

Tracking down the origin of NWP model uncertainty : coarse-graining studies

Glenn Shutts* and Alfons Callado Pallares \diamond

* ECMWF, Shinfield Park, Reading, RG2 9AX, United Kingdom

\diamond Spanish Meteorological Agency (AEMET), Barcelona, Spain

ABSTRACT

Current implementations of the perturbed parametrization tendency method for representing uncertainty rely on *ad hoc* assumptions about its magnitude and its spatial and temporal correlation scales. Ideally one would use observational data to ascertain the statistical character of parametrization tendency errors and use the resulting probability distribution functions to devise and calibrate the perturbed tendency approach. The reality is that observations rarely have the coverage, representativity and accuracy to form a useful comparison with model data. A less satisfactory alternative is to use high resolution modelling to provide a ‘truth’ simulation and then compare this with an equivalent but lower resolution simulation. Tendency fields from both simulations are coarse-grained to a resolution compatible with the assumed horizontal correlation scale in the perturbed tendency method and the bias-corrected differences between them are used to quantify statistical uncertainty. Early results using the ECMWF IFS forecasts appear to show that the variance of the coarse-grained tendency differences is proportional to the tendency in the lower-resolution forecast. However the current perturbed parametrization tendency scheme at ECMWF assumes that the *standard deviation* of the perturbations is proportional to the tendency itself. Probability distribution functions of the high-resolution model tendency, sub-sampled by narrow ranges of the low-resolution model tendency, seem to be consistent with an underlying Poisson process.

Comparing model vorticity tendencies in matching high and low resolution forecasts has shown that the difference (i.e. ‘error’) implies direct kinetic energy input into fairly low wavenumbers (i.e. between wavenumber 10 and 30). Moreover, the globally-averaged coarse-grained kinetic energy input at any instant is always positive whereas that computed from the stochastic backscatter scheme is highly noisy and fluctuates about zero. This noisiness is presumably due to the random phases of waves comprising the pattern generator. Tests with the anti-diffusive algorithm called *vorticity confinement* show its potential as a backscatter scheme.

1 Introduction

A variety of algorithms exist currently for representing random model error in ensemble forecasting systems. First and foremost, they were designed to correct the lack of spread in the ensemble rather than to act as an improved physical description of the atmosphere. Houtekamer et al (1996) introduced uncertainty into the design of their ensemble prediction system by using different combinations of physical parametrization schemes; representations of the orography, and randomly-perturbed surface fields (e.g. sea surface temperature, roughness length and albedo). Buizza et al (1999) considered uncertainty in physical parametrization as resulting from its underlying statistical nature and experimented with tendency perturbations generated by multiplying the total parametrization tendency by a number drawn randomly from a uniform probability distribution function ranging from 0.5 to 1.5. For best impact on skill, the random numbers were assigned to 10° latitude/longitude boxes and refreshed every 4 time steps. Such an approach could be justified on the grounds that parametrization schemes are designed to provide ensemble-average tendencies and statistical fluctuation is often sufficiently large that the instantaneous tendency is substantially different from the ensemble mean - even when averaged over a 10° latitude-longitude box (Palmer, 2001). A revised version of the perturbed tendency scheme was introduced into the operational EPS in 2009 (the Stochastic Perturbed Parametrization Tendency or

SPPT scheme) which uses a spectral pattern generator based on a first-order auto-regressive model to provide the tendency multiplier. The pattern generator is composed of three individual patterns, each with its own spatial and temporal decorrelation scale (see Leutbecher, 2011). A similar spectrally-based, perturbed tendency scheme has been in use at Environment Canada since 2007 (Charron et al, 2010).

Another approach used at the Met Office, perturbs the values of a subset of parametrization parameters that are poorly constrained by theory and/or measurement (e.g. the convective entrainment rate or ice crystal fall speed; see Bowler et al, 2008). Each member, running with its own fixed set of perturbed parameters, acts as a different deterministic model (the approach taken by the Hadley centre's 'Quantifying Uncertainty in Model Predictions' (QUMP) project; see Rougier et al, 2009). In the context of medium-range ensemble prediction, the Met Office Global and Regional Ensemble Prediction System (MOGREPS) uses this 'Random Parameters' technique but in contrast to QUMP, allows time variation of the parametrization parameters using an auto-regressive process with six hour decorrelation time. The justification for allowing time variation without space variation is not clear and it implies that the convective entrainment rate parameter varies synchronously at all points across the globe. It seems to mix the ideas of 'knowledge uncertainty' with 'statistical uncertainty'.

The Stochastic Kinetic Energy Backscatter (SKEB) algorithm (Shutts, 2005) is used - albeit with different implementations - at ECMWF, Met Office and Environment Canada (Berner et al, 2009; Bowler et al, 2009, and Charron et al, 2010 respectively). In contrast to perturbed tendencies or Random Parameters, backscatter aims to represent a missing physical process - upscale energy transport - and counteracts spurious dissipative processes in the model such as numerical diffusion. It is assumed that this upscale energy transport is manifest as random vorticity perturbations injected into the flow, with a magnitude dictated by a local background kinetic energy dissipation rate. Importantly, the phase of these vorticity perturbations is assumed to be independent of the model flow and so the local kinetic energy input can be positive or negative and it is only in an ensemble-average or long-term mean sense that the energy input is positive.

In order to put some constraints on the adjustable parameters in the above stochastic physics schemes and validate some of their assumption, coarse-graining high-resolution 'truth' simulations provides a useful source of information. Such simulations could be of the cloud-resolving model type for which convection parametrization is not necessary, or they could be high-resolution numerical weather prediction model forecasts. Shutts and Palmer (2007) used the Met Office cloud-resolving model to generate data for coarse-graining studies and computed probability distributions of the coarse-grained diabatic heating when sub-sampled according to different ranges of parametrized convective heating (using a convective parametrization on the coarse-grained fields). Their results provided tentative support for the Buizza et al (1999) perturbed tendency scheme although the fluctuations in diabatic heating did not vanish at zero parametrized heating as they should do in such multiplicative noise schemes.

One advantage of using NWP model forecasts rather than limited-area cloud-resolving model simulations is that of realism (e.g. global domains with initial state provided by data assimilation). Global NWP model forecasts with horizontal gridlengths of the order of 1 km would ideally be compared with forecasts made at current operational resolution i.e. gridlengths ≥ 20 km so that the truth simulation contains no convection or gravity wave drag parametrization. This would hopefully provide the statistical information required to better calibrate the perturbed tendency schemes.

2 Coarse-graining with the ECMWF IFS

The coarse-graining technique used here consists of making two 12-hour forecasts from the same start time; one at T1279 horizontal resolution and the other at T159 resolution and both using 91 vertical levels. Tendency contributions coming from parametrized radiation, convection, large-scale condensa-

tion, turbulent mixing, non-orographic gravity wave drag and their total are output every hour on model levels and then spatially-filtered in spectral space using the quasi-Gaussian filter described in Weaver and Courtier (2001) with filter scales of either 250 or 500 km (dynamical tendencies are shown for reference too). The resulting coarse-grained tendency fields on a particular model level are subtracted to give a ‘model error’ field. The pattern field making the dominant contribution to SPPT has a spatial decorrelation scale of 500 km and a decorrelation time of 6 hours and so, to filter out short time scale tendency fluctuations, the tendencies are filtered in time using a triangular filter centred on T+6 hours. The mathematical representation of the coarse-graining process is described by:

$$\hat{f}^k(\lambda, \phi) = \sum_{k=1}^{11} \sum_{m=-N}^{m=N} \sum_{n=|m|}^N f_{m,n}^k W_{m,n}^k Y_n^m(\lambda, \phi) \quad (1)$$

with

$$W_{m,n}^k = \exp \left[-\frac{1}{2} \left(\frac{R_f}{a} \right)^2 n(n+1) \right] \frac{1}{6} \left(1 - \frac{|k-6|}{6} \right) \quad (2)$$

where \hat{f}^k is the smoothed representation of the field $f^k(\lambda, \phi)$; Y_n^m is the spherical harmonic associated with zonal wavenumber m and order n ; $f_{m,n}^k$ is the coefficient in the spherical harmonic expansion of $f^k(\lambda, \phi)$; R_f is the smoothing scale; a is the Earth’s radius; λ and ϕ are the longitude and latitude respectively; k is the number of hours into the forecast. A potential problem with subtracting tendency fields from the two different resolution forecasts is that the model levels do not coincide exactly in physical height. This is mainly due to the different orographic height specifications and the dependence of model level pressure on surface pressure (for levels located in the troposphere at least). Various checks, such as masking out the land points, suggest that this model height discrepancy is not important here.

Results to be shown are based on 24 pairs of forecasts initialized at 12 Z for the 20th day of each month in 2009 and 2010. The aim is to compute the standard deviation of the T159 tendency ‘error’ (i.e. the difference in the coarse-grained T159 and T1279 tendencies) from these 24 forecasts for different ranges of T159 tendency. In this way, the functional relationship between standard deviation and the mean tendency can be computed and compared with the proportionality assumption upon which SPPT is based.

Figure 1 shows standard deviation of the T159 400 hPa temperature tendency error versus the mean temperature tendency for individual contributions from the physical parametrizations (and their total) together with the dynamical (advection) term for reference. In relation to SPPT, the total physics tendency (black) curve does not support the assumption that the standard deviation of the error is proportional to the mean with its decreasing slope with increasing mean temperature tendency and local non-zero minimum between -1 and -2 K per day. This displacement of the minimum away from zero is caused by the radiative contribution to the total temperature tendency which has its lowest error at around -1 K per day. This reflects the fact that radiative temperature tendency differences between the coarse-grained T159 and T1279 forecasts are smallest where cloud is absent and those regions typical cool at around -1 to -2 K per day. The dotted, straight line in Figure 1 shows the relationship between standard deviation and mean assumed in SPPT. Although not of the right form compared to the total tendency curve, the overall magnitude of the tendency perturbations assumed in SPPT is supported by the coarse-graining calculation.

In contrast to the total temperature tendency, both the convective and explicit condensation temperature tendencies have a cusp minimum in standard deviation at zero mean tendency. However, the corresponding red and green lines are not straight as SPPT would require but imply a proportionality between the variance of the error and the mean tendency as can be seen from the (dashed and dash-dot) curve fits. It is interesting that the condensation tendency is considerably more uncertain than the convective tendency. The contributions of turbulent dissipation and non-orographic gravity wave drag are essentially




Figure 1: Standard deviation of coarse-grained, temperature tendency error (T1279-T159) at 400 hPa coming from the parametrization of convection, radiation, vertical turbulent mixing, non-orographic gravity wave drag (see the key) plotted with respect to the bin-averaged T159 temperature tendency. Also included are the explicit condensation temperature tendency and the total model physics tendency (denoted as ‘physics’). ‘Dynamics’ refers to the dynamical (or advective) temperature tendency. The dashed lines with a $\sqrt{(|mean|)}$ factor in the key reference are curve fits for which the variance is proportional to the mean tendency. The dotted line shows the linear relationship between standard deviation and mean tendency assumed in the dominant pattern of SPPT. The variable-sized dots indicate the centre of the bins into which the T159 tendency is sub-sampled and the size of the dots is a logarithmic indicator of the number of points in the corresponding bin. The model tendencies come from 24 forecasts made on the 20th day of each month in 2009 and 2010.




Figure 2: as for Figure 1 except for q - the specific humidity




Figure 3: as for Figure 1 except for u - the east-west wind component




Figure 4: as for Figure 1 except for surface precipitation rate alone (red curve) and uses a set of 32 forecasts between 15 October and November 15 2009. The green curve shows the curve representing a standard deviation proportional to 0.2 of the mean precipitation rate. Units : mm/day

negligible here. The standard deviation in the dynamical temperature tendency is much less variable over this range mean tendencies and lowest between 1 and 2 K per day. One might speculate that this displacement of the minimum to positive values of tendency is the balancing counterpart of the radiative error minimum and could be associated with subsidence warming.

Figure 2 shows the standard deviation/mean tendency relations for specific humidity and reveals very similar functional dependencies as for temperature tendency. For instance the red curve, which represents the convective drying, shows that the standard deviation of the tendency error is roughly proportional to the square root of the tendency magnitude. The dotted straight line is (again), a poor fit to the black curve which shows the uncertainty in the total physics tendency. The most uncertain of the specific humidity tendency contributions comes from explicit condensation - as for temperature tendency. Also of interest is the qualitative similarity of dynamical tendency curve to that of the other processes with a sharp minimum at zero mean tendency.

This approximate proportionality of the error variance to the mean can even be seen in the eastward momentum component at 400 hPa (Figure 3) although perhaps not for the vertical diffusion tendency. The dynamical momentum tendency has much larger values generally and the pink curve only represents a limited portion of the entire range.

Finally, it is interesting to look at the model precipitation error at T159 with truth being defined by coarse-grained T1279 precipitation rates. The calculation is essentially the same with model tendency being replaced by surface precipitation rate. From a hydrological viewpoint, this is similar to determining the standard deviation of the T159 forecast error in its representation of the rainfall averaged over a particular catchment area. Figure 4 shows this error as a function of the mean precipitation rate together with the line that represents a 20% error. As with the other physical parametrization tendencies,

the slope of this curve decreases with increasing precipitation rate. Since surface precipitation rate is a measure of vertically-integrated diabatic heating due to condensation, the steeper slope of the 400 hPa temperature tendency due to condensation in Figure 1 suggests some cancellation of error in the vertical.

3 Backscatter by Vorticity Confinement

Using a streamfunction forcing function that depends on model dissipation rate, stochastic backscatter schemes aim to represent the random component of upscale energy flux that is suppressed or absent from the near or sub-gridscale flow. The streamfunction forcing can be assumed to represent chaotic fluctuations in unresolved Reynolds stresses. The need for a systematic, flow-dependent upscale energy flux has arisen in many contexts over the years such as the anticipated potential vorticity method of Sadourny and Basdevant (1985). Another simple algorithm which has many applications in computational fluid dynamics is Vorticity Confinement (VC). Its main use has been as a device for preventing the dissipation of the energy in vortices close to the gridscale in the simulation of three-dimensional flows. ‘Confinement’ here refers to its tendency to counteract the diffusive spreading of vorticity with an opposing pseudo-advective transport term in the vorticity equation. A brief review of VC is now presented and its role as a form of deterministic backscatter supported by energy calculations using T159 IFS forecasts.

Steinhoff et al (2005) define two different forms of vorticity confinement and we shall be concerned with the type 1 scheme. Steinhoff and Underhill (1994) define VC as the second term on the right-hand side of the following momentum equation:

$$\frac{D\mathbf{V}}{Dt} + \nabla\phi = \mu\nabla^2\mathbf{V} + \varepsilon\hat{\mathbf{n}}_* \times \zeta \quad (3)$$

where

$$\hat{\mathbf{n}}_* = \frac{\nabla\eta}{|\nabla\eta|}$$

and $\zeta = \nabla \times \mathbf{V}$ with $\eta = -|\zeta|$; $\nabla\phi$ is the pressure gradient and $\mu\nabla^2\mathbf{V}$ represents the effect of diffusion. For the purposes of this discussion, the diffusion term will represent the effect of numerical approximation. For NWP model implementation, \mathbf{V} will be the horizontal velocity vector ($= \mathbf{V}_H$), ζ will be the vorticity measured on a model coordinate surface and $\hat{\mathbf{n}}_*$ will be the unit vector pointing in the direction of $\nabla_H|\zeta|$ where the subscript H on the nabla denotes horizontal differentiation on the model surface. The equation of motion (including now the Coriolis force) becomes:

$$\frac{D\mathbf{V}_H}{Dt} + f\mathbf{k} \times \mathbf{V}_H + \nabla\phi = \mu\nabla^2\mathbf{V}_H + \varepsilon\hat{\mathbf{n}} \times |\zeta|\hat{\mathbf{k}} \quad (4)$$

with

$$\hat{\mathbf{n}} = \frac{\nabla_H\zeta}{|\nabla_H\zeta|}$$

noting that the vorticity confinement term has been redefined to avoid differentiation of a modulus. This effectively creates a horizontal force that lies along the contours of vorticity (on a model level) and has a magnitude proportional to the modulus of the vorticity. The force is rotated clockwise by a right-angle from the direction of the vorticity gradient so that on average it will circulate with the same sense as the tangential wind in a vortex. Forming an equation for ζ by taking $\mathbf{k} \cdot \text{Curl}$ of eq.(4) can be shown to give:

$$\frac{\partial\zeta}{\partial t} + \nabla \cdot \{(\zeta + f)\mathbf{V} + \mathbf{s}|\zeta|\} = \mu\nabla^2\zeta \quad (5)$$




Figure 5: Spectral power distribution of the globally-averaged, 250 hPa kinetic energy input derived from coarse-graining (green line); from the vorticity confinement force in T159 forecasts (red line; $\epsilon = 0.8$), and from SKEB (blue line; backscatter ratio = 0.095). In the cases of VC and SKEB, a streamfunction forcing field was obtained from their associated forcing functions and the spectral coefficients of their spherical harmonic expansions are multiplied by the corresponding coefficients in the spectral representation of the model's vorticity field. The resulting products are summed over zonal wavenumber for each value of n . In the case of the coarse-graining results, the streamfunction forcing used to compute energy input was obtained from the difference in streamfunction tendencies in T1279 and T159 forecasts. Each of the three types of energy input were calculated from forecasts started at 12Z each day in the period October 15 until November 14 2009. All forecasts were made with CY36R4 of the ECMWF IFS. The magenta line shows the effect on the SKEB energy input calculation of extending the number of forecasts included until March 1st 2010 and taking forecast fields every 3 hours between T+6 hrs and T+24 hrs.

where the vector $\mathbf{s} = \epsilon \mathbf{n}$ acts as a velocity field. Alternatively one can write eq.(5) as

$$\frac{\partial \zeta}{\partial t} + \nabla \cdot \{ \zeta + f \} V = \nabla \cdot [\mu \nabla \zeta + \mathbf{s} |\zeta|] = \nabla \cdot \left[\nabla \zeta \left(\mu - \frac{\epsilon |\zeta|}{|\nabla \zeta|} \right) \right] \quad (6)$$

from which it can be seen that $\epsilon |\zeta| / |\nabla \zeta|$ acts like a diffusivity term that opposes the explicit diffusion term.

The VC momentum forcing term of eq.(4) has been coded in the ECMWF IFS and tested in deterministic forecasts and in EPS forecasts. Its action therefore, is to provide a gentle force that tends to be in the same direction as the flow (exactly true for circular vortices) - thereby imparting energy to the flow and helping to offset numerical energy loss.

In order to provide some justification for the use of vorticity confinement and SKEB, an estimate of the spectral distribution of backscatter energy input is computed from T+3 hour u and v tendencies at 250 hPa (dynamics+physics) in a set of T159 and T1279 forecasts starting from the same analyses.

These gridpoint tendencies are transformed to spectral tendencies of vorticity and divergence and the difference between the T1279 and T159 forecast vorticity tendencies is regarded as the ‘error’ that one would wish to parametrize using algorithms like SKEB and VC. Multiplying the spectral coefficients of this tendency error by the spectral coefficients of the vorticity field in the T159 forecast and dividing by the total wavenumber squared $n(n+1)/a^2$ gives the energy input to that spectral mode. Summing these modal energy tendencies over zonal wavenumber, for each value of n , gives the spectral distribution of kinetic energy input that is missing from the T159 forecast i.e. the backscatter energy input. Likewise, the energy input spectrum due to SKEB and vorticity confinement can be computed and compared with that implied by the vorticity tendency differences (see Figure 5). On a cautionary note, it is important to recognize that these T1279-T159 vorticity tendencies differences may not necessarily be due to negative viscosity effects associated with quasi-2D upscale energy cascades but could, for instance, result from the effects of horizontal divergence on the vorticity field.

The first thing to note is that the input of kinetic energy by SKEB is very uniform across wavenumber except at low wavenumbers where 31 forecast fields is an insufficient sample size for the energy input to be representative of the ensemble mean. The magenta line (labelled ‘SKEB extended’) represents the SPBS energy input from 966 fields as opposed to the 31 used to compute the blue line and shows positive energy input at low wavenumbers. In stark contrast, the energy input implied by the T1279-T159 vorticity tendency difference is very different and subject to much less noise. The dominant energy input is in the wavenumber band 10 to 40 with a maximum at the lower wavenumber end near $n = 12$ and markedly less for $n < 10$. Fortunately, the vorticity confinement scheme (with $\epsilon = 0.8$) has, for much of the wavenumber range, a remarkably similar energy input spectrum to that computed from T1279-T159 differences. Only for low wavenumbers ($n < 5$) is there a big difference, and here VC acts as an energy sink. The tendency of VC to sharpen vorticity gradients appears to be responsible for this downscale energy cascade at very low wavenumbers. It should be noted that the energy input deduced from T1279-T159 vorticity tendency differences is sensitive to the forecast range and T+3 hrs was chosen to minimize spin-up effects but avoid too much forecast divergence. In spite of this, the shape of its spectral distribution does not vary much up to T+6 hrs. $\epsilon = 0.8$ was therefore chosen to give a best match between VC energy input and that implied by T1279-T159 forecast differences.

Since the kinetic energy input by VC so closely matches the difference seen between T1279 and T159 forecasts, one might expect that a deterministic forecast could be made more accurate by implementing VC. To examine this possibility, a set of 31 forecasts at T95 resolution have been made, with and without VC, and their global-mean r.m.s. error has been computed with respect to analyses for selected variables. It was found that the r.m.s. forecast error in 500 hPa height could be reduced by up to 2% in the first 3 to 4 days of the forecast with the inclusion of VC. Thereafter, the r.m.s. error of the forecasts with VC became worse than the no-VC forecasts, presumably because of the increased cyclone intensity. Choosing a forecast model resolution as low as T95 obviously helps to expose the anti-diffusive benefits of VC but one could question the need for such a term at current EPS resolution (T639 at ECMWF). Even so, seasonal and climate forecasting could clearly benefit from such an effect and some supporting evidence for this has been given at this workshop (the presentation by Antje Weisheimer which showed impacts on seasonal time-scale biases and blocking frequency, and the poster presentation of Claudio Sanchez which showed the positive impact of VC in Hadley Centre climate runs). The reader should be cautioned that these are early results and require more analysis.

A stochastic variant of vorticity confinement has been developed in which ϵ is made to vary in space and time using the SPPT pattern generator. Since ϵ represents an upgradient advective velocity, the effect of its spatial variation would be to cause distortions to the shape of vorticity blobs as well as changes to their intensity. Vorticity perturbations ‘manufactured’ this way - from the actual vorticity field - have the potential to be more realistic than perturbations that only know about the model flow indirectly through the SKEB dissipation rate calculation. Tests of the stochastic version of VC with the ECMWF EPS have shown its potential to increase spread and probabilistic skill at T159 resolution (for which the EPS is

substantially under-spread). It remains to be demonstrated that stochastic vorticity confinement could be a suitable replacement for SKEB. An alternative may be to use VC (without stochasticity) as the deterministic component of backscatter and allow SKEB to be the random component.

4 Summary

Forecasts made with the ECMWF IFS at T1279 resolution have been used as ‘truth’ relative to matching T159 forecasts. Differences in model tendencies between them have been used to infer some statistical information on the component of model error that derives from resolution. Of particular interest is the dependence of its size on the magnitude of the tendency itself since the SPPT scheme assumes that the standard deviation of the error is proportional to the mean tendency. In the above analysis, model tendencies are coarse-grained to a common resolution using a spectral smoothing process. Early results show that the *variance* of the model error is proportional to the mean tendency rather than the standard deviation. Evidently, the associated probability distribution function resembles a Poisson distribution, for which the variance equals the mean.

Vorticity confinement has been introduced as a kind of ‘deterministic backscatter’ and shown to provide a similar spectral distribution of energy input to that seen in the mean difference between T1279 and T159 forecasts. By contrast, SKEB provides a rather uniform spectral distribution of energy input and exhibits a variability at low wavenumbers that cannot be seen in the differences between T1279 and T159 forecasts. This suggests that a key assumption in SKEB (that the phases of the waves making up the pattern are independent of the actual flow) is invalid. A stochastic variant of vorticity confinement could provide an alternative algorithmic basis for backscatter in which vorticity perturbations are effectively created from small rearrangements of the vorticity distribution at any instant.

References

- Berner J., Shutts G. J., Leutbecher M., Palmer T. N. (2009) A Spectral Stochastic Kinetic Energy Backscatter Scheme and Its Impact on Flow-Dependent Predictability in the ECMWF Ensemble Prediction System. *J. Atmos. Sci.*, **66**, 603626.
- Bowler N. E., Arribas A., Mylne K. R., Robertson K. B. and Beare S. E. (2008) The MOGREPS short-range ensemble prediction system. *Q. J. R. Meteorol. Soc.*, **134**, 703-722.
- Bowler N. E., Arribas A., Beare S. E. Mylne K. R. and Shutts G. J. (2009) The local ETKF and SKEB: Upgrades to the MOGREPS short-range ensemble prediction system. *Q. J. R. Meteorol. Soc.*, **135**, 767 - 776.
- Buizza R., Miller M. and Palmer T. N. (1999) Stochastic representation of model uncertainty in the ECMWF Ensemble Prediction System, *Q. J. R. Meteorol. Soc.*, **125**, 2887-2908.
- Charron M., Pellerin, G. Spacek L., Houtekamer P. L., Gagnon N., Mitchell H. L. and Michelin L. (2010) Toward Random Sampling of Model Error in the Canadian Ensemble Prediction System. *Mon. Wea. Rev.*, **138**, 1877-1901.
- Houtekamer P. L., Lefaivre L., Derome J., Ritchie H., Mitchell H. L., (1996) A System Simulation Approach to Ensemble Prediction. *Mon. Wea. Rev.*, **124**, 1225-1242.
- Leutbecher M. (2011) Stochastic tendency perturbations for NWP ensembles. *ECMWF workshop entitled ‘Representing Model Uncertainty and Error in Weather and Climate Prediction’*, June 20-24 2011.
- Palmer T. N. (2001) A nonlinear dynamical perspective on model error: A proposal for non-local

stochastic-dynamic parametrization in weather and climate prediction models, *Q. J. R. Meteorol. Soc.*, **127**, 279304.

Rougier J., Sexton D. M., Murphy J. M. and Stainforth D. (2009) Analyzing the climate sensitivity of the HadSM3 climate model using ensembles from different but related experiments. *J. Clim.*, **22**, 3540-3557.

Sadourny R. and Basdevant C. (1985) Parameterization of Subgrid Scale Barotropic and Baroclinic Eddies in Quasi-geostrophic Models: Anticipated Potential Vorticity Method. *J. Atmos. Sci.*, **42**, 1353-1363.

Shutts G. J. (2005) A kinetic energy backscatter algorithm for use in ensemble prediction systems. *Q. J. R. Meteorol. Soc.*, **131**, 3079 - 3102.

Shutts G. J. and Palmer T. N. (2007) Convective Forcing Fluctuations in a Cloud-Resolving Model: Relevance to the Stochastic Parameterization Problem. *J. Climate*, **20**, 187202.

Steinhoff J. and Underhill D. (1994) Modification of Euler Equations for Vorticity Confinement : Application to the Computation of Interacting Vortex Rings. *Physics of Fluids*, **6**, 2738.

Steinhoff J., Lynn N. and Wang L. (2005) Computation of high Reynolds number flows using vorticity confinement : I. Formulation. *UTSI Preprint, University of Tennessee Space Institute, Tullahoma, TN*.

Weaver A. and Courtier P. (2001) Correlation modelling on the sphere using a generalized diffusion equation. *Q. J. R. Meteorol. Soc.*, **127**, 1815-1846.

