Ocean Modelling 88 (2015) 38–53

Contents lists available at ScienceDirect

Ocean Modelling

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

Optimisation of an idealised ocean model, stochastic parameterisation of sub-grid eddies

Fenwick C. Cooper *, Laure Zanna

Atmospheric, Oceanic and Planetary Physics, Department of Physics, University of Oxford, Oxford OX1 3PU, UK

ARTICLE INFO

Article history:
Received 26 May 2014
Received in revised form 15 December 2014
Accepted 21 December 2014
Available online 27 February 2015

Keywords:
Optimisation
Stochastic
Parameterisation
Sub-grid
Eddy

ABSTRACT

An optimisation scheme is developed to accurately represent the sub-grid scale forcing of a high dimensional chaotic ocean system. Using a simple parameterisation scheme, the velocity components of a 30 km resolution shallow water ocean model are optimised to have the same climatological mean and variance as that of a less viscous 7.5 km resolution model. The 5 day lag-covariance is also optimised, leading to a more accurate estimate of the high resolution response to forcing using the low resolution model.

The system considered is an idealised barotropic double gyre that is chaotic at both resolutions. Using the optimisation scheme, we find and apply the constant in time, but spatially varying, forcing term that is equal to the time integrated forcing of the sub-grid scale eddies. A linear stochastic term, independent of the large-scale flow, with no spatial correlation but a spatially varying amplitude and time scale is used to represent the transient eddies. The climatological mean, variance and 5 day lag-covariance of the velocity from a single high resolution integration is used to provide an optimisation target. No other high resolution statistics are required. Additional programming effort, for example to build a tangent linear or adjoint model, is not required either.

The focus of this paper is on the optimisation scheme and the accuracy of the optimised flow. However the forcing can provide insights in the design of deterministic and stochastic parameterisations. In the present study, we found that the stochastic parameterisation correcting the model variance is associated with the spatial pattern of eddy-decorrelation timescales rather than the spatial pattern of the amplitude of the variance. The method can be applied in future investigations into the physical processes that govern barotropic turbulence and it can perhaps be applied to help understand and correct biases in the mean and variance of a more realistic coarse or eddy-permitting ocean model. The method is complementary to current parameterisations and can be applied at the same time without modification.

1. Introduction

Due to the limitations of finite computational power, current numerical methods are not capable of accurately resolving the ocean circulation down to the viscous scale. Since there exists no universal sub-grid scale turbulence model that can close for all unresolved quantities (Reynolds stresses, turbulent fluxes, etc.) ad hoc representations are required, and state of the art numerical models exhibit serious differences and inaccuracies in their climatologies (e.g. Flato et al. (2013, Section 9.4.2)). The simplest approach to parameterise sub-grid scale processes is to dissipate any small-scale motion while simultaneously stabilising the model. This is typically achieved by employing an eddy diffusivity designed, for example, to improve spectral characteristics near the grid-scale (e.g. Smagorinsky (1963), Leith (1967)), or by using a diffusive integration scheme (e.g. Ritchie (1988)). Another approach is to mimic the physical processes in the real ocean. For example, mesoscale eddies in the ocean interior tend to rearrange fluid parcels along isopycnals (constant density surfaces) which leads to the widely implemented Gent–McWilliams parameterisation scheme in the tracer equations (Gent et al., 1990). Such approaches to find the sub-grid momentum or buoyancy forcing are often based upon the time-mean effect of the sub-grid scale forcing upon the large scale flow as diagnosed by comparing a low resolution model with measurements, or a high resolution integration. The approximate functional form of the sub-grid momentum or buoyancy forcing in terms of the grid scale flow of a turbulent system may be found using high resolution integrations (e.g. Achatz and Branstator (1999)), using, for example, a polynomial fit. A stochastic term

* Corresponding author. Tel.: +44 7531 290 687.
E-mail address: Fenwick@LittleStick.com (F.C. Cooper).

http://dx.doi.org/10.1016/j.ocemod.2014.12.014
1463-5003/

may be used to represent the fit residual (e.g. Wilks (2005), Frederiksen and Kepert (2006), Zidikheri and Frederiksen (2009), Kitsios et al. (2013), Arnold et al. (2013)) or realistic variance (e.g. Hasselmann (1976), Buizza et al. (1999), Palmer (2001), Berloff (2005)). The deterministic and stochastic sub-grid forcing can be derived from theoretical considerations (e.g. Kraichnan (1959), Herring and Kraichnan (1972), Frederiksen and Davies (1997), Holm et al. (1999), Marshall et al. (2012), Grooms and Majda (2013), Mana and Zanna (2014)), although such an approach can be practically difficult to implement (Foiaş et al., 2001; Mana and Zanna, 2014).

In many cases one or more parameters that govern the strength of these schemes must be chosen with limited guidance from theory. Trial and error comparison of model output, as a function of parameter values, with ocean data, is one method often referred to as "tuning". Tuning up to around five or six parameters is possible in principle with a simple parameter search, however larger numbers of parameters require an "expert" opinion, independence from each other, or an acceptance that the optimal values will not be found. The problem is that to explore each direction of a parameter space of dimensionality d across n different climatologies requires n^d points to be evaluated. Given a high ($d \gg 5$) dimensional vector $\mathbf{p}$ of parameters, evaluation of the entire parameter space is not practical and in order to optimise anything we are forced to define an objective target to optimise for, or in other words a cost function $G(\mathbf{p})$ to minimise. From an initial guess $\mathbf{p}_0$ a direction to change $\mathbf{p}$ may be given by the gradient of the cost function

$$\mathbf{p}_1 = \mathbf{p}_0 - \left.\frac{\partial G(\mathbf{p})}{\partial \mathbf{p}}\right|_{\mathbf{p}=\mathbf{p}_0} \delta p. \tag{1}$$

Here $\mathbf{p}_1$ is an improved estimate of the optimal parameters in comparison with $\mathbf{p}_0$ and δp is a small positive constant with the appropriate units. The process can be iterated until no further optimisation is possible. Accurate estimation of $\partial G/\partial \mathbf{p}$ can be difficult, requiring for example tangent linear and adjoint models to be integrated. Implementation of adjoint models for ocean circulation problems has been achieved for sensitivity analysis and data assimilation capabilities (e.g. Marotzke et al. (1999), Moore et al. (2004)) as has optimisation of the eddy-buoyancy sub-grid parameters from the climatological mean state (Ferreira et al., 2005). However there are still some unresolved issues for large-scale chaotic systems. Firstly the programming effort is substantial leading to the development of semi-automatic differentiation packages for this purpose (e.g. Giering (1999), Heimbach et al. (2005)). Secondly if a system has a stochastic element the problem of optimising stochastic parameterisation has not, to the author's knowledge, been considered. Finally, although the adjoint approach is useful for short time optimisation in ocean (e.g. Gebbie et al. (2006), Mazloff et al. (2010), Balmaseda et al. (2013)) and atmosphere (e.g. Kalnay et al. (1996), Dee et al. (2011)) state estimation, it is not currently capable of optimising for the long time climate averages of a chaotic system (e.g. Lea et al. (2000), Eyink et al. (2004)) and approximations are required. Some attempts to solve this problem in a slightly different context include the methods of Abramov and Majda (2009) applied to climate response, who use the full non-linear model for the short time gradient estimate and a Gaussian model approximation for longer times, and Wang et al. (2014) who uses a modified adjoint algorithm to stabilise the gradient estimation algorithm. Fortunately an estimate of $\partial G/\partial \mathbf{p}$ does not need to be particularly accurate for the purposes of optimisation. It is merely required to follow a trajectory in parameter space that eventually leads toward the optimum and to tend to zero as the optimum is approached. Therefore we have the opportunity to optimise with a much simpler criteria if a very approximate direction of $\partial G/\partial \mathbf{p}$ can be found. This is the approach of the present paper. In our case, with the climate change problem in mind, the goal is accurate optimisation of the climatological mean and variance and approximate optimisation of the response of the system to a forcing, using a "truth" as the optimisation target.

1.1. The mean

Current state of the art ocean models exhibit a different climatological mean state to that observed in the real ocean (Flato et al., 2013). For example, the poor representation of eddy-mean flow processes leads to unrealistic western boundary currents (Gulf Stream and Kuroshio) responsible for large sea surface temperature biases (Large and Danabasoglu, 2006). Their predictions are therefore approximations about a different climatological mean point in state space to that of reality. To account for such deviations from the observed climatology, post integration bias correction is sometimes applied (e.g. Stockdale (1997)). A more accurate approach would be to have a model that has the correct climatological mean state in the first place. This can be achieved for example by adding a spatially varying, but constant in time, parameter to the right hand side of the governing equations (Achatz and Branstator, 1999). This spatially-varying time-independent parameter represents the contribution to the climatological mean of all of the sub-grid processes that are not included in the basic low resolution model minus any biases introduced by incorrect additional terms, such as high viscosity. The size of the improvement in accuracy relative to post integration bias correction can be important. For example, in a coupled ocean–atmosphere model some studies suggest that the mean location of the ocean boundary currents have an important impact upon atmospheric dynamics (e.g. Kirtman and Vecchi (2011), Scaife et al. (2011)). The ocean bias therefore has the potential to cause atmospheric bias that may be difficult to correct post integration.

1.2. The variance

Often, due to artificially high viscosity in a dynamical ocean model and the lack of sub-grid variability, the variance of the prognostic variables is underestimated. Without a time dependent external forcing such as the seasonal cycle, one can often obtain a steady state in very low-resolution ocean models, where time derivatives of all prognostic variables are equal to zero. In non-eddying ocean models, any effect of the variance due to eddies is therefore reduced or missing. The fluctuations brought about by resolving the eddies in an ocean model can potentially lead to additional dynamical regimes being explored (e.g. Palmer (2001), Palmer and Weisheimer (2011)) and important processes such as eddy saturation (Straub, 1993; Munday et al., 2013) or jet rectification (Berloff, 2005; Waterman et al., 2012; Waterman and Hoskins, 2013). In addition, the lack of variance between the members of an ensemble of model integrations contributes to over confidence, in a statistical sense, in model predictions. For these reasons we consider it desirable for our model climatological variance, and hence the turbulent eddy kinetic energy, to be as close as possible to the measured ocean variance. Moreover, since the correlations of a turbulent system decay in time, we would like the correlations of any parameterised source of variance to also decay after some time. The simplest approach is to add a stochastic term, with a spatially varying amplitude and time scale, to the right hand side of the governing equations. In this paper we require that the parameters governing such a process ensure that the model's climatological variance is as accurate as possible, relative to the "truth".

1.3. The response to forcing

For climate change experiments or for seasonal forecasts, the accuracy of a model's response to forcing is an important characteristic. Unfortunately, unlike in the case of the mean and the variance, we can not always directly measure the true response of a system to a particular forcing and compare it with our low resolution model response. However we would like to be able to predict in advance the change in the climatological mean and variance due to forcing. A forcing change may be, for example, an increase in the concentration of carbon dioxide or in wind stress over the Southern Ocean (e.g. Fyfe et al. (2007)).

The fluctuation–dissipation theorem (e.g. Marconi et al. (2008)) guarantees, under very general assumptions, that the sensitivity to a small forcing is equal to a time integral of the correlation of some function of the prognostic variables. Thus fluctuation–dissipation theorem provides a method by which we can relate the response of a complex system to variables that we can measure. In the case of a single variable linear stochastic system the sensitivity is simply the time integration of its autocorrelation. For application to a general circulation model see Gritsun and Branstator (2007). Unfortunately the response according to its autocorrelation function when assuming a linear stochastic model may not be accurate (Cooper and Haynes, 2015) and the true response of a chaotic nonlinear system to forcing is not straightforward to evaluate using the fluctuation–dissipation theorem (Cooper and Haynes, 2011; Cooper et al., 2013). Despite those limitations, one can reasonably assume that a model will respond in a more similar manner to the truth it represents, the more the autocorrelation functions of the respective systems are alike. We therefore require a parameterisation that pushes the model to have an autocorrelation function to be close to truth. Given that a stochastic term is characterised by an amplitude and a time scale, it can therefore be used to adjust both the variance and lag-covariance of a model towards the true climatological values. In our case the stochastic term is represented in a discretized model by a system of linear stochastic ordinary differential equations

$$d\xi = \mathbf{B}\xi dt + \mathbf{V}\sqrt{\mathbf{D}}d\mathbf{w}. \tag{2}$$

Here $\mathbf{B}$ is a constant matrix, restricted to have all negative eigenvalues so that the system is stable, ξ is a vector of model grid points, each element of which is to be added to the right hand side of the discretized governing equations, $\mathbf{w}$ denotes a vector of independent Wiener or white noise processes with unit variance. The covariance of the stochastic term is set by the matrix $\mathbf{Q}$ with the matrix of eigenvectors of $\mathbf{Q}$ denoted by $\mathbf{V}$ and diagonal matrix of eigenvalues by $\mathbf{D}$. Changing the values of $\mathbf{B}$ and $\mathbf{Q}$ allows us to set the integral of the autocorrelation to approximately the true value, and hence improve the accuracy of the model's response to a small forcing, see Fig. 1.

1.4. A simple parameterisation scheme

In summary we propose adding a spatially-dependent but constant in time forcing term (Section 1.1) and a linear stochastic term (Sections 1.2 and 1.3) to the equations governing a turbulent fluid. The aim is to reproduce some of the statistics of a low viscosity fluid with a high viscosity numerical integration. In the present work, we focus on finding the idealised mean sub-grid barotropic eddy forcing terms in an ocean double gyre configuration. The forcing is limited to the momentum sub-grid forcing in a shallow-water model, unlike the study of Ferreira et al. (2005) which tackled the eddy buoyancy forcing in a primitive equation model. In this study, geostrophic balance ensures that a constant forcing of the momentum equations acts to correct both velocity and sea surface height biases. The approach is similar to a standard linear relaxation where prognostic variables are linearly forced towards some basic state (e.g. in the configuration described in Held and Suarez (1994)). However, in our case the variable being relaxed, in a stochastic sense, is independent of the unparameterised prognostic variables. Berloff (2005) applies a similar scheme with parameters diagnosed from single high and low resolution integrations to a quasi-geostrophic system. The novel approach used in the present work is to optimise the values of these parameters and to develop a method that is applicable to complex models involving more than one prognostic variable. We optimise a low resolution model in such a way as to improve its characteristics when compared with the truth as defined by observations or a high resolution

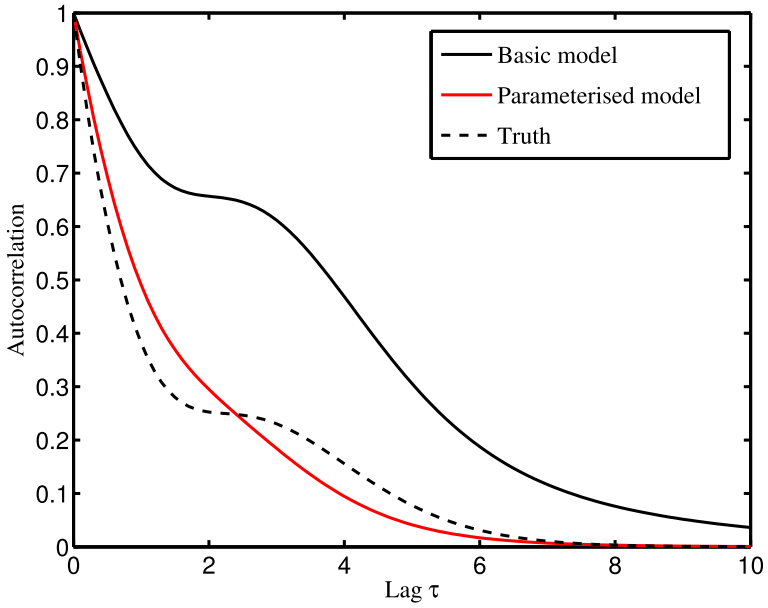

Fig. 1. In this hypothetical illustration, the integral of the autocorrelation function of a model without parameterisation, is much larger than the integral of the autocorrelation function of the true system. Therefore we might expect that the model's climatological response to forcing is somewhat greater than the true climatological response to forcing. After adding a linear stochastic parameterisation with an appropriate choice of parameters, the integral of the autocorrelation function of the model and the truth is identical, presumably leading to a more similar climatological response.

integration. Our scheme involves empirically derived parameters and therefore we refer to it as a parameterisation scheme. However given the number of parameters, three for each grid cell, that the scheme captures the bulk effect of unspecified sub-grid scale processes as opposed to mimicking specific processes and the fact that no dependence on the large-scale flow or external parameters is found, some readers may not consider it to be a parameterisation scheme in the conventional sense. Calling it a bias correction scheme may be more appropriate, albeit in a free running model.

Our choice of an idealised barotropic ocean gyre as test model is primarily motivated by the more general, and to the authors knowledge unsolved, optimisation problem for the long time statistics of a high dimensional chaotic system. We hope that our particular configuration, described in section 2, contains no special features that make the optimisation problem solvable in only this case. However further investigations would be required to check the applicability in other physical configurations such as a periodic channel model and moving from a two dimensional, to a three dimensional fluid.

Integrations performed using a shallow water model are detailed in Section 2 and the optimisation method in Section 3. In Section 4, we compare the high and low resolution integrations without parameterisation with the optimised low resolution system and investigate the effect of the parametrisation scheme.

2. The truth and model integrations

Our choice of an idealised ocean model is motivated by the ocean double gyre setup of Berloff (2005) although rather than integrate the quasi-geostrophic equation, we are integrating the shallow water equations with a linear free surface. To reduce the computational cost, our system is barotropic and therefore omits the baroclinic modes and their interaction with the mean flow. Although the resulting setup is idealised it is sufficient to illustrate our parameterisation scheme.

The integration domain (Fig. 2) is a 3840 km by 3840 km rectangular box located on a mid-latitude beta plane with a flat bottom. The flow is forced by a (time) constant surface wind which varies latitudinally (see Fig. 2) but is uniform in longitude. It is dissipated by a viscous term whose magnitude depends upon the resolution of the model grid. The ocean floor and horizontal boundaries are free slip. Berloff's domain was uniformly 3 km deep, however for the magnitude of forcing chosen, a 3 km deep barotropic system exhibits extremely long timescales at low resolution. Such long timescales makes the experiments rather expensive. A simple solution to reduce the cost is to increase the turbulent nature of the flow that in turn reduces the correlation times present and makes estimation of the autocorrelation computationally feasible. The system is made more turbulent by reducing the ocean depth to a constant 500 m and increasing the wind forcing by a factor of three. Unfortunately this more chaotic system might resemble even less a realistic ocean gyre but nevertheless should not affect the main conclusions.

The equations governing the time evolution of the prognostic variables, zonal velocity u (x direction), meridional velocity v (y direction) and sea surface height η are

$$\frac{\partial u}{\partial t}+u\frac{\partial u}{\partial x}+v\frac{\partial u}{\partial y}-fv+g\frac{\partial \eta}{\partial x}-\kappa\nabla^2 u=\frac{1}{\rho_0 h}(\alpha(y)+F_u(x,y)+\xi_u(x,y,t)), \tag{3}$$

$$\frac{\partial v}{\partial t}+u\frac{\partial v}{\partial x}+v\frac{\partial v}{\partial y}+fu+g\frac{\partial \eta}{\partial y}-\kappa\nabla^2 v=\frac{1}{\rho_0 h}(F_v(x,y)+\xi_v(x,y,t)), \tag{4}$$

$$\frac{\partial \eta}{\partial t}+h\left(\frac{\partial u}{\partial x}+\frac{\partial v}{\partial y}\right)=0, \tag{5}$$

where the acceleration due to gravity $g = 9.81\ \mathrm{m\,s^{-2}}$, density of water $\rho_0 = 999.8\ \mathrm{kg\,m^{-3}}$, depth $h = 500$ m and the Coriolis parameter $f = f_0 + \beta y$ with $\beta = 2\times10^{-11}\ \mathrm{m^{-1}s^{-1}}$ and $f_0 = 4.46\times10^{-5}\ \mathrm{s^{-1}}$. $\alpha(y)$ represents a zonal wind forcing of the form

$$\alpha(y)=\alpha_0\left[\cos\left(\frac{2\pi(y-L/2)}{L}\right)+2\sin\left(\frac{\pi(y-L/2)}{L}\right)\right], \tag{6}$$

where the domain width in each direction $L = 3840$ km with $0\leqslant y\leqslant L$, $0\leqslant x\leqslant L$ corresponding to a domain bounded between latitudes 17.8° N and 56.3° N. The constant $\alpha_0 = 0.12$ Pa. The model Eqs. (3)–(5) are discretized onto a uniform Cartesian Arakawa C-grid (Arakawa and Lamb, 1977), initialised to zero or with low amplitude random noise, and integrated using the MITgcm (Marshall et al., 1997).

In the parameterised low resolution system the constant in time but spatially varying forcing $F_u(x,y)$ and $F_v(x,y)$, represented by the vectors $\mathbf{f}_u$ and $\mathbf{f}_v$, are found by optimisation. $\xi_u(x,y,t)$ and $\xi_v(x,y,t)$ are stochastic terms, represented by $\boldsymbol{\xi}_u$ and $\boldsymbol{\xi}_v$. $\boldsymbol{\xi}_u$ and $\boldsymbol{\xi}_v$ would, in a fairly general case, be governed by Eq. (2) with an appropriate

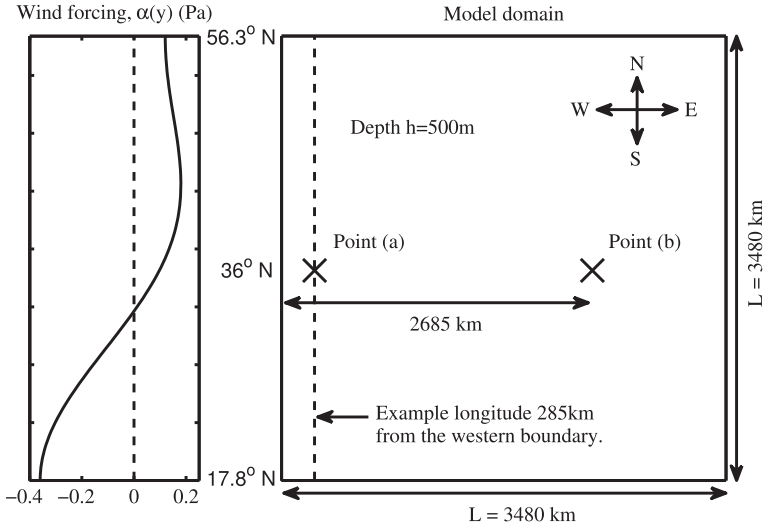

Fig. 2. The wind forcing profile and integration domain, see text. The autocorrelation of the sea surface height, η at points (a), (b) is presented in Fig. 3. The climatological mean sea surface height and eddy kinetic energy at the longitude indicated is presented in Figs. 6 and 9 respectively.

choice of $\mathbf{B}_u$, $\mathbf{B}_v$, $\mathbf{Q}_u$ and $\mathbf{Q}_v$ found via the optimisation. We consider the simpler case where all elements of ξ_u and ξ_v are independent of each other (uncorrelated in space) and the equations for their time evolution reduce to

$$\mathrm{d}\xi_{ui} = b_{ui}\xi_{ui}\mathrm{d}t + \sqrt{q_{ui}}\mathrm{d}w_{ui} \tag{7}$$

and

$$\mathrm{d}\xi_{vi} = b_{vi}\xi_{vi}\mathrm{d}t + \sqrt{q_{vi}}\mathrm{d}w_{vi} \tag{8}$$

for $i = 1 \ldots d$ where d is the number of grid cells in the integration and w_{ui} and w_{vi} are white noise processes with unit variance. In (7) and (8) the variables are the elements of the vector quantities $\boldsymbol{\xi}_u = (\xi_{u1}, \xi_{u2}, \ldots)$, $\boldsymbol{\xi}_v = (\xi_{v1}, \xi_{v2}, \ldots)$ and the constants $\mathbf{b}_u = (b_{u1}, b_{u2}, \ldots)$, $\mathbf{b}_v = (b_{v1}, b_{v2}, \ldots)$, $\mathbf{q}_u = (q_{u1}, q_{u2}, \ldots)$ and $\mathbf{q}_v = (q_{v1}, q_{v2}, \ldots)$ represent the diagonals of $\mathbf{B}_u$, $\mathbf{B}_v$, $\mathbf{Q}_u$ and $\mathbf{Q}_v$ respectively. Test integrations show that the addition of $\xi_u(x, y, t)$ and $\xi_v(x, y, t)$ in the coarse resolution set-up, resulting from Eqs. (7) and (8), do not cause the model to have a systematic long term drift.

The high resolution "truth" has a grid spacing of $\Delta x = \Delta y = 7.5$ km corresponding to 512 by 512 grid cells, a viscosity of $\kappa = 10\ \mathrm{m}^2\ \mathrm{s}^{-1}$ and $f_u = f_v = \xi_u = \xi_v = 0$. The truth is integrated for 10^4 days after discarding a spin up of 10^3 days. 10^4 days is chosen so as to provide estimates of the climatological mean, variance and lag-covariance with sufficient accuracy. A grid spacing of $\Delta x = \Delta y = 30$ km corresponding to 128 by 128 grid cells, a viscosity of $\kappa = 470.23\ \mathrm{m}^2\ \mathrm{s}^{-1}$ integrated for the same time is used to represent a low resolution "model" of the "truth". The viscosities are chosen for stability to yield a Munk layer width M_w along the western boundary of at least three grid cells, $M_w = \pi(\kappa/\beta)^{1/3} > 3\Delta x$. Throughout this paper we refer to high 7.5 km resolution integrations as the truth and low 30 km resolution integrations as the model. The initial model without additional constant or stochastic forcing is referred to as the unparameterised model and the optimised model with the additional constant and stochastic forcing is referred to as the parameterised model. To compare the model with the truth, linearly interpolated values of the truth integration at the locations of the low resolution model grid cell variables are used.

3. The optimisation algorithm

For the discretized u and v fields, dropping the subscript for simplicity, the parameters $\mathbf{f}$ (representing a time-independent spatially varying forcing), $\mathbf{b}$ and $\mathbf{q}$ (governing the respective timescale and amplitude of the stochastic process), are estimated using an iterative process. These vectors are initialised to zero, $\mathbf{f}^0 = 0$, $\mathbf{b}^0 = 0$ and $\mathbf{q}^0 = 0$ where the superscript indicates an iteration number n. For each iteration, the low resolution model is integrated. After an initial spin up period, the model climatological mean of u or v at each grid point, denoted $\mathbf{m}^n_{\mathrm{model}}$, is estimated by integrating over a sufficiently long time (in our case 10^4 days). The model climatological mean is compared with the true climatological mean vector $\mathbf{m}_{\mathrm{truth}}$ and $\mathbf{f}$ is updated as follows

$$\mathbf{f}^{n+1} = \mathbf{f}^n + \left(\mathbf{m}_{\mathrm{truth}} - \mathbf{m}^n_{\mathrm{model}}\right)\delta f, \tag{9}$$

where δf is a suitably small positive constant. A similar procedure is followed for updating $\mathbf{q}$

$$\mathbf{q}^{n+1} = \mathbf{q}^n + \left(\boldsymbol{\sigma}^2_{\mathrm{truth}} - \left(\boldsymbol{\sigma}^n_{\mathrm{model}}\right)^2\right)\delta q, \tag{10}$$

where $\boldsymbol{\sigma}^n_{\mathrm{model}}$ is the standard deviation of the low resolution parameterised model u or v at each grid point measured at the n'th iteration, $\boldsymbol{\sigma}_{\mathrm{truth}}$ is the standard deviation of the true system and δq is again a suitably small positive constant. The minimum value that $\mathbf{q}$ can obtain is clipped at zero, corresponding to no stochastic forcing at the relevant grid point.

As mentioned in Section 1.3, we wish to set $\mathbf{b}$ such that the integral of the lag-covariance (or equivalently the integral of the autocorrelation function), of the model and the truth are the same. From data, estimates of the integral of the lag-covariance are not as accurate as estimates of the lag-covariance at a particular lag (Cooper and Haynes, 2015). Therefore we aim to choose a single lag τ and optimise $\mathbf{b}$ so that the model and the truth have the same autocorrelation at this lag. The idea is that if the model and the truth have the same autocorrelation at this lag, they also approximately have the same integral of the autocorrelation. The autocorrelation of the stochastic term is an exponential decay. It turns out that choosing a lag that is too small leads to an overestimate of the integrated lag-covariance because the initial decay in the truth autocorrelation function is slower than exponential (DelSole, 2000). On the other hand the uncertainty in the autocorrelation as a percentage increases with lag, so a lag that is too large leads to a large uncertainty in the decorrelation time. By plotting the autocorrelation of a selection of the high resolution variables against lag (as in Fig. 3) over exponentials with various decay constants, we estimate that the two curves meeting at a lag $\tau = 5$ days gives a reasonable, but imperfect, exponential approximation to the autocorrelation. This is similar to the value that would be obtained by Berloff (2005), see their Fig. 4. Fig. 3 demonstrates that the autocorrelation varies across the domain, so as an alternative to our choice of $\tau = 5$ days, it might be more reasonable to assume that the appropriate τ to use is found when the autocorrelation first reaches 0.6 or some other reasonable value.

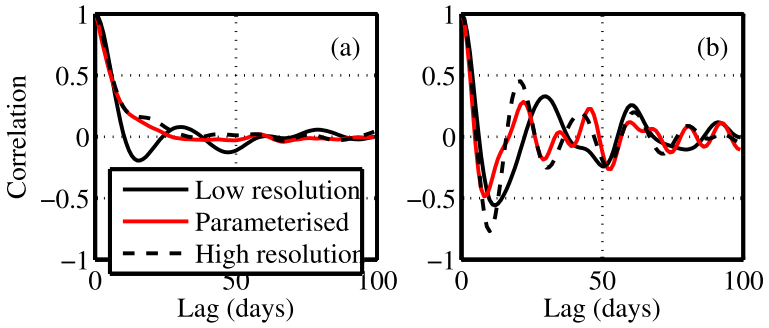

Fig. 3. The autocorrelation of the sea surface height of each low resolution model and the high resolution truth at 36° latitude, (a) 285 km and (b) 2685 km from the western boundary, see Fig. 2. Point (a) is right in the centre of the most turbulent part of the domain where the largest value of the mean eddy kinetic energy is found (compare with Berloff (2005), Fig. 4(a)). Point (b) is in a region dominated by wave activity, the autocorrelation function is more oscillatory with the oscillations taking longer to decay (compare with Berloff (2005), Fig. 4(b)).

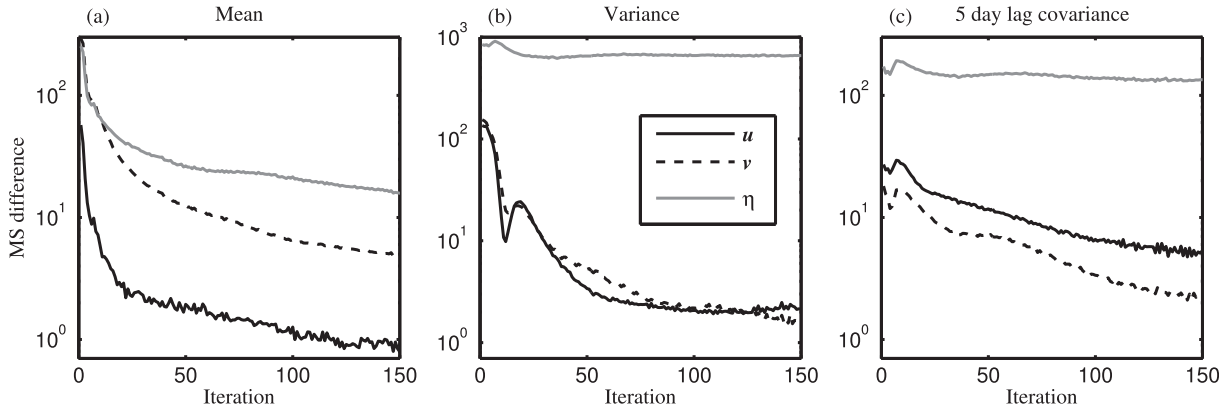

Fig. 4. The mean squared (MS) difference between the parameterised model and high resolution truth climatological (a) mean, (b) variance and (c) 5 day lag covariance for the zonal velocity u (solid), meridional velocity v (dashed) and sea surface height η (grey). Note the logarithmic y axis. The unparameterised model represents the first iteration.

It can be shown, see Appendix A, that for linear systems, $-1/b_i$ is approximately linearly related to the lag covariance introduced to the system, where b_i is element i of $\mathbf{b}$. Therefore the procedure for updating $\mathbf{b}$ is given by

$$\frac{-1}{b_i^{n+1}} = \frac{-1}{b_i^{n}} + \left(c_{i,\text{truth}}(\tau) - c^n_{i,\text{model}}(\tau)\right)\delta b \tag{11}$$

for $i = 1 \ldots d$. Here $c_{i,\text{truth}}(\tau)$ and $c^n_{i,\text{model}}(\tau)$ represent the lag τ covariance, independent at each grid point, for the respective truth and n'th model iteration. δb once again is a sufficiently small positive constant. $-1/b_i$ is constrained to be always greater than 600 s. Each element of $\mathbf{b}$ is therefore always negative and the stochastic systems (7) and (8) are guaranteed to be stable.

The constants $\delta f = 0.1\ \text{kg m}^{-2}\ \text{s}^{-1}$, $\delta q = 0.1\ \text{kg}^2\ \text{m}^{-4}\ \text{s}^{-3}$, and $\delta b = 5.0\ \text{m}^{-2}\ \text{s}^3$, were found by trial and error to lead to convergence. If they are too small, the algorithm is too slow to converge and if they are too large this Euler type method is unstable. The model's initial condition at each iteration of the optimisation is the final state taken from the end of the previous iteration. Using this method it was found that it was only necessary to discard a spin up of 500 days at the start of each iteration. When the difference between the truth and model climatological state at a grid point becomes smaller than the uncertainty in the mean, variance and 5 day lag covariance, further optimisation is not possible. The integration length governs this uncertainty. Longer (and more computationally expensive) integrations than the 10^4 days used here is found to lead to faster convergence per iteration and the optimisation converges to a more accurate state.

We do not make use of additional assumptions, such as divergence free forcing and further investigations, beyond our current scope, are necessary to understand the importance of various conservation laws. The integrations reported here did not exhibit any obvious long term drifts.

4. Results

In our barotropic double gyre configuration, the parameters control the sub-grid eddy momentum forcing. The optimisation algorithm finds a slightly improved set of parameters at each iteration. With the first point corresponding to the unparameterised model, the mean squared differences between the high resolution and parameterised climatological mean, variance and 5 day lag covariance are plotted with a logarithmic y axis in Fig. 4. Using this metric, the difference between the parameterised and true climatology has been substantially reduced when compared with the unparameterised low resolution system. After 150 iterations the mean u, v and η fields are continuing to improve; however the variance of all three fields seems to have reached a plateau. The 5 day lag covariance of the u and v fields are still improving while for η it seems to have also plateaued. The minimal improvement in η, which has no sub-grid forcing applied to its governing equation, is examined in more detail in Section 4.2.

4.1. The climatological mean

After 150 iterations, Fig. 5 shows that the climatological mean state of the low resolution parameterised model (panels (d)–(f)) is significantly closer to the high resolution mean (panels (g)–(i)) when compared with the low resolution model without parameterisation (panels (a)–(c)). The large values of u, v and η near the western boundary are correctly reproduced and the magnitude and shape of the zonal barotropic jet has been improved. Taking a single longitude of the climatological mean sea surface height 285 km from the western boundary, Fig. 6 demonstrates that the parameterised model adequately mimics the truth. For example at 40° latitude, 285 km from the western boundary, η is 51.3, −53.9 and −57.8 cm for the unparameterised, parameterised and truth integrations respectively. Note the change in sign. The relatively small difference between the high resolution and parameterised integrations, spread unevenly over the whole domain, stem from the fact that the system is chaotic. Estimation of the climatological mean of a chaotic system is subject to some error, proportional to $1/\sqrt{t_{\max}}$ where $t_{\max}$ is the integration length. Therefore increasing $t_{\max}$ reduces this error. Constant forcing of the u and v fields leads to optimisation of the η field because of approximate geostrophic balance. Additional forcing of η by adding a spatially dependent term, $F_\eta(x, y)$, to (5) introduces additional sources and sinks of mass, but may also lead to small improvements because neither geostrophic balance or the numerical scheme used is exact.

4.2. The climatological variance

Figs. 7 and 8 show that the low resolution unparameterised model (panels (a)–(c)) has a much lower climatological variance and 5 day lag covariance than the true high resolution model (panels (g)–(i)). The optimisation has succeeded in increasing the variance and 5 day lag covariance of the low resolution parameterised model to be much closer to that of the high resolution system for the u and v fields (panels (d) and (e)). Both the variance and 5 day lag covariance in the η field is not so well reproduced (compare panels (f) and (i)). The 5 day lag autocorrelation, Fig. 3,

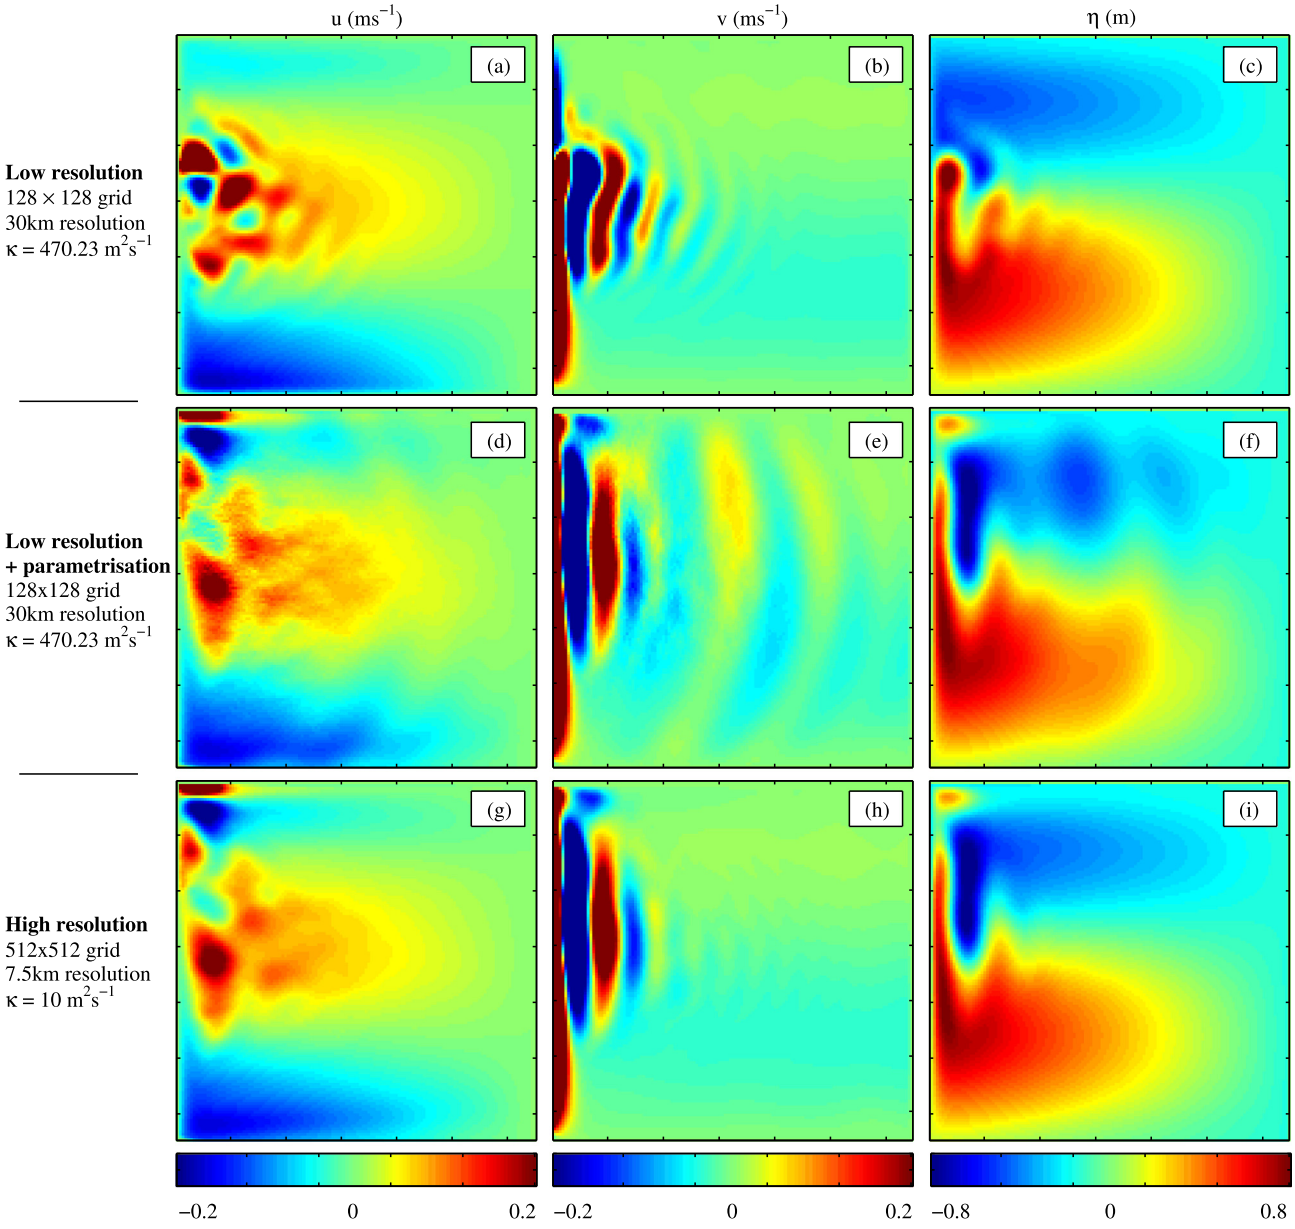

Fig. 5. The climatological mean zonal velocity u (left), meridional velocity v (middle) and sea surface height η (right) for the low resolution (top), low resolution parameterised at iteration 150 (middle), and high resolution model (bottom). κ is the viscosity parameter.

indicates that the long persistence of the low resolution model, relative to the high resolution model, has been reduced. Decomposing the velocities into their mean and varying components, $u = \bar{u} + u'$ and $v = \bar{v} + v'$, the eddy kinetic energy, defined as $\frac{1}{2}\Delta x\, \Delta y\, h\, \rho_0(u'^2 + v'^2)$, is proportional to the variance in the u and v fields. Taking the climatological eddy kinetic energy at a single longitude, 285 km from the western boundary, Fig. 9 demonstrates that the parameterised model is again a good representation of the truth. For example at 40° latitude, 285 km from the western boundary, within a 30 km grid cell it is 0.60×10^{14} J, 2.32×10^{14} J and 2.34×10^{14} J for the unparameterised, parameterised and truth integrations respectively.

η is well approximated by the two dimensional stream function, ψ, defined by $u = -\partial\psi/\partial y$ or $v = \partial\psi/\partial x$. ψ, and hence η, can therefore be approximated as a spatial integral over u or v. The noise terms ξ_u and ξ_v are uncorrelated in space, so their integral over the domain tends to be small. A large positive ξ_u at one latitude is likely to be cancelled out by negative values of ξ_u at other latitudes. Thus the contribution of ξ_u and ξ_v to the variability of η turns out to be relatively small. To get the correct variability of the η field a stochastic term would need to be added to the right hand side of the equation for η, (5). Alternatively, including spatial correlation (represented by the off diagonals of $\mathbf{Q}$), when integrating Eqs. (7) and (8) would increase the variance of the spatially integrated ξ_u and ξ_v. Therefore the variance in η would also increase. Either of these options requires more investigation and are beyond our current scope.

4.3. The constant forcing parameters

The parameterisation terms added to the equations for u and v found by the optimisation at iteration 150 are shown in Fig. 10. The constant sub-grid momentum forcing term necessary to maintain the recirculating gyre, F_u, applied to the equation for u (panel

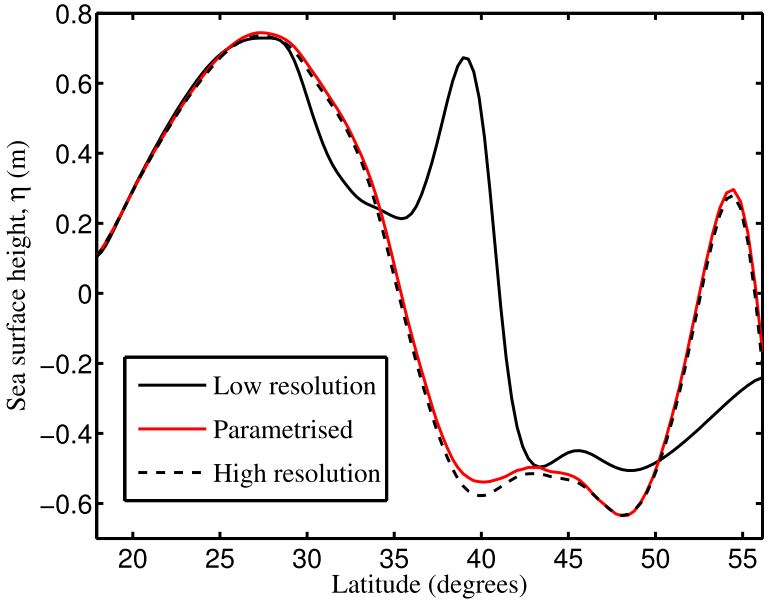

Fig. 6. The climatological mean sea surface height η at a single longitude 285 km from the western boundary, see Fig. 2.

(a)) is dominated by a large positive value, denoting forcing fluid from west to east, in the north-west corner of the domain. There are however patches of forcing with amplitude changing rapidly over small length scales, close to that of the grid scale, in parts of the turbulent region.

This may be due either to slow convergence of the optimisation algorithm (perhaps the true forcing is smooth) or due to uncertainty in the climatological mean stemming from the finite integration length. Perhaps the small scales in the forcing structure are necessary as suggested in other studies (e.g. Kraichnan (1976), Zidikheri and Frederiksen (2010), Mana and Zanna (2014) and references therein). This was tested by smoothing the iteration 150 F_u and F_v fields using a nearest neighbour average for each grid cell. In Fig. 4a at iteration 150 the MS difference between the parameterised model and the truth for $\bar{u}$, $\bar{v}$ and $\bar{\eta}$ was 0.93 $m^2\,s^{-2}$, 4.8 $m^2\,s^{-2}$ and 16 m^2 respectively. By comparison the respective MS errors with the smoothed forcing applied was 3.0 $m^2\,s^{-2}$, 20 $m^2\,s^{-2}$ and 24 m^2. The smoothing reduced the quality of the climatological mean and demonstrates that at least some small-scale features are needed for an optimal solution.

The constant forcing of the v field, F_v (Fig. 10, panel (d)), is very strong (northwards) for a single grid box alongside the western boundary. For example at 40° latitude $F_v = 1.25$ Pa. It reaches a maximum of 1.32 Pa at 36.6° latitude which compares with a maximum, over the remainder of the domain, of 0.62 Pa at a point in the red (colour saturated) region close to, but not always alongside, the western boundary. In the main body of the domain there are alternate bands of northward and southward forcing approximately aligned with the western boundary. A possible explanation for the two strongest bands closest to the western boundary is that they extend the meridional flow patterns, and hence the flow boundary separation point, northwards. In the low resolution model without parameterisation these flow patterns stop half way along the domain, compare with Fig. 5 panel (b). Fig. 5 also shows some differences between the low resolution parameterised and the high resolution models far from the western boundary (compare panels (e) and (h)). Perhaps these differences are due to a finite integration time giving a climatological mean that is not perfectly resolved. The shape of these differences reflects the linear Rossby waves present in this region and the optimisation algorithm may be trying to correct for these differences but inadvertently magnifying them. On the other hand, the forcing elsewhere may lead to a correction far from the western boundary that the optimisation is in turn trying to correct. Longer integrations or more iterations may resolve this issue.

Now that the values of F_u and F_v have been found, it is interesting to consider which physical processes set the pattern of the forcing. In the barotropic double gyre experiment, we expect the non-linear (Reynolds stresses) and viscous terms to dictate the mean sub-grid eddy forcing. Examination of these terms in the truth and low resolution experiment indicate that forcing of the form

$$F_u \approx \gamma_1 \nabla^2 \overline{u_T} \quad \textit{and} \quad F_v \approx \gamma_1 \nabla^2 \overline{v_T}, \tag{12}$$

might be a reasonable approximation. Here γ_1 is a constant, the over bar indicates the time mean and the subscript T indicates that data from the truth integration is used. We choose $\gamma_1 = 4 \times 10^9$ kg s^{-1} in order to approximately match the amplitude of the optimised forcing shown in Fig. 10. In the present model configuration F_u and F_v then correspond to the time mean of an additional eddy viscosity that corrects the mean flow. Forcing the model with these terms, leads to the separation point of the jet being further north and encouraging changes to the structure of the eastwards jet. However the strength and pattern is significantly different from the high resolution truth suggesting that if this form of forcing plays a role, non-linear feedbacks are important. These results in combination with the stochastic forcing mentioned at the end of Section 4.4, are presented in Fig. 11. Performing identical integrations without the stochastic forcing altered the variances but did not have a large impact upon the mean response.

4.4. The stochastic forcing parameters

The variance $\mathbf{q}_u$ of the noise in the stochastic system governing ξ_u is relatively large at a few grid points in the north-west corner of the domain and along the northern and southern boundaries (Fig. 10, panel (b)). There is also noise in the vicinity of the high resolution eddy activity, but there are regions in the centre of the domain and all along the western boundary where variance of ξ_u is zero. $\mathbf{q}_v$ on the other hand (Fig. 10, panel (e)) is large close to the northern part of the western boundary, small or zero along the northern and southern boundaries, small in the region of eddy

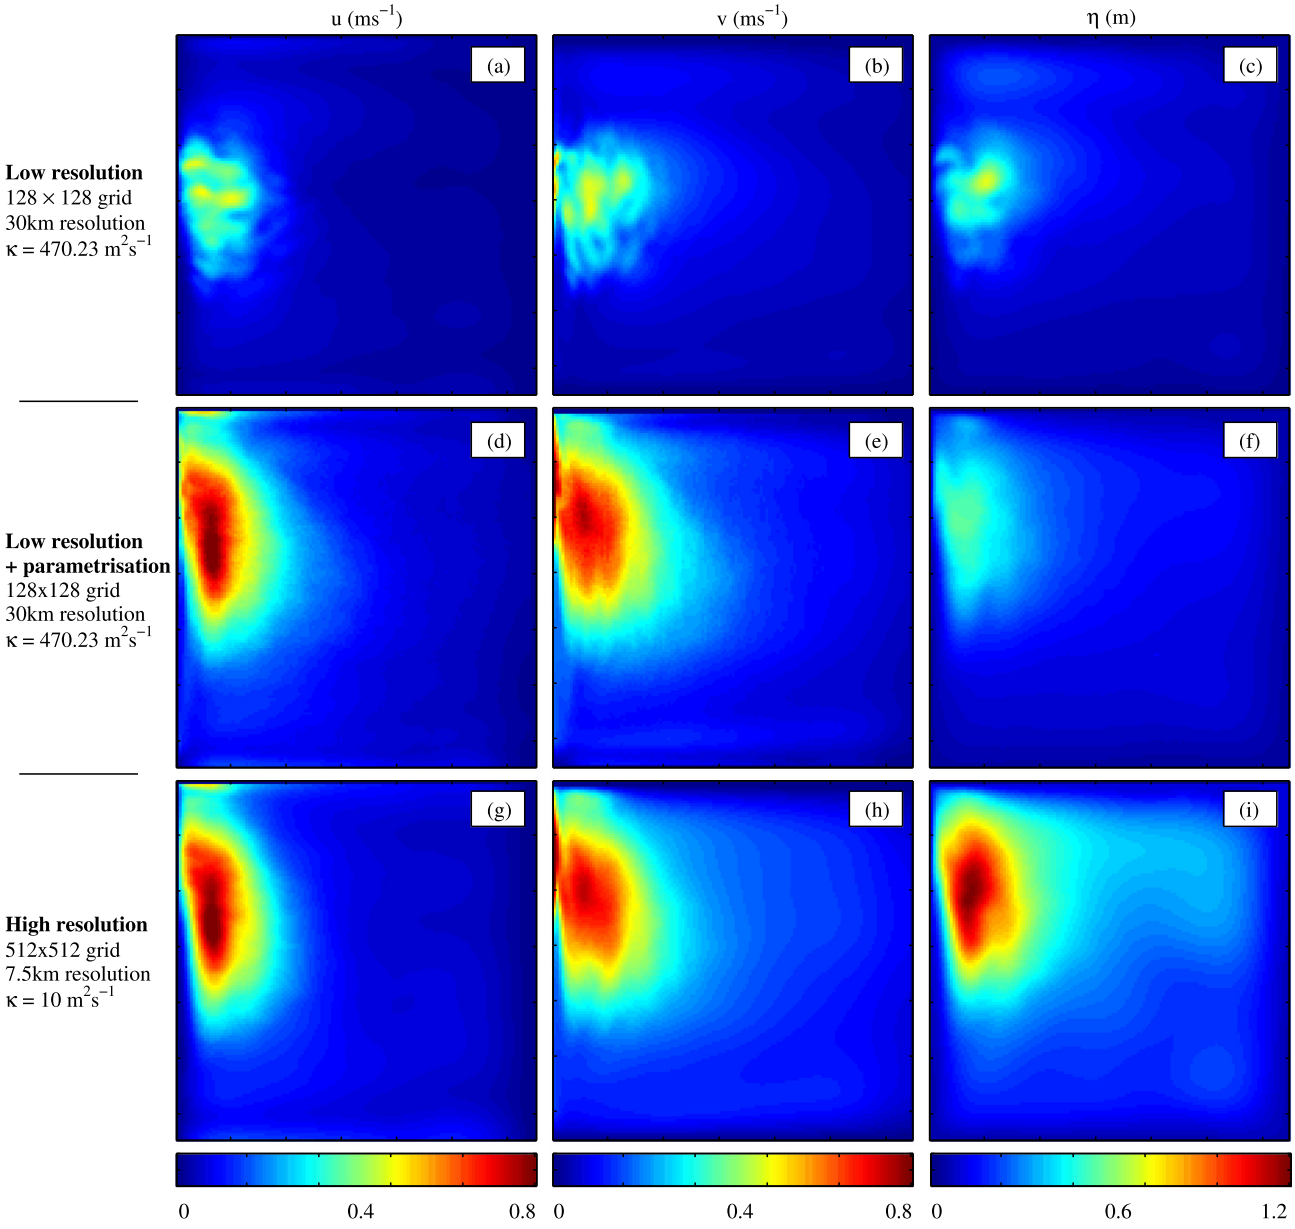

Fig. 7. The climatological standard deviation of the zonal velocity u (left), meridional velocity v (middle) and sea surface height η (right) for the low resolution (top), low resolution parameterised at iteration 150 (middle), and high resolution model (bottom). κ is the viscosity parameter.

activity and large in the region of linear wave activity. Also there is an oval region at the centre of eddy activity where $\mathbf{q}_v$ is zero.

A measure of the effective forcing of the stochastic term is given by $\sqrt{-\mathbf{q}_u/\mathbf{b}_u}$ and $\sqrt{-\mathbf{q}_v/\mathbf{b}_v}$ with units of Pascals and the division denotes the element wise division of each element of $\mathbf{q}$ by the corresponding element of $\mathbf{b}$. The square root is also taken element wise. Thus in addition to $\mathbf{q}_u$ and $\mathbf{q}_v$, the time scale of ξ_u and ξ_v determined by $\mathbf{b}_u$ and $\mathbf{b}_v$ needs to be taken into account. Panels (c) and (f) of Fig. 10 show that the largest stochastic forcing is located in the region of high resolution eddy activity. In this region, the amplitude of the stochastic term is small, but the time scale, of around 10–12 days, is relatively long when compared with the time scale in the rest of the domain, ~ 5 h. In Berloff (2005) stochastic forcing was applied using a first and second order auto-regressive process. For various experiments their first order process was given a spatially uniform time scale between 3.3 and 30 days.

To further try to disentangle the overlapping effects of the four additional forcing terms F_u, F_v, ξ_u and ξ_v, a set of experiments was performed where a subset of these terms was set to zero and the optimisation was performed again. These experiments, the results of which are summarised in Table 1, show that using the full set of four terms leads to the most optimal climatological means and variances. With this experimental configuration the constant forcing terms F_u and F_v are largely responsible for correcting the climatological mean and the stochastic terms, ξ_u and ξ_v, are mainly responsible for correcting the climatological variance and 5 day lag covariance. In addition, it is crucial that both the amplitude and time scale of the stochastic term is allowed to vary spatially.

In situations where an optimisation cannot practically be performed, an approximation to the optimal forcing may be useful. Comparing Figs. 10b with 10c and 10e with 10f, indicates that the form of $\mathbf{q}_u$ and $\mathbf{q}_v$ is quite complex and is often relatively large in regions where the effective forcing ($\sqrt{-q_i/b_i}$) is small. We might

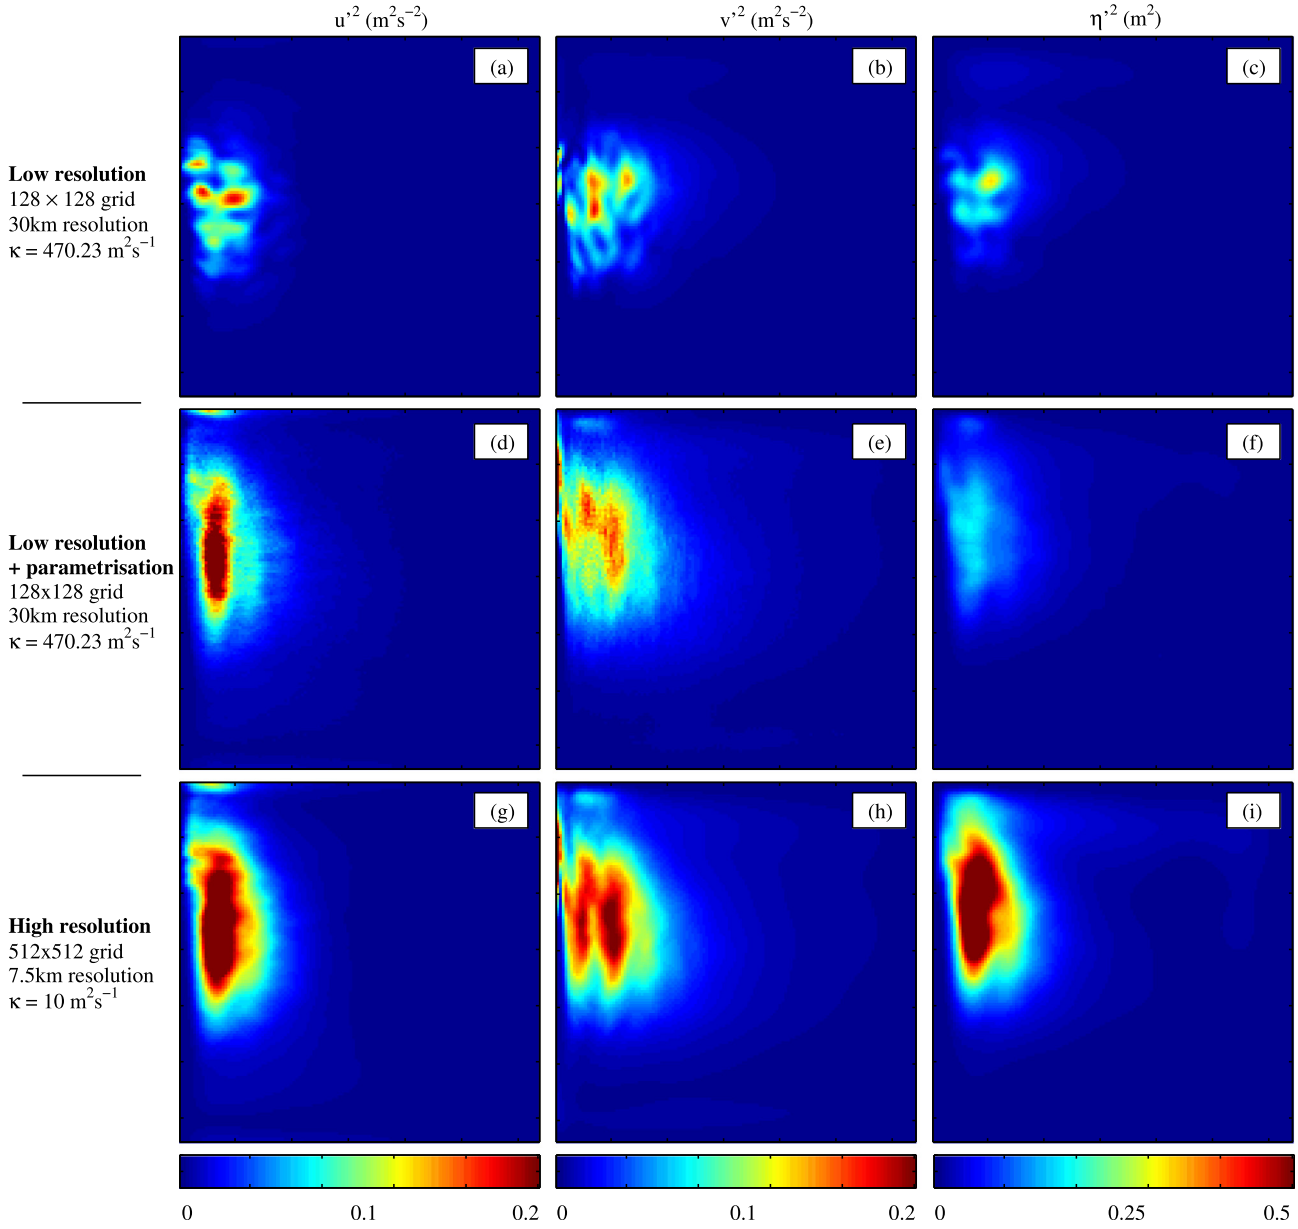

Fig. 8. The climatological 5 day lag covariance of the zonal velocity u (left), meridional velocity v (middle) and sea surface height η (right) for the low resolution (top), low resolution parameterised at iteration 150 (middle), and high resolution model (bottom). κ is the viscosity parameter.

therefore expect that the effective forcing is largely set by the time scale parameters $\mathbf{b}_u$ and $\mathbf{b}_v$. A simple approximate relation at each grid point is

$$b_{ui} = \frac{\gamma_2}{\frac{1}{2}\rho_0\left|\overline{u'^2_{iT}} - \overline{u'^2_{iM}}\right|} \quad \text{and} \quad b_{vi} = \frac{\gamma_2}{\frac{1}{2}\rho_0\left|\overline{v'^2_{iT}} - \overline{v'^2_{iM}}\right|} \tag{13}$$

where γ_2 is an undetermined constant, the subscripts M and T indicate data from the respective model and truth integrations, the over line indicates the time mean and the prime indicates the time varying part from the standard Reynolds decomposition, $u = \bar{u} + u'$ and $v = \bar{v} + v'$. As a test, we choose that $q_{ui}^{1/2} = q_{vi}^{1/2} = 0.004$ Pa s$^{-1/2}$ in the region defined by the western and northern boundaries and 2100 km east and 2940 km south of these boundaries, $q_u = q_v = 0$ elsewhere, b_u and b_v are given by (13) with $\gamma_2 = 2 \times 10^{-4}$ Pa s^{-1} and F_u and F_v are given by (12). The climatological statistics of the system integrated with these parameters are shown in Fig. 11. Comparing the plots in Fig. 7 with 11d–f, indicates that the variance of this test system has been improved relative to the low resolution model without parameterisation. It provides some evidence that the time scale of the sub grid forcing is important, rather than a particular spatial form of the amplitude. Unfortunately the 5 day lag covariance of the test system is poor, compare Fig. 8 with 11g–i.

Equations similar to (12) and (13), may yield a quick indication of where forcing might be applied in a parameterisation strategy. However Fig. 11 demonstrates that using Eqs. (12) and (13) to parameterise the sub-grid scale flow is much less accurate than performing an optimisation. Also with the spurious lag covariance shown in Fig. 11, we would expect additional error in its response to forcing.

4.5. The response to a change in the wind

Arguably, the most desirable property of a climate model, or a seasonal forecast model, is an accurate response to changes in the applied forcing. As mentioned in Section 1.3 and Fig. 1, the

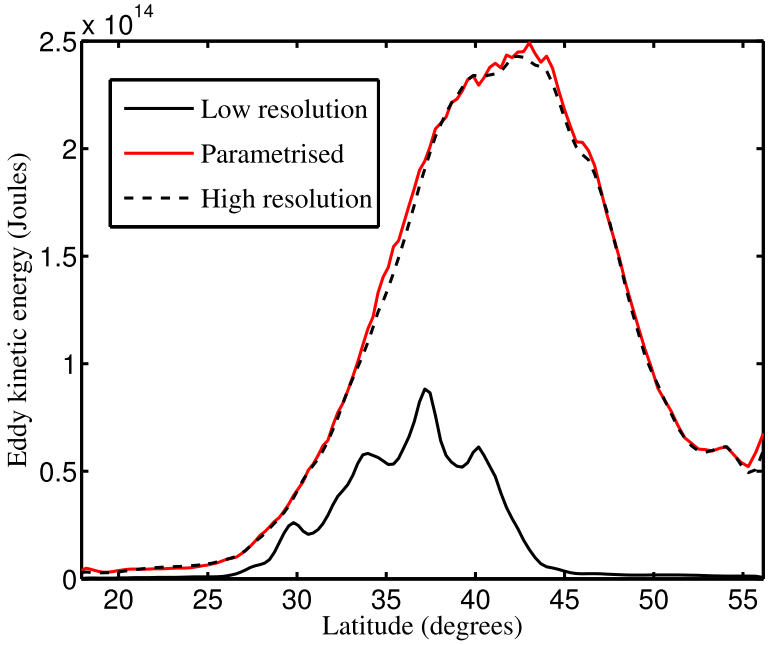

Fig. 9. The climatological eddy kinetic energy within each 30 km grid cell at a single longitude 285 km from the western boundary, see Fig. 2.

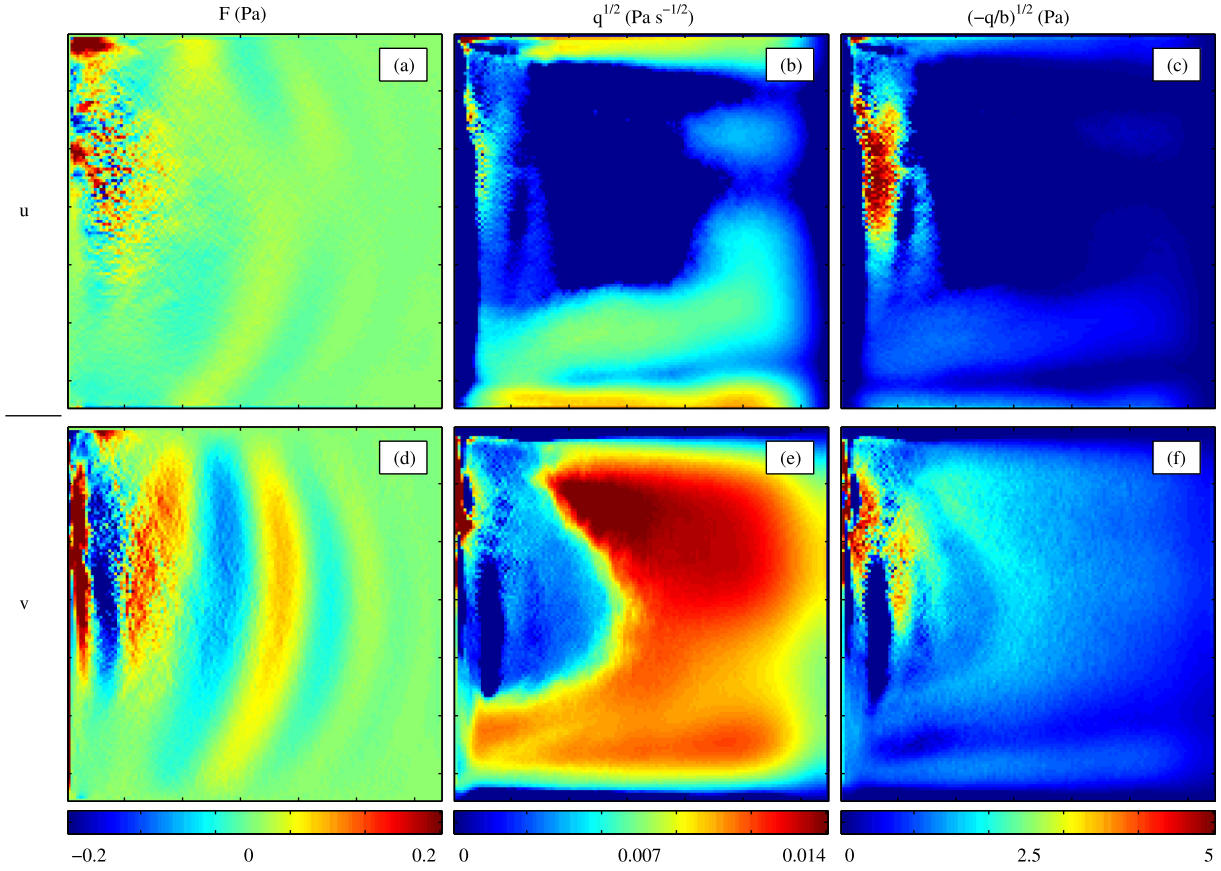

Fig. 10. Iteration 150 of the optimised values for the zonal u (top) and meridional v (bottom) velocity forcing parameters $\mathbf{f}$ (left) and stochastic parameters $\mathbf{q}$ (middle) and $(-\mathbf{q}/\mathbf{b})^{1/2}$ (right) where the division and square root is taken separately for each vector element. The quantity $(-\mathbf{q}/\mathbf{b})^{1/2}$ has Pascals as its units and is a measure of the strength of the stochastic forcing.

fluctuation–dissipation theorem suggests that a possible consequence of optimising the timescales in the model is to improve the response to forcing. Statistics related to the timescales, namely the variance and 5 day lag covariance, have been somewhat improved. In this section we describe the impact upon the model's forced response.

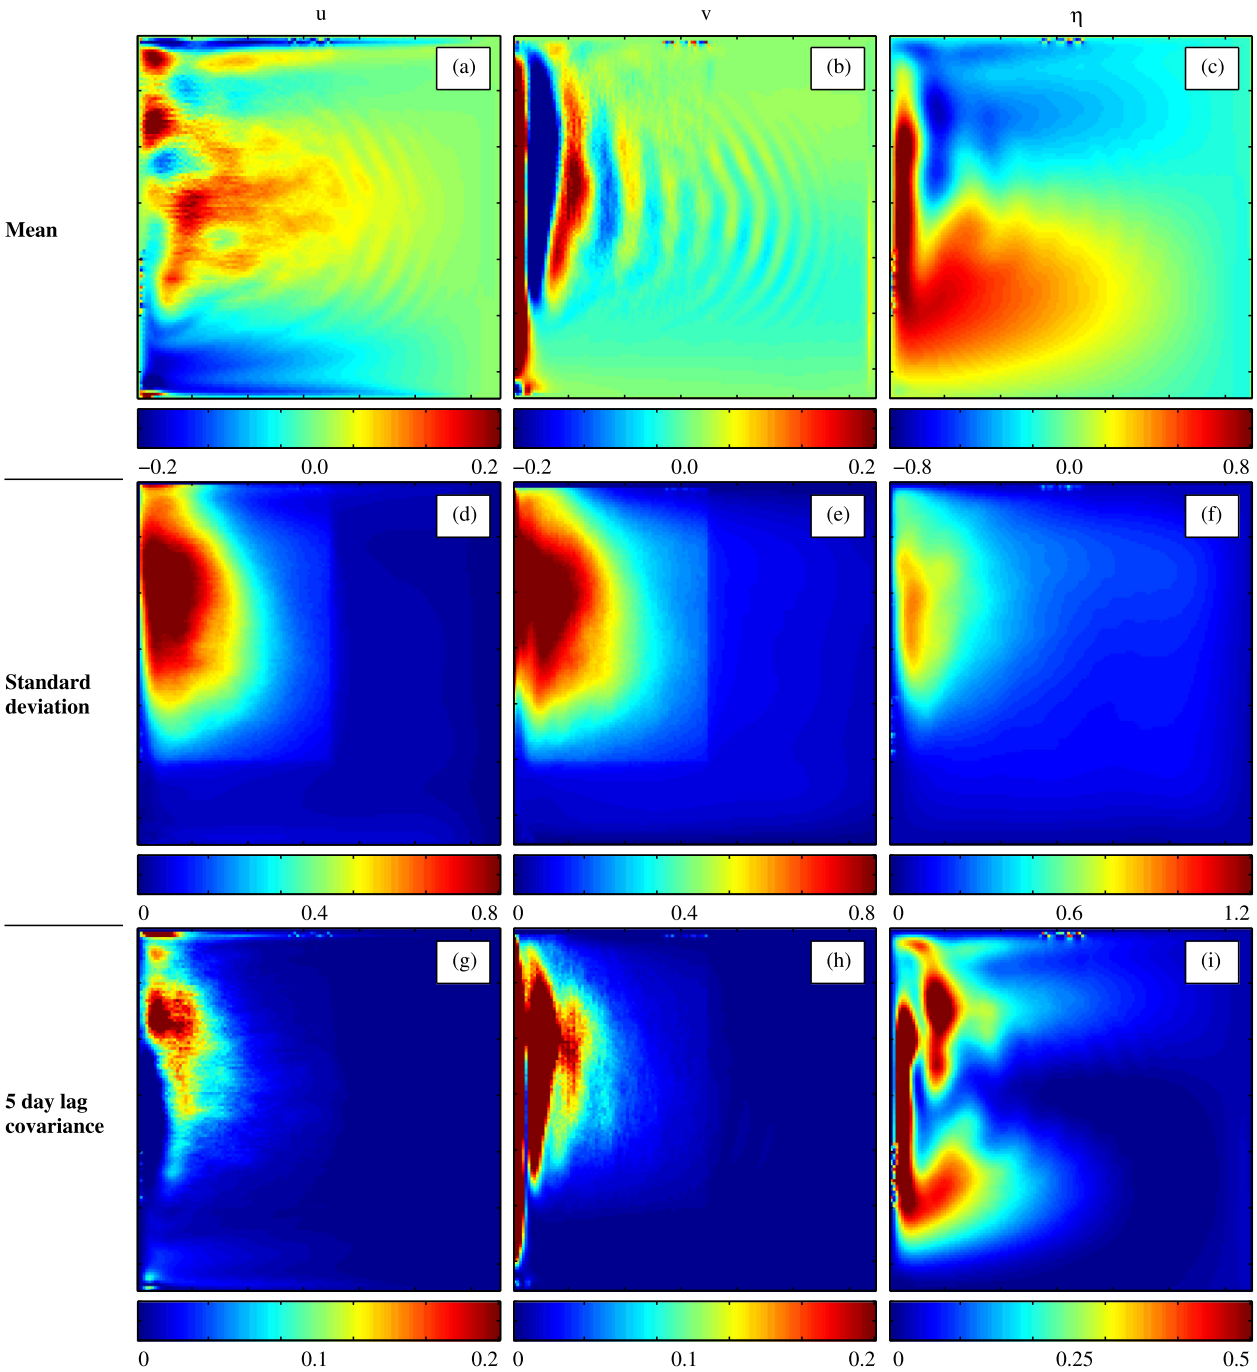

Fig. 11. Statistics of a low resolution integration with $F_u = \gamma_1 \nabla^2 \bar{u}_T$, $F_v = \gamma_1 \nabla^2 \bar{v}_T$, see the end of Section 4.3. In addition $q_u^{1/2} = q_v^{1/2} = 0.004$ Pa s$^{-1/2}$ in the region of high variance and b_u and b_v are given by (13), see the end of Section 4.4. Statistics of u (left), v (middle) and η (right) including the mean (top), standard deviation (middle) and 5 day lag covariance (bottom) are plotted. To enable comparison, the colour axis for each plot is the same as that used in Figs. 5 (for the mean), 7 (for the standard deviation) and 8 (for the lag covariance).

To test how the climatological mean of a particular configuration responds to forcing, it is useful to apply several magnitudes of forcing to check for linearity of the response, and to perform multiple independent integrations with different initial conditions to check for the uncertainty in the response. Therefore for five values of the wind forcing amplitude $\alpha_0 = 0.10,\ 0.11,\ 0.12,\ 0.13$ and 0.14 Pa, an integration of ten ($N = 10$) ensemble members with random initial conditions was performed. Fig. 12 shows the ensemble mean climatological mean response of each system to a small change to the forcing. Uncertainty in the response, quantified as the ensemble standard deviation of the climatological mean multiplied by $2/\sqrt{N}$, peaked at around 10–15 percent in regions of large response. Given the variance in the ensemble members, non-linearity of the response as a function of forcing amplitude was undetectable.

Fig. 12 shows that the response to forcing of the high resolution u field (panel (g)) is dominated by the gyre in the north-west corner of the domain (cf. Fig. 5, panel (g)). The low resolution u field without parameterisation (panel (a)) has a strong response close to the centre of the western boundary. The parameterised system successfully reduces this response to an amplitude more similar to the high resolution system but it fails to introduce the response

Table 1
A summary of integrations performed using different combinations of parameterisation terms. The left column indicates the terms included and optimised for in Eqs. (3) and (4). Terms not included in the optimisation are equal to zero. The other two columns indicate the quality of the parameterised model climatological mean and variance as optimisation progresses.

Forcing terms used	Climatological mean	Climatological variance
F_u	The mean u field converges relatively quickly, the v and η fields converge relatively slowly	No improvement in the amplitude or shape of the variance
F_v	The mean v field converges relatively quickly, the u and η fields converge relatively quickly and then plateau at a relatively poor value of the mean squared error	No improvement in the amplitude or shape of the variance.
$F_u,\ F_v$	The mean $u,\ v$ and η fields converge relatively quickly	No improvement in the amplitude or shape of the variance
$F_u,\ \xi_u$	The mean u field converges relatively quickly, the v and η fields converge relatively slowly	The variance in u converges relatively quickly, and v plateaus at a relatively poor value of the mean squared error. The 5 day lag covariance in u and v converges. No convergence in η
$F_v,\ \xi_v$	The mean $u,\ v$ and η fields converge relatively quickly.	The variance in v converges relatively quickly, and u plateaus at a relatively poor value of the mean squared error. The 5 day lag covariance in u converges and v converges relatively slowly. No convergence in η
$\xi_u,\ \xi_v$	No improvement in the amplitude or shape of the mean	The u and v variance fields converge relatively quickly. No convergence in η
$d\xi_{ui} = \sqrt{q_{ui}}dw_{ui}\ d\xi_{vi} = \sqrt{q_{vi}}dw_{vi}$ with various configurations of F_u and F_v	The mean is optimised depending upon the inclusion of F_u and F_v as above	No improvement in the amplitude or shape of the variance
$F_u,\ F_v,\ \xi_u,\ \xi_v$ (full parameter set)	Fastest convergence and most accurate	Fastest convergence and most accurate. No convergence in η

to the north-west corner. It is plausible that the location of the circulation in the low resolution unparameterised model is simply further north in the high resolution model, and since our parameterisation is local it is not able to reproduce such a non-local change. The response in the v field is located all along the western boundary in the high resolution system (Fig. 12, panel (h)), whereas the response in the low resolution system without parameterisation (panel (b)) stops half way and has a higher frequency pattern in the eastwards direction close to the centre of the domain. The low resolution system with parameterisation reproduces the pattern of the high resolution response in v quite well (compare panels (e) and (h)) although the amplitude is lower. However in the north, once again the response is not captured. A similar story applies to the η field (panels (c), (f) and (i)). The mean squared difference between the low resolution parameterised and high resolution $u,\ v$ and η fields is 0.606, 0.500 and 0.650 times the mean squared difference between the low resolution unparameterised and high resolution fields respectively. Values below one indicate that for each field, according to this measure, the parameterisation has improved the response.

5. Discussion

We have successfully optimised a high dimensional spatially varying linear stochastic parameterisation scheme for barotropic sub-grid eddy turbulence. By finding the mean sub-grid eddy forcing term, our method is successful at reproducing the climatological mean of a high resolution idealised shallow water ocean gyre using low resolution integrations. By finding the spatially varying amplitude and timescale of a local stochastic sub-grid forcing term, our method is also successful at reproducing the climatological variance and 5-day lag-covariance of the velocity variables. The response to forcing of the low resolution parameterised test system was significantly closer to the true high resolution response (estimated using the mean squared difference) than the default system without parameterisation. Only data from one integration of the high resolution system was used for optimisation, so by substituting sufficient measurements, or reanalysis data, representing the real world, this method is potentially useful in climate change experiments. The time evolution of both the low and high resolution test systems is chaotic.

The key to our method is optimisation. In a system with many parameters it is impossible to explore the entire space defined by them and a direction to optimise in must be assumed. This is normally found using a tangent linear approach. We make the simpler assumption that the direction to optimise in is given by the difference between low and high resolution climatological states. When measurements or a high resolution integration is available, optimisation allows us to avoid complications such as having to fit the sub-grid model as in Achatz and Branstator (1999), or spectral representations of the flow as in Frederiksen and Kepert (2006). We also avoid making assumptions present in less empirical theories that are either difficult to implement in reality (e.g. Kraichnan (1959), Mana and Zanna (2014)), or do not apply in practical cases. The result is a simple and accurate method that can be applied without modification to state of the art ocean models.

The measurements that we require are not too demanding. We require for example the velocity measured at a single point in space at a specified time. We do not require the tendencies that are used in Frederiksen and Kepert (2006), Achatz and Branstator (1999) and Achatz et al. (2013). We require sufficient measurements to be able to estimate the climatological mean, variance and lag-covariance in a region over some period of time.

In pioneering work, using a tangent linear model and its adjoint, Ferreira et al. (2005) optimise the eddy stresses of a $4^\circ \times 4^\circ$ resolution global ocean model to obtain a model with an accurate climatological mean temperature as defined by observations (Levitus and Boyer, 1994). As is the case with all sufficiently coarse resolution ocean models, their global model was integrated with a viscosity sufficient to damp away all of the chaotic eddies. In the absence of a time varying forcing all time differentials are equal to zero and their system is not chaotic. This enables the use of an adjoint method to optimise the climatology. By contrast, the low resolution model we use in this study is chaotic. Unfortunately, for climate problems, adjoint methods like that applied by Ferreira et al. (2005) cannot easily cope with chaotic systems (see e.g. Lea et al. (2000) and Eyink et al. (2004)). It is therefore difficult to extend them to higher resolution models which are chaotic due to having lower viscosity.

Our simple approach requires many iterations (150 in our test case) of a low resolution model integration. Ferreira et al. (2005) use 120 iterations of their forward and adjoint model. Given the different model configurations, number of degrees of freedom and optimisation tools, it remains difficult to assess which method is computationally cheaper. We applied a basic Euler method at each iteration step so it may be possible to reduce the number of iterations by using a higher order approach. Additionally the

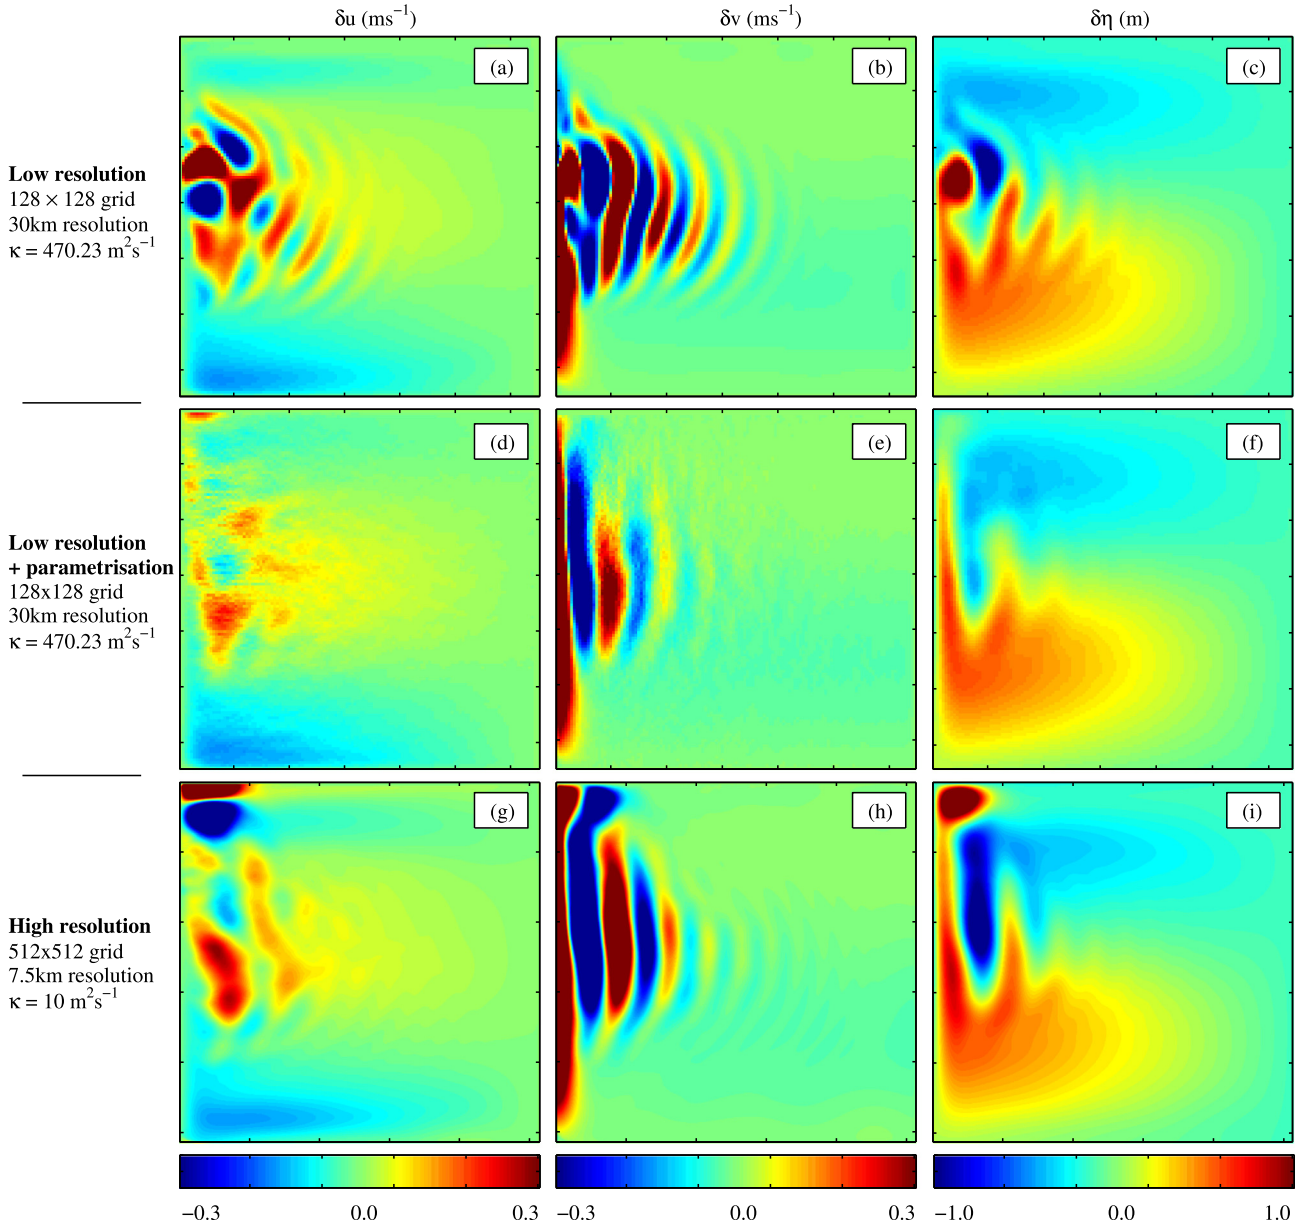

Fig. 12. The response to a small change to the wind forcing, in the zonal velocity u (left), meridional velocity v (middle) and sea surface height η (right) for the low resolution (top), low resolution parameterised at iteration 150 (middle), and high resolution model (bottom). κ is the viscosity parameter. The quantities shown are the ensemble mean climatological mean from the $\alpha_0 = 0.13$ Pa minus the ensemble mean climatological mean from the $\alpha_0 = 0.11$ Pa integrations.

direction that we push our parameterisation vectors is not optimal, therefore it may also be possible to reduce the number of iterations by more accurately estimating this direction using a modified adjoint method (Wang et al., 2014), or by starting the optimisation with parameters defined by a fit to the high resolution statistics as in Achatz and Branstator (1999) and Frederiksen and Kepert (2006), rather than the low resolution climatology.

Since our shallow water ocean gyre test system exhibits a high degree of geostrophic balance and is well approximated by the equivalent single layer quasi-geostrophic system, one may expect that optimisation of a single prognostic variable, u, v or η, would be sufficient. However this was found not to be the case. The variability of the sea surface height was not forced, or well reproduced in the low resolution model. Thus there is scope for improvement by including a stochastic term in its governing equation. The fluctuation–dissipation theorem guarantees that the response to a forcing is related to the underlying variability, so this may also improve the response of the parameterised model. When estimating the forced response, we neglected any change in the sub-grid parameters as a result of the forcing. So the low resolution parameterised response estimate can potentially be improved by including these changes as estimated using the fluctuation–dissipation theorem (Achatz et al., 2013). A complementary approach is to include optimisation of the correlation between grid points of the stochastic variables.

For objective measures of the quality of any model of a physical system, measurements of the system to be represented are absolutely necessary. In our case and in those of Berloff (2005), Achatz and Branstator (1999), Achatz et al. (2013), Frederiksen and Kepert (2006) and Zidikheri and Frederiksen (2009), the measurements are represented by values taken from a higher resolution integration. In this paper, in addition to assessing the

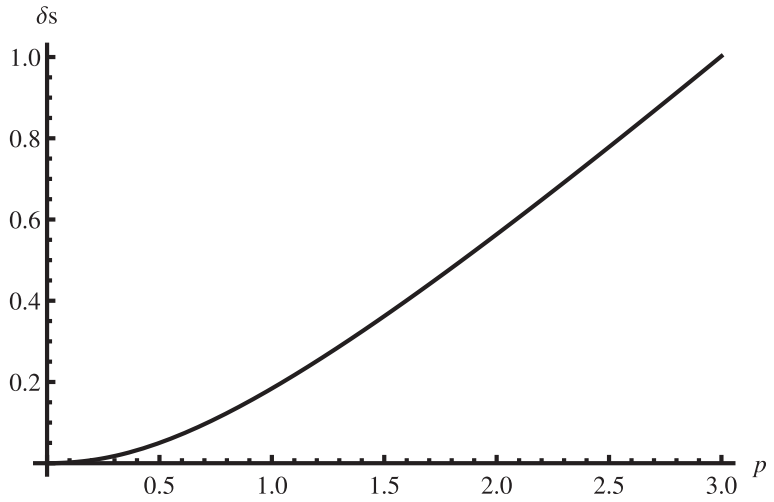

Fig. A.13. The lag covariance δs added to a linear model by a linear stochastic term with a timescale p measured in units of time. Eq. (A.4) with $a = -1,\ g = -1$ and $\tau = 1$.

quality of our parameterisation, the high resolution integration is used as an optimisation target. For a more realistic ocean model, interpolated measurements of the real ocean, or reanalysis data, would be used instead and the high resolution integration is therefore not required. For example Ferreira et al. (2005) use 1994 world ocean atlas data and more up to date reanalysis is available for sectors of the ocean, e.g. Mazloff et al. (2010). We have not developed a fully self-consistent theory of ocean turbulence so it is impossible for us to derive the values of our high dimensional parameterisation vector in advance. In regions where we have no measurements (or theory), assumptions must be made. For example we may assume that the parameters are all zero and the default model is the best, or that parameters are the same as in a similar region of ocean, or even that the parameters are given by experiments with a high resolution regional model. Another approach is to assume some given form of the forcing. As an example we considered the hypothesis that the constant forcing is given by the viscous terms in the model, (multiplied by an unknown constant). In our model setup, a simple relation between the eddy kinetic energy difference and the timescales of the stochastic forcing, was tested. Our results suggest that the impact of the stochastic forcing is mainly determined by the spatial pattern of these timescales rather than the spatial pattern of the amplitude of the variance. Using our optimisation method, the numerical values of the parameters change if the model parameters (density, resolution etc.) or boundary conditions (e.g. location of continents) change. The solution to this problem, for those who wish to make use of a particular model in a situation for which the appropriate parameters have not been found, is to apply the optimisation algorithm detailed here to find a new set of optimal parameters.

Our scheme is very simple but makes use of six parameters and two variables for each model grid cell. Although this large number of parameters is not a practical disadvantage, it is not particularly elegant. Through understanding of the ocean system (Gent et al., 1990), or turbulence in general (Holm et al., 1999), it may be possible to find a scheme of similar quality using fewer parameters. Since we have been able to estimate the mean sub-grid eddy forcing and have the amplitude and timescale of a variable forcing term (in our case stochastic), our approach may help with insights into a more developed theory. One advantage of our method is that it does not conflict with other parameterisations. If necessary it can be used in conjunction with the parameterisation schemes already present in complex ocean models both at coarse or eddy-permitting resolution.

Acknowledgements

This work was funded by UK NERC Grant NE/J00586X/1. We thank the three referees for their input which undoubtedly led to improvement of this paper. Thanks to David Munday and Miroslaw Andrejczuk for a huge amount of help with the MITgcm and thanks to James Maddison, Peter Düben, PierGianLuca Porta Mana, Mark Forshaw, and Tim Palmer for the many discussions that contributed to the ideas in this paper.

Appendix A. Influence of a stochastic parameterisation upon a linear system

Assume that a single element one dimensional model, x, is represented by a linear stochastic system of the form

$$\mathrm{d}x = ax\,\mathrm{d}t + \sqrt{b}\,\mathrm{d}w \tag{A.1}$$

where w represents Gaussian white noise with zero mean and unit variance and a and b are constants. We wish to add an additional term r to the right hand side as a parameterisation giving

$$\mathrm{d}x' = ax'\,\mathrm{d}t + \sqrt{b}\,\mathrm{d}w + r\,\mathrm{d}t \tag{A.2}$$

where r is given by

$$\mathrm{d}r = cr\,\mathrm{d}t + \sqrt{g}\,\mathrm{d}w'.$$

w' is also a white noise term with the same properties as w and c and g are constants.

Using the fact that for a d dimensional linear system $\mathbf{C}(\tau) = \exp(\mathbf{B}\tau)\mathbf{C}(0)$ where $\mathbf{C}(\tau)$ is the systems lag τ covariance matrix, $\mathbf{B}$ is a constant matrix and we use the matrix exponential, it can be shown that the lag τ variance difference, δs, between (A.1) and (A.2) is given by

$$\delta s = \frac{cg\ \exp(a\tau) - ag\ \exp(c\tau)}{2a^3c + 2ac^3}, \quad c < 0. \tag{A.3}$$

Substituting the variable $p = -1/c$ into (A.3) gives

$$\delta s = \frac{gp^2(\exp(a\tau) + ap\ \exp(-\tau/p))}{2a^3p^2 - 2a}, \quad p > 0 \tag{A.4}$$

Then over some region of $p + \delta p$, for $\tau > 0,\ \delta s$ is approximately proportional to p, so changing p in an iterative step (11) will lead to a reasonable change in the lag covariance, see Fig. A.13.

References

Abramov, R.V., Majda, A.J., 2009. New algorithms for low frequency climate response. J. Atmos. Sci. 66, 286–309.

Achatz, U., Branstator, G., 1999. A two-layer model with empirical linear corrections and reduced order for studies of internal climate variability. J. Atmos. Sci. 56, 3140–3160.

Achatz, U., Löbl, U., Dolaptchiev, S.I., 2013. Fluctuation–dissipation supplemented by nonlinearity: a climate-dependent subgrid-scale parameterization in low-order climate models. J. Atmos. Sci. 70, 1833–1846.

Arakawa, A., Lamb, R., 1977. Computational design of the basic dynamical processes of the UCLA general circulation model. Meth. Comput. Phys. 17, 173–265.

Arnold, H.M., Moroz, I.M., Palmer, T.N., 2013. Stochastic parameterizations and model uncertainty in the Lorenz 96 system. Philos. Trans. R. Soc. A 371, 20110479.

Balmaseda, M.A., Mogensen, K., Weaver, A.T., 2013. Evaluation of the ECMWF ocean reanalysis system ORAS4. Quart. J. R. Meteorol. Soc. 139, 1132–1161.

Berloff, P.S., 2005. Random-forcing model of the mesoscale oceanic eddies. J. Fluid Mech. 529, 71–95.

Buizza, R., Miller, M., Palmer, T.N., 1999. Stochastic representation of model uncertainties in the ECMWF ensemble prediction system. Quart. J. R. Meteorol. Soc. 125 (560), 2887–2908.

Cooper, F.C., Haynes, P.H., 2011. Climate sensitivity via a nonparametric fluctuation–dissipation theorem. J. Atmos. Sci. 68 (5), 937–953.

Cooper, F.C., Haynes, P.H., 2015. Assessment of the fluctuation–dissipation theorem as an estimator of the tropospheric response to forcing. Quart. J. R. Meteorol. Soc. (Submitted).

Cooper, F.C., Esler, J.G., Haynes, P.H., 2013. Estimation of the local response to a forcing in a high dimensional system using the fluctuation–dissipation theorem. Nonlinear Processes Geophys. 20, 239–248.

Dee, D.P., Uppala, S.M., Simmons, A.J., Berrisford, P., Poli, P., Kobayashi, S., Andrae, U., Balmaseda, M.A., Balsamo, G., Bechtold, P.B.P., Beljaars, A.C.M., van de Berg, L., Bidlot, J., Bormann, N., Delsol, C., Dragani, R., Fuentes, M., Geer, A.J., Haimberger, L., Healy, S.B., Hersbach, H., Hlm, E.V., Isaksen, L., Kllberg, P., Khler, M., Matricardi, M., McNally, A.P., Monge-Sanz, B.M., Morcrette, J.-J., Park, B.-K., Peubey, C., de Rosnay, P., Tavolato, C., Thpaut, J.-N., Vitart, F., 2011. The ERA-interim reanalysis: configuration and performance of the data assimilation system. Quart. J. R. Meteorol. Soc. 137, 553–597.

DelSole, T., 2000. A fundamental limitation of Markov models. J. Atmos. Sci. 57, 2158–2168.

Eyink, G.L., Haine, T.W.N., Lea, D.J., 2004. Ruelle's linear response formula, ensemble adjoint schemes and Lévy flights. Nonlinearity 17, 1867–1889.

Ferreira, D., Marshall, J., Heimbach, P., 2005. Estimating eddy stresses by fitting dynamics to observations using a residual-mean ocean circulation model and its adjoint. J. Phys. Oceanogr. 35 (10), 1891–1910.

Flato, G., Marotzke, J., Abiodun, B., Braconnot, P., Chou, S., Collins, W., Cox, P., Driouech, F., Emori, S., Eyring, V., Forest, C., Gleckler, P., Guilyardi, E., Jakob, C., Kattsov, V., Reason, C., Rummukainen, M., Flato, G., Marotzke, J., Abiodun, B., Braconnot, P., Chou, S., Collins, W., Cox, P., Driouech, F., Emori, S., Eyring, V., Forest, C., Gleckler, P., Guilyardi, E., Jakob, C., Kattsov, V., Reason, C., Rummukainen, M., 2013. Evaluation of climate models. In: Stocker, T.F., Qin, D., Plattner, G.-K., Tignor, M., Allen, S.K., Boschung, J., Nauels, A., Xia, Y., Bex, V., Midgley, P.M., Stocker, T.F., Qin, D., Plattner, G.-K., Tignor, M., Allen, S.K., Boschung, J., Nauels, A., Xia, Y., Bex, V., Midgley, P.M. (Eds.), Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Foiaş, C., Holm, D.D., Titi, E.S., 2001. The Navier–Stokes-alpha model of fluid turbulence. Phys. D 152–153, 505–519.

Frederiksen, J.S., Davies, A.G., 1997. Eddy viscosity and stochastic backscatter parameterizations on the sphere for atmospheric circulation models. J. Atmos. Sci. 54 (20), 2475–2492.

Frederiksen, J.S., Kepert, S.M., 2006. Dynamical subgrid-scale parameterizations from direct numerical simulations. J. Atmos. Sci. 63 (11), 3006–3019.

Fyfe, J.C., Saenko, O.A., Zickfeld, K., Eby, M., Weaver, A.J., 2007. The role of poleward-intensifying winds on southern ocean warming. J. Clim. 20, 5391–5400.

Gebbie, G., Heimbach, P., Wunsch, C., 2006. Strategies for nested and eddy-permitting state estimation. J. Geophys. Res. 111, C10073.

Gent, P.R., McWilliams, J.C., 1990. Isopycnal mixing in ocean circulation models. J. Phys. Oceanogr. 20 (1), 150–155.

Giering, R., 1999. Tangent linear and Adjoint Model Compiler. Users Manual 1.4 (TAMC Version 5.2), MIT/JPL/FastOpt, gie:99, <http://www.autodiff.com/tamc/>.

Gritsun, A., Branstator, G., 2007. Climate response using a three-dimensional operator based on the fluctuation–dissipation theorem. J. Atmos. Sci. 64, 2558–2575.

Grooms, I., Majda, A.J., 2013. Efficient stochastic superparameterization for geophysical turbulence. Proc. Natl. Acad. Sci. U.S.A. 110, 4464–4469.

Hasselmann, K., 1976. Stochastic climate models, Part I. Tellus 28, 473–485.

Heimbach, P., Hill, C., Giering, R., 2005. An efficient exact adjoint of the parallel MIT general circulation model, generated via automatic differentiation. Future Gen. Comput. Syst. (FGCS) 21, 1356–1371.

Held, I.M., Suarez, M.J., 1994. A proposal for the intercomparison of the dynamical cores of atmospheric general circulation models. Bull. Am. Meteor. Soc. 75, 1825–1830.

Herring, J.R., Kraichnan, R.H., 1972. Comparison of some approximations for isotropic turbulence. In: Rosenblatt, M., Van Atta, C. (Eds.), Statistical Models and Turbulence, Lecture Notes in Physics, vol. 12. Springer-Verlag, Berlin, pp. 148–194.

Holm, D.D., Kouranbaeva, S., Marsden, J.E., Ratiu, T., Shkoller, S., 1999. A nonlinear analysis of the averaged Euler equations. arXiv preprint chao-dyn/9903036.

Kalnay, E., Kanamitsu, M., Kistler, R., Collins, W., Deaven, D., Gandin, L., Iredell, M., Saha, S., White, G., Woollen, J., Zhu, Y., Leetmaa, A., Reynolds, R., Chelliah, M., Ebisuzaki, W., Higgins, W., Janowiak, J., Mo, K.C., Ropelewski, C., Wang, J., Jenne, R., Joseph, D., 1996. The NCEP/NCAR 40-year reanalysis project. Bull. Am. Meteor. Soc. 77, 437–471.

Kirtman, B., Vecchi, G.A., 2011. Why climate modelers should worry about atmospheric and oceanic weather. In: Chang, C.-P., Ding, Y., Lau, N.-C., Johnson, R.H., Wang, B., Yasunari, T. (Eds.), The Global Monsoon System: Research and Forecast, vol. 5. World Scientific, pp. 511–524 (Chapter 29).

Kitsios, V., Frederiksen, J.S., Zidikheri, M.J., 2013. Scaling laws for parameterisations of subgrid eddy–eddy interactions in simulations of oceanic circulations. Ocean Modell. 68 (8), 88–105.

Kraichnan, R.H., 1959. The structure of isotropic turbulence at very high Reynolds numbers. J. Fluid Mech. 5 (4), 497–543.

Kraichnan, R.H., 1976. Eddy viscosity in two and three dimensions. J. Atmos. Sci. 33 (8), 1521–1536.

Large, W.G., Danabasoglu, G., 2006. Attribution and impacts of upper-ocean biases in CCSM3. J. Clim. 19, 2325–2346.

Lea, D.J., Allen, M.R., Haine, T.W.N., 2000. Sensitivity analysis of the climate of a chaotic system. Tellus 52A, 523–532.

Leith, C.E., 1967. Diffusion approximation to inertial energy transfer in isotropic turbulence. Phys. Fluids 10, 1409–1416.

Levitus, S., Boyer, T., 1994. World ocean atlas 1994 volume 4: Temperature. NOAA Atlas NESDIS 4, U.S. Department of Commerce, Washington, D.C.

Mana, P.P., Zanna, L., 2014. Toward a stochastic parameterization of ocean mesoscale eddies. Ocean Modell. 79, 1–20.

Marconi, U.M.B., Puglisi, A., Rondoni, L., Vulpiani, A., 2008. Fluctuation–dissipation: response theory in statistical physics. Phys. Rep. 461, 111–195.

Marotzke, J., Giering, R., Zhang, K.Q., Stammer, D., Hill, C., Lee, T., 1999. Construction of the adjoint MIT ocean general circulation model and application to Atlantic heat transport sensitivity. J. Geophys. Res. 104 (C12), 29529–29547.

Marshall, J., Adcroft, A., Hill, C., Perelman, L., Heisey, C., 1997. A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers. J. Geophys. Res. 102, 5753–5766.

Marshall, D.P., Maddison, J.R., Berloff, P.S., 2012. A framework for parameterizing eddy potential vorticity fluxes. J. Phys. Oceanogr. 42 (4), 539–557.

Mazloff, M.R., Heimbach, P., Wunsch, C., 2010. An eddy-permitting southern ocean state estimate. J. Phys. Oceanogr. 40, 880–899.

Moore, A.M., Arango, H.G., Di Lorenzo, E., Cornuelle, B.D., Miller, A.J., Neilson, D.J., 2004. A comprehensive ocean prediction and analysis system based on the tangent linear and adjoint of a regional ocean model. Ocean Modell. 7 (1–2), 227–258.

Munday, D.R., Johnson, H.L., Marshall, D.P., 2013. Eddy saturation of equilibrated circumpolar currents. J. Phys. Oceanogr. 43 (3), 507–532.

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. R. Meteorol. Soc. 127 (572), 279–304.

Palmer, T.N., Weisheimer, A., 2011. Diagnosing the causes of bias in climate models – why is it so hard? Geophys. Astrophys. Fluid Dyn. 105, 351–365.

Ritchie, H., 1988. Application of the semi-lagrangian method to a spectral model of the shallow water equations. Mon. Weather Rev. 116, 1587–1598.

Scaife, A.A., Copsey, D., Gordon, C., Harris, C., Hinton, T., Keeley, S., O'Neill, A., Roberts, M., Williams, K., 2011. Improved atlantic winter blocking in a climate model. Geophys. Res. Lett. 38, L23703.

Smagorinsky, J., 1963. General circulation experiments with the primitive equations, I, the basic experiment. Mon. Weather Rev. 91, 99–164.

Stockdale, T.N., 1997. Coupled ocean–atmosphere forecasts in the presence of climate drift. Mon. Weather Rev. 125, 809–818.

Straub, D.N., 1993. On the transport and angular momentum balance of channel models of the antarctic circumpolar current. J. Phys. Oceanogr. 23 (4), 776–782.

Wang, Q., Hu, R., Blonigan, P., 2014. Least squares shadowing sensitivity analysis of chaotic limit cycle oscillations. J. Comput. Phys. 267, 210–224.

Waterman, S., Hoskins, B.J., 2013. Eddy shape, orientation, propagation, and mean flow feedback in western boundary current jets. J. Phys. Oceanogr. 43 (8), 1666–1690.

Waterman, S., Jayne, S.R., 2012. Eddy-driven recirculations from a localized transient forcing. J. Phys. Oceanogr. 42 (3), 430–447.

Wilks, D.S., 2005. Effects of stochastic parameterizations in the Lorenz 96 system. Quart. J. R. Meteorol. Soc. 131, 389–407.

Zidikheri, M.J., Frederiksen, J.S., 2009. Stochastic subgrid parameterizations for simulations of atmospheric baroclinic flows. J. Atmos. Sci. 66, 2844–2858.

Zidikheri, M.J., Frederiksen, J.S., 2010. Stochastic modelling of unresolved eddy fluxes. Geophys. Astrophys. Fluid Dyn. 104 (4), 323–348.