastic–dynamic parametrization in weather and climate prediction models. *Quart. J. Roy. Meteor. Soc.*, **127**, 279–304, doi:10.1002/qj.49712757202.

——, and P. D. Williams, 2008: Introduction. Stochastic physics and climate modelling. *Philos. Trans. Royal Soc. London*, **366A**, 2421–2427, doi:10.1098/rsta.2008.0059.

Petersen, M. R., S. J. Williams, M. E. Maltrud, M. W. Hecht, and B. Hamann, 2013: A three-dimensional eddy census of a high-resolution global ocean simulation. *J. Geophys. Res. Oceans*, **118**, 1759–1774, doi:10.1002/jgrc.20155.

Porta Mana, P., and L. Zanna, 2014: Toward a stochastic parameterization of ocean mesoscale eddies. *Ocean Modell.*, **79**, 1–20, doi:10.1016/j.ocemod.2014.04.002.

Randall, D. A., and Coauthors, 2007: Climate models and their evaluation. *Climate Change 2007: The Physical Science Basis*, S. Solomon et al., Eds., Cambridge University Press, 589–662.

Schmittner, A., and A. J. Weaver, 2001: Dependence of multiple climate states on ocean mixing parameters. *Geophys. Res. Lett.*, **28**, 1027–1030, doi:10.1029/2000GL012410.

Scott, R. B., 2003: Predictability of SST in an idealized, one-dimensional, coupled atmosphere–ocean climate model with stochastic forcing and advection. *J. Climate*, **16**, 323–335, doi:10.1175/1520-0442(2003)016<0323:POSIAI>2.0.CO;2.

Shaffrey, L. C., and Coauthors, 2009: U.K. HiGEM: The new U.K. High-Resolution Global Environment Model—Model description and basic evaluation. *J. Climate*, **22**, 1861–1896, doi:10.1175/2008JCLI2508.1.

Shutts, G. J., and T. N. Palmer, 2007: Convective forcing fluctuations in a cloud-resolving model: Relevance to the stochastic parameterization problem. *J. Climate*, **20**, 187–202, doi:10.1175/JCLI3954.1.

Smith, R. S., 2012: The FAMOUS climate model (versions XFXWB and XFHCC): Description update to version XDBUA. *Geosci. Model Dev.*, **5**, 269–276, doi:10.5194/gmd-5-269-2012.

——, J. M. Gregory, and A. Osprey, 2008: A description of the FAMOUS (version XDBUA) climate model and control run. *Geosci. Model Dev.*, **1**, 53–68, doi:10.5194/gmd-1-53-2008.

Srokosz, M. A., and H. L. Bryden, 2015: Observing the Atlantic meridional overturning circulation yields a decade of inevitable surprises. *Science*, **348**, 1255575, doi:10.1126/science.1255575.

Sura, P., and C. Penland, 2002: Sensitivity of a double-gyre ocean model to details of stochastic forcing. *Ocean Modell.*, **4**, 327–345, doi:10.1016/S1463-5003(02)00008-2.

Treguier, A.-M., and Coauthors, 2014: Meridional transport of the salt in the global ocean from an eddy-resolving model. *Ocean Sci.*, **10**, 243–255, doi:10.5194/os-10-243-2014.

Williams, P. D., 2005: Modelling climate change: The role of unresolved processes. *Philos. Trans. Royal Soc. London*, **363A**, 2931–2946, doi:10.1098/rsta.2005.1676.

——, 2012: Climatic impacts of stochastic fluctuations in air–sea fluxes. *Geophys. Res. Lett.*, **39**, L10705, doi:10.1029/2012GL051813.

——, P. L. Read, and T. W. N. Haine, 2003: Spontaneous generation and impact of inertia–gravity waves in a stratified, two-layer shear flow. *Geophys. Res. Lett.*, **30**, 2255, doi:10.1029/2003GL018498.

——, T. W. N. Haine, and P. L. Read, 2004: Stochastic resonance in a nonlinear model of a rotating, stratified shear flow, with a simple stochastic inertia–gravity wave parameterization. *Nonlinear Processes Geophys.*, **11**, 127–135, doi:10.5194/npg-11-127-2004.

——, E. Guilyardi, G. Madec, S. Gualdi, and E. Scoccimarro, 2010: The role of mean ocean salinity in climate. *Dyn. Atmos. Oceans*, **49**, 108–123, doi:10.1016/j.dynatmoce.2009.02.001.

——, M. J. P. Cullen, M. K. Davey, and J. M. Huthnance, 2013: Mathematics applied to the climate system: Outstanding challenges and recent progress. *Philos. Trans. Roy. Soc. London*, **371A**, 20120518, doi:10.1098/rsta.2012.0518.

Williamson, D. L., 1999: Convergence of atmospheric simulations with increasing horizontal resolution and fixed forcing scales. *Tellus*, **51A**, 663–673, doi:10.1034/j.1600-0870.1999.00009.x.

——, 2008: Convergence of aqua-planet simulations with increasing resolution in the Community Atmospheric Model, version 3. *Tellus*, **60A**, 848–862, doi:10.1111/j.1600-0870.2008.00339.x.

Zavala-Garay, J., A. M. Moore, C. L. Perez, and R. Kleeman, 2003: The response of a coupled model of ENSO to observed estimates of stochastic forcing. *J. Climate*, **16**, 2827–2842, doi:10.1175/1520-0442(2003)016<2827:TROACM>2.0.CO;2.