Research Article **Open Access**

Nick Verheul*, Jan Viebahn, and Daan Crommelin

Covariate-based stochastic parameterization of baroclinic ocean eddies

https://doi.org/10.1515/mcwf-2017-0005
Received August 25, 2017; accepted November 21, 2017

Abstract: In this study we investigate a covariate-based stochastic approach to parameterize unresolved turbulent processes within a standard model of the idealised, wind-driven ocean circulation. We focus on vertical instead of horizontal coarse-graining, such that we avoid the subtle difficulties of horizontal coarse-graining. The corresponding eddy forcing is uniquely defined and has a clear physical interpretation related to baroclinic instability. We propose to emulate the baroclinic eddy forcing by sampling from the conditional probability distribution functions of the eddy forcing obtained from the baroclinic reference model data. These conditional probability distribution functions are approximated here by sampling uniformly from discrete reference values. We analyze in detail the different performances of the stochastic parameterization dependent on whether the eddy forcing is conditioned on a suitable flow-dependent covariate or on a time-lagged covariate or on both. The results demonstrate that our non-Gaussian, non-linear methodology is able to accurately reproduce the first four statistical moments and spatial/temporal correlations of the stream function, energetics, and enstrophy of the reference baroclinic model.

Keywords: Stochastic parameterization, baroclinic eddies, flow-dependent parameterization

MSC: 65C20, 60H35, 62G09, 76M35

1 Introduction

1.1 Background and motivation

The large-scale ocean circulation is strongly influenced by mesoscale turbulent eddies [26, 37]. Baroclinic instability is the primary generating mechanism for mesoscale eddies in oceanic flows [38, 44, 46]. How accurately the impact of baroclinic instability is represented in ocean models depends on the accuracy of the baroclinic eddy forcing that appears in the equations of motion. Mesoscale ocean eddies exist on spatial scales roughly between $O(10\,\text{km})$ and $O(100\,\text{km})$. Therefore, global climate models need grid resolutions smaller than $O(10\,\text{km})$ in their ocean component in order to directly resolve these turbulent motions. Due to computational limitations, such high resolution is still infeasible in current climate models, and the effects of turbulent eddies must be parameterized. Parameterizations here are understood to be parametric mathematical models to be applied to an ocean model with a spatial resolution that leaves eddy forcing partly unresolved.

***Corresponding Author: Nick Verheul:** Centrum Wiskunde & Informatica (CWI), Science Park 123, 1098XG Amsterdam, The Netherlands, E-mail: n.verheul@cwi.nl
Jan Viebahn: Centrum Wiskunde & Informatica (CWI), Science Park 123, 1098XG Amsterdam, The Netherlands, E-mail: j.p.viebahn@cwi.nl
Daan Crommelin: Centrum Wiskunde & Informatica (CWI), Science Park 123, 1098XG Amsterdam, The Netherlands. And KdV Institute for Mathematics, University of Amsterdam, Science Park 105, 1098XG Amsterdam, The Netherlands, E-mail: daan.crommelin@cwi.nl

Mesoscale eddy parameterizations are commonly formulated in a deterministic way, typically based on the Gent–McWilliams (GM) parameterization [14, 15, 20, 49]. Deterministic eddy parameterizations represent an approximation of the integrated effects of the unresolved processes in terms of the resolved scale flow. These parameterizations are motivated by the idea that the properties of the unresolved scale processes can be uniquely represented by the resolved scale states. However, in practice, the resolved states are associated with many possible unresolved states [8, 17]. Therefore, deterministic parameterizations can, at best, provide an ensemble-mean representation of the unresolved scale processes. To overcome the limitations of deterministic parameterizations, atmospheric research has in recent years started to focus on stochastic parameterizations [8, 18, 36, 40]. Examples in an atmospheric context relevant to our work include Markov Chain models to represent atmospheric convection [10, 12, 32]. Stochastic eddy parameterizations are a more recent development in oceanic research. Grooms and Majda [24] developed a new approach combining elements from superparameterization and stochastic parameterizations applicable to quasi-geostrophic turbulence. Cooper and Zanna [9] posed a linear stochastic term that stochastically parameterizes transient eddies in an idealized barotropic ocean gyre model. They suggested an efficient search method along parameter space that optimizes their parameters with respect to a reference climatological mean, variance, and decorrelation time scales of the horizontal flow velocities.

In the current work we explore how to use the novel data-driven stochastic methodology posed in Verheul and Crommelin [47] for eddy parameterization. The main assumption for our parameterization is that sample data from a 'high-resolution' ocean model is available. We use this sample data to approximate conditional probability distribution functions (CPDFs) for the mesoscale eddy forcing. By conditioning on appropriately chosen covariates, i.e. model variables that correlate significantly with the mesoscale eddy forcing, we define a flow-dependent parameterization that samples directly from the CPDFs. The main goal of this parameterization is to drive a reduced ocean model in such a way that the resulting stochastic model is able to emulate the dynamics produced by the full model.

Typically, in such a reduced model the vertical discretization of the 'high-resolution' ocean model is preserved, and the horizontal grid resolution is reduced, see e.g. [6, 29, 41]. While such a horizontal coarse-graining set-up preserves some of the properties induced by the vertical stratification, a clear definition of the associated mesoscale eddy forcing is difficult both numerically and physically. Numerically, one is faced with nontrivial filtering options [41, 48] that subtly change the definition of the eddy forcing. In turn, physical interpretations of such eddy forcings are non-transparent to a certain extent because horizontally coarse-grained terms mix partly resolved processes of both barotropic and baroclinic nature. To avoid such concerns, we focus instead on a 'vertical coarse-graining' set-up which preserves horizontal grid resolution, but reduces the vertical discretization [21, 28]. This less common approach enables a clear and unambiguous definition of the eddy forcing with the clear physical interpretation related solely to the baroclinic nature of the flow. The 'vertical coarse-graining' allows us to focus fully on the development of our stochastic methodology in this spatially extended setting without being detracted by the aforementioned concerns.

In this study, we aim to drive a reduced ocean model with flow-dependent as well as spatially and temporally correlated stochastic forcing. Recent related work using flow-dependent stochastic parameterizations include the following examples. Using the stochastic approach of [19], Kitsios et al. [33] parameterizes subgrid eddy-eddy interactions as a combination of deterministic eddy viscosity and stochastic backscatter eddy viscosity. Furthermore, in [33] they formulated scaling laws for the respective coefficients dependent on the resolution, enstrophy flux, Rossby radius, and energy range. Jansen and Held [29] modeled the amplitude of the flow-dependent energy source due to backscatter forcing with simple spatially uncorrelated Gaussian white noise. By combining a standard hyperviscous closure with this stochastic term they successfully returned the energy otherwise spuriously dissipated, as is typical for hyperviscous closures. Finally, some recent studies have focused on increasing the efficiency of superparameterization, an extremely expensive computational approach to parameterization that embeds high-resolution simulations in grid cells of low-resolution large scale simulations [22, 23, 31]. Majda and Grote [35] proposed a framework that models the small-scale dynamics with quasilinear stochastic partial differential equations, which was later implemented with success in Grooms and Majda [25] for a one-dimensional turbulent system. However, the feedback to the large scales was effectively non-stochastic in this implementation. Grooms and Majda [24] successfully used unidirec-

tional plane waves in random directions for efficient computation of the flow-dependent Fourier integrals that determine the stochastic feedback to the large scales.

Most relevant to our proposed methodology are the studies by Berloff [4, 6], and by Zanna and colleagues [41, 52]. The goal in these studies is to model the spatio-temporal correlations of the ocean flow, a goal that we share here. Moreover, Zanna and colleagues employ a stochastic methodology based on the use of a covariate, again similar to what is proposed here. Regarding the former, Berloff [6] showed that the temporal correlations of a diagnosed eddy forcing can be reproduced by forcing a 'non-eddy-resolving' stationary double gyre ocean model with a simple flow-independent but spatially varying autoregressive process. This approach showed good results in reproducing the desired statistics in the stochastic model, however it required many parameters to be estimated. In Berloff et al. [4] this methodology was extended to model spatio-temporal correlations in a coupled ocean-atmosphere model. The low-frequency coupled variability in this system gave novel non-stationary statistical properties of the reference ocean eddy forcing. These properties were modeled in the stochastic forcing by introducing flow-dependency in the variance of the forcing (dependent on the baroclinicity of the ocean flow).

In Porta Mana and Zanna [41] it was proposed to reproduce spatio-temporal correlations by sampling from reference values of the eddy forcing. To achieve this, the CPDFs for the eddy forcing were approximated with Gaussian distributions, conditioned on a suitable covariate. The stochastic and deterministic feedback to a double gyre quasi-geostrophic ocean model using this covariate were explored in Zanna et al. [52]. This parameterization drastically improved the mean state and variability of the ocean state. While similar in design philosophy to our work here, we note some important differences.

Firstly, Zanna and colleagues [41, 52] develop their methodology in a set-up of horizontal coarse-graining, with the related difficulties discussed earlier in this introduction. Secondly, the covariate specified in [41, 52] is based on temporal tendencies of the vorticity, however these tendencies are in turn dependent on the unresolved eddy forcing. To close the system, the temporal tendencies must be approximated. By contrast, in the approach developed here we use resolved model variables and lagged self-conditioning to formulate our parameterization. Thirdly, this lagged self-conditioning allows us to explicitly represent temporal correlations in the parameterization. This feature is implicitly included only with respect to the sampling interval in [41, 52]. Fourthly, whereas in [41, 52] the CPDFs for the eddy forcing are approximated with Gaussian distributions, we assume no underlying distribution. Instead, we sample directly from the CPDFs as described by the sample data. Fifthly, and lastly, while the covariate used in [41, 52] is motivated physically as well as justified numerically, the parameterization concerns a single covariate. Therefore, all dynamical effects on the ocean flow are attributed to that one covariate. We aim to make the dynamical effects of our covariates intuitive and transparent by using multiple simple covariates. This allows us to perform sensitivity analyses, as well as compare between two-way coupled simulations with different configurations to illustrate the differences between covariates.

The presentation of this work is as follows: in the remainder of this section we present a formal problem description. In Section 2 we define the physical ocean model. The stochastic model and methodology are detailed in Section 3. Finally, different choices for stochastic models and the accompanying results are discussed in Sections 4 and 5, respectively.

1.2 Problem description.

An ocean-climate model consists of coupled partial differential equations (PDEs) resulting from a set of conservation laws [11, 38]. In an abstract description of an ocean model a state vector $\boldsymbol{u}$ is evolved over time in response to some external forcing $\mathcal{F}$, a linear operator $\mathcal{L}\boldsymbol{u}$ and some non-linear operator $\mathcal{N}(\boldsymbol{u})$. Without loss of generality, we consider the state vector to consist of two orthogonal modes $\boldsymbol{u} := (u_0, u_1)$. The coupled PDEs with quadratic $\mathcal{N}$ can then be written as:

$$\partial_t u_0 = \mathcal{F}_0 + \mathcal{L}_0 u_0 + \mathcal{N}_{00}(u_0, u_0) + \mathcal{N}_{01}(u_0, u_1)$$
$$\partial_t u_1 = \mathcal{F}_1 + \mathcal{L}_1 u_1 + \mathcal{N}_{11}(u_1, u_1) + \mathcal{N}_{10}(u_0, u_1),$$

where $\mathcal{N}_{00}, \mathcal{N}_{11}$ indicate the nonlinear self-interaction of each mode, and $\mathcal{N}_{01}, \mathcal{N}_{10}$ represent the nonlinear coupling of the different modes.

Next, we consider a reduced ocean model where only the variable u_0 is evolved. To distinguish it from the u_0 in the coupled model above, we denote it as $\widetilde{u}_0$ in the reduced model in the following. Without parameterization to compensate for the missing term $\mathcal{N}_{01}$, the dynamics of the reduced model can differ significantly from the dynamics of u_0 in the full, coupled model. The stochastic approach explored in this study aims to remedy this shortcoming by driving the reduced model with a stochastic process $\widetilde{\boldsymbol{R}}$, i.e. to define a stochastic reduced model:

$$\partial_t \widetilde{u}_0 = \mathcal{F}_0 + \mathcal{L}_0 \widetilde{u}_0 + \mathcal{N}_{00}(\widetilde{u}_0, \widetilde{u}_0) + \widetilde{R}.$$

The main objective of our work is to choose the stochastic process $\widetilde{\boldsymbol{R}}$, the so-called stochastic eddy forcing, in such a way that both the long-term statistical behavior and the physical properties of $\widetilde{u}_0$ resemble those of u_0. Hence, the criteria that we use to assess the accuracy of our stochastic reproduction are the field's first four statistical sample moments, the autocorrelations, and spatial covariances and correlations, as well as the energetics and enstrophy (see Section A.4 for formal definitions of these quantities).

2 Physical model

In this study, we consider a standard model of idealised ocean dynamics, namely, quasi-geostrophic (QG), potential-vorticity (PV) equations in a classical double-gyre configuration (see e.g. Vallis [46]). The fluid-dynamic model describes idealised, wind-driven midlatitude ocean circulation with prescribed density stratification in a flat-bottom square basin with north-south and east-west boundaries. We employ a set-up in which the vertical discretisation is done by projection onto the two leading vertical eigenmodes (see e.g. [16, 21, 28]), i.e. the barotropic mode and the first baroclinic mode. The potential vorticity conservation for the barotropic (baroclinic) mode stream function ψ_0 (ψ_1) with rigid lid vertical boundary conditions then reads:

$$\begin{aligned}
&\partial_t q_0 + J(\psi_0, q_0) + R + \beta \partial_x \psi_0 = A_H \nabla^4 \psi_0 + \frac{\partial_x \tau^y - \partial_y \tau^x}{\overline{\rho} H}\,, \\
&\partial_t q_1 + J(\psi_1, q_0) + J(\psi_0, q_1) + \epsilon_{111} J(\psi_1, q_1) + \beta \partial_x \psi_1 = A_H \nabla^4 \psi_1 + \frac{\gamma(\partial_x \tau^y - \partial_y \tau^x)}{\overline{\rho} H}\,,
\end{aligned} \tag{2M}$$

where $J(f, g) := (\partial_x f)(\partial_y g) - (\partial_y f)(\partial_x g)$, the relative PVs are given by $q_0 = \nabla^2 \psi_0$ and $q_1 = \nabla^2 \psi_1 - \lambda^{-2} \psi_1$, respectively, and R is defined as:

$$R(x, y, t) := J(\psi_1(x, y, t), q_1(x, y, t)). \tag{1}$$

The term R acts as the feedback of the baroclinic mode on the barotropic mode, and is interpreted as baroclinic eddy forcing term in this study. In the following, we denote the 2-mode model with 2M, and we refer to the first equation of 2M with a stochastic representation of R as the stochastic 1-mode model (S1M, see Section 3), and to the first equation of 2M with R set to zero as the deterministic 1-mode model (D1M):

$$\partial_t q_0 + J(\psi_0, q_0) + \beta \partial_x \psi_0 = A_H \nabla^4 \psi_0 + \frac{\partial_x \tau^y - \partial_y \tau^x}{\overline{\rho} H}\,. \tag{D1M}$$

The code for above deterministic models is part of OMUSE [39] and is available at the OMUSE project website: https://bitbucket.org/omuse. In our numerical model simulations, the flow is driven at the surface by the asymmetric double-gyre zonal wind stress (as e.g. in [6]),

$$\tau^x(y) = \tau_0 \left[\cos\left(\frac{2\pi(y - L/2)}{L} \right) + 2 \sin\left(\frac{\pi(y - L/2)}{L} \right) \right], \qquad \tau^y = 0\,,$$

where $\tau_0 = 0.05$ Nm^{-2}, and $L = 4000$ km is the size of the square basin with $0 \le x, y \le L$. The first internal Rossby radius of deformation, λ, represents a length scale of baroclinic eddies and is set to $\lambda = 50$ km, a typical value for the midlatitude ocean circulation. We use an eddy-resolving horizontal resolution

of 10 km in our numerical simulations with a correspondingly small lateral viscosity coefficient, $A_H = 100$ m^2s^{-1}, as well as free-slip boundary conditions. Furthermore, we use typical values for the mean ocean depth, $H = 4000$ m, the mean density of sea water, $\overline{\rho} = 1000$ kgm^{-3}, and the meridional variation of the Coriolis parameter, $\beta = 1.8616 \times 10^{-11} m^{-1}s^{-1}$. Finally, we consider the idealized case of constant stratification such that $\epsilon_{111} = 0$ and $\gamma = \sqrt{2}$ (see e.g. Hua and Haidvogel [28]). All numerical model simulations in this work have a spin-up time of 30 years and an integration length of another 30 years. See Table 1 for an overview of parameter values used in this study.

Figure 1 shows snapshots as well as the temporal averages $\mu(H\psi_0)$ and standard deviations std$(H\psi_0)$ of the barotropic stream function ψ_0 in Sverdrup (1 Sv = 10^6 m^3s^{-1}) for 2M and D1M, respectively. The statistical quantities are calculated from simulation results stored on a weekly basis. For both models the time-mean flow (see Figures 1b and 1e) consists of the southern (subtropical) and northern (subpolar) gyres that fill about 2/3 and 1/3 of the basin, respectively, which is consistent with the wind stress pattern. In the eastern part of the basin the flow is characterized by the linear Sverdrup balance which leads to essentially identical time-mean flow fields for both 2M and D1M. Near the western boundary, on the other hand, narrow boundary currents close the circulation and nonlinear terms are significant. The magnitude and the meridional extension of the time-mean western boundary currents are significantly larger for 2M than for D1M. In terms of fluctuations, for both models the basin can be partitioned into the more energetic western boundary region, characterized by strong vortices, and the less energetic eastern part, dominated by the planetary waves (see [7] for details). However, in addition to the strengthened time-mean flow, the variability is significantly more pronounced as well for 2M, as visible in both snapshots (Figures 1a and 1d) and standard deviation fields (Figures 1c and 1f). In particular, significantly larger and stronger vortices are present at the western boundary. Variability is also dominant in the rest of the basin, whereas for D1M the instantaneous flow pattern in the eastern part largely resembles the time-mean flow pattern. Finally, Figures 2a and 2b show corresponding time series of kinetic energy in joules (1 J = 1 kgm^2s^{-2}) and enstrophy in kgs^{-2} which again demonstrate larger mean values and stronger variability for 2M than for D1M.

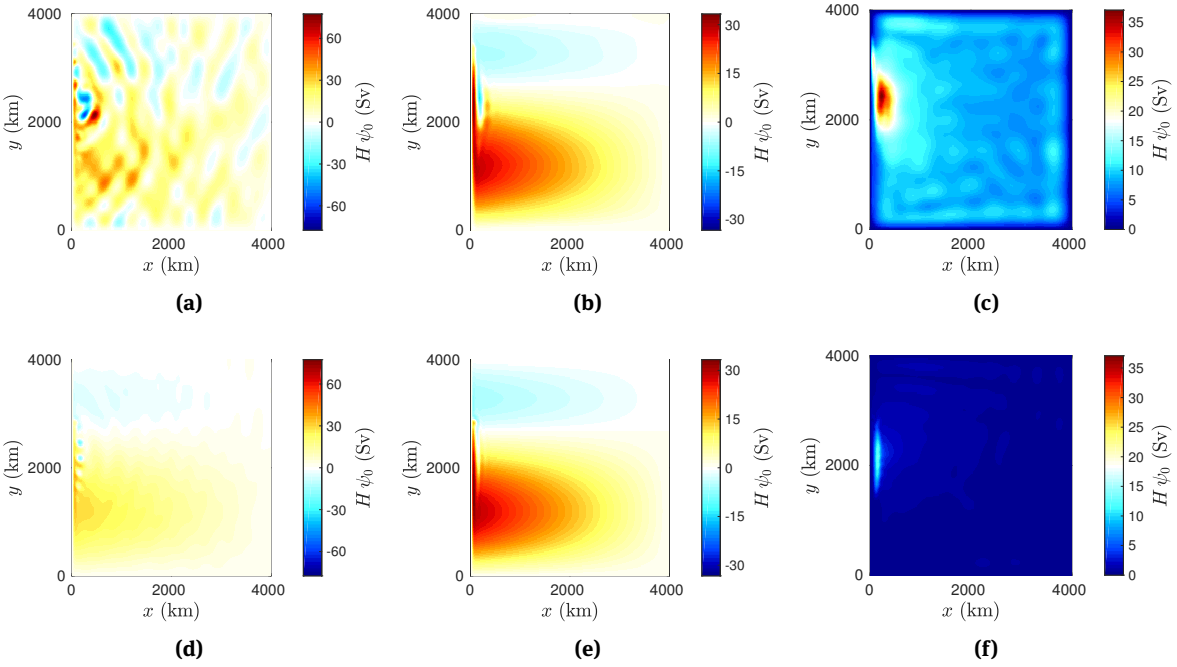

Figure 1: Properties of ψ_0 in 2M/D1M: (a)/(d) snapshot, (b)/(e) pointwise temporal mean, (c)/(f) pointwise temporal standard deviation.

We consider 2M as a minimal model that captures the main barotropic and baroclinic processes of interest, as well as the interactions between these dynamical processes. Notwithstanding, it is clear that 2M is strongly idealized, because of e.g. the assumption of QG dynamics, the idealized square basin geometry and the coarse vertical discretization with only two vertical modes. The model could be made more realistic, e.g. by increasing the number of vertical modes, by including a vertically dependent stratification, or by applying different boundary conditions. In particular, the relatively small eastward jet extension (related to the boundary conditions and stratification, see Figure 1) is not very realistic. However, the model allows for straightforward implementation of our stochastic modeling approach, as we discuss in the next section. Therefore, we consider it an appropriate test model for the purposes of developing the stochastic methodology from Verheul and Crommelin [47] in the setting of a spatially extended model.

As already mentioned in the introduction, in this study we focus on model reduction by vertical coarse-graining instead of horizontal coarse-graining. We aim to formulate a reduced model for the barotropic stream function and vorticity, with a stochastic representation of the baroclinic feedback R. Here, it amounts to reducing the number of degrees of freedom by a factor of 2, a modest reduction compared to what can be achieved with the more commonly pursued horizontal coarse-graining. We point out that it is straightforward to generalize to a higher number of vertical modes, by changing the definition of R in (1) to the sum of nonlinear feedback terms on the barotropic mode, $\sum_i J(\psi_i, q_i)$. For a recent investigation into the roles of the different individual baroclinic modes and their interaction, see Shevchenko and Berloff [43]. In our context the entire effect of baroclinicity is reduced to the one baroclinic eddy forcing of the barotropic mode (for any number of baroclinic modes in the reference model). Consequently, our approach is formally unaffected by more baroclinic modes in the reference model. Finally, Shevchenko and Berloff [43] reports mainly quantitative changes due to the inclusion of more baroclinic modes with the overall eddy effects remaining qualitatively similar. We speculate that the same will hold for the baroclinic eddy forcing and its stochastic parameterization.

More importantly, with vertical coarse-graining the full (baroclinic) model and the reduced model are both discretized on the same horizontal grid, so that R purely represents the effects of baroclinic instability. In this way, we avoid the conceptual difficulties of filtering (coarse-graining) fields that are discretized on a lattice, as also discussed in the introduction. R has a clear physical interpretation and we do not have to disentangle physical effects of unresolved processes from, e.g, grid transfer effects and reduced accuracy of discretized differential operators.

2.1 Spatial structure and restriction of the eddy forcing

Figures 3a and 3b show the temporal average $\mu(R)$ and standard deviation std(R) of the eddy forcing R as diagnosed from 2M. Both fields exhibit the same order of magnitude and are essentially confined to a narrow band at the western boundary. Within this region, $\mu(R)$ exhibits a dipole structure with negative values in the southern half and positive values towards the north. The two local extrema in $\mu(R)$ correspond to two local maxima in std(R).

The spatial structure of R suggests that it might be sufficient to model R within only a subdomain at the western boundary instead of the entire basin. In order to test this, we performed a 'truncated' 2-mode model simulation, T2M, which is identical to 2M except that R is set to zero outside the western boundary region WB = $[10, 490] \times [500, 3490]$ km. Figure 2 also shows the corresponding time series of enstrophy and the kinetic energy for T2M. The mean levels as well as both the short and long-term variability of enstrophy and kinetic energy are similar for T2M and 2M (the same holds when comparing the spatial fields shown in Figures 1a–1c with corresponding T2M fields, not shown) which indicates that T2M and 2M essentially produce the same flow dynamics for our model set-up. Consequently, we will focus on modeling R restricted to WB in the following; this has the practical advantage of reducing the volume of data that must be handled in our stochastic modeling of R (see Section 3.4).

Table 1: Parameter settings for all modal ocean numerical models (2M, D1M, S1M)

Parameter	Explanation	Value
α	Robert–Asselin filter parameter	0.1
β	rate of Coriolis change	1.8616×10^{-11} m^{-1}s^{-1}
A_H	lateral viscosity coefficient	10^2 m^{2}s^{-1}
$\overline{\rho}$	mean density	10^3 kgm^{-3}
H	ocean depth	4×10^3 m
τ_0	magnitude of wind-forcing	5×10^{-2} Nm^{-2}
ϵ_{111}	triple interaction coefficient	0
$\phi_1(z=0)$	surface eigenfunction	$\sqrt{2}$
λ	first Rossby radius of deformation	5×10^4 m
Δx	horizontal grid spacing x-direction	10^4 m
Δy	horizontal grid spacing y-direction	10^4 m
L_x	horizontal domain size x-direction	4×10^6 m
L_y	horizontal domain size y-direction	4×10^6 m
N_x	number of grid points x-direction	401
N_y	number of grid points y-direction	401
Δt	integration time step	1.8×10^3 s
δt	sampling interval	1.8×10^3 s
T_c	conditioning interval	9.43488×10^8 s
T_s	spin-up time	9.43488×10^8 s
T	integration time	9.43488×10^8 s
N_B	number of bins per conditioning variable	5

3 Stochastic model

The goal of this work is to formulate a stochastic process $\widetilde{\boldsymbol{R}}$ that emulates the eddy forcing $\boldsymbol{R}$ (see (1)). Adding this stochastic eddy forcing to D1M results in the stochastic 1-mode model (S1M):

$$\partial_t \widetilde{q}_0 + J(\widetilde{\psi}_0, \widetilde{q}_0) + \widetilde{R} + \beta \partial_x \widetilde{\psi}_0 = A_H \nabla^4 \widetilde{\psi}_0 + \frac{\partial_x \tau^y - \partial_y \tau^x}{\overline{\rho} H}. \tag{S1M}$$

Recall that throughout this work we compare variables from deterministic models (e.g. ψ_0) with their counterparts in a stochastic model (e.g. $\widetilde{\psi}_0$), and we use the same symbols for both but emphasize the difference with a tilde.

3.1 Conditioning procedure

To close the system S1M, a model that describes the temporal evolution of $\widetilde{R}$ is needed. We model $\widetilde{R}$ as a stochastic process, following the approach discussed in Verheul and Crommelin [47]. This approach is a form of resampling, in which $\widetilde{R}$ is sampled uniformly from conditioned observed values of R. However, whereas in [47] we considered a situation in which R was a scalar quantity, here we are dealing with a spatially extended system in which R is a two-dimensional field. Therefore, we extend the method from [47] to a multi-dimensional setting, and apply it pointwise to sample the field $\widetilde{R}$. In this extension, we preserve the modular design philosophy behind the stochastic methodology, as well as the ability to represent non-linear and non-Gaussian behavior.

Our stochastic methodology makes use of sample data $(\boldsymbol{\Psi}, \boldsymbol{Q}, \boldsymbol{R})$ from the full model 2M. As follows from (1), let us write $R^n := J(\psi_1(x, y, n\Delta t), q_1(x, y, n\Delta t))$ and $\boldsymbol{R} = (R^0, \ldots, R^N)$ to denote the time series of R resulting from a 2M simulation. Let $R^n(i, j)$ denote the eddy forcing value in grid point (i, j) at time n. The stochastic

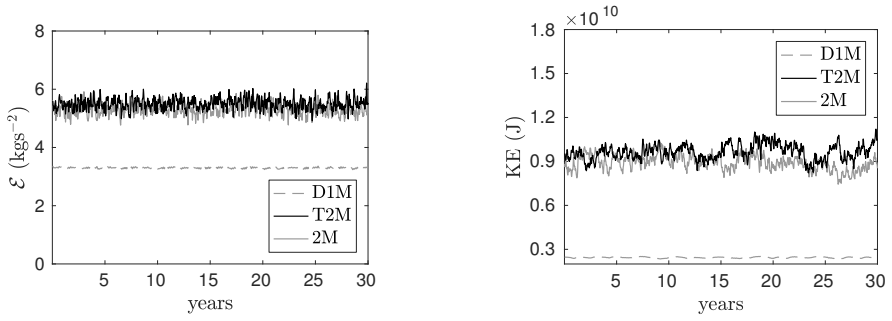

Figure 2: Comparison of scalar physical properties between the deterministic models: (a) enstrophy, (b) kinetic energy.

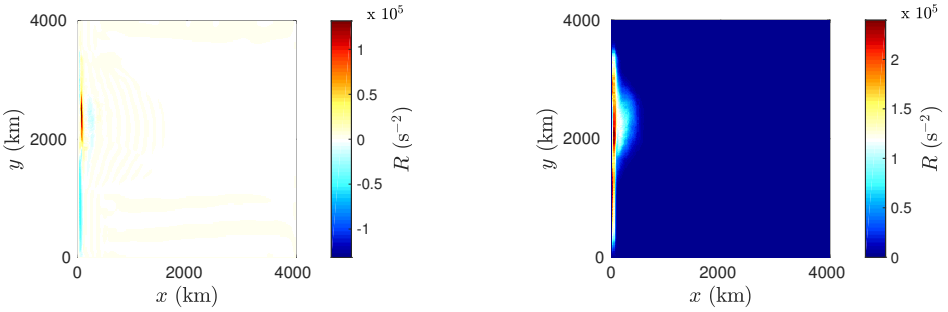

Figure 3: Properties of R in the reference 2M: (a) pointwise temporal mean, (b) pointwise temporal standard deviation.

forcing $\widetilde{R}$ in S1M is then sampled pointwise from sample data R conditioned on two types of covariates: time-lagged R-values $\widetilde{R}^{n-l_\theta(i,j)}(i,j)$ and a flow-dependent model variable $C(\widetilde{\psi}_0^n)(i,j)$:

$$\widetilde{R}^{n+1}(i,j) \sim R^{n+1}(i,j) \mid (R^{n-l_\theta(i,j)}(i,j) = \widetilde{R}^{n-l_\theta(i,j)}(i,j), C(\psi_0^n)(i,j) = C(\widetilde{\psi}_0^n)(i,j)), \quad \forall i,j \in \mathrm{WB},$$

or for short:

$$\widetilde{R}^{n+1}(i,j) \sim R^{n+1}(i,j) \mid (\widetilde{R}^{n-l_\theta(i,j)}(i,j), C(\widetilde{\psi}_0^n)(i,j)), \quad \forall i,j \in \mathrm{WB}, \tag{2}$$

where $l_\theta(i,j) + 1$ denotes the conditioning time lag for grid point (i,j) (see Section 3.2), and $C(\widetilde{\psi}_0^n)(i,j)$ is a function of $\widetilde{\psi}_0^n$ (see Section 3.3). Also, the stochastic forcing is generated only in the region WB and considered zero outside of this region (see Section 2.1).

Intuitively, the formulation in (2) says that the stochastic values are sampled from precisely those reference R-values that occurred in 2M when certain model variables in 2M matched the values of the corresponding variables in S1M. The distributions $R^{n+1}(i,j) \mid (\widetilde{R}^{n-l_\theta(i,j)}(i,j), C(\widetilde{\psi}_0^n)(i,j))$ are usually not known, therefore we approximate them with a simple binning method (see Section 3.4).

An important detail that we highlight is that the sampling interval used in the conditioning procedure (2) equals the integration time step (half an hour, see Table 1). This allows our parameterization a level of detail not usually seen in stochastic climate models, but this does come at a cost in the form of considerable memory requirements.

Crucially, the temporal evolution of $\widetilde{\boldsymbol{R}}$ is governed by sampling from the CPDFs in (2). Obviously, the set of conditioning model variables could be chosen arbitrarily large, making this methodology easily extendable. But for the sake of simplicity, and the test cases we discuss here, we consider at most one of each covariates: a single function $C(\psi_0)$ and a single lagged value of R. The set of conditioning variables is denoted $\{R^{n-l_\theta}, C(\psi_0))\}$. Different choices for $l_\theta(i,j)$ and $C(\psi_0^n)(i,j)$ give different sampling methods, and, in turn, different discrete stochastic processes $\widetilde{R}$. These stochastic processes together with S1M define several stochastic

models. In the following we discuss several sets of conditioning model variables, and we denote each stochastic model by their variable choices, i.e. S1M-$R[l_\theta]$ C. For example, if one chooses $l_\theta(i,j) = 0$, $\{C(\psi^n)(i,j)\} = \emptyset$, then S1M-$R[0]$ describes the 1-mode model S1M forced by a simple time-correlated stochastic process (see (2)).

3.2 Time-lagged covariate

An important criterion for our stochastic simulations is the reproduction of the autocorrelations of ψ_0 exhibited in the full model 2M. To reproduce the temporal correlations of R we condition the CPDFs in (2) on temporally lagged values of R, i.e. $R^{n-l_\theta(i,j)}$ for some $l_\theta(i,j) \geq 0$ (see (2)). The choice for lag times relates to an interesting, but different, question entirely: if one could sample the stochastic term $\widetilde{R}^{n+1}$ from the conditioned probability distribution $R^{n+1} \mid \widetilde{R}^n, \ldots, \widetilde{R}^{n-L}$, how large does L need to be to accurately reproduce temporal correlations shown in ($\boldsymbol{\Psi}$, $\boldsymbol{Q}$, $\boldsymbol{R}$)? This question can be phrased intuitively as: how much information of the history of the stochastic process is sufficient for our conditioning procedure. While this is an interesting problem, we consider this investigation outside of the scope of this paper, and take a heuristic approach.

We condition the CPDFs in (2) on a single lagged $\widetilde{R}$. We consider temporal decorrelation for each grid point to be the time at which the autocorrelation function first drops below the threshold $\theta = e^{-1}$. Figure 4 shows that the decorrelation time of R varies widely over the grid, i.e. anywhere from a day near the western boundary to 10 weeks around the eastward jet. Therefore, we expect the need to choose a pointwise lagged time l_θ in the conditioning procedure (2), i.e. to define a spatially dependent lag time $l_\theta(i,j)$. In Sections 4.1 and 4.3 we show results from stochastic simulations conditioning with lag times constant over the grid, as well as with spatially varying $l_\theta(i,j)$.

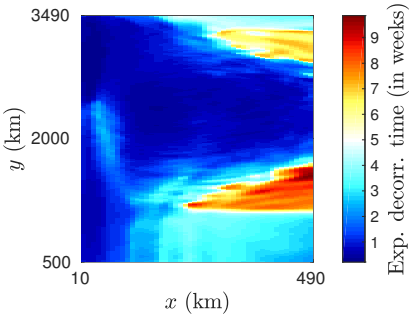

Figure 4: Pointwise decorrelation time (in days) of R over the region WB.

3.3 Flow-dependent covariate

Ideally, one conditions the CPDFs in (2) on resolved covariates, i.e. resolved model variables (in both D1M and 2M) that correlate strongly with the eddy forcing. The existence of such model variables in any multiscale model is not guaranteed. However, several studies have investigated and proposed candidates in related climate models (e.g. Porta Mana and Zanna [41]). In this section we choose a set of covariate candidates from the different terms constituting the PV budget. Expressing the eddy forcing R in 2M in terms of the other model variables gives:

$$R = -\partial_t \nabla^2 \psi_0 - J(\psi_0, q_0) - \beta \partial_x \psi_0 + A_H \nabla^4 \psi_0 + \frac{\partial_x \tau^y - \partial_y \tau^x}{\overline{\rho} H}. \tag{3}$$

This means that the eddy forcing R is expressed directly as a linear combination of each of the terms in the right-hand-side equation (3). However, unlike [41], we don't consider $\partial_t \nabla^2 \psi_0$ as a covariate candidate, because this is an unresolved term dependent on R itself. Each of the resolved terms in (3) are natural candi-

dates for our covariate selection procedure. Therefore, we define the set of covariate candidates V as:

$$V := \{\psi_0, q_0, \beta\partial_x\psi_0, J(\psi_0, q_0), A_H\nabla^4\psi_0, \frac{\partial_x\tau^y - \partial_y\tau^x}{\bar{\rho}H}\}.$$

We use linear regression analysis and standard Pearson coefficient plots to assess these candidate covariates. For the regression analysis all model candidate variables in V are considered the regressors, and R the response variable, respectively. The r^2-value (not to be confused with the eddy forcing R) for a given linear regression is a statistical quantity for the percentage of the response variable's variability that is 'explained' by the covariates. While a high r^2-value indicates a good regression fit, it by no means guarantees the best covariate. Therefore, we make the following observations associated with r^2-values only to compare different sets of covariates, and not to 'prove' quality.

Our pointwise regression analysis shows that from all variables in the set V, it is the Jacobian $J(\psi_0, q_0)$ that explains most of the variability of R. This can be seen by comparing two different regressions, first between the entire set V and R, second between $J(\psi_0, q_0)$ and R. Figures 5a and 5b show the pointwise r^2-values for these two regressions. Comparing the two different plots, one sees near identical r^2-values. This strongly indicates that the other candidates provide hardly any additional information for the regression. However, while the Jacobian shows the highest r^2-values, we note that it is far from a perfect predictor, as the central and eastern regions of WB remain badly represented.

Let us next compare the point wise Pearson correlation coefficients $p_{X,Y} = \mathrm{Cov}(X, Y)(\mathrm{std}_X\mathrm{std}_Y)^{-1}$, where $\mathrm{Cov}(X, Y)$ denotes the covariance between X and Y, and std_X and std_Y the standard deviations of X and Y respectively. Comparing each candidate in V coupled with R, we confirm the previous assessments that $J(\psi_0^n, q_0^n)$ correlates significantly with R, see Figure 5c. The highest correlation found between the Jacobian and R is located within the region WB described in Section 2.1. Other candidates in V show either significantly lower correlation to R or significantly lower r^2-values (not shown).

In addition to the statistical analysis above, our intuitive understanding for why the Jacobian is the most suitable covariate is that the Jacobian is the only resolved scale-coupling term in (3), i.e. term that is dependent on both small and large scale vortices. Therefore, despite the relatively low r^2-values in large parts of WB, $J(\psi_0, q_0)$ is our choice for covariate out of the tested candidates, and we choose $C(\widetilde{\psi}_0^n)(i, j) := J(\widetilde{\psi}_0^n(i, j), \widetilde{q}_0^n(i, j))$ in the following sections.

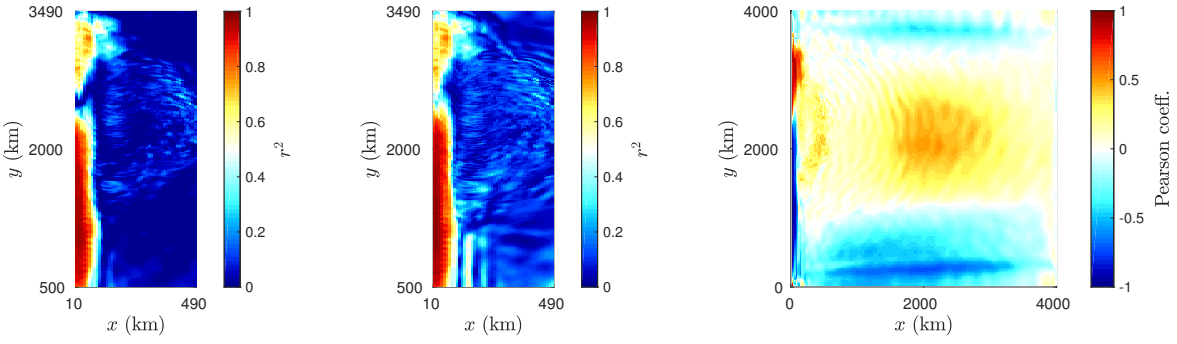

Figure 5: Covariate selection criteria. r^2-values for each grid point in $[10, 490] \times [500, 3490]$ km for pointwise regression analysis: (a) with regressor $J(\psi_0, q_0)$, (b) with regressors V. (c) Pearson coefficient $J(\psi_0, q_0)$ per grid point.

3.4 Sampling from empirical distribution

We choose the set of conditioning variables to contain either one or both of the time-lagged R-values $\widetilde{R}^{n-l_\theta(i,j)}(i, j)$ (see Section 3.2) and the Jacobian $J(\widetilde{\psi}_0^n(i, j), \widetilde{q}_0^n(i, j))$ (see Section 3.3). Following the set-up

from Verheul and Crommelin [47], we apply an equidistant binning procedure to approximate the CPDFs in conditioning procedure (2). We refer to uniformly drawing samples from these bins as sampling from the empirical distribution. To establish the binning associated with the chosen covariates, the range between the minimum and maximum of each covariate is independently partitioned in N_B equidistant intervals, with the outer intervals considered half-open. This partitioning results in disjoint bins, each of which describes a set of $R^{n+1}(i,j)$-values. Through this discretization, one obtains a mapping from conditioning variables to sets of $R^{n+1}(i,j)$-values. See Appendix A.2 for further technical details on the binning procedure.

Because our stochastic sampling procedure (2) acts pointwise, excessive artificial spatial roughness can arise in the generated fields $\widetilde{R}$. To prevent this phenomenon, we apply Gaussian spatial smoothing to the stochastic fields, as detailed in Appendix A.3. This is an ad hoc way to promote spatial smoothness, but one we consider adequate for the scope of this work. We leave a more systematic way to generate smooth fields for a future study.

A limitation of discrete sampling methods is that there is no predetermined way of handling situations in which the values of the conditioning variables are outside of the ranges exhibited in the sample data. We refer to bins outside of the ranges of the sample data as empty bins. When conditioning on both covariates, the bins are projected onto the linear trend $\widehat{J}$ between $R^{n-l_\theta(i,j)}(i,j)$ and $J(\psi_0^n(i,j), q_0^n(i,j))$ to more efficiently use the available bins (see Appendix A.1).

Furthermore, during simulations of the S1M, the Jacobian is removed from the conditioning variables whenever it goes outside the range in the sample data. The same is done at grid points where the correlation between J and R is low in the sample data (see Figure 5c).

We point out the computational efficiency of this sampling procedure. To evolve our stochastic model for the eddy forcing over time, we only need to calculate from which bin to sample for each grid point, and then draw a uniform random sample from that bin from memory. For comparison, in, for example, the approach from Cooper and Zanna [9], also relying on availability of high-resolution data, a system of linear stochastic ordinary differential equations (SDE) must be evolved at each model time step, involving six parameters and two variables per model grid point (for the model grid used here, this would amount to an SDE with 320000 variables). Besides the computational cost of integrating the model in time, the cost of constructing the stochastic eddy forcing model (i.e. the 'training phase') can be substantial. For our approach, it involves simple binning of data, with negligible computational cost. By contrast, the approach from [9] requires an expensive optimization procedure involving many reduced model runs.

3.5 Emulated stochastic eddy forcing

Similar to the investigations in [41], let us verify our stochastic methodology before coupling the stochastic process for R back to the reduced barotropic model. We use the output from a 2M simulation (i.e. sample data $\boldsymbol{J}$ and $\boldsymbol{R}$) to generate a so-called offline 'emulation' of the stochastic process. For the purposes of such an emulated process, let us choose $l = 0$:

$$\widetilde{R}^{n+1}(i,j) \sim R^{n+1}(i,j) \mid (\widetilde{R}^n(i,j), J(\widetilde{\psi}_0^n(i,j), \widetilde{q}_0^n(i,j))), \quad \forall i,j \in \mathrm{WB}.$$

This stochastic term is then compared to the same reference eddy forcing sample data. While not directly testing our ultimate goal of driving a reduced model with stochastic forcing, this poses an interesting question in itself: can our methodology reproduce the statistical properties of the reference eddy forcing when the input to our conditional sampling method is known to be 'correct'? Thus, this test should be considered a verification of the consistency of our procedure, rather than a validation. As intuition would suggest, even this simple verification can fail if, e.g, the number of bins is chosen too small or if the conditioning variables are not effective predictors.

Individual snapshots over time of these emulated stochastic fields show very little error. The long-term statistics are shown in Figure 6, where one can see an accurately reproduced mean and standard deviation for the spatially filtered emulated stochastic forcing. We note that the spatial smoothing (see Section A.3) decreases the standard deviation of $\widetilde{R}$ somewhat, compared to R. This should not be surprising, given that

the variability of $\widetilde{R}$ is artificially smoothed out. However, based on our experiments, we consider the benefit from the spatial cohesion and smoothness of $\widetilde{R}$ as a result of this spatial smoothing more significant than the disadvantage of its decreased standard deviation.

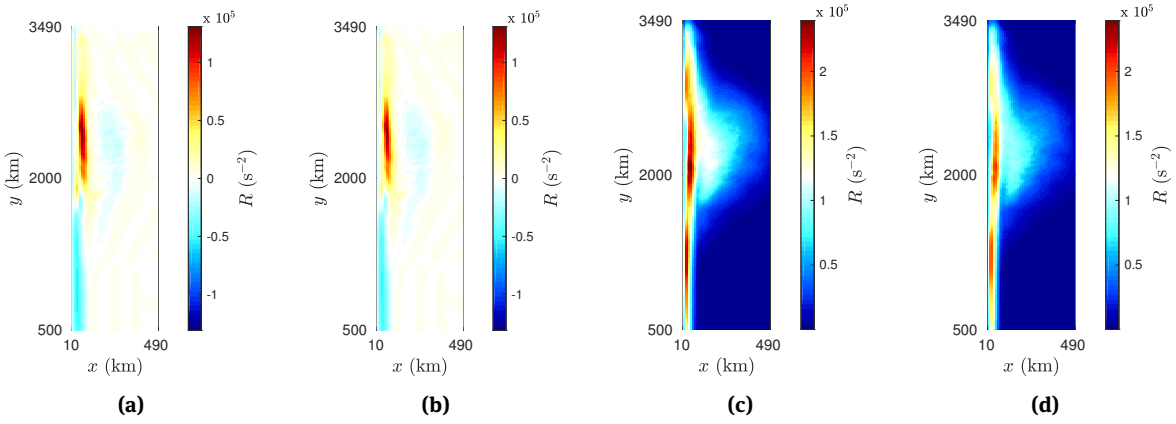

Figure 6: Comparison between pointwise temporal mean (a)-(b) and standard deviation (c)-(d) for the reference R (a) and (c) versus the filtered emulated $\widetilde{R}$ (b) and (d).

4 Results

The natural point of comparison between S1M and 2M is the evolution of the barotropic modes $\widetilde{\psi}_0$ and ψ_0. Therefore, we assess our stochastic parameterization by inspecting how well $\widetilde{\psi}_0$ reproduces the physical and statistical properties of ψ_0 in the reference 2M. The specific quantities we compare are: the enstrophy, kinetic energy, and energy transfer related to R (physical), and the statistical moments, and spatial and temporal correlations (statistical). See Appendix A.4 for a reference of formal definitions for each of these.

4.1 One-way coupling

Here, we consider a flow-independent sampling method, i.e. we consider only the time-lagged eddy forcing $R^{n-l_\theta(i,j)}$ as covariate and condition one-dimensionally: $R^{n+1} \mid \widetilde{R}^{n-l_\theta(i,j)}$. Let us start with a spatially constant $l_\theta(i,j) = l$:

$$\widetilde{R}^{n+1}(i,j) \sim R^{n+1}(i,j) \mid \widetilde{R}^{n-l}(i,j) \quad \forall i,j \in \mathrm{WB}. \tag{R[l]}$$

While we will use such simulations to highlight the reasons that spatially dependent lag times are desired, they also result in some interesting observations. Let us discuss these results for three different choices for l (recall the choice for half an hour integration time-step, see Section 2): half an hour (S1M-$R[\frac{1}{2}\mathrm{h}]$), a day (S1M-$R[1\mathrm{d}]$), and three days (S1M-$R[3\mathrm{d}]$).

These three reduced model simulations reproduce the mean barotropic stream function very well (not shown), with somewhat better results for longer time lags. The standard deviation of these barotropic stream functions (not shown) are a major improvement over D1M (Figure 1f), but are still significantly smaller than in 2M (Figure 1c). Figure 7a shows that the mean enstrophy (A.10) is also reproduced accurately for each of the lag time choices with a maximal error of 11% for S1M-$R[3\mathrm{d}]$ and an error of only 0.9% for S1M-$R[1\mathrm{d}]$. For each of the lag options the enstrophy's variability is slightly overestimated compared to the 2M reference diagnostics. Both the kinetic energy (A.11) and the energy exchange term in watts ($1\ \mathrm{W} = 1\ \mathrm{Js}^{-1}$, A.12) are significantly underestimated compared to 2M, as shown in Figures 7b and 7c, respectively. The mean kinetic

energies in Figure 7b deviate between 20% and 30% from 2M, in addition to showing long excursions from their means, unlike what the reference diagnostic exhibits. The energy transfer means deviate between 2% (for S1M-R[1d]) and 29% (for S1M-R[3d]), but show standard deviations that are between 18% and 26.7% off from the reference values. However, we do consider these results to be remarkably good, considering the straightforward methodology that generated them.

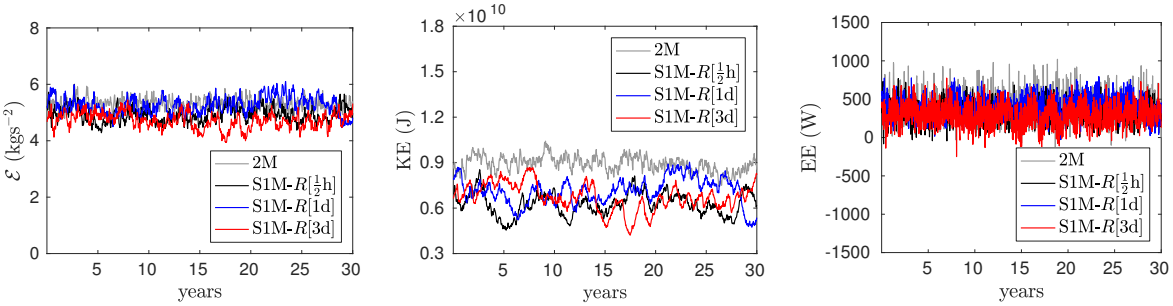

Figure 7: Comparison of scalar physical properties between 2M and S1M-$R[l]$ for different lags l: (a) enstrophy, (b) kinetic energy, (c) energy exchange term.

Figure 8 shows that similar plots result from spatially dependent lag times $l_\theta(i,j)$ (chosen equal to the time at which the autocorrelation function (ACF, A.9) first crosses some threshold θ):

$$\widetilde{R}^{n+1}(i,j) \sim R^{n+1}(i,j) \mid \widetilde{R}^{n-l_\theta(i,j)}(i,j) \quad \forall i,j \in \text{WB}. \qquad (R[l_\theta])$$

.

All means and deviations of the stochastic barotropic stream function improve significantly, a trend illustrated in Table 2. The temporal means of the kinetic energy plots in Figure 8b improve compared to Figure 7b to an error between 9% and 16%. Aside from this result, Figure 8 indicates that spatially dependent lag times in S1M-$R[l_\theta]$ fail to significantly improve the diagnostic physical results. However, these results instead illustrate the limitations of flow-independent conditioning methods $R[l]$ and $R[l_\theta]$, as supported by later results in Section 4.3.

To illustrate the reproduction of temporal correlations with a scalar quantity we first consider the ACFs in the two grid points [200, 2390] km and [440, 3190] km in two different dynamical areas in region WB. The ACF plots in Figures 11 and 12 show that, unlike the energetics and enstrophy, a significant improvement can be seen when comparing the S1M-$R[l]$ simulations (Figure 11) with the S1M-$R[l_\theta]$ simulations (Figure 12). This should not be surprising given that the constant lag times chosen are shorter than the $l_\theta(i,j)$-values, even for $\theta = 0.9$. This means that the information added to the stochastic process by the process history is relatively insignificant, i.e. the lagged R-values are not significantly decorrelated.

The spatial covariances here are represented by the covariances between $\widetilde{\psi}_0$ in a central grid point and its surrounding grid points (we again choose both [200, 2390] km and [440, 3190] km as the two example points). Figures 10c and 10d show that the S1M-$R[\frac{1}{2}h]$ very accurately reproduces the spatial structure of the reference covariances in Figures 10g and 10h, unlike the D1M covariances in Figures 10a and 10b. Note the significantly smaller magnitude for the barotropic references in Figures 10a and 10b, illustrating the strong improvement by each of the stochastic simulations over D1M. Quantitatively, however, the spatial covariances in these grid points (Figures 10c and 10d) are significantly underestimated compared to Figures 10g and 10h.

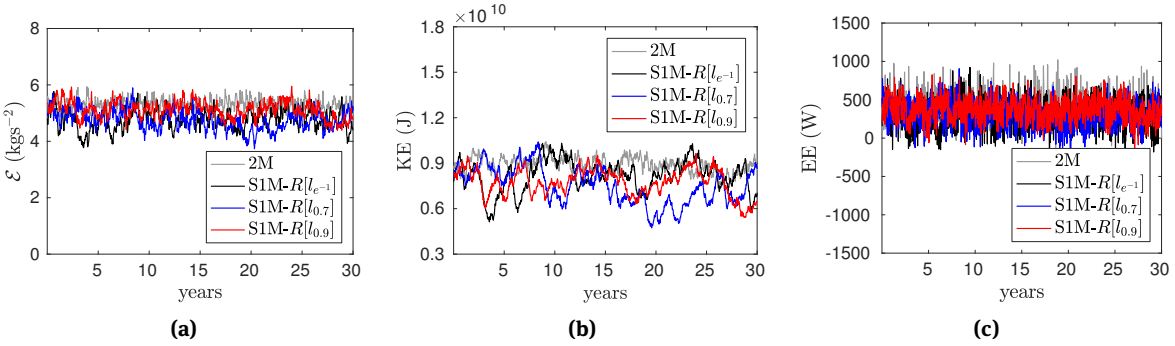

Figure 8: Comparison of scalar physical properties between 2M and S1M-$R[l_\theta]$ for different spatial lag patterns l_θ: (a) enstrophy, (b) kinetic energy, (c) energy exchange term.

4.2 Two-way coupling: single conditioning variable

The results from Section 4.1 already show promise for the suggested methodology, but could use improvement when it comes to the scalar physical quantities and spatial covariances, see Figures 10c and 10d. We expect both aspects to improve by conditioning on the Jacobian, which simultaneously adds flow-dependency and implicitly represented spatial correlations to neighboring grid points (as discussed in Section 3.3):

$$\widetilde{R}^{n+1}(i,j) \sim R^{n+1}(i,j) \mid J(\widetilde{\psi}_0^n(i,j), \widetilde{q}_0^n(i,j)) \quad \forall i,j \in \text{WB}. \tag{J}$$

Similar to the results in the previous section, S1M-J reproduces the mean barotropic stream function very well (not shown). However, it also significantly overestimates the standard deviation (not shown). Figure 9 shows that the comparisons to the flow-independent methods in Section 4.1 are unfavorable. The S1M-J simulation severely overestimates all of the previously considered physical quantities.

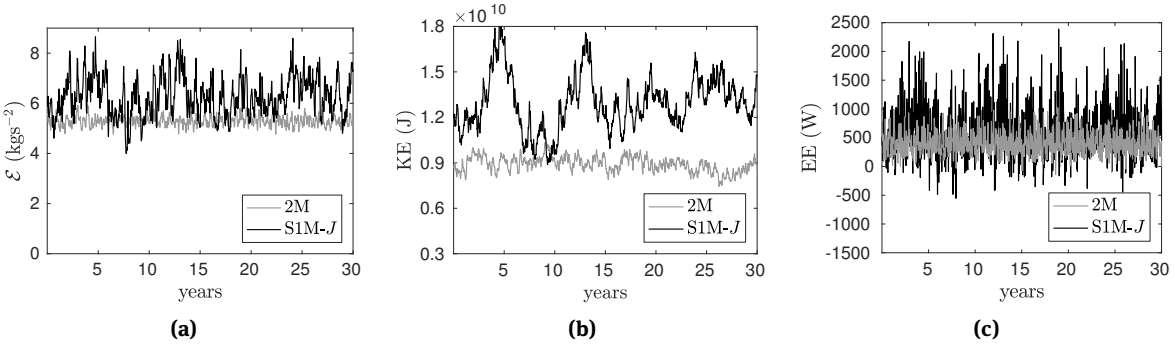

Figure 9: Comparison of scalar physical properties between 2M and S1M-J: (a) enstrophy, (b) kinetic energy, (c) energy exchange term (with wider range on y-axis).

Comparing the ACF plots of $\widetilde{\psi}_0$ between S1M-J and S1M-$R[l]/R[l_\theta]$ simulations tells a similar story. Figures 11 and 12 show that the ACFs are reproduced much more accurately when conditioning on time lagged values for R. This difference is to be expected, given that the flow-dependent conditioning (J) does not involve the process' history.

In contrast to the covariance plots for D1M (Figures 10a and 10b), the covariance plots for S1M-J (Figures 10e and 10f) show the same spatial structure of the covariances shown for 2M (Figures 10g and 10h). However, whereas the covariances for S1M-$R[\frac{1}{2}\text{h}]$ are significantly underestimated, the covariances for S1M-J are significantly overestimated.

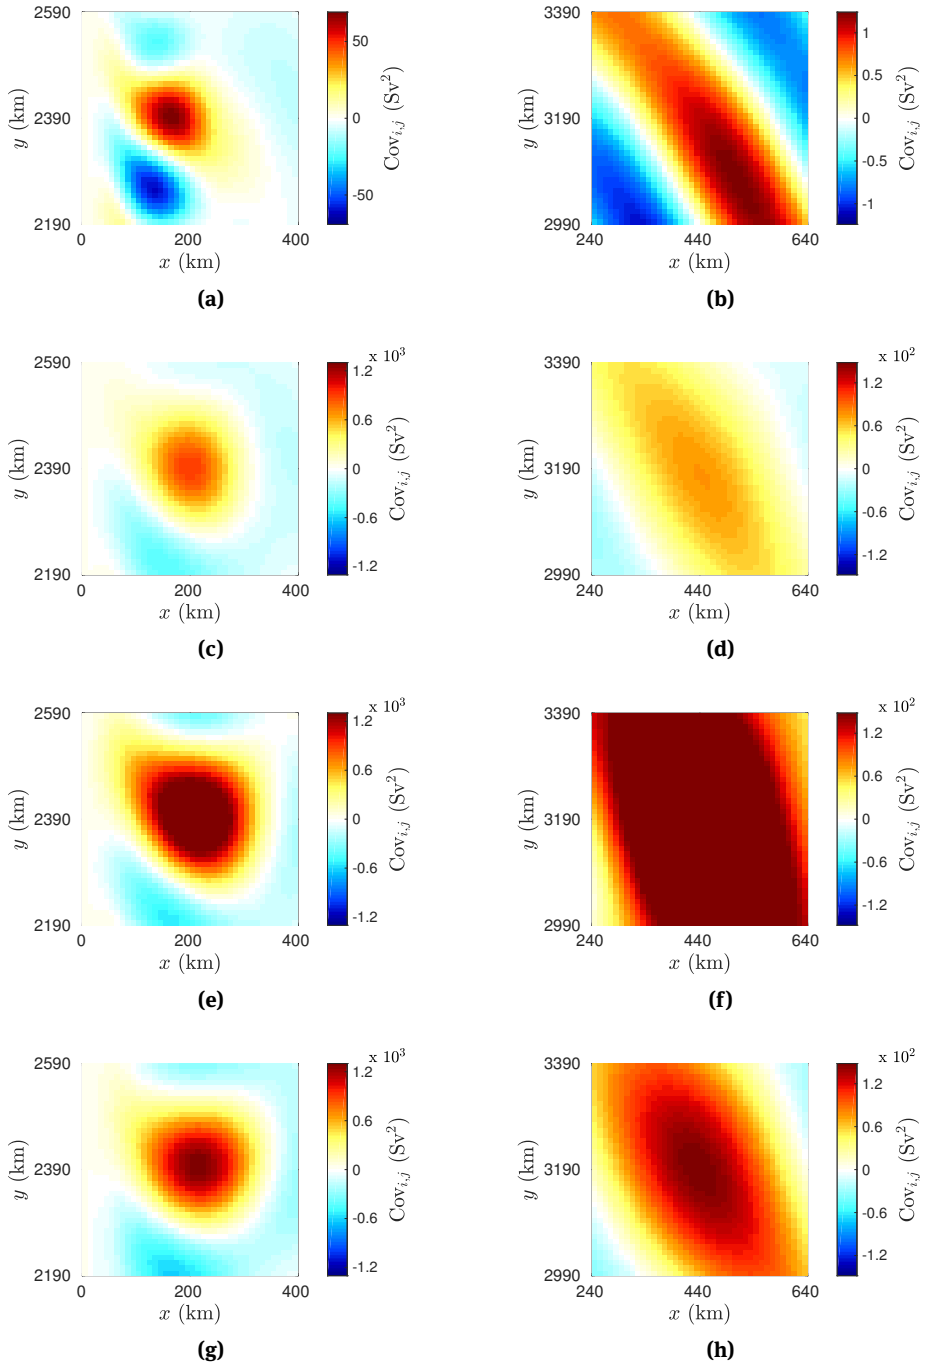

Figure 10: Covariance plots for ψ_0 between grid point [200, 2390] km/[440, 3190] km and neighbouring grid points (see (A.13)) for each of the following models: (a)/(e) D1M, (b)/(f) S1M-$R[\frac{1}{2}\text{h}]$, (c)/(g) S1M-J, (d)/(h) 2M.

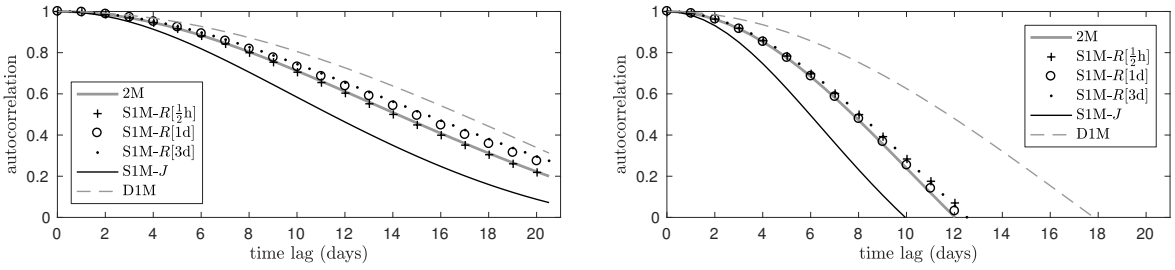

Figure 11: Comparison of ACFs of ψ_0 between D1M, 2M, S1M-J, and S1M-$R[l]$ for different lags l for grid point: (a) [200, 2390] km, (b) [440, 3190] km.

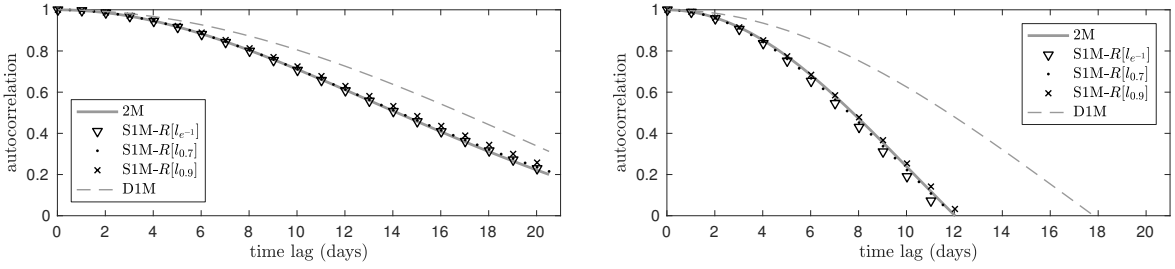

Figure 12: Comparison of ACFs of ψ_0 between D1M, 2M, and S1M-$R[l_\theta]$ for different lag patterns l_θ for grid point: (a) [200, 2390] km, (b) [440, 3190] km.

The overall conclusion from these tests is then that the Jacobian leads to overestimated amplitudes for most considered diagnostic criteria, i.e. enstrophy, energetics, spatial covariances, and standard deviation of ψ_0 (see Table 2), whereas the autocorrelations are underestimated. This further emphasizes our assessment that the Jacobian is far from a perfect predictor (as briefly discussed in Section 3.3). However, this flow-dependent spatially correlated driving force, albeit too erratic as sole conditioning variable, can improve the previously discussed flow-independent results.

4.3 Two-way coupling: double conditioning variables

By combining the aspects of the tests described in Sections 4.1 and 4.2, one arrives at the two-fold conditioning procedure, as described in (2), where $C(\psi_0^n)(i,j) = \widehat{J}(\psi_0^n(i,j), q_0^n(i,j))$ ($\widehat{J}$ denotes the linearly fitted J, see Section A.1). Similar to the tests in Section 4.1, let us first consider the simulations that condition on lagged R-values with constant lag l:

$$\widetilde{R}^{n+1}(i,j) \sim R^{n+1}(i,j) \mid (\widetilde{R}^{n-l}(i,j), \widehat{J}(\widetilde{\psi}_0^n(i,j), \widetilde{q}_0^n(i,j))), \quad \forall i,j \in \mathrm{WB}. \tag{$R[l]\widehat{J}$}$$

The motivation for the sampling method ($R[l]\widehat{J}$) is to combine the benefits of both conditioning variables, i.e. the spatial structure of the Jacobian, and temporal correlations from lagged R-values, respectively.

Similar to previous simulations, all the tested simulations reproduce the mean barotropic stream function excellently (not shown), but the S1M-$R[4h]\widehat{J}$ and S1M-$R[3d]\widehat{J}$ simulations overestimate the standard deviation (not shown). While an immediate improvement over the S1M-$R[l]$ simulations in Section 4.1 can be seen, the results with the S1M-$R[l]J$ model are quite sensitive to the choice of l. This is illustrated in Figure 13, where

physical results from a S1M-$R[4\text{h}]\widehat{J}$ simulation are shown in addition to the same lag-time choices discussed in Section 4.1. In Figure 13 the enstrophy and energetics are plotted for the various spatially constant lag times. On the one hand, the simulations S1M-$R[\frac{1}{2}\text{h}]\widehat{J}$ and S1M-$R[1\text{d}]\widehat{J}$ result in major improvements over the physical diagnostics resulting from S1M-$R[l]$ (Figure 7) and S1M-J (Figure 9). Specifically, the mean of the kinetic energy only deviates 2.9% and 1.7% from 2M's reference for S1M-$R[\frac{1}{2}\text{h}]\widehat{J}$ and S1M-$R[1\text{d}]\widehat{J}$, respectively. On the other hand, the S1M-$R[4\text{h}]\widehat{J}$ model performs overall worse than S1M-$R[l]$ (Figure 8).

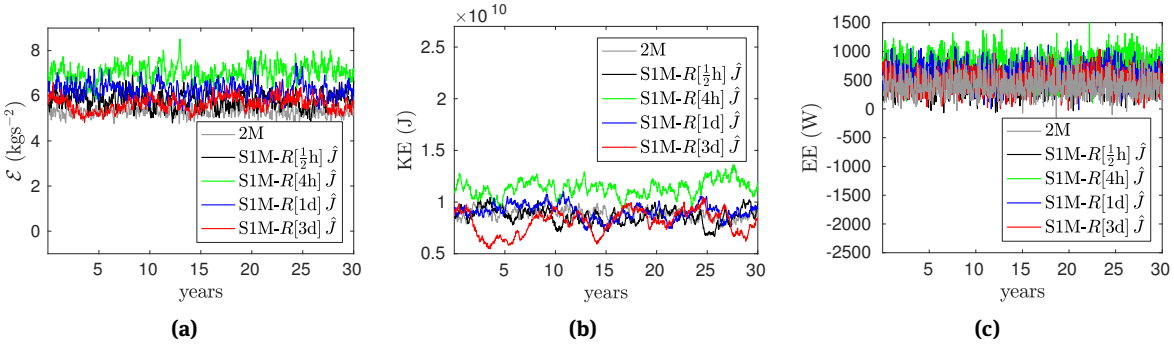

(a) (b) (c)

Figure 13: Comparison of scalar physical properties between 2M and S1M-$R[l]\widehat{J}$ for different lags l: (a) enstrophy, (b) kinetic energy (with wider range on y-axis), (c) energy exchange term.

The sensitivity discussed above stems from the choice for a constant lag l because, as discussed in Section 3.2, decorrelation times of R vary widely between grid points. Instead of using the spatially constant lag times, we use spatially variable lag times based on the decorrelation time scales of the eddy forcing $R(i,j)$. As in Section 4.1, this spatially variable $l_\theta(i,j)$ is chosen equal to the time lag at which the ACF for $R(i,j)$ first crosses some threshold θ:

$$\widetilde{R}^{n+1}(i,j) \sim R^{n+1}(i,j) \mid (\widetilde{R}^{n-l_\theta(i,j)}(i,j), \widehat{J}(\widetilde{\psi}_0^n(i,j), \widetilde{q}_0^n(i,j))), \quad \forall i,j \in \text{WB}. \qquad (R[l_\theta]\widehat{J})$$

The results for S1M-$R[l_\theta]\widehat{J}$ are shown in Figure 14, for several values of θ. With $\theta = 0.9$, enstrophy and energy exchange are too high. We hypothesize that with $\theta = 0.9$, the lagged R-values are still very strongly correlated, so that they add little information and the conditioning is dominated by the Jacobian. As a result, S1M-$R[l_{0.9}]\widehat{J}$ suffers from similar errors as S1M-J (see Section 4.2).

The results with $\theta = 0.7$ and $\theta = e^{-1}$ are overall very good, with diagnostics in Figure 14 close to those of the reference model 2M. We focus here on S1M-$R[l_{0.7}]\widehat{J}$, however results for S1M-$R[l_{e^{-1}}]\widehat{J}$ are highly comparable. For $\theta = 0.7$ the mean of the enstrophy, kinetic energy, and energy exchange terms are all reproduced excellently, with an error of 3.5%, 0.9%, and 9.2%, respectively. Additionally, the standard deviation of the energy exchange term is also within 4.7% of the reference, proving another significant improvement over the previously tested approaches.

We note that the standard deviation of the kinetic energy is too high for all S1M-$R[l_\theta]\widehat{J}$. This is caused by the limited spatial dependency in our sampling method ($R[l_\theta]\widehat{J}$), which can lead to forcing fields that are spatially less smooth than in the 2M reference model. This increased spatial roughness affects local gradients and thereby the kinetic energy (see (A.11)). Despite this shortcoming, S1M-$R[l_{0.7}]\widehat{J}$ performs well by all other criteria.

By preserving the temporal information provided by the lagged R-values, the ACFs for 2M in $[200, 2390]$ km and $[440, 3190]$ km are reproduced almost exactly, as shown in Figure 15. This drastically improves on the autocorrelations reproduced by S1M-J and S1M-$R[l]$ (Figure 11), and equals the best results obtained with pure temporal stochastic parameterizations S1M-$R[l_\theta]$ (Figure 12).

Let us next consider ACFs more comprehensively by focusing on the entire WB region. Consider the exponential decorrelation time scales, i.e. the time lag at which the ACF for the barotropic stream function first

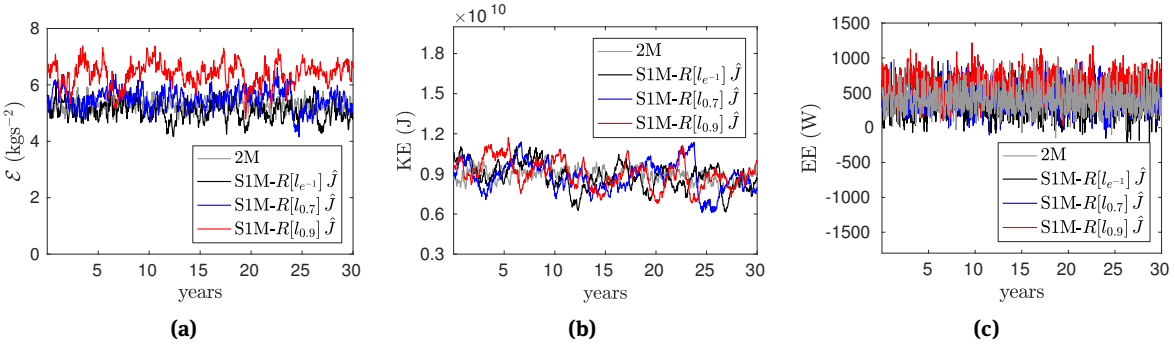

Figure 14: Comparison of scalar physical properties between 2M and S1M-$R[l_\theta]\widehat{J}$ for different spatial lag patterns l_θ: (a) enstrophy, (b) kinetic energy, (c) energy exchange term.

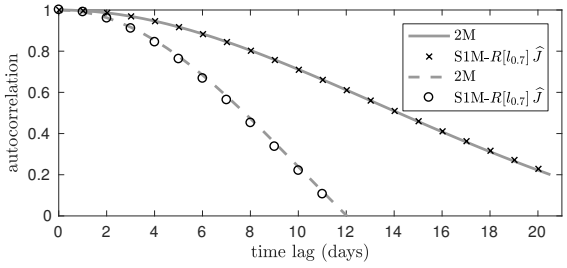

Figure 15: Comparison of the ACFs of ψ_0 between the two grid points [200, 2390] km (line and crosses) and [440, 3190] km (dashed and circles) between 2M and S1M-$R[l_{0.7}]\widehat{J}$

dips below e^{-1}, pointwise over the whole western boundary region (WB). Figure 16 shows the drastic differences of exponential decorrelation time scales between D1M and 2M; the significantly weaker vortices near the western boundary present in D1M cause much longer decorrelation time scales across the WB region.

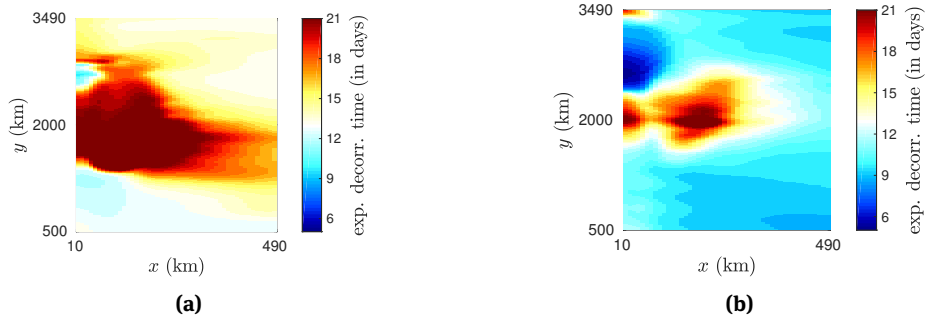

Figure 16: Pointwise exponential decorrelation time (in days) of ψ_0 in each of the following models: (a) D1M, (b) 2M.

Consistent with the ACFs in Figures 11a and 11b, S1M-J preserves little of the temporal history (Figure 17a), as its ACFs decorrelate much faster over the whole WB region. However, this result already presents a strong improvement over decorrelation time scales exhibited in D1M (Figure 16a). Significant improvements come once again when we observe the results for simulations that condition on lagged R-values: qualitatively similar patterns for both S1M-$R[\frac{1}{2}\mathrm{h}]$ (Figure 17b) and S1M-$R[l_{0.7}]\widehat{J}$ (Figure 17c). These patterns approximate the reference 2M (Figure 16b) very well. None of the tested simulations are able to adequately reproduce the high decorrelation time scales in the $[10, 40]\times[1800, 2300]$ km region. This may be caused by boundary effects, or by the dynamic complexities of the gyre's detachment point at this location. We do, however, emphasize that the main outcome from these results is that our goal of improving the stochastic model for $\widetilde{R}$ by combining

spatial and temporal information (through conditioning on J and $R[l_\theta]$, respectively) is accomplished with the spatially dependent sampling method ($R[l_\theta]\widehat{J}$).

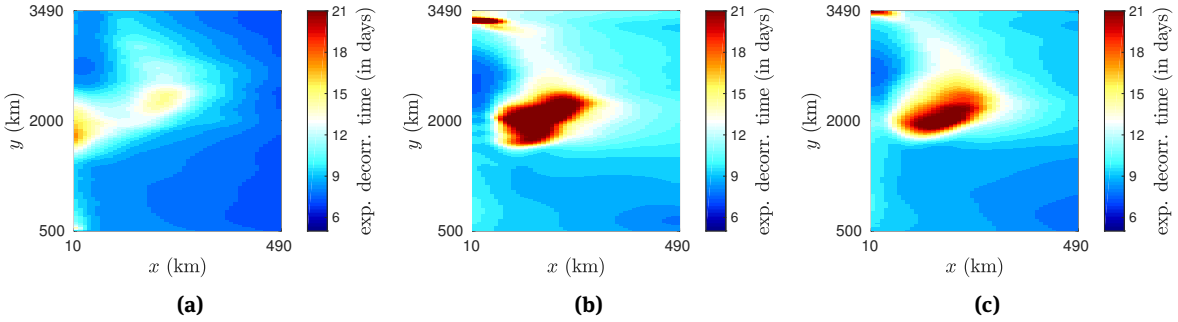

Figure 17: Pointwise exponential decorrelation time (in days) of $\widetilde{\psi}_0$ in each of the following models: (a) S1M-J, (b) S1M-$R[\frac{1}{2}\text{h}]$, (c) S1M-$R[l_{0.7}]\widehat{J}$.

Table 2 shows the expectation and maximum of the absolute errors ϵ_i over the grid for the first four statistical moments, e.g. $\epsilon_1 = \mu(H\psi_0) - \mu(H\widetilde{\psi}_0)$. Multiple errors for S1M-$R[\frac{1}{2}\text{h}]\widehat{J}$ are higher than their corresponding values for either S1M-$R[\frac{1}{2}\text{h}]$ or S1M-J, illustrating the previously discussed difficulties with spatially constant lag times. However, the S1M-$R[l_{0.7}]\widehat{J}$ simulation does not solve this completely, as the error ϵ_1 is shown to be worse than for S1M-$R[l_{0.7}]$. One sees that adding conditioning variables for our discrete sampling method does not guarantee a universal improvement to the statistical quantities, because the added conditioning variable J does correlate strongly with the eddy forcing in only part of the WB region. Aside from the first statistical moment, the S1M-$R[l_{0.7}]\widehat{J}$ simulation gives the best overall results for the statistical moments, consistent with our motivations and results, described earlier in this section. The most drastic improvements with respect to the D1M are found in the errors of the standard deviation, reducing the error by an order of magnitude. Crucially, these results are further emboldened by observing the spatial fields of the statistical moments of $\widetilde{\psi}_0$. Figures 18a and 18b show that the reference mean and standard deviation of ψ_0 (Figures 1b and 1c) are indeed extremely well reproduced by S1M-$R[l_{0.7}]\widehat{J}$.

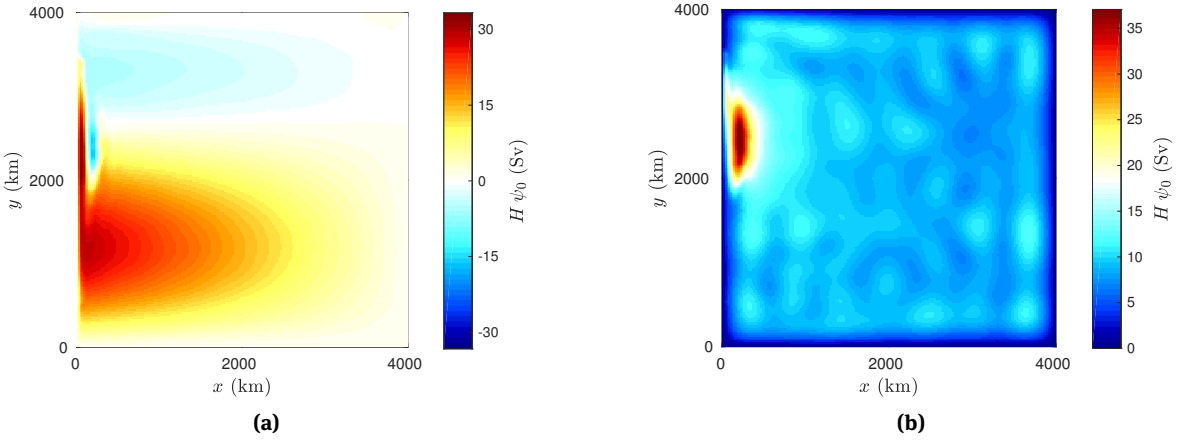

Figure 18: Pointwise temporal statistical moments of $\widetilde{\psi}_0$ in S1M-$R[l_{0.7}]\widehat{J}$: (a) mean, (b) standard deviation.

Figure 19 shows that, besides improving the energetics and statistical moments of the system, S1M-$R[l_{0.7}]\widehat{J}$ is also able to reproduce spatial covariances present in 2M, as can be seen by comparing Figure

Table 2: Mean and maximum absolute errors of the first four statistical moments ϵ_s of the barotropic stream functions of stochastic simulations with different sets of conditioning variables.

conditioning	E ϵ_1 (10^{-1} Sv)	max ϵ_1 (10^{1} Sv)	E ϵ_2 (Sv)	max ϵ_2 (10^{1} Sv)	E ϵ_3 (10^{-2})	max ϵ_3	E ϵ_4 (10^{-1})	max ϵ_4 (10^{1})
D1M	3.33	2.82	8.19	3.33	7.16	3.31	3.81	3.00
S1M-$R[\frac{1}{2}\text{h}]$	1.95	1.18	2.80	1.20	5.56	3.15	2.12	3.04
S1M-$R[l_{0.7}]$	1.42	0.71	1.00	0.97	5.24	3.22	2.73	3.02
S1M-J	2.05	1.03	5.50	1.81	6.13	4.38	2.40	3.01
S1M-$R[\frac{1}{2}\text{h}]\widehat{J}$	1.75	0.95	0.77	1.39	6.02	2.83	2.02	3.01
S1M-$R[l_{0.7}]\widehat{J}$	1.64	0.83	0.72	0.77	5.45	2.94	1.92	2.97

19 with Figures 10g and 10h. This is a significant improvement over both S1M-$R[\frac{1}{2}\text{h}]$ (Figures 10c and 10d) and S1M-J (Figures 10e and 10f). Given that both S1M-$R[\frac{1}{2}\text{h}]$ and S1M-J reproduced the spatial patterns of the covariances qualitatively well, this quantitative improvement is most likely due to the more accurately reproduced standard deviation of the barotropic stream field (see Table 2).

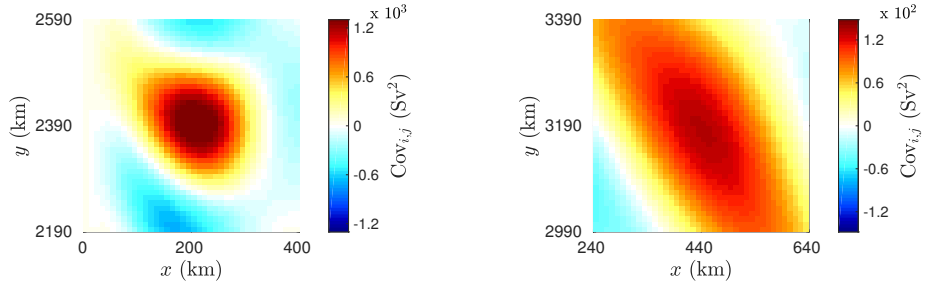

Figure 19: Covariance plots for ψ_0 for S1M-$R[l_{0.7}]\widehat{J}$ between grid point [200, 2390] km/[440, 3190] km) in (a)/(b) and neighbouring grid points.

To support these claims, let us investigate spatial correlations over the region [10, 690]×[1700, 3200] km, i.e. the region of the domain where the standard deviation of the barotropic stream is significant (see Figure 1c). Similar to the spatial covariances, for each grid point in the region we compute the correlations between ψ_0 in this central grid point and its surrounding grid points, see (A.14). Contrasted with the reference correlations in 2M, the mean absolute correlation errors in grid point (i, j) for a stochastic model are given by:

$$g(i,j) = P(i,j)^{-1} \sum_{\substack{i',j'>0 \\ i-20\le i'\le i+20 \\ j-20\le j'\le j+20}} |\,\mathrm{Corr}(\psi_0(i,j),\psi_0(i',j')) - \mathrm{Corr}(\widetilde{\psi}_0(i,j),\widetilde{\psi}_0(i',j'))\,|, \tag{4}$$

where $P(i, j)$ denotes the number of grid points (i', j') over which the summation in (4) runs (the summation cannot run over grid points that exceed the boundaries of the full domain), and ψ_0 and $\widetilde{\psi}_0$ are the barotropic stream functions of 2M and of the stochastic model, respectively.

These mean absolute correlation errors are shown in Figure 20 for each of the highlighted test models. The figure shows that the two grid points chosen to illustrate the covariances in Figure 10 are representative for the globally reproduced correlations. Importantly, Figure 20 shows that all stochastic models improve significantly on the spatial correlations as reproduced by the barotropic reference D1M. Additionally, one sees that, besides overestimating the magnitude of the covariances (Figures 10e and 10f), S1M-J (Figure 20b) re-

produces the spatial patterns of the correlations with less accuracy than both S1M-$R[\frac{1}{2}\text{h}]$ (Figure 20c) and S1M-$R[l_{0.7}]\widehat{J}$ (Figure 20d). These latter two reproduce the spatial correlations of the reference 2M with a similarly high accuracy. This indeed indicates that the improvements to the spatial covariances by S1M-$R[l_{0.7}]\widehat{J}$ (Figures 19a and 19b) are due to the standard deviation of the flow being better resolved, and is less likely attributed to the spatial correlations.

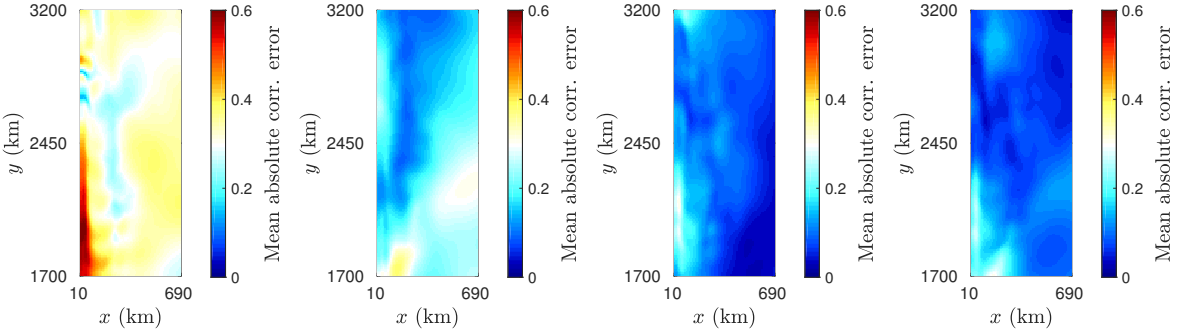

Figure 20: Mean absolute correlation error (4) of $\widetilde{\psi}_0$ in each of the following models: (a) D1M, (b) S1M-J, (c) S1M-$R[\frac{1}{2}\text{h}]$, (d) S1M-$R[l_{0.7}]\widehat{J}$.

Concluding, we have successfully introduced flow-dependency into the stochastic parameterization. In many respects the flow-independent parameterization S1M-$R[l_{0.7}]$, discussed in Section 4.1, already shows promising results. By introducing the Jacobian into the conditioning S1M-$R[l_{0.7}]\widehat{J}$ we have further improved almost all of the considered physical and statistical criteria posed in Section 1.2.

5 Summary and discussion

In this study we investigated a covariate-based stochastic approach to parameterize unresolved processes within a standard model of the idealised, wind-driven ocean circulation. We considered the reduction of the reference 2-mode baroclinic model (2M) to a 1-mode barotropic model (D1M). The reduced D1M lacks the baroclinic feedback given by the baroclinic eddy forcing R. We developed a stochastic model for R, and coupled it to the 1-mode model to obtain a stochastic 1-mode barotropic model (S1M). With a suitable stochastic model for R, S1M is able to mimic the behavior of the barotropic mode in the reference 2M closely.

We focused on vertical instead of horizontal coarse-graining, such that all considered models (2M, D1M, S1M) are discretized on the same high-resolution horizontal grid. Hereby, we avoid the subtle difficulties of horizontal coarse-graining (e.g. choices of filter and grid transformation), and can fully focus on the stochastic model formulation. The corresponding eddy forcing R is uniquely defined and has a clear physical interpretation solely related to the baroclinic nature of the flow.

The stochastic parameterization of the eddy forcing R is based on a covariate-approach recently developed in Verheul and Crommelin [47] within a scalar set-up. Here we construct a pointwise spatial extension of the covariate-approach such that it can be applied to a spatially extended ocean model. More precisely, in S1M the eddy forcing R is modeled as a spatially extended stochastic process $\widetilde{R}$. Sample data from a 2M reference simulation for both the eddy forcing R and the resolved model variables is assumed to be given for our approach. The stochastic term $\widetilde{R}$ is sampled uniformly from conditional probability distribution functions (CPDFs) approximated over the available sample data, i.e. sampled from the so-called conditional empirical distributions. The CPDFs are approximated with a simple binning procedure. $\widetilde{R}$ is then conditioned on both the most suitable flow-dependent covariate, which turned out to be the resolved nonlinear advection field,

and past states of $\widetilde{R}$ itself, inducing adequate temporal correlations. For our specific ocean model, the volume of employed data could be significantly reduced by limiting the stochastic forcing to the western boundary sub-region of the grid, which in diagnostic tests proved to reproduce the flow dynamics of 2M.

In order to evaluate the performance of the stochastic parameterization, we compared the reference and reduced stochastic models (i.e. 2M and S1M) with respect to a range of physical and statistical criteria. These criteria are the first four statistical moments, the autocorrelation function, spatial covariances and correlations of the stream function, kinetic energy, energy conversion, and the enstrophy.

The results show that the flow-dependent covariate and the past states of $\widetilde{R}$ each contribute in their own way to the stochastic model. Conditioning on the past states of $\widetilde{R}$ introduces temporal consistency (Section 4.1). The flow-dependent covariate introduced an energetic driving force and improved the spatial cohesion of the system, but by itself is an insufficient predictor for the eddy forcing (Section 4.2).

Conditioning both on the flow-dependent covariate and on past $\widetilde{R}$-values resulted in a good model for R, so that S1M reproduced many characteristics of the 2M reference model very well (Section 4.3). In order to achieve this it turned out to be crucial to account for the spatial dependence of the decorrelation time scale of R.

Overall our final results with S1M show a large improvement over D1M, as well as very close similarity to the reference 2M. Additionally, the results provide further proof of concept of our methodology, extending the approach from Verheul and Crommelin [47] to a spatially extended setting. Because the empirical distribution is discrete and needs only one sample point per bin to sample from, it is a method robust to small sample sets. However, practical limitations can pose challenges, because of large amount of data involved for spatially extended systems. This led us to develop straightforward methods to use sample data more effectively, as well as to rethink ways to handle sparsely populated bins (Section 3.4). With a large amount of data, efficient storage and access of the data becomes more important for our methodology. In this study, optimization of memory usage and data storage was not our priority, however we point out that several straightforward optimizations are possible for our approach, e.g. reduced sampling interval, fast database-based lookup tables, and interpolated sample data. We leave these for future study. Furthermore, we emphasize the computational efficiency of the methodology developed here, as our methodology requires only straightforward calculations to determine what piece of memory to read out, i.e. what bin to sample from.

In future work we intend to explore methods to remedy these limitations. The approach proposed here can be further extended to improve the representation of spatial structures and correlations in the parameterized eddy forcing fields by, e.g, exploring additional types of flow-dependent covariates (e.g. based on energetics). A more thorough sensitivity analysis of the conditioning time lags of $\widetilde{R}$ would also be helpful and insightful for the covariate selection. Furthermore, we are interested in exploring continuous approximation methods for the CPDFs, as well as other continuous stochastic parameterizations.

Acknowledgement: The authors would like to thank Henk A. Dijkstra for insightful discussions. This research is funded by the Netherlands Organization for Scientific Research (NWO) through the Vidi project "Stochastic models for unresolved scales in geophysical flows".

References

[1] Arakawa, A. (1966). Computational design for long-term numerical integration of the equations of fluid motion: Two-dimensional incompressible flow. part i. *J. Comput. Phys.*, 1(1):119–143.

[2] Asselin, R. (1972). Frequency filter for time integrations. *Mon. Weather Rev.*, 100(6):487–490.

[3] Berloff, P. (2015). Dynamically consistent parameterization of mesoscale eddies. part i: Simple model. *Ocean Modell.*, 87:1–19.

[4] Berloff, P., Dewar, W., Kravtsov, S., and McWilliams, J. (2007). Ocean eddy dynamics in a coupled ocean–atmosphere model. *J. Phys. Oceanogr.*, 37(5):1103–1121.

[5] Berloff, P. S. (2005a). On dynamically consistent eddy fluxes. *Dyn. Atmos. Oceans*, 38(3):123–146.

[6] Berloff, P. S. (2005b). Random-forcing model of the mesoscale oceanic eddies. *J. Fluid Mech.*, 529:71–95.

[7] Berloff, P. S., McWilliams, J. C., and Bracco, A. (2002). Material transport in oceanic gyres. part i: Phenomenology. *J. Phys. Oceanogr.*, 32(3):764–796.
[8] Berner, J., Achatz, U., Batte, L., Bengtsson, L., Cámara, A. d. l., Christensen, H. M., Colangeli, M., Coleman, D. R., Crommelin, D., Dolaptchiev, S. I., et al. (2017). Stochastic parameterization: Toward a new view of weather and climate models. *Bull. Am. Meteorol. Soc.*, 98(3):565–588.
[9] Cooper, F. C. and Zanna, L. (2015). Optimisation of an idealised ocean model, stochastic parameterisation of sub-grid eddies. *Ocean Modell.*, 88:38–53.
[10] Crommelin, D. and Vanden-Eijnden, E. (2008). Subgrid-scale parameterization with conditional Markov chains. *J. Atmos. Sci.*, 65(8):2661–2675.
[11] Dijkstra, H. A. (2005). *Nonlinear physical oceanography: a dynamical systems approach to the large scale ocean circulation and El Nino*, volume 28. Springer Science & Business Media.
[12] Dorrestijn, J., Crommelin, D. T., Siebesma, A. P., and Jonker, H. J. (2013). Stochastic parameterization of shallow cumulus convection estimated from high-resolution model data. *Theor. Comput. Fluid Dyn.*, 27(1-2):133–148.
[13] Durran, D. R. (1999). *Numerical methods for wave equations in geophysical fluid dynamics*.
[14] Eden, C., Jochum, M., and Danabasoglu, G. (2009). Effects of different closures for thickness diffusivity. *Ocean Modell.*, 26(1):47–59.
[15] Farneti, R. and Gent, P. R. (2011). The effects of the eddy-induced advection coefficient in a coarse-resolution coupled climate model. *Ocean Modell.*, 39(1):135–145.
[16] Flierl, G. R. (1978). Models of vertical structure and the calibration of two-layer models. *Dyn. Atmos. Oceans*, 2(4):341–381.
[17] Franzke, C. L., O'Kane, T. J., Berner, J., Williams, P. D., and Lucarini, V. (2015). Stochastic climate theory and modeling. *Wiley Interdiscip. Rev. Clim. Change*, 6(1):63–78.
[18] Frederiksen, J. S. and Davies, A. G. (1997). Eddy viscosity and stochastic backscatter parameterizations on the sphere for atmospheric circulation models. *J. Atmos. Sci.*, 54(20):2475–2492.
[19] Frederiksen, J. S. and Kepert, S. M. (2006). Dynamical subgrid-scale parameterizations from direct numerical simulations. *J. Atmos. Sci.*, 63(11):3006–3019.
[20] Gent, P. R. and McWilliams, J. C. (1990). Isopycnal mixing in ocean circulation models. *J. Phys. Oceanogr.*, 20(1):150–155.
[21] Ghil, M., Feliks, Y., and Sushama, L. (2002). Baroclinic and barotropic aspects of the wind-driven ocean circulation. *Phys. D*, 167(1):1–35.
[22] Grabowski, W. W. (2001). Coupling cloud processes with the large-scale dynamics using the cloud-resolving convection parameterization (crcp). *J. Atmos. Sci.*, 58(9):978–997.
[23] Grabowski, W. W. and Smolarkiewicz, P. K. (1999). Crcp: A cloud resolving convection parameterization for modeling the tropical convecting atmosphere. *Phys. D*, 133(1):171–178.
[24] Grooms, I. and Majda, A. J. (2013). Efficient stochastic superparameterization for geophysical turbulence. *Proc. Natl. Acad. Sci. USA*, 110(12):4464–4469.
[25] Grooms, I. and Majda, A. J. (2014). Stochastic superparameterization in a one-dimensional model for wave turbulence. *Commun. Math. Sci*, 12(3):509–525.
[26] Hallberg, R. and Gnanadesikan, A. (2006). The role of eddies in determining the structure and response of the wind-driven southern hemisphere overturning: Results from the modeling eddies in the southern ocean (meso) project. *J. Phys. Oceanogr.*, 36(12):2232–2252.
[27] Haltiner, G. J. and Williams, R. T. (1980). *Numerical Prediction and Dynamic Meteorology*. Wiley, 2nd edition.
[28] Hua, B. and Haidvogel, D. (1986). Numerical simulations of the vertical structure of quasi-geostrophic turbulence. *J. Atmos. Sci.*, 43(23):2923–2936.
[29] Jansen, M. F. and Held, I. M. (2014). Parameterizing subgrid-scale eddy effects using energetically consistent backscatter. *Ocean Modell.*, 80:36–48.
[30] Kalnay, E. (2003). *Atmospheric modeling, data assimilation and predictability*. Cambridge university press.
[31] Khairoutdinov, M., Randall, D., and DeMott, C. (2005). Simulations of the atmospheric general circulation using a cloud-resolving model as a superparameterization of physical processes. *J. Atmos. Sci.*, 62(7):2136–2154.
[32] Khouider, B., Biello, J., Majda, A. J., et al. (2010). A stochastic multicloud model for tropical convection. *Commun. Math. Sci.*, 8(1):187–216.
[33] Kitsios, V., Frederiksen, J., and Zidikheri, M. (2013). Scaling laws for parameterisations of subgrid eddy–eddy interactions in simulations of oceanic circulations. *Ocean Modell.*, 68:88–105.
[34] Lilly, D. K. (1965). On the computational stability of numerical solutions of time-dependent non-linear geophysical fluid dynamics problems. *Mon. Weather Rev.*, 93(1):11–25.
[35] Majda, A. J. and Grote, M. J. (2009). Mathematical test models for superparametrization in anisotropic turbulence. *Proc. Natl. Acad. Sci. USA*, 106(14):5470–5474.
[36] Majda, A. J., Timofeyev, I., and Vanden-Eijnden, E. (2003). Systematic strategies for stochastic mode reduction in climate. *J. Atmos. Sci.*, 60(14):1705–1722.
[37] McWilliams, J. C. (2008). The nature and consequences of oceanic eddies. *Ocean Modeling in an Eddying Regime*, pages 5–15.
[38] Olbers, D., Willebrand, J., and Eden, C. (2012). *Ocean dynamics*. Springer Science & Business Media.

[39] Pelupessy, I., Van Werkhoven, B., Van Elteren, A., Viebahn, J., Candy, A., Zwart, S. P., and Dijkstra, H. (2017). The oceanographic multipurpose software environment (omuse v1. 0). *Geosci. Model Dev.*, 10(8):3167.
[40] Plant, R. and Craig, G. C. (2008). A stochastic parameterization for deep convection based on equilibrium statistics. *J. Atmos. Sci.*, 65(1):87–105.
[41] Porta Mana, P. and Zanna, L. (2014). Toward a stochastic parameterization of ocean mesoscale eddies. *Ocean Modell.*, 79:1–20.
[42] Robert, A. J. (1966). The integration of a low order spectral form of the primitive meteorological equations. *J. Meteorolog. Soc. Jpn.*, 44:237–245.
[43] Shevchenko, I. and Berloff, P. (2017). On the roles of baroclinic modes in eddy-resolving midlatitude ocean dynamics. *Ocean Modell.*, 111:55–65.
[44] Storch, J.-S. v., Eden, C., Fast, I., Haak, H., Hernández-Deckers, D., Maier-Reimer, E., Marotzke, J., and Stammer, D. (2012). An estimate of the lorenz energy cycle for the world ocean based on the storm/ncep simulation. *J. Phys. Oceanogr.*, 42(12):2185–2205.
[45] Thomas, J. W. (2013). *Numerical partial differential equations: finite difference methods*, volume 22. Springer Science & Business Media.
[46] Vallis, G. K. (2006). *Atmospheric and oceanic fluid dynamics: fundamentals and large-scale circulation*. Cambridge University Press.
[47] Verheul, N. and Crommelin, D. (2016). Data-driven stochastic representations of unresolved features in multiscale models. *Commun. Math. Sci.*, 14(5):1213–1236.
[48] Viebahn, J. (2017). Toward a stochastic closure of the baroclinic energy pathway. (in preparation).
[49] Viebahn, J. and Eden, C. (2010). Towards the impact of eddies on the response of the southern ocean to climate change. *Ocean Modell.*, 34(3):150–165.
[50] Williams, P. D. (2009). A proposed modification to the robert-asselin time filter. *Mon. Weather Rev.*, 137(8):2538–2546.
[51] Young, J. A. (1968). Comparative properties of some time differencing schemes for linear and nonlinear oscillations. *Mon. Weather Rev.*, 96(6):357–364.
[52] Zanna, L., Mana, P. P., Anstey, J., David, T., and Bolton, T. (2017). Scale-aware deterministic and stochastic parametrizations of eddy-mean flow interaction. *Ocean Modell.*, 111:66–80.

A Appendix

A.1 Linear fit details

The conditioning procedure as defined in Section 3.1 is defined pointwise over the grid, therefore the number of sample points for each approximated CPDF in (2) is limited by the length of the time series ($T_c/\delta t$, see Table 1). This proved somewhat problematic in our experiments, because a straightforward equidistant binning (see Appendix A.2) did not always result in sufficiently filled bins in the case where (2) is conditioned on both the covariate $J(\widetilde{\psi}_0^n, \widetilde{q}_0^n)$ and the time lagged $\widetilde{R}^{n-l_\theta(i,j)}$. Therefore, in the case where we condition on both, we project the binning on the linear trend between $J(\widetilde{\psi}_0^n, \widetilde{q}_0^n)$ and $\widetilde{R}^{n-l_\theta(i,j)}$. This linear fit allows us to make more efficient use of the number of bins, i.e. to more evenly partition the sample data over the number of bins.

Let us define:

$$\bar{J}(\psi_0(i,j), q_0(i,j)) = k(i,j) + m(i,j)R(i,j), \tag{A.1}$$

to be the linear trend between $J(\psi_0(i,j), q_0(i,j))$ and $R(i,j)$ (see (1)), where $[k(i,j), m(i,j)] = (\boldsymbol{X}^T\boldsymbol{X})^{-1}\boldsymbol{X}^T\boldsymbol{y}$, and $\boldsymbol{X}(i,j) = [1, (\boldsymbol{R})^T]$ a $(T/\delta t \times 2)$-tensor, and $\boldsymbol{y}(i,j) = (\boldsymbol{J})^T$ a $(T/\delta t \times 1)$-columnvector. Then, by subtracting this linear trend from the Jacobian values we get:

$$\widehat{J}(\psi_0(i,j), q_0(i,j)) = J(\psi_0(i,j), q_0(i,j)) - \bar{J}(\psi_0(i,j), q_0(i,j)),$$

and using the conditioning set $\{\widetilde{R}^{n-l_\theta(i,j)}(i,j), \widehat{J}(\widetilde{\psi}_0^n(i,j), \widetilde{q}_0^n(i,j))\}$ results in the linearly fitted sampling procedure:

$$\widetilde{R}^{n+1}(i,j) \sim R^{n+1}(i,j) \mid \widetilde{R}^{n-l_\theta(i,j)}(i,j), \widehat{J}(\widetilde{\psi}_0^n(i,j), \widetilde{q}_0^n(i,j)) \text{ for } i \in \{1,\dots,N_x\}, j \in \{1,\dots,N_y\} \tag{A.2}$$

A.2 Empirical distribution details

Here we elaborate on the general description of the empirical distribution that is found in Section 3.4. The equidistant binning method independently partitions the range between the minimum and maximum of each of the covariates into N_B intervals α_b. Each of the intervals α_b describes a set ρ_b of R^{n+1}-values. Let $\mathcal{U}(A)$ denote a uniform distribution over the elements of the set A, i.e. if $a' \sim \mathcal{U}(A)$ then $\forall a \in A : P(a' = a) = |A|^{-1}$. The conditional sampling method (2) is then approximated by the empirical approach:

$$\widetilde{R}^{n+1}(i,j) \sim \mathcal{U}(\rho_b(i,j)), \text{ where } b : (\widetilde{R}^n(i,j), \widehat{J}(\widetilde{\psi}_0^n(i,j), \widetilde{q}_0^n(i,j)) \in \alpha_b(i,j). \tag{A.3}$$

This empirical approach is a prototypical example of a discrete sampling method. An obvious limitation of such methods is that they can exclusively sample from observed sample values. Therefore, the stochastic model has no predetermined way of handling situations in which the values of the conditioning variables are outside of the ranges exhibited in the sample data. With respect to the equidistant binning procedure, see Section 3.4, this situation manifests as empty bins, i.e. a bin $\alpha_b(i,j)$ for which $\rho_b(i,j)$ is empty. The likelihood of the conditioning variable's values to arrive at an empty bin is predominantly determined by three factors: the number of conditioning variables D, the number of bins per conditioning variable N_B, and the limitations on the amount of sample data. With each of the aforementioned factors, this is a non-trivial issue. We solve this problem by linking each of the empty bins to a nonempty bin in the training phase of the simulation. Because the covariate $\widehat{J}(\psi_0, q_0)$ is least likely to hold predictive qualities for the sampling procedure (as explained in Section 3.4), the predictive quality of this covariate is considered negligible in empty bins. And thus, the conditioning in empty bins is considered 1-dimensional, i.e. in empty bins we condition only on $R^{n-l(i,j)}$, which is, by construction, guaranteed to be in range of sample data.

A.3 Spatial smoothing

With the pointwise conditioning procedure (see Section 3.1) the spatial correlations of R are only modeled implicitly by conditioning on the Jacobian term $J(\psi_0, q_0)$. As discussed in Section 1, finding more rigorous ways to explicitly reproduce spatial correlations will stay a topic for an immediate future project. We use a more heuristic method to introduce spatial smoothness and remedy the spatially uncorrelated fields here. Namely, we use a truncated Gaussian smoothing filter $G(x,y)$:

$$G'(x,y) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2+(y-\mu)^2}{2\sigma^2}}, \qquad G(x,y) = \frac{G'(x,y)}{\sum_{x,y} G'(x,y)}, \tag{A.4}$$

where we choose a basic 3×3 filter $G(x,y)$ with $\mu = 0$ and $\sigma = 0.5$, where $x, y \in \{-1, 0, 1\}$. Illustrating snapshots are shown in Figure 21, where one can see the spatially smoothed snapshot more closely resembling the resolved R, both in structure and spatial smoothness.

As previously mentioned, this smoothing is an added heuristic method to strengthen the spatial correlations between neighboring grid points. But more importantly, because the stochastic procedure, as defined in (2), works completely pointwise over the grid, this Gaussian filter will smooth out unwanted spatial roughness. Otherwise, the spatial roughness would artificially 'add energy' into the system, spuriously exciting the system beyond realistic goals.

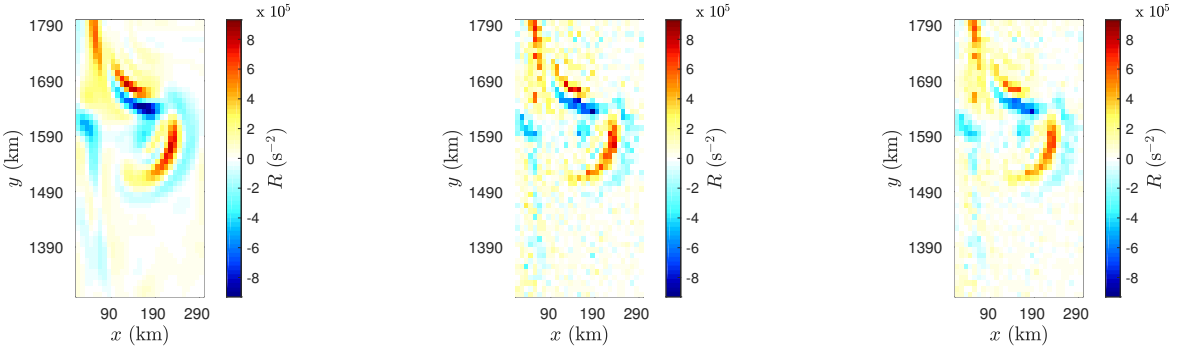

Figure 21: A comparison between snapshots of the resolved R (left), the unfiltered emulated stochastic forcing (middle), and the filtered emulated stochastic forcing (right).

A.4 Definitions of diagnostic criteria

- The first four statistical sample moments: mean, standard deviation, skewness, and kurtosis over the model variables ψ_0 (pointwise in physical space):

$$\mu(\psi_0) = \mathrm{E}(\psi_0^i) \tag{A.5}$$

$$\mathrm{std}(\psi_0) = \left(\mathrm{E}((\psi_0^i)^2) - (\mathrm{E}(\psi_0^i))^2\right)^{1/2} \tag{A.6}$$

$$\gamma(\psi_0) = \mathrm{E}\left[(\psi_0^i - \mathrm{E}(\psi_0^i))^3\right](\mathrm{Var}(\psi_0^i))^{-3/2} \tag{A.7}$$

$$\mathrm{Kurt}(\psi_0) = \mathrm{E}\left[(\psi_0^i - \mathrm{E}(\psi_0^i))^4\right](\mathrm{Var}(\psi_0^i))^{-2} \tag{A.8}$$

- The autocorrelation function (ACF) with time lag l over the model variable ψ_0:

$$\mathrm{ACF}_l(\psi_0) = \mathrm{E}\left[(\psi_0^i - \mathrm{E}(\psi_0^i))(\psi_0^{i+l} - \mathrm{E}(\psi_0^i))\right](\mathrm{Var}(\psi_0^i))^{-1} \tag{A.9}$$

- The mean, standard deviation, and the variability (both short-term and long-term variability) of the time-dependent scalar quantities enstrophy $\mathcal{E}$, kinetic energy KE, and energy exchange/transfer EE as horizontally integrated over grid A:

$$\mathcal{E} = \frac{\overline{\rho}H}{2}\int_A q_0^2, \tag{A.10}$$

$$\mathrm{KE} = -\frac{\overline{\rho}H}{2}\int_A (\nabla\psi_0)^2, \tag{A.11}$$

$$\mathrm{EE} = \overline{\rho}H\int_A \psi_0 J(\psi_1, q_1). \tag{A.12}$$

- The covariance $\mathrm{Cov}(\psi_0)$ between grid points (i, j) and (i', j') for model variable ψ_0:

$$\mathrm{Cov}(\psi_0(i,j), \psi_0(i',j')) = \mathrm{E}([\psi_0(i,j) - \mathrm{E}(\psi_0(i,j))][\psi_0(i',j') - \mathrm{E}(\psi_0(i',j'))]) \tag{A.13}$$

- Similar to the covariance above, the correlation $\mathrm{Corr}(\psi_0)$ between grid points (i, j) and (i', j') for model variable ψ_0:

$$\mathrm{Corr}(\psi_0(i,j), \psi_0(i',j')) = \frac{\mathrm{Cov}(\psi_0(i,j), \psi_0(i',j'))}{\mathrm{std}(\psi_0(i,j))\mathrm{std}(\psi_0(i',j'))} \tag{A.14}$$

A.5 Numerical integration details

The numerical implementation uses the standard centered difference scheme $(\Delta x)^{-n}\,\delta_x^n f(x,y,t)$ of order $O((\Delta x)^2)$ to approximate the n-th order spatial derivatives [45] in 2M, where:

$$\delta_x^n f(x,y,t) = \begin{cases} \sum_{i=0}^{n} (-1)^i \binom{n}{i} f\left(x + \left(\frac{n}{2} - i\right)\Delta x, y, t\right) & \text{if } n \text{ even}, \\ \sum_{i=0}^{n} (-1)^i \binom{n}{i} \frac{1}{2}\left[f\left(x + \left(\frac{n+1}{2} - i\right)\Delta x, y, t\right) + f\left(x + \left(\frac{n-1}{2} - i\right)\Delta x, y, t\right)\right] & \text{if } n \text{ odd}, \end{cases} \tag{A.15}$$

and analogously for partial derivatives to y.

We choose the Arakawa stencil $J_1 = \frac{1}{3}(J^{++} + J^{+\times} + J^{\times+})$ to discretize the Jacobian terms for its conservation properties, this Arakawa scheme satisfies finite difference analogs of energy and mean squared vorticity conservation laws [1].

To integrate D1M, 2M, and S1M over time, let us consider the second-order centered-difference leapfrog scheme [13, 27, 30], which is used widely in weather and climate models. The leapfrog time stepping method admits a well-documented spurious computational mode that manifests as spurious oscillations between even and odd time steps that amplify during nonlinear simulations [13, 34, 51]. This phenomenon is referred to as time-splitting. The Robert–Asselin filter [2, 42] is an oft-used time-filtering solution (for a list of examples in climate models see, e.g, Williams [50]) that dampens this computational mode and thus controls the time-splitting instability. The combination of the leapfrog scheme with the Robert–Asselin filter results in the Asselin-leapfrog scheme:

$$\psi^{n+1} = \overline{\psi}^{n-1} + 2\Delta t\, F(\psi^n), \tag{A.16}$$

where $F(\psi)$ denotes the right-hand side terms in the PDE governing the temporal evolution of ψ, and the bar notation in (A.16) denotes the Robert–Asselin time-filter:

$$\overline{\psi}^n = \psi^n + \alpha(\overline{\psi}^{n-1} - 2\psi^n + \psi^{n+1}), \tag{A.17}$$

where the parameter α denotes the Robert–Asselin filter strength.

We further impose the free-slip condition $n_\perp \cdot (\nabla \circ \nabla)\psi_m = 0$ along the lateral boundaries, where $n_\perp$ denotes the horizontal unit vector normal to the boundary, and $\circ$ is the Hadamard product, defined for matrices A and B with same dimensions by $(A \circ B)_{ij} = (A)_{ij}(B)_{ij}$.

All our choices for model parameters are shown in Table 1. The discussed discretization methods result in the following discretized integration scheme for the 2-mode referemce model 2M:

$$\begin{aligned} &\psi_0^{n+1} = \overline{\psi}_0^{n-1} + 2\Delta t\, S_0(\psi_0^n, \psi_1^n), \qquad \psi_1^{n+1} = \overline{\psi}_1^{n-1} + 2\Delta t\, S_1(\psi_0^n, \psi_1^n), \\ &\text{where } \nabla^2 S_0(\psi_0^n, \psi_1^n) = -J(\psi_0^n, q_0^n) - R^n - \beta\partial_x\psi_0^n + A_H\nabla^4\psi_0^{n-1} + \frac{\partial_x\tau^y - \partial_y\tau^x}{\overline{\rho}H} \\ &\text{where } \nabla^2 S_1(\psi_0^n, \psi_1^n) = -J(\psi_1^n, q_0^n) - J(\psi_0^n, q_1^n) - \epsilon_{111}J(\psi_1^n, q_1^n) - \beta\partial_x\psi_1^n + \\ &\qquad A_H\nabla^4\psi_1^{n-1} + \frac{\phi_1^n(z=0)(\partial_x\tau^y - \partial_y\tau^x)}{\overline{\rho}H}, \end{aligned} \tag{A.18}$$

where numerical solutions to the Poisson's equations for $\nabla^2 S_m$ are found with the Intel MKL Poisson solver. The viscosity terms above are taken from the previous time-step for numerical stability purposes.

The discretized 2M (A.18) is initialized with the conditions $\overline{\psi}_0^{n=-1} = \overline{\psi}_1^{n=-1} = \psi_0^{n=0} = \psi_1^{n=0} = 0$ and $q_0^0 = q_1^0 = 0$. To equilibrate any baroclinic simulation starting from rest, a sufficient spin-up period needs to be performed. For all our experiments, we choose a 30 year spin-up time. The integration time-step Δt is chosen to be half an hour for both the barotropic and baroclinic models, these choices are made to find comparison with similar set ups (e.g. [3, 5, 41]).

Applying the pointwise sampling method (2) to evolve $\widetilde{R}$ over time, and numerically integrating S1M results in the following numerical integration scheme for the discretized stochastic 1-mode model:

$$\begin{aligned} &\widetilde{\psi}_0^{n+1} = \overline{\widetilde{\psi}}_0^{n-1} + 2\Delta t\, S(\widetilde{\psi}_0^n), \qquad \widetilde{R}^{n+1} \sim R^{n+1} \mid (R^{n-l_\theta}, C(\psi_0))) = (\widetilde{R}^{n-l_\theta}, C(\widetilde{\psi}_0)) \\ &\text{with } \nabla^2 S(\widetilde{\psi}_0^n) = -J(\widetilde{\psi}_0^n, \widetilde{q}_0^n) - \widetilde{R}^n - \beta\partial_x\widetilde{\psi}_0^n + A_H\nabla^4\widetilde{\psi}_0^{n-1} + \frac{\partial_x\tau^y - \partial_y\tau^x}{\bar{\rho}H}, \end{aligned} \tag{A.19}$$

where the initial conditions are chosen to start from an equilibrated baroclinic state before starting the 30 years spin-up for the stochastic simulation: $\overline{\widetilde{\psi}}_0^{-1} = \overline{\psi}_0^{M-1}$, $\overline{\widetilde{\psi}}_1^{-1} = \overline{\psi}_1^{M-1}$, $\widetilde{\psi}_0^0 = \psi_0^M$, $\widetilde{\psi}_1^0 = \psi_1^M$ and $\widetilde{q}_0^0 = q_0^M$, $\widetilde{q}_1^0 = q_1^M$.