SOME INTERPRETATIONS OF STOCHASTIC PHYSICAL PARAMETRIZATIONS

Glenn Shutts

Met Office, Bracknell, England

Summary: Current parametrizations used in weather and climate models aim to represent ensemble-average fluxes of various sub-gridscale process elements such as convective clouds. In any flow realization, contributions to the flux averaged over the filter scale of a model are likely to vary enormously in space and time implying that the effect of the sub-gridscale physical process will be different from that of the corresponding parametrization. Some arguments will be given to support the introduction of a stochastic element to current dynamical parametrizations to represent this uncertainty.

1 INTRODUCTION

The aim of this lecture is to provide some physical motivation for the introduction of a stochastic element into current dynamical parametrizations such the convection and gravity wave drag parametrization schemes. Here the term 'stochastic' refers to the existence of variables in parametrization schemes whose value is obtained by random selection in accordance with some probability distribution function e.g. a number selected in the range 0 to 1 with the same probability of finding the number between 0 and 0.1 as in finding it in the range 0.1 to 0.2 or any other arbitrary interval of width of 0.1 . These 'random numbers' reflect the true unpredictability of physical parameters at or near the gridscale due to processes that are absent from the model or poorly-resolved. It is a sobering fact that many of the most important physical processes in numerical weather prediction are only partially resolved e.g. the drag exerted on mesoscale mountain ranges such as the Swiss Alps. The near gridscale response generated by such processes might only be considered to be known in some statistical sense.

The idea of adding random numbers to a weather forecast model in the hope of obtaining improved predictive skill seems far-fetched. Clearly the magnitude of these stochastic modifications must lie within some bounds related to the expected 'model error' associated with grid-box means. This naturally leads to the question of what the grid-box mean represents. Since finite difference methods assume smoothness at the gridscale so that the truncation error may be tolerably small, the grid-box mean represents some kind of weighted-spatial average over several grid boxes. Near the surface where parametrized fluxes dominate, gridpoint temperatures and winds are better regarded as parameters in energy and momentum budget formulae (respectively) rather than actual wind values. For instance in hilly terrain it would not be meaningful to compare a local 10 metre wind measurement with a gridpoint wind value due to the direct influence of hills (e.g. hill top speed-up, lee-side wakes of low wind speed etc).

Another example which will be discussed in greater detail later on in this paper is the effect of buoyant convection where groups of cumulonimbus clouds have their own local circulations that may be comparable in magnitude with the synoptic-scale flow as represented on a forecast model grid. These deep cloud systems ($\sim$ 10 km) are not, in the context of current forecast models, predictable and whilst parametrization schemes seek to represent their mean transport properties, one can anticipate large time fluctations of vertical fluxes of mass, heat, moisture and momentum on grid-box scales. The standard deviation of the mean convective mass flux within a model grid-box depends on the average number of independent convective updraughts (N) contained within the grid-box and would tend to decrease as $1/\sqrt{N}$. Even at current operational NWP resolutions, the assumption that N is large is invalid : the real convective mass fluxes averaged over the scale of model grid-boxes (or a somewhat larger filter scale) have substantial instantaneous departures from the mean.

Whilst we know from chaos ideas that a computer weather forecast cannot be regarded as deterministic, one tends to feel uncomfortable with the idea of undermining what determinism we do have by introducing random numbers into the forecast model formulation. As a device for increasing the forecast differences

between between individual members of an ensemble prediction system, we are probably happier (Buizza et al, 1999). Is it possible though, that parametrizing the stochastic variability in near-gridscale motions might have a small but positive impact on the forecast system through many data assimilation cycles ? We have yet to discuss how mimicking the observable statistical fluctuations in convective transports and gravity wave drag found might translate into improved forecasts and so at this stage, the random number exercise looks merely cosmetic. In order to give some credibility to the stochastic forcing technique, let us now review the application of 'stochastic backscatter' in large eddy simulation models where some success has been achieved.

2 Stochastic backscatter in LES

The Large Eddy Simulation (LES) strategy is now a well-established approach to modelling turbulence (Mason, 1994). It consists of combining of a standard numerical model for evolving the Navier-Stokes equations in time with a sub-filter scale model based on a turbulent viscosity closure. The truncation scale of the explicit model is imagined to be sufficiently fine that the bulk of the kinetic energy of the turbulence is resolved, with only a small fraction (e.g. 5% or less) lying below the filter scale. When this condition is met it would seem plausible that the exact nature of the sub-grid model would not significantly affect the turbulence properties and this has been shown to be the case (Mason, 1994). Whilst this basic requirement can easily be met for 'free turbulence' away from rough boundaries, it does inevitably fail in boundary layers as one approaches a 'wall' (to use the engineering terminology). This is because the characteristic turbulent eddy scale varies with distance from the wall and ultimately becomes smaller than the mesh size close to the wall. Here the sub-grid model must carry all of the momentum transfer - consistent with the total viscous shearing stress and pressure drag exerted at the wall surface. Experience with the LES approach shows that the most well-known of turbulent boundary layer properties – the 'log-law' – is not properly represented with an error in the form of enhanced shears at a height where the sub-grid model hands over to the explicit model (Fig.1 (a)). It appears that at this transition scale (the near-gridscale of the explicit model) there is insufficient resolution for the explicit turbulent eddies to carry the bulk of the momentum transport and so the mean shear increases so that sub-grid model can compensate for this.

Fig. 1 (b) shows the height variation of the non-dimensional shear $\frac{\partial u}{\partial z}\frac{kz}{u_*}$ taken from LES simulations using the Smagorinsky-Lilly turbulent viscosity closure where u is a horizontal wind component, z is the height above the boundary, u_* is the friction velocity and k is von Karman's constant. The bulge in the profile near the surface represents a departure from logarithmic behaviour which should be characterised by a value of unity for the non-dimensional shear. The apparent underestimation of the shear in the upper part of the profile is exaggerated by the smallness of the shear (u_*/kz) used in the non-dimensionalization. Mason and Thomson (1992) show that this spurious non-logarithmic behaviour is corrected with the introduction of random, near-gridscale stress fluctuations. Their stochastic forcing assumes statistically isotropic stress fluctuations whose local values are independent of the filter scale and given by a Gaussian probability distribution function. The magnitude of the stress fluctuations was chosen in their work so that the rate of supply of kinetic energy to the resolved scales of the model is about equal to the kinetic energy dissipation rate.

The effect of stochastic forcing on the wind and wind shear profiles is also shown by the solid lines in Fig. 1 (a) and (b) respectively. The horizontally-averaged wind is systematically too strong above a height of 100 m without a stochastic backscatter parametrization and is much closer to logarithmic form with its introduction. The non-dimensional shear is now fairly height-independent and close to unity (Fig. 1 (b)). Another benefit of stochastic backscatter is the removal of a 'spin-up' problem in the transition to turbulence where the unforced LES exhibits an initial dramatic overshoot in turbulent kinetic energy (Fig. 2). This results from the inability of the model to generate turbulence until a critical threshold of boundary layer shear is reached. Backscatter apparently enables the LES to release turbulent kinetic energy from the intense near-surface shear by providing energy to the near-gridscales eddies.

Whilst the effect of stochastic backscatter appears to be beneficial and has a fairly sound physical rationale, one does not approach the problem in the same way that an NWP modeller would parametrize deep convection. Specifically, there is not clear picture of the turbulent structures that contribute to sending energy to larger scales that matches a meteorologists understanding of the structure of an organized convective

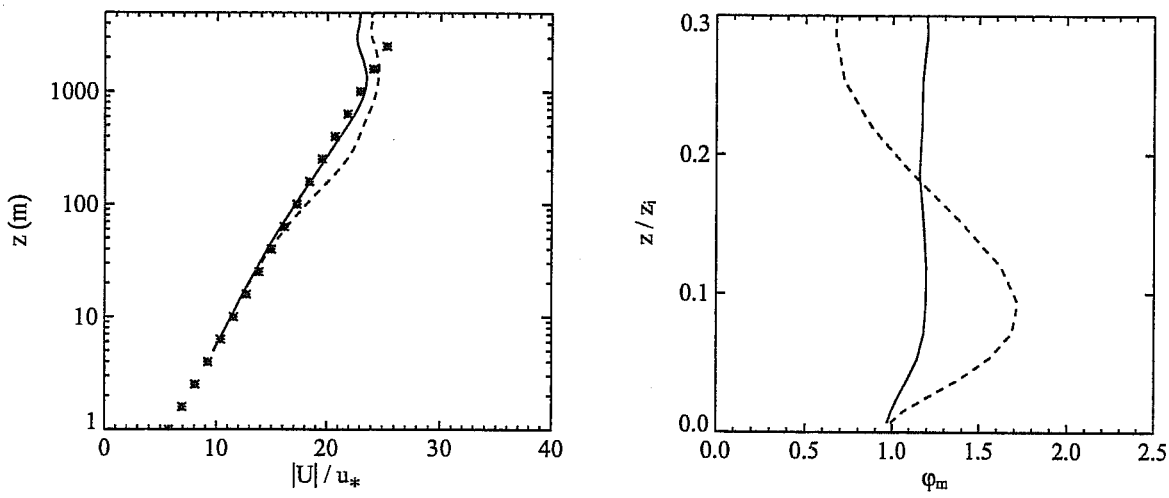

Figure 1: (a) Plot of the non-dimensionalized, horizontal and time averaged wind speed $|\mathbf{V}|/u_*$ versus height from large eddy model simulations carried out by A. Brown . The dashed (solid) line shows the wind profile in a run without (with) backscatter and the asterisks indicate the log-law. (b) shows plots of the non-dimensional wind shear $\varphi_m = \frac{kz}{u_*}\left|\frac{\partial \mathbf{V}}{\partial z}\right|$(horizontal and time average) versus non-dimensional height z/z_i where z_i is the boundary layer depth. The line styles represent backscatter and no-backscatter as in (a).

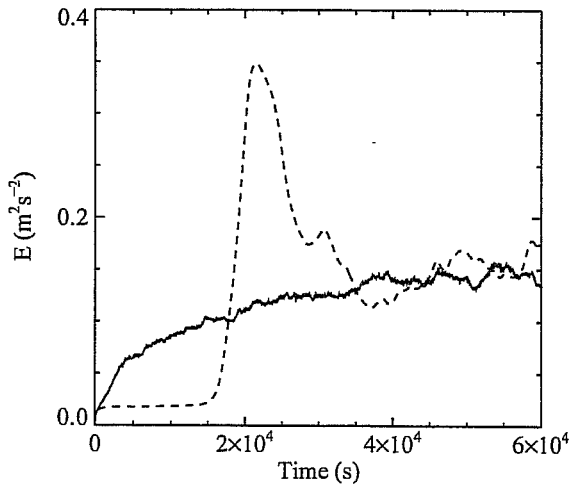

Figure 2: The volume-averaged turbulent kinetic energy in the large eddy model as a function of time; dashed line, without backscatter; solid line, with backscatter.

storm. Perhaps it doesn't matter because in attempting to parametrize a statistical process we should not necessarily be concerned with the structure of individual process elements. Ultimately though, one feels that some understanding of the near-wall processes may at least help to optimize the way stochastic forcing is formulated.

The key process that causes the cascade of energy to small scales in turbulence is the stretching of vortex tubes which can, with great rapidity, send energy from the eddy scale to the viscous scale. In a tangle of vortex tubes we suppose that, on average, vortex tubes are more likely to be stretched than compressed. Nevertheless, a small fraction of vortex tubes will be *compressed* with consequent upscale transfer of energy. Does the backscatter parametrization represent the contribution to the near-gridscale of kinetic energy produced by compressing sub-gridscale vortex tubes ? The problem is that the 'tangle of vortex tubes' model is just a convenient idealization for thinking about the 'sinews' of turbulence : it does not help much in understanding the dynamics of the wall region.

The phenomenology of near-wall turbulence has received much attention in the last decade or so and in particular the recognition that highly-organized flow structures erupt from the wall, interact with the background shear and regenerate i.e. the hairpin vortices. Although their dynamics has yet to be clarified they appear to be a manifestation of an explosive instability resulting from the upward buckling of densely-packed vortex lines at the surface. The anti-parallel vortex filaments that form at the 'loop' are then subjected to strong background shear and self-advect away from the wall forming long trailing 'hairpin-like' legs. Visualization of these 'burst events', as they are often called, is quite difficult (Haidari and Smith, 1994). The process of vorticity transfer away from the wall in these burst events resembles flow separation and may require direct numerical simulation to model it i.e. models that have sufficiently high resolution that the eddies are fully resolved and the sub-grid model is simply the parametrization of viscous diffusion. In marked contrast to the steady diffusive transfer of vorticity away from the surface in the laminar layer, burst events appear to eject vorticity from the boundary and through the viscous layer. Could it be that the energy backscatter process used in LES is crudely representing this violent and highly intermittent vorticity transfer process that characterizes near-wall turbulence generation ? It should be noted that this surface instability view of turbulence generation is not the only explanation of the flow structures seen in laboratory simulations. Hunt and Morrison (2000) take a 'top-down' view of neutral, very high Reynolds number turbulent boundary layers wherein eddies in the middle layer impinge on the boundary from above and cause vorticity ejections through interaction with the near surface flow.

If stochastic backscatter is addressing the absence of burst events in an LES model one might wonder whether the form which the stochastic forcing function takes matters. In the work of Mason and Thomson (1992) and Mason and Brown (1994) the stochastic forcing of the momentum equation is obtained by spatially smoothing a three-dimensional vector field defined by random number components (whose means are zero) and then taking the *Curl* giving a non-divergent vector field. This vector forcing field - taken to be the flow acceleration - is scaled so that the net rate of energy input is consistent with a certain scaling argument. Unlike the construction of NWP model parametrization schemes, no attempt is made to design a forcing function that reflects known properties of the sub-grid flow structures (e.g. hairpin vortices). This policy may be considered to be the result of insufficient understanding of those sub-grid motions or a deliberate decision to avoid introducing empirically-derived knowledge into the model. Whether the use of information on hairpin vortex structures could add to the realism of the model's turbulent statistics is questionable. In any case a turbulence purist might regard this as 'cheating'. Parametrization of dynamical processes in weather forecast models often begins with understanding of individual parametrizable elements (e.g. a convection cloud) and attempts to infer the statistical properties of a group of such elements.

3 Convective backscatter

3.1 Mesoscale Convective Systems

Arguably the most important source of near-gridscale kinetic energy relevant to NWP models is that due to deep convective storms and mesoscale convective systems. Convective systems on smaller scales such as boundary layer convection (e.g. fair weather cumulus over land and cold airstream cumulus over the ocean) are ineffective in passing their kinetic energy to the larger balanced scales of motion. Convective

parametrization schemes usually implicitly assume that all kinetic energy released in clouds is also dissipated within them. Their effect is primarily to address the stabilzation of the atmosphere through subsidence forced by convective mass transfer and modifying the large-scale humidity field. In presence of considerable moist convective instability (characterized by large values of the Convective Available Potential Energy, CAPE), convection can transfer appreciable fractions of released kinetic into quasi-balanced mesoscale flow structures . This type of organized convection is known as a Mesoscale Convective System (MCS) and is characterized by an upper-level outflow extending over hundreds of kilometres in the form of an elliptical cloud shield as viewed from space. The MCS may comprise many updraught cores - often organized into lines – and feeding the trailing stratiform cloud system. Precipitation from this mesoscale cloud shield drives a descending in flow through the evaporational cooling (together with both sublimation and melting of snow).

In order to appreciate the way an MCS interacts with the large-scale flow, it is important to regard its principal effect as the redistribution of mass in a rotating, stratified flow. From this viewpoint, we are dealing with a class of geostrophic adjustment problem for which the potential temperature of air parcels is reset on passage through the storm system. If we were also to assume that the convective circulation was axisymmetric, then the angular momentum of any toroidal air parcel would be conserved. Knowledge of the potential temperature and angular momentum of air parcels coupled with mass continuity is sufficient to determine a unique balanced flow state (Shutts et al, 1988). In these idealized mass transfer calculations, the balanced endstate takes the form of a upper-level anticyclonic lens of convected mass and a lower level cyclonic vortex of much smaller radius . The model suffices principally to convey the idea that deep penetrative convection in a rotating stratified flow feels the effect of both the inertial stiffness imparted by rotation and the resistance to vertical displacement implied by gravity in a thermally-stratified environment. Spreading of the outflow tends to be limited by angular momentum conservation though departure from symmetry can break this constraint. Indeed this symmetry breaking may be necessary for the persistent intensification of hurricanes.

Whilst the previous model may appear overly restrictive, it illustrates a more general view of the interaction of convection with its environment. This requires that we treat convective updraughts and downdraughts as agencies for redistributing mass in what otherwise is a quasi-balanced flow in the sense required of the potential vorticity invertibility principle (Hoskins et al, 1985). Deep convection lies outside of this balanced flow view yet we may regard the mass transfer it achieves as an effective mass source/sink to the balanced flow. Here we regard the governing equations used in NWP models as quasi-balanced even though they are capable of supporting inertia-gravity wave motions. In the presence of explicit mass sources and sinks, Ertel's theorem for the potential vorticity conservation becomes

$$\frac{Dq}{Dt} = -qM_f/\rho$$

where $q = \rho^{-1}\zeta \cdot \nabla\theta$ (ρ is the density, ζ is the absolute vorticity and θ is the potential temperature) and M_f is defined by

$$\frac{D\rho}{Dt} + \rho\nabla \cdot \mathbf{V} = M_f$$

(Shutts and Gray, 1997).

Gray et al (1998) show how the balanced flow that remains after deep convection in a cloud-resolving model can be parametrized by the introduction of an explicit mass forcing term into the anelastic continuity equation used by the model. If M_c is the mass convected in a single cloud system then Shutts and Gray (1994) argue that the residual balanced energy is proportional to M_c^2 in 2D flows whereas in three-dimensional flows Gray et al (1998) suggest the dependence is $M_c^{5/3}$. Since the amount of kinetic energy released scales as $CAPE \times M_c$, there exists in principle, an upper bound on the convective mass flux where all of the kinetic energy released goes into creating a residual balanced flow structure (i.e. anticyclonic lens and lower level meso-vortex or shearline front; see also Shutts, 1987 for the adjustment problem in a baroclinic environment). In their two-dimensional case, Gray et al (1998) found that up to 60% of the released kinetic energy could be captured in balanced flow associated with convectively-generated potential vorticity anomalies. For isolated storms in three-dimensions this reduces to about 30%. The efficiency of balanced flow production by convective storms is far greater in an MCS than in a group of isolated cumulonimbus clouds carrying the

same total mass flux. For instance let E_{MCS} be the balanced flow energy generated by an MCS and given by the relation

$$E_{MCS} = \alpha M_C^{5/3}$$

where α is a constant. If this mass flux is equally divided amongst n isolated storms then the total energy E' is given by

$$E' = \alpha n \left(\frac{M_c}{n}\right)^{5/3} = \frac{E_{MCS}}{n^{2/3}} \tag{1}$$

and if $n = 10$, the amount of balanced flow energy generated by a group of ten isolated thunderstorms is $10^{2/3} \sim 4.6$ times less than in an MCS.

As stated earlier, convection parametrizations implicitly assume that the kinetic energy released in up-draughts and downdraughts is dissipated within the model grid-box. Potential vorticity modification to an airstream containing convection is implied by the vertical profile of parametrized convective heating though this should, in principle, occur on a horizontal scale that matches the scale of the instability. Because MCS's are near-gridscale phenomena their growth implies an inject of kinetic energy at that scale.

At current NWP global model horizontal mesh sizes, mesoscale convective systems cannot be resolved and so their effect must be handled within the convective parametrization scheme. Clearly the parametrized response, which only acts directly on thermodynamic field, will not address the near-gridscale potential vorticity anomalies generated by convection. Whilst there is some topographic control over the location at which MCS's are triggered, there is sufficient unpredictability to warrant treating the process as stochastic. Computer animations of satellite images taken over continental regions in summer show the sporadic and explosive growth of MCS's and other large convective clouds – at least with respect to the time scale of depression formation. It is often possible to discern the rapid anticyclonic rotation of the upper outflow cloud shield as it evaporates, consistent with the dilution of potential vorticity substance (in the terminology of Haynes and McIntyre, 1990).

Of all the sub-grid physical processes currently parametrized, deep convection seems to be the one most suitable for stochastic representation. This is not simply because of its impulsive and unpredictable nature but also the fact that mechanisms for the upscale cascade of energy to the explicitly-resolved, *balanced* scales of motion have been identified. A stochastic parametrization of MCS activity could be triggered in regions of large CAPE and might, for instance, backscatter a certain fraction of the kinetic energy generated by buoyancy forces within the clouds (i.e. the product of the mean convective mass flux and CAPE).

Recent unpublished case studies of the effect of stochastic forcing due to MCS's by M. Gray at the Met Office go further and use dynamical understanding of the effect of MCS's to force a specific potential vorticity anomaly structure, albeit varying in scale and intensity. Using the assumption that the MCS potential vorticity anomalies are dominated by the associated vertical vorticity rather than stability anomaly, he forces the Met Office Unified Model (UM) with axisymmetric wind patterns that represent the upper anticyclonic lens and mid-level meso-vortex.

Specifically, Gray defines a characteristic upper cloud shield radius a given by

$$a = c\sqrt{CAPE}/\max(|f|, f_0)$$

where $f_0 = 2 \times 10^{-5}\ s^{-1}$, c is a dimensionless stochastic parameter set by the convective precipitation rate P_{con} i.e.

$$c = 0.1 P_{con} R$$

where R is a number selected at random from the range $0 \rightarrow 1$ with a flat probability distribution and P_{con} is in units of mm/hr. The absolute vorticity of an outflow cloud shield of thickness d_a is assumed to be zero - consistent with the imposition of an anticyclonic vortex with tangential velocity V given by

$$V(r) = \begin{cases} -\frac{1}{2} f r & \text{for} \quad r < a \\ -\frac{f}{2}\frac{a^3}{r^2} & \text{for} \quad r \geq a \end{cases} .$$

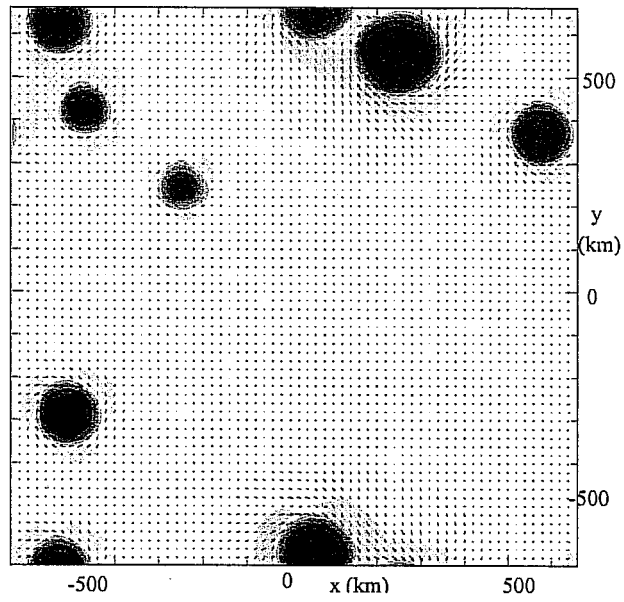

Figure 3: An example of the vorticity and wind fields associated with a group of convective clouds (at the upper outflow level) obtained from Gray's proposed stochastic convective backscatter scheme.

The accompanying meso-vortex, with depth d_v, is taken to be of the form

$$V(r) = \begin{cases} \frac{1}{2} f \frac{r}{n} \frac{d_a}{d_v} & \text{for} \quad r < na \\ \frac{f}{2} \frac{a^3}{r^2 n} \frac{d_a}{d_v} & \text{for} \quad r \geq na \end{cases}$$

where na is the radius of the meso-vortex. An exclusion algorithm is implemented which removes selected vortices to ensure that the nearest neighbour lies at least 2 vortex radii away (using the larger of the two vortices) . Prior to implementing the scheme in the operational UM, this stochastic convective backscatter algorithm was tested in the Met Office's Large Eddy Model. Fig. 3 shows a particular realization of the vorticity and wind fields created by the scheme.

In Gray's sensitivity tests using the UM, $d_v = 2$ km and $d_a = 1.5$ km, with the meso-vortex being centred 100 mb above the freezing level and the upper cloud shield centred 50 mb below the cloud top height - as diagnosed by the convective parametrization scheme. The u and v increments were computed based on a time scale of three hours to add the above vortices and introduced via the convective momentum transport subroutine.

Out of five case studies upon which Gray's preliminary work is based, three were during the northern hemisphere summer, one was a southern hemisphere summer case, and one was an equinoctial case. The impacts of the stochastic convective backscatter were measured in terms of the root-mean square (rms) difference of two different backscatter runs (based on different random numbers) relative to a control run with no backscatter. It was found that the global-mean rms differences of the sea-level pressure and geopotential height at 500 and 200 mb at day five of the forecast are general quite small (e.g. 0.29 to 1.57 mb for the sea-level pressure) and mainly coming from the southern hemisphere in spite of the fact that most of the convective events were diagnosed in tropical regions. The influence of the MCS vortices was greatest near the mid-latitude baroclinic zones of the southern hemisphere where small changes to the orientation of low pressure troughs may have a large impact on subsequent development, principally the phase though . In a few of the cases the intensity of tropical storms was altered and in one case a parametrized MCS formed the seed for tropical storm initiation !

At this stage there remain some questions regarding the implementation of the convective backscatter, particularly with respect to its performance over land. For instance, the intensity of vorticity forcing is heavily dependent on the UM's convective precipitation rate which is unlikely to be reliable in the conditions in which MCS's form. The MCS itself may not encapsulated a broad enough range of deep convective

processes which classify as stochastic convective events. Satellite image animations of oceanic stormtracks in wintertime give the impression that active frontal zones contain much embedded deep convective structure, often leading to, or manifesting itself as mesoscale cloud leafs or comma cloud systems. An important and enigmatic aspect of this is the role of slantwise instability and the presumption that this may be released in slantwise convection. A fundamental difficulty with this concept is its two-dimensional basis and the assumption that air conserves absolute momentum during slantwise frontal ascent. In spite of weaknesses in the theoretical underpinning of this process there exists some compelling evidence from numerical modelling and satellite imagery that the rapid growth of some frontal cloud shields is a manifestation of slantwise convection (Dixon et al, 2001). It seems likely that these slantwise convective plumes are still beyond the resolution of current global NWP models and as such represent a potential candidate for stochastic parametrization.

Another approach to the implementation of stochastic physical parametrization has been taken by Buizza et al (1999) who multiply the tendencies of all physical parametrizations in the ECMWF forecast system by a height-independent number (r) selected randomly from a range centred on 1 with uniform probability. The value of r used in any forecast run was chosen from one of three ranges; $r \in [0, 2]$, $r \in [0.5, 1.5]$ and $r \in [0.75, 1.25]$. The same random number is used for all grid-boxes within latitude/longitude squares of side D degrees and for time period T hours. They experimented with values of $D = 1, 5$ and 10 and $T = 0.75, 3$ and 12 hours. In general, the range $r \in [0.5, 1.5]$ and values $D = 3$ degrees, $T = 3$ hrs gave the most desirable impact - increasing the spread of an ensemble whilst decreasing the ensemble-mean rms error. Since the dominant effect in this stochastic physics package probably comes from the convective heating it can, to some extent, be compared with stochastic convective backscatter though acting in a quite different way - modifying the local buoyancy field rather than the vorticity field. Buizza et al (1999) are able confirm, using TOGA-COARE data, that the gross effect of the stochastic forcing on the precipitation rate is not unreasonable. Even so, it would be interesting to see the effect of the stochastic forcing on the levels of inertia-gravity wave energy explicitly resolved by the forecast model.

From a numerical techniques standpoint, it is desirable that the stochastic forcing functions are fairly smooth at the gridscale to avoid generating spurious two-gridlength waves. On the other hand the whole point of the stochastic forcing method is to force the near-gridscale motions which are only partially resolved such as the MCS. A further numerical concern is the tendency of convective parametrization schemes to excite gridscale gravity waves due to their column-based design. As the horizontal resolution of forecast models increases, so the gridscale variability of parametrized convective heating becomes more effective in destroying local thermal wind balance and generating two-gridlength gravity waves with highly inaccurate phase speeds (e.g. zero). Spurious gridscale motions can then interact with the convective parametrization scheme in a way that maintains and spreads them. Unless measures are taken to ensure spatial smoothness of parametrized convective forcing, it could be argued that this noisy gridscale behaviour of existing parametrization implementations already has a stochastic character.

An interesting alternative approach to generating and imposing realistic patterns of convective forcing onto near-gridscale fields is through the cellular automaton model (Palmer, 1997; Palmer, 2001). Palmer envisages the modelling of organized mesoscale convective structure on a cellular automaton grid where each grid cell exists in an 'on' or 'off' state prescribed according to a set of rules chosen to reflect those physical processes which determine the time scale and space organization of convection. The 'on' states may be associated with cells containing convective activity and the probability of their occurrence may be linked to the CAPE diagnosed from the forecast model. The persistence of on-states may be loosely associated with the degree of 'clumping' of existing cells so that groups of adjacent 'on-state' cells would behave like MCS's. Patterns of convective organization generated this way could then be used to impose structure on convective parametrization increments or form the basis of a stochastic convective backscatter scheme as described above.

Lin and Neelin (2000) discuss the beneficial effects of a stochastic convective parametrization scheme that works by perturbing a measure of CAPE in the Betts/Miller parametrization scheme (Betts and Miller, 1986) with a random noise term formulated as a first-order Markov process. When used in a general circulation model they find noticeable improvements to the tropical intraseasonal variability for certain plausible noise amplitudes and autocorrelation timescales. As is often the case with the use of stochastic parametrization schemes it is not always clear why the effects are beneficial even if, at the cosmetic level, the stochastic modifications seems reasonable.

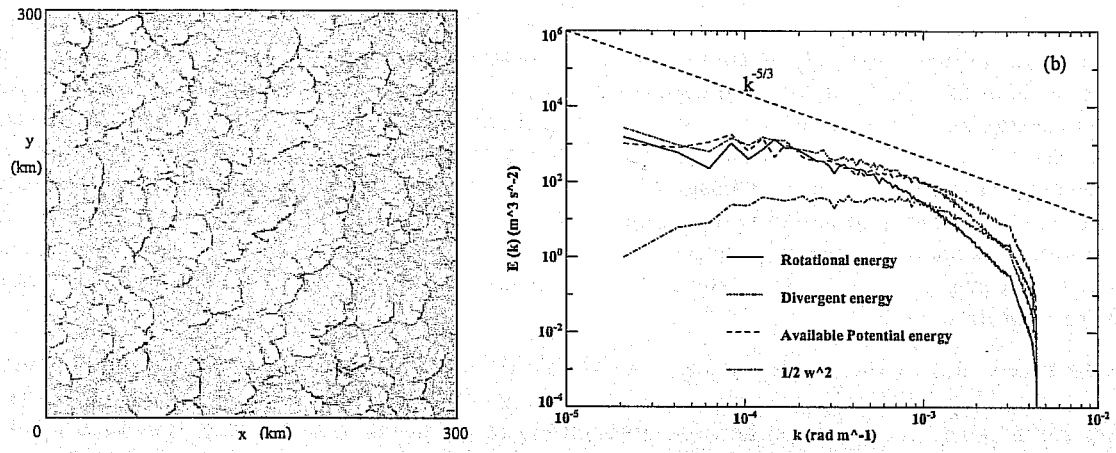

Figure 4: (a) Vertical velocity (w) at a height of 458 m in a simulation of deep convection forced by imposed cooling after about one day. The shading represents a range of w extending from -3.1 ms^{-1} to $+3.1$ ms^{-1} with dark shading being upward motion. (b) Energy spectra at a height of 7.6 km after about one day. The horizontal wind is split into rotational and divergent contributions and the spectral energy density of each is plotted along with the contribution of vertical motion to the kinetic energy, and the available potential energy. The straight dashed line indicate the slope one would see if the energy density followed the $k^{-5/3}$ inertial sub-range of two-dimensional turbulence.

Whilst the vortical motion associated with MCS's is fairly weak (e.g. ~ 5 ms^{-1}), other forms of convectively-driven vortices are much stronger e.g. polar lows and hurricanes. For many climate models, these vortices are unresolvable and even in current NWP models with gridlengths $\sim$ 50 km they are only partially resolved. This raises the question of possible benefits from parametrizing their stochastic effect.

3.2 Upscale energy cascade

The process of potential vorticity forcing by MCS's described above was made possible by background rotation. In the absence of Coriolis forces it is still possible for convective clouds (and other small to mesoscale forcing agencies) to export kinetic energy to large scales. Lilly (1983) hypothesized that cumulonimbus anvil outflows might behave as collapsing two-dimensional wakes whose kinetic energy then preferentially spreads to larger scales in the manner of two-dimensional or 'stratified' turbulence. He postulated that only a small fraction of the total kinetic energy released was required to account for the observed $k^{-5/3}$ energy spectrum in the atmospheric mesoscale (Gage and Nastrom, 1986). In an attempt to verify this mechanism and to quantify the influence of the Coriolis effect, Vallis et al (1997) used a high resolution numerical model to study the approach to statistical convective equilibrium under the influence of uniform imposed cooling and fixed surface temperature. The model had a square horizontal domain of side 300 km and resolved with a uniform 1 km grid which is sufficiently high to permit the growth of fairly realistic deep convective cloud fields. In runs with a uniform geostrophic wind imposed and sufficient vertical resolution near the surface to permit a crude frictional boundary layer, low-level convective clouds organize into 'cuspy lines' separated by a distance of about 60 km whereas the deeper clouds tend to lie at various points along the lines (Fig. 4 a). The line structure arises as convective downdraughts spread laterally in the presence of boundary layer wind shear. The relative flow approaches each line, and feeds the updraught, from the southeast. These linear structures are absent from the mid-troposphere upwards and the w field there consists mainly of a collection of near-gridscale updraught cores lying close to the convection lines below.

A breakdown of the spectral distribution of energy at a height of 7.6 km is given in Fig. 4 (b), close to the mean level of neutral buoyancy for the updraughts. From this one might conclude that the spectral slope was quite shallow and not very different from the observed $k^{-5/3}$ spectrum at the mesoscale. However there are several aspects of the energy spectrum and its evolution that are difficult to reconcile with the energy-cascading inertial sub-range of two-dimensional turbulence e.g. larger divergent energy than rotational,

inefficient upscale spectral energy transfer rate (Vallis et al, 1997). Shutts and Gray (1999) argued that the spectral slope was a direct result of the convective line structure and embedded convective updraught cells. At the neutral buoyancy level, positive horizontal divergence tends to be concentrated at the updraught cores with convergence occurring at the convective line separation scale. The result is a shallow spectral slope lying between k^{-1} and k^{-2} . The mechanism of energy transfer to large scales does not appear to be due to an energy cascade in the way envisaged in $2D$ turbulence. It relates more to those processes that determine the fractional area occupied by updraughts (typically about 1%) and the convective line spacing. Perhaps this upscale energy transfer is simply the transition from the convective scale to the mesoscale : air ascends in updraughts on the 1 km scale and then spreads out into cloud anvils over tens of kilometres (similarly for downdraughts).

It should be noted that in the above simulations, MCS's do not form spontaneously and the convective organization resembles that seen in cold airstreams over the sea in satellite images. This mode of convection is less effective in generating balanced flow features related to potential vorticity anomalies than the MCS mechanism but may be the source of an upward flux of gravity wave energy above the cloud tops. As regards stochastic parametrization, the scheme due to Buizza et al (1999) would probably suffice to capture the time and space variability of the subsidence warming averaged over groups of grid-boxes.

4 Gravity wave drag parametrization

4.1 Orographic form drag error

The stochastic character of convection in an NWP context derives from instability of flow and the consequent sensitivity to initial conditions and topographic forcing. Internal gravity waves forced by hills and mountains (and their associated wave drag) would appear to more predictable. Given an upstream wind profile, the orographic height field and a suitable upper boundary condition, even linear theory can provide surprisingly realistic lee wave fields (Vosper, 2001). However for mountains sufficiently high that the parameter U/Nh is less than unity (where U is the approaching low-level flow speed; N is the buoyancy frequency and h is the mountain height), flow blocking sets in and the amount of gravity wave energy radiated, and wave drag accrued, is far less certain. This uncertainty is amplified when U/Nh is of order unity for which hydraulic behaviour is possible and nonlinear resonances increase the sensitivity of the drag to small flow changes (Miranda and Valente, 1997). The existence of wind and temperature profiles that trap lee waves (i.e. when the Scorer parameter decreases with height in the troposphere) allows linear resonance to occur and accounts for the ubiquitous lee wavetrains seen in satellite imagery. This is also a potential source of sensitivity to the large-scale flow that might make the prediction of wave drag probabilistic (Bretherton, 1969). Like the problem of convection parametrization, we seek the statistical effect of many parametrizable elements and so wave drag schemes based on single mountain flow responses may be unrepresentative (Welch et al, 2001). Since one mountain may lie in the wake of another, it is doubtful whether all mountains meeting the criteria for a hydraulic flow response would actually excite such a response.

Another source of uncertainty in the parametrization of gravity wave drag lies in the incomplete specification of the sub-grid orography. Parametrization schemes tend to use elementary statistical properties of the sub-gridscale orography such as the variance of the orographic height about the mean and the mean-square orographic slopes e.g. Gregory et al, 1998 or attempt to model the orography as a particular idealized form e.g. elliptical mountains Lott and Miller, 1997; Scinocca and McFarlane, 2000. However it is quite possible to imagine sub-gridscale orographic configurations that have identical orographic height variance but a very different drag response. Vosper (1996) considered the effect on the drag of wave interference between two neighbouring mountains. In cases of strong wave trapping the trapped lee wavetrain of the upstream mountain may cause constructive or destructive interference with the downstream mountain depending on the resonant wavelength and the mountain separation. In linear theory at least, it would be possible to arrange for the wavetrain due to the upstream mountain to annihilate that due to the downstream mountain or, if the mountain separation was changed by half a wavelength, to double its amplitude. Such dependencies would not exist for a parametrization scheme based on the orographic variance alone and serve to emphasize the inherent error - even if the dynamical model was perfect.

Whilst there is no shortage of reasons why drag formulae have errors that might be stochastic in nature,

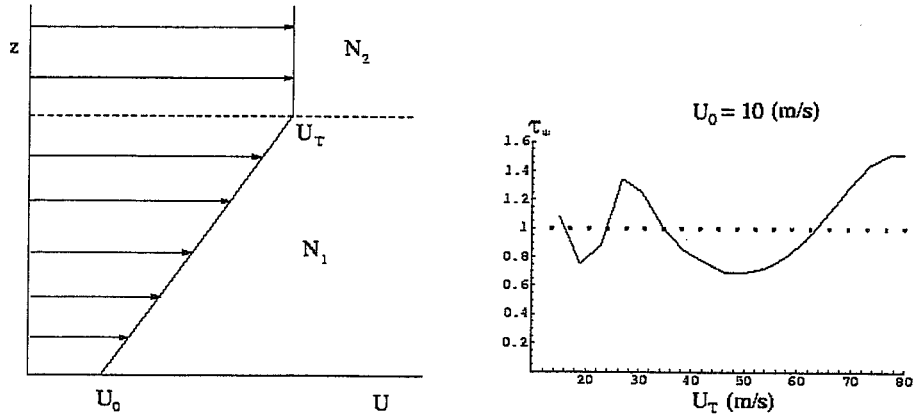

Figure 5: (a) idealized two-layer flow representing the mean properties of the troposphere and stratosphere in a model of flow over a circular hill defined by eq. 2 with $a = 40$ km. The wind is uni-directional and determined by a surface wind $U_0 = 10\ ms^{-1}$ and a variable tropopause wind U_T with constant shear in the troposphere and zero shear above. The buoyancy frequency $N_1 = 10^{-2}\ s^{-1}$ in the troposphere and $N_2 = 2 \times 10^{-2}\ s^{-2}$ in the stratosphere. (b) the total wave drag τ due to the flow shown in (a) and plotted versus U_T. The drag is normalized with respect to the value for a constant wind (i.e. given by eq. 3).

physical idealizations are often made that obscure important dependencies. For instance, the surface gravity wave drag formula, which is based on the linear theory of uniform, stratified flow over a hill, is often assumed to be accurate in the presence of realistic tropospheric vertical wind shear - as if it behaved like an aerodynamic surface drag law. For an isolated circular hill of height $h(r)$ given by

$$h(r) = \frac{h_0}{(1 + r^2/a^2)^{3/2}} \tag{2}$$

where h_0 is the hill top height, r is the radial distance from the hill top and a is a half-width parameter, then the net wave drag (D) on the hill predicted by linear theory for a uniform flow of speed U and buoyancy frequency N is

$$D = \frac{\pi}{4} N U_0 a h_0^2. \tag{3}$$

Fig. 5 (a) specifies an idealized, two-layer flow through the parameters U_0, U_T, N_1 and N_2 and (b) shows the dependence of the normalized wave stress on U_T The wave field and stress were obtained using linear theory, enforcing a radiation boundary condition at the layer interface. It can be seen that the surface wave drag varies by up to 50% from its constant wind value. The variations are caused by partial internal reflection due to the stability discontinuity at the tropopause and the increasing wind with height. In principle, this dependence of the surface drag on the mean tropospheric wind shear could be accounted for in a more advanced surface drag law but the complexity of real profiles of U and N suggests that this may not be worthwhile. Perhaps therefore, this uncertainty may also be regarded stochastic in nature.

4.2 Gravity wave breaking - where does it happen ?

When a strong geostrophic flow develops over a mountain, an internal gravity wave field extending upwards into the stratosphere forms within an hour or so. In idealized two-dimensional models this make take the form of an upwardly-directed 'beam' of wave energy if the hill is sufficiently wide and the wave response lies in the hydrostatic regime. Narrower hills elicit a non-hydrostatic response which may lead to wave trapping if N/U decreases with height. These trapped lee wavetrains may spread laterally for hundreds of kilometres with little amplitude attenuation. As shown by Shutts (1998), the form of orographic gravity wave fields is much more complex for isolated hills in flows which vary in direction with height. In these situations,

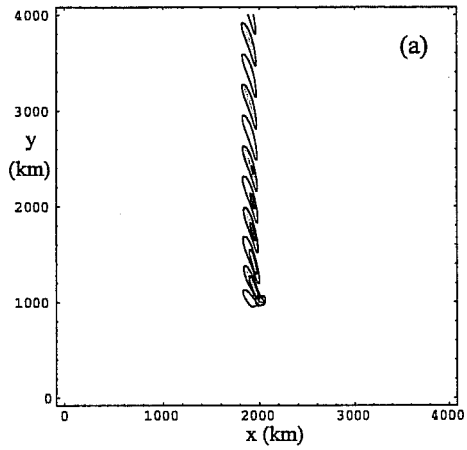

Figure 6: (a) u' field at $z = U_0/\Lambda$for flow over the isolated mountain profile given by eq.(2) with $h_0 = 500\ m$ and $a = 50$ km. The flow is horizontally-uniform and has a height dependence given by $\mathbf{V} = (U_0 - \Lambda z, V_0)$ where U_0 and V_0 are constants (set here to 10 ms^{-1}) and the shear $\Lambda = \sqrt{10} \times 10^{-3}\ s^{-1}$. The buoyancy frequency is set to $10^{-2}\ s^{-1}$. **N.B.** The downstream amplitude decay of the wavetrain results from adding a small positive imaginary number ϵ to the frequency. This numerical device was used to avoid singularities in the Fourier integral solution and permit a solution that does not 'wrap-around'. The magnitude of ϵ can be reduced as the resolution in the Fourier integral is increased and the amplitude decay scale of the wavetrain varies inversely with ϵ.

depending on the degree to which the wind backs or veers with height, the wave energy is dispersed laterally through three-dimensional critical layers implied by the existence of wavevectors at right angles to the local wind. The effect on the vertical flux of wave pseudo-momentum is analogous to detrainment in convective mass flux schemes but here is controlled by the rate of backing or veering of the wind with height. Unlike two-dimensional critical level problems where the wind speed goes to zero at the critical level, here the wind speed remains finite at the point where wave packets have zero vertical group speed. The result is that the energy density of mountain wave beams is depleted smoothly in height without the type of local singular build-up that characterizes the 2D problem and leads to wavebreaking.

Another process that tends to force laterally spreading of wave energy is the effect of rotation when waves have their intrinsic frequency doppler-shifted towards f - the Coriolis frequency. Because the vertical group speed again tends to zero at finite wind speed i.e. when $\left|\frac{\mathbf{V_h}.\mathbf{K}}{f}\right| = 1$ where $\mathbf{V_h}$ is the horizontal wind vector and $\mathbf{K}$ is the wavevector, wave packets get advected away from the wave source. Since the previous condition is met at different heights according to the wavevector magnitude and direction, the effect is to defocus the build-up of wave pseudo-momentum and energy and cause a downstream inertia wavetrain (Eliassen, 1969; Shen and Lin, 1999; Shutts, 2001a). In the non-rotating study of Shutts (1998), it was deduced that the wave energy density (E) in the wavetrains that trail away from the mountain decreases with distance but the local vertical wavenumber (m) increases with distance. Since a characteristic perturbed wind shear is $mE^{1/2}$, it must increase with distance from the mountain and so eventually lead to wavebreaking where the 'wave' Richardson number $N^2/(m^2E) \sim 1$. Shutts (1998) estimated that this might typically be of the order of 10,000 km (though obviously dependent on the strength of orographic forcing). Recent calculations of mountain wave fields for directionally-sheared flows that include rotation suggest, like the two-dimensional case studied by Eliassen (1969), an inertia wavetrain whose amplitude does not decrease downstream (Shutts, 2001 b; see Fig. 6) . This conclusion was also reached by Broutman et al (2001) using ray tracing methods. The implication is that wavebreaking in the trailing wavetrains should occur closer to the mountain than without rotation.

Since the relative group velocity of quasi-inertia waves is small compared with typical atmospheric flow speeds, the inertia wavetrain behaves like a flow 'streakline' and in any real flow this will not be a straight line as in Fig. 6 due to synoptic timescale flow variability. However, wave theory tells us that the horizontal wavevector will be refracted by synoptic scale wind shear and the absolute frequency will change due to time dependence in the flow. Changes to the wavevector and frequency will affect the group velocity and

may permit wave packets to continue their upward journey. Clearly, the ray trajectories emanating from mountain ranges in real, time-dependent and spatially-varying flows may be very complex and the location at which wave breaking occurs could be highly variable. It would be interesting to know what percentage of the wave pseudo-momentum generated by a mountain $\sim$ 100 km wide travels in wave packets more than 1000 km without wavebreaking. This view is to be contrasted with the extremely simplistic assumption used in current orographic gravity wave drag parametrization that the upward transmission of wave pseudo-momentum is a one-dimensional process that obeys the Eliassen-Palm theorem when the wave stress is below its saturation value. The bottom line is that the all-important gravity wavebreaking process that leads to irreversible flow acceleration/deceleration is not as accurate or deterministic as its conventional formulation might suggest.

4.3 Gravity wave soup

Physical mechanisms responsible for gravity wave generation in the troposphere are numerous and include topographic forcing, convection, frontogenesis, geostrophic adjustment in jetstreams and shearing instability. The range of intrinsic frequencies spanned by these processes is bounded from below by f and from above by N, with topographic forcing and convection being at the upper end, and frontogenesis and geostrophic adjustment being at the lower end. The relative contribution of these processes to the global-mean, upward flux of pseudo-momentum is not known but one suspects that mountain waves dominate the budget. As noted earlier, wave properties will be changed as the waves propagate through spatially and temporally-varying winds and stratification and the size of this effect will depend on the lifespan of a wave packet. Waves with high intrinsic frequency and fast vertical group speeds will propagate quickly into the stratosphere and be dissipated in a wave breaking event before the synoptic or mesoscale variability of the flow has had time to act on the wavevector and absolute frequency. In contrast, waves that are forced at low intrinsic frequency, or suffer doppler-shifting to low intrinsic frequency, will have low vertical group speeds and may spend long periods in the presence of horizontal background shear and time variability in the wind. The properties of the waves may be scrambled so that stationary waves become travelling waves and long waves become short waves etc. Little is known of this scrambling process though it is possible to see evidence of it, such as the refraction of trapped lee wave cloud bands near the low-level jet of an approaching cold front.

In middle atmosphere gravity wave drag parametrization schemes it is usually assumed that waves originating from the troposphere have been sufficiently scrambled that it is no longer meaningful to associate them with any specific forcing mechanism. The different forms of gravity wave packet originating from different sources are assumed to have blended into a 'gravity wave soup' with certain universal spectral properties (e.g. an m^{-3} fall off in the spectral energy density at large vertical wavenumber, m) The amplitude of this tail in the vertical wavenumber spectrum of energy is set by simple wave saturation arguments that depend only on the local buoyancy frequency and m. The height-independence of this large m tail has been observed over many density scale heights and it is this property that drives the wave pseudo-momentum flux deposition process in the Warner-McIntyre (**WM**) gravity wave drag parametrization scheme (Warner and McIntyre, 2001). Currently, parametrization schemes like that of **WM** are column-based and have no geographical dependence or association with the underlying meteorology (e.g. areas of deep convection).

The total sub-gridscale gravity wave energy is not a routinely observable quantity in the atmosphere and since it is likely to vary substantially with horizontal position and in time, a stochastic approach could be justified. For instance, a random number term could be added to the pseudo-momentum flux magnitudes at the 'launch height' (with azimuthal dependence). Indeed a stochastic gravity wave stress model was proposed by Dunkerton (1982) in which the phase speed or total pseudo-momentum of a wave event was selected randomly in accordance with a probability distribution function. Clearly there are many ways in which randomness could be introduced into gravity wave parametrization schemes.

5 Summary and Discussion

I have tried to motivate the use of stochastic physical parametrizations based on the phenomenology of certain near-gridscale processes that are absent in numerical models of fluid flow. One can identify three distinct strategies for implementing this technique. In the first of these (the 'sea-of-noise' approach) one considers

the details of the partially-resolved fluid flow (i.e. coherent structures) as irrelevant and provides the model with a randomly-contrived forcing at the near-gridscale that conforms with some very basic constraints like flow non-divergence and smoothness in time. Here, these flow corrections are designed to represent an assumed 'sea-of-noise' that may be the result of some turbulent energy/enstrophy cascade or, for instance, a saturated spectrum of inertia-gravity waves. The second approach (noisy parametrization) consists of adapting existing parametrization schemes to include some stochastic variability as typified by the Buizza et al (1999) scheme. The scope for doing this is very wide, ranging from randomly adjusting the outputs of some parametrization scheme to randomly varying some physical parameter within the scheme (e.g. diffusion coefficients). The third approach (the coherent structure method) is to use fluid dynamical understanding to devise specific forcing functions that describe the direct, near-gridscale effect of some unresolved, organized structure. Unlike conventional parametrizations this strategy involves forcing these flow structures directly into the shortest resolvable scales of a model rather than changing the value of some model variable at a single gridpoint. This is a recognition of the fact that the parametrizable elements may be bigger than the grid-box dimensions but still not explictly resolvable by the model. Gray's stochastic convective backscatter scheme is a good example of the 'coherent structure' approach to stochastic forcing.

One practical aspect that needs to be raised concerns the numerical algorithms implemented in conventional parametrization schemes and their potential for noisy, bad behaviour at the near-gridscale. Because most parametrization schemes are applied column-wise in NWP models, neighbouring grid columns are treated as dynamically independent. This may result in large amounts of gridscale noise as thermal wind balances are destroyed by, for instance, convective heating being applied to one grid column but not its neighbours. For gridlengths of a few hundred kilometres this is not a big problem but for current operational global model resolutions (gridlengths of $\sim$ 60 km or less) there is the potential for noisy behaviour at the near-gridscale *without deliberate inclusion of a stochastic character to the relevant parametrization scheme.* For convection parametrization schemes, this bad behaviour follows from the inappropriate use of schemes based on an ensemble of deep convective clouds in models where the grid-box size is approaching that of an individual cloud system. The use of 'switches' (characterized by 'IF' tests in the code) exacerbates the problem by helping to excite spurious near-gridscale gravity waves. It is doubtful whether one could justify the use of switches in the representation of a cloud ensemble in any case. The existence of bad numerical behaviour in some parametrization schemes may therefore be unintentionally acting as stochastic forcing in the above 'noisy parametrization' sense though this is hardly an endorsement for sloppy coding and poorly-crafted parametrization methods !

There is a need to understand why stochastic forcing schemes often have a beneficial impact. Like any conventional parametrization scheme, model improvements may follow for the wrong reason and block future progress with other parametrization schemes whose improved scientific basis may lead to worse results. For instance, the Buizza et al (1999) scheme may lead to an increase in the energy of explicitly-resolved, upward-radiating gravity waves which may then partake in wave-mean flow interactions such as those that drive the quasi-biennial oscillation in the tropical stratosphere. Improvements to the tropical wind structure might then lead to better mid-latitude planetary Rossby wave structure at these levels and ultimately the entire climatology of the model. Alternatively, the effect of stochastic physics may be to source an upscale cascade of energy in the tropical atmosphere as a two-dimensional turbulence process, leading to enhanced low frequency variability at zonal wavenumbers 1 to 3, or perhaps an increase in 'easterly wave' activity. There is a danger that the introduction of a stochastic character to existing parametrizations might merely generate a few more tuning parameters in what is already a practice that is laden with simple model idealizations and heavy empiricism.

It is interesting to speculate on the likely future developments in parametrization complexity as operational model gridlength approach about 10 km. Whilst one would expect to see substantial growth in the area of cloud microphysical parametrization, it less clear what degree of complexity should be required of the dynamical parametrizations such as gravity wave drag and deep convection. There is much to be said for 'keeping it simple' and reducing the complexity of these parametrizations scheme in the run-up to sub-20 km grids. Perhaps a combination of the resolution afforded by a 10 km grid together with a simple stochastic convective forcing function could perform just as well as a complicated and cleverly-contrived convective parametrization scheme ? The philosophy of parametrizing ensembles of convective clouds will clearly be inappropriate for a 10 km grid. Even though realistic convective clouds cannot be represented on such a coarse grid it may be true that the important statistical transfer properties effected by deep convection are reasonably well described by near-gridscale convective structures excited by a stochastic forcing function.

This would seem to be an important area for future research using LES models.

6 ACKNOWLEDGEMENTS

I thank Andy Brown for clarifying the role of stochastic backscatter in LES and supplying Figs 1 and 2. Also I would like to thank Tom Allen for pointing out the literature on hairpin vortices and other boundary layer topics, together with David Thomson and Mike Gray who made useful comments on the text.

REFERENCES

Betts, A. K. & Miller, M. J. 1986 , A new convective adjustment scheme. part ii: Single column tests using gate wave, bomex, atex and arctic air-mass data sets, *Q.J.R. Meteorol. Soc.* **125**, 2887–2908.

Bretherton, F. P. 1969 , Momentum transport by gravity waves, *Q.J.R. Meteorol. Soc.* **95**, 213–243.

Broutman, D., Rottman, J. W. & Eckermann, S. 2001 , A hybrid method for wave propagation from a localized source, with application to mountain waves, *Q.J.R. Meteorol. Soc.* **127**, 129–146.

Buizza, R., Miller, M. J. & Palmer, T. N. 1999 , Stochastic representation of model uncertainties in the ecmwf ensemble prediction system, *Q.J.R. Meteorol. Soc.* **125**, 2887–2908.

Dixon, R., Browning, K. A. & Shutts, G. J. 2001 , The relation of symmetric instability and upper-level potential vorticity anomalies to the observed evolution of cloud heads, *Q.J.R. Meteorol. Soc.* **127**, to appear.

Dunkerton, T. J. 1982 , Stochastic parameterization of gravity wave stresses, *J. Atmos. Sci.* **39**, 1711–1725.

Eliassen, A. 1969 , On meso-scale mountain waves on the rotating earth, *Geofysica Publikasjoner* **27**, 1–15.

Gage, K. S. & Nastrom, G. D. 1986 , Theoretical interpretation of atmospheric wavenumber spectra of wind and temperature observed by commercial aircraft during gasp, *J. Atmos. Sci.* **43**, 729–739.

Gray, M. E. B., Shutts, G. J. & Craig, G. C. 1998 , The role of mass transfer in describing the dynamics of mesoscale convective systems, *Q.J.R. Meteorol. Soc.* **124**, 1183–1207.

Gregory, D., Shutts, G. J. & Mitchell, J. R. 1998 , A new gravity-wave-drag scheme incorporating anisotropic orography and low-level wave breaking: impact upon the climate of the uk meteorological office unified model., *Q.J.R. Meteorol. Soc.* **124**, 463–493.

Haidari, A. H. & Smith, C. A. 1994 , The generation and regeneration of single hairpin vortices, *J. Fluid Mechs.* **277**, 135–162.

Haynes, P. H. & McIntyre, M. E. 1990 , On the conservation and impermeability theorems for potential vorticity, *J. Atmos. Sci.* **47**, 2021–2031.

Hoskins, B. J., McIntyre, M. E. & Robertson, A. W. 1985 , On the use and significance of isentropic potential vorticity maps, *Q.J.R. Meteorol. Soc.* **111**, 877–946.

Lilly, D. K. 1983 , Stratified turbulence and the mesoscale variability of the atmosphere, *J. Atmos. Sci.* **40**, 749–761.

Lott, F. & Miller, M. J. 1997 , A new subgrid-scale orographic drag parametrization: its formulation and testing, *Q.J.R. Meteorol. Soc.* **123**, 101–127.

Mason, P. J. 1994 , Large-eddy simulation: A critical review of the technique, *Q.J.R. Meteorol. Soc.* **120**, 1–26.

Mason, P. J. & Brown, A. R. 1994 , The sensitivity of large-eddy simulations of turbulent shear flow to subgrid models, *Boundary Layer Meteorol.* **70**, 133–150.

Mason, P. J. & Thomson, D. J. 1992 , Stochastic backscatter in large-eddy simulations, *J. Fluid Mechs.* **242**, 51–78.

Miranda, P. M. A. & Valente, M. A. 1997 , Critical level resonance in three-dimensional flow past isolated mountains., *J. Atmos. Sci.* **54**, 1574–1588.

Palmer, T. N. 1997 , On parametrizing scales that are only somewhat smaller than the smallest resolvable scales, with application to convection and orography, *ECMWF Workshop proceedings on 'New insights and approaches to convective parametrization'* pp. 328–337.

Scinocca, J. F. & McFarlane, N. A. 2000 , The parametrization of drag induced by stratified flow over anisotropic orography, *Q.J.R. Meteorol. Soc.* **126**, 2353–2393.

Shen, B.-W. & Lin, Y.-L. 1999 , Effects of critical levels on two-dimensional back-sheared flow over an isolated mountain ridge on an f plane, *J. Atmos. Sci.* **56**, 3286–3302.

Shutts, G. J. 1987 , Balanced flow states resulting from penetrative, slantwise convection, *J. Atmos. Sci.* **44**, 3363–3376.

Shutts, G. J. 1998 , Stationary gravity-wave structure in flows with directional wind shear, *Q.J.R. Meteorol. Soc.* **124**, 1421–1442.

Shutts, G. J. 2001*a* , Inertia-gravity wave and neutral eady wavetrains forced by directionally-sheared flow over isolated hills, *J. Atmos. Sci.* **58**, submitted.

Shutts, G. J. 2001*b* , A linear model of back-sheared flow over an isolated hill in the presence of rotation, *J. Atmos. Sci.* **58**, in press.

Shutts, G. J. & Gray, M. E. B. 1996 , A balanced flow perspective of the effect of deep convection, *ECMWF Workshop proceedings on 'New insights and approaches to convective parametrization'* pp. 269–287.

Shutts, G. J. & Gray, M. E. B. 1999 , Numerical simulations of convective equilibrium under prescribed forcing, *Q.J.R. Meteorol. Soc.* **125**, 2767–2787.

Shutts, G. J., Booth, M. & Norbury, J. 1988 , A geometric model of balanced, axisymmetric flows with embedded penetrative convection, *J. Atmos. Sci.* **45**, 2609–2621.

Vallis, G. K., Shutts, G. J. & Gray, M. E. B. 1997 , Balanced mesoscale motion and stratified turbulence forced by convection, *Q.J.R. Meteorol. Soc.* **123**, 1621–1652.

Vosper, S. 1996 , Gravity wave drag on two mountains, *Q.J.R. Meteorol. Soc.* **122**, 993–999.

Vosper, S. 2001 , Development and testing of a high resolution mountain-wave forecasting system, *Meteorol. Appl.* p. submitted.

W-B.Lin, J. & Neelin, J. D. 2000 , Influence of a stochastic moist convective parametrization on tropical variability, *Geophys. Res. Lett.* **27**, 3691–3694.

Warner, C. D. & McIntyre, M. E. 2001 , An ultrasimple spectral parametrization for nonorographic gravity waves, *J. Atmos. Sci.* **58**, to appear.

Welch, W., Smolarkiewicz, P., Rotunno, R. & Boville, B. A. 2001 , The large-scale effects of flow over periodic mesoscale topography, *J. Atmos. Sci.* **58**, 1477–1492.